DOCTOR OF PHILOSOPHY

Detecting plagiarism in the forensic linguistics turn

Rui Sousa Silva

2013

Aston University

Detecting Plagiarism in the Forensic Linguistics Turn

Rui Manuel Sousa Silva

M.A.

Aston University

January, 2013

Aston University

Title: ‘Detecting Plagiarism in the Forensic Linguistics Turn.’

Name: Rui Manuel Sousa Silva

Degree: Doctor of Philosophy

Year of submission: 2013

Synopsis:

This study investigates plagiarism detection, with an application in forensic contexts. Two types of data were collected for the purposes of this study. Data in the form of written texts were obtained from two Portuguese Universities and from a Portuguese newspaper. These data are analysed linguistically to identify instances of verbatim, morpho-syntactical, lexical and discursive overlap. Data in the form of survey were obtained from two higher education institutions in Portugal, and another two in the United Kingdom. These data are analysed using a 2 by 2 between-groups Univariate Analysis of Variance (ANOVA), to reveal cross-cultural divergences in the perceptions of plagiarism. The study discusses the legal and social circumstances that may contribute to adopting a punitive approach to plagiarism, or, conversely, reject the punishment. The research adopts a critical approach to plagiarism detection. On the one hand, it describes the linguistic strategies adopted by plagiarists when borrowing from other sources, and, on the other hand, it discusses the relationship between these instances of plagiarism and the context in which they appear. A focus of this study is whether plagiarism involves an intention to deceive, and, in this case, whether forensic linguistic evidence can provide clues to this intentionality. It also evaluates current computational approaches to plagiarism detection, and identifies strategies that these systems fail to detect. Specifically, a method is proposed to translingual plagiarism.
The findings indicate that, although cross-cultural aspects influence the different perceptions of plagiarism, a distinction needs to be made between intentional and unintentional plagiarism. The linguistic analysis demonstrates that linguistic elements can contribute to finding clues for the plagiarist’s intentionality. Furthermore, the findings show that translingual plagiarism can be detected by using the method proposed, and that plagiarism detection software can be improved using existing computer tools.

Keywords: intention and intentionality, punitive turn, computational forensic linguistics, translingual plagiarism, linguistic evidence.

To my mother and to my late father.

To the Little Prince.

On ne voit bien qu'avec le cœur. L'essentiel est invisible pour les yeux.

– Antoine de Saint-Exupéry, *Le Petit Prince*

Acknowledgments

I would like to extend my gratitude to Tim Grant for his effort in guiding me for the last four years. Special thanks are in order for the challenging discussions on philosophy and morals, and most of all for his continuing support and encouragement.

I am also grateful to Malcolm Coulthard, for his swift reply to my first e-mail, for his invitation to attend the launch of the Centre for Forensic Linguistics, for encouraging me to come to Aston, and for all the inspirational discussions.

I would like to thank Maria Teresa Turell for her warm welcome to the ForensicLab, at the Universitat Pompeu Fabra. I cannot thank Maite enough for her scientific and personal support, for making me feel at home in Barcelona, for being an example of courage and strength, and especially for giving me another reason to keep going back.

I would also like to express my gratitude to Belinda Maia, whom I met at a young age, for teaching me, a long time ago, what *Bright Eyes* taught me much later: that 'language just happened, it was never planned.' I am really grateful to her for her continuing encouragement before and during my PhD, for her friendship, and for sacrificing her time to proofread my thesis.

I would also like to extend my gratitude to my examiners, Edward Finegan and Anne Wheeler, for the stimulating intellectual discussion at the viva and for their insightful comments and suggestions.

To my colleagues at Aston, and all the people I have met throughout my PhD, for all the insightful discussions, which, one way or another, influenced my work. I would also like to thank all the colleagues and friends at the University of Porto, not only for encouraging my work, but also for providing me with most of my data.

I would also like to express my thanks and appreciation to Silvana, my fantastic loving partner and friend, whom I have always truly admired, for the inspirational

discussions, for her support throughout (often) troubled times, for understanding my absence from home, for her smile, and especially for teaching me that cities also 'have a soul and body and life' (of course, she meant London, but she hardly knew Birmingham by then)! To our cats, Leo and Leya, for showing their unconditional love, and for keeping me company throughout most of my writing.

To my friends, thanks for their continuing support. I cannot name them all, so will thank Nicci for making me feel at home at Aston, for showing me around, for the long lunches, for the drinking, and for the 'numerical mystery tours'.

To Rui, my very best friend, for being with me through the most troubled times, for making me believe that true friends are always at hand; I cannot thank him enough for everything he has taught me, for always being there for me, for his understanding, for his extraordinary friendship, and for the inspirational discussions about... life!

Most of all, to my family, especially to my mum, Lucinda, and my late dad, Zeferino, for all their financial and emotional support, but most of all for their unconditional pride. Their effort has always been an inspiration to me. To my sisters, Margarida and Emília, for their continuing support, and especially to São, who is not only a sister, but a true friend. Her courage, her strength and her kindness are an extraordinary example. To my nephews and nieces, especially to Ana and to Inês, for their witty remarks, and for making me believe that this world can be a fairer world.

Finally, my thanks to Birmingham, where I will always feel at home!

This work was supported by grant SFRH/BD/47890/2008 FCT-Portugal, co-financed by POPH/FSE.

UNIÃO EUROPEIA
Fundo Social Europeu

QUADRO DE REFERÊNCIA ESTRATÉGICO NACIONAL
PORTUGAL 2007.2013

POPH

Contents

List of Tables

List of Figures

List of Abbreviations

ANOVA – Analysis of Variance

CDA – Critical Discourse Analysis

DPA – Discriminative Power Analysis

ETS – Education Testing Service

HEA – Higher Education Academy

IL – Interlanguage

IR – Information Retrieval

JISC – Joint Information Systems Committee

MT – Machine Translation

NL – Native Language

NLP – Natural Language Processing

NUS – National Union of Students

PLA – Principle Component Analysis

POS - Part-of-speech

SL – Source Language

SMT – Statistical Machine Translation

SVM – Support vector machine

TL – Target Language

UCAS – British Universities and Colleges Admissions Service

UI – User Interface

VLE – Virtual Learning Environment

Chapter 1

Introduction: Investigating Plagiarism

1.1 Introduction

> *Un simple imitateur est un estomac ruiné qui rend l'aliment comme il le reçoit: un plagiaire est un faussaire.*
>
> — Voltaire

Plagiarism has attracted the attention of diverse disciplines over the centuries, including that in recent years of forensic linguistics / language and the law. In fact, it was not until recently that linguists initiated their investigation into plagiarism in forensic contexts, either to explain why a certain instance represents plagiarism or otherwise, or to improve the detection methods and procedures available. However, the legality of unattributed textual reuse has been questioned over time, and this quote, attributed to Voltaire, partly reflects that legal nature; in particular, it describes the plagiarist as a *faussaire*, a deceiver whose acts, consisting of reusing the product of someone else's work without acknowledgment, and consequently passing it off as their own, are counterfeit. Incapable of properly addressing and managing their sources, the imitator is therefore *an estomac ruiné*, a 'ruined stomach' that is condemned to 'vomiting' the ingested nourishment in an identical, yet altered form that only slightly resembles the original. The plagiarist is consequently a deceiver, whose actions, possibly illegal, are absolutely immoral.

The social, legal and moral implications of plagiarism for the plagiarist's reputation are as serious today as they seemed to be in Voltaire's own time. As recent cases covered by the media demonstrate, charges of plagiarism have the potential to bring the plagiarist under public trial, and eventually affect irreversibly someone's career.

Recent cases of this sort involved a lecturer of a Portuguese higher education institute, who resigned in 2010 as a result of accusations that she had plagiarised her doctoral thesis. Later, in 2011, the German Defence Minister Karl-Theodor zu Guttenberg (temporarily) renounced his doctorate title and eventually resigned following accusations of plagiarism. In 2012, the Romanian Prime Minister Victor Ponta was accused of having plagiarised large portions of his doctoral thesis, and faced pressure to resign.

These cases show that even when produced in academic contexts, plagiarism can and does have implications out of the academy; but the consequences of plagiarism are not limited to academic work. A few years ago a journalist of the Portuguese quality newspaper *Público* was accused of plagiarising from *Nature*, the *New Scientist* and the *Wikipedia*. She continued working for the newspaper after she admitted to having plagiarised and made a public apology. More recently, *The Independent* journalist Johann Hari was suspended after he was found to have plagiarised news articles, and in August the renowned *CNN* and the *Time* magazine journalist Fareed Zakaria was suspended after he admitted to having plagiarised a *New Yorker* article. Zakaria was later readmitted after the magazine screened all his opinion articles, but his reputation was certainly marred.

These social implications resulting from the 'public judgement' of cases such as those of Guttenberg, Hari or Ponta, as well as Zakaria, suggest, not the least, that plagiarism is both immoral and an act of deception; ultimately, however, it is the illegal nature of the text reuse, the product of knowingly concealing a fact or misrepresenting aspects of the truth, that accounts for its consideration as a fraudulent activity, requiring the adoption of legal action. The public and legal judgements of plagiarism, however, bring along with them two other problems. Firstly, most cases of plagiarism – even those of an inherently 'social' nature – have legal and moral implications, and have the potential to debilitate or even destroy someone else's professional career and damage their life, thus requiring more than a frivolous, light-hearted accusation. Secondly, in the public sphere any insufficiently founded accusations can easily slip into a case of defamation, in which in a reversal of roles, the suspect plagiarist becomes the plaintiff, seeking for compensation for financial and moral damage.

One of the fundamental problems of plagiarism, as we shall see in the next chapters, is therefore investigating it and finding evidence that a certain text was plagiarised. In cases of barefaced, straightforward lifting, where the text is copied and

pasted literally, word-for-word, a simple comparison of the suspect text against the known original can be sufficient to demonstrate that the two texts are uncommonly similar, or even identical. Most instances are not, however, as simple as verbatim lifting, involving sophisticated strategies of text manipulation; and even in cases of literal unacknowledged borrowing, apparently clear cases can be easily challenged by circumstances such as the amount of text that is reused, what type of alterations (if any) have been made, which one was produced and published first, and the surrounding conditions granting access to the texts, to name just a few.

As a result of the technological developments of the last decades, access to more information made virtually any text in any part of the world more susceptible to plagiarising. And although, luckily, these same developments also facilitated the detection of those instances, necessity being what it is – 'the mother of invention' – new strategies are developed almost every day, with plagiarists having to devise new ways of plagiarising without being easily detected. From a plagiarist's strategic point of view, this meant making alterations to the text, by deleting, editing or adding to the original, replacing words or phrases, changing the word order, rephrasing and paraphrasing the original, or even more complex, translating from another source. Consequently, in parallel with the limitations (even the failure) of simple computational text matching techniques in detecting more sophisticated instances of plagiarism, as will be shown in chapter 7, new methods of plagiarism detection had to be developed, sometimes computational, but usually based on more advanced linguistic methods that are necessary to compare the suspect text against the potential original to then conclude whether the text has been plagiarised.

The type of linguistic analysis developed and used by forensic linguists, in particular, has demonstrated good results in plagiarism detection. Previous and ongoing research into forensic linguistics, while approaching linguistic analysis within the scope of the interaction between language and the law, has demonstrated that the likelihood that two texts may have been produced independently can often be determined accurately, and this data can be used essentially as an investigative tool, and frequently as linguistic evidence.

It adds to this that, in the academy, as in research contexts, plagiarism is an unacceptable practice, as institutional approaches to plagiarism demonstrate. Students are required to properly acknowledge their sources, lecturers and tutors are increasingly more demanding with students, academic journals are growingly intolerant to

'unacknowledged' or 'improper' text reuse, university administrations (although not always for the right reasons) encourage and cooperate with investigations of plagiarism and, maybe as a consequence, the general public seems to cooperate. Paradoxically, some cases of textual reuse remain unchallenged, even when obviously reusing considerable portions of text from other sources, on the grounds that they are texts of a particular genre, often formulaic, to which the requirements of original authorship, and correspondingly proper and due acknowledgement, do not apply. This is the case of some documents of legal or technical nature that circulate around the world, unacknowledged and unattributed. This dichotomic overview of plagiarism as fraudulent behaviour or, conversely, as legitimate text reuse suggests a straightforward approach to detecting plagiarism that is unrealistic and oversimplistic. On the contrary, as will be discussed throughout this thesis, plagiarism is a complex web of concepts, perceptions, understandings, and sometimes even competing discourses which are difficult to disentangle.

1.1.1 Rationale for the Research

Plagiarism is a growing concern among authors, publishers, intellectual property rights owners and the academic community all over the world, as the multitude of policies, rules and regulations demonstrate. And although it is several centuries old, it has gained public attention and social visibility in recent years, in no small part due to the media coverage dedicated to high profile cases like the ones discussed in the previous paragraphs. However, possibly due to to newsworthy criteria, other instances of plagiarism, including academic plagiarism, remain restricted to the academy, hardly attracting media attention, and despite being equally or even more important than other cases, they remain less visible.

Nevertheless, plagiarism in the academy remains a serious issue. It is possible that, every day, thousands, perhaps millions of students worldwide, with access to millions of resources in hundreds of different languages, plagiarise when writing academic essays. Speculatively, only a small fraction of these are detected, and the issue brought to their attention. This represents a problem of effectiveness in the detection procedure. Only a systematic, rather than spontaneous, analysis of all academic work permits a comprehensive coverage. The main drawback to this systematic approach is that it requires both technical capacities and, correspondingly, sufficient resources. If, on the one hand, a detection software needs to be capable of detecting plagiarism

strategies, on the other hand powerful computer resources, with processing speed, are required to systematically analyse all work submitted for assessment – not to mention the competent human resources required to interpret the plagiarism reports. On the other hand, the list of strategies used to plagiarise can be comprehensive, but is not exhaustive. Plagiarism is inherently a creative act, so in the same way that anti-virus software is updated everyday with new virus definitions and new computer viruses are being invented and released immediately after, so should plagiarism detection software be regularly updated with new 'plagiarism definitions' to cater for the sophisticated, newly developed plagiarism strategies.

This scenario is ideal, but not realistic. Firstly, unlike computer viruses, that always cause some sort of problem (even those operating in 'silent mode'), instances of plagiarism do not cause any (immediate) problems if they pass undetected. Secondly, once methods are developed to detect the most straightforward forms of plagiarism, such as those cases involving reuse of similar or identical text, more in-depth linguistic research is needed to describe more sophisticated strategies, which usually require linguistic knowledge before they can be implemented computationally. Unfortunately, the linguistics equivalent of computer scientists standing behind their desks working on the newly found cases is very rare, if they exist at all.

Detection being the most fundamental element of the fight against plagiarism, as some plagiarism policies seem to suggest, everything else is accessory. This view is not altogether wrong if we consider that the single most important element is placing in the 'hands' of the machine the systematic screening of similar or identical instances that we, as humans with limited processing capacities, could not possibly afford to do. It would then be a task for the 'human detector' to analyse the instances identified by the software, and subsequently determine the outcome of the identified instance, whether that means, pending a determination of the seriousness of the text reuse, subjecting the student to disciplinary action, failing their assignment or the whole module, or simply asking the student to correctly attribute the sources and resubmit. But how often are plagiarising students referenced for further training, or the reasons for their plagiarism investigated further, in order to identify and act on the problem issues? Speculating that the answer to these questions is 'not very often', the effectiveness of the detection procedure can be challenged.

As will be demonstrated in chapter 6, detection software currently performs fairly well, despite the foreseen room for constant improvement, especially in order to ac-

commodate new strategies, as will be argued in chapter 7. However, unless its main aim is to identify the instance of lifting, assess its seriousness and punish the plagiarist accordingly – a punitive aim that appeals to many, as the survey presented in chapter 3 will illustrate – detection systems risk being ineffective. Understanding why a student plagiarises, whether they lack the necessary academic writing skills, the knowledge of academic writing conventions, or simply training in writing academically, or whether, knowing how to write academically, they act out of carelessness or intention to deceive, are all elements of the plagiarism equation. All of them contribute, not only to determining the severity of the punishment, but also to act on possible problems and permit the adoption of corrective measures early enough in the student's career to guarantee a proactive intervention, as will be argued below. Otherwise, by not acting over the main issues and correcting them, these problems may be perpetuated to an extent that risks overloading the processing capacity of computers, thus endangering the detecting ability of the software, and threatening its sustainability.

The rationale behind this project is to contribute to the development of research on plagiarism detection, while simultaneously considering existing laws, rules and regulations, but also the diverse, varying perceptions of plagiarism. This research approaches the problem of plagiarism detection from a *forensic linguistic perspective*. Its fundamental assumption is that, as Finegan (1993) demonstrates, the linguistics expert can only attend to matters within his/her expertise, and therefore can only express an opinion based on linguistic evidence. However, in the case of plagiarism detection, this linguistic evidence can also take into account circumstances such as the students' academic writing skills or training – or their lack thereof –, the degree of awareness of the plagiarist, and their understanding of plagiarism. Consideration of these elements not only contributes to explaining the linguistic evidence found but, more importantly, the linguistic evidence can help to explain the strategies used, and, ultimately, allow the linguistics expert to do his/her job: to 'render an opinion' (Finegan, 1993: 185). As Finegan ultimately recommends, 'a careful expert should endeavour to imagine other stories, other contexts for the facts presented, perhaps competing versions of the truth' (Finegan, 2009: 272).

The following section 1.1.2 presents the research questions that this study attempts to address, given the rationale explained.

1.1.2 Research Questions

This research aims to identify, from a forensic linguistics perspective, the linguistic patterns and strategies used across different instances of plagiarism. It attempts to contribute to improving detection methods and techniques and, accordingly, indicating directions in computational forensic linguistics for a more effective implementation of those techniques. In particular, this research focuses on plagiarism detection in Portuguese (although these detection methods can also be of use to other languages) from a forensic linguistics perspective, with an application to computational linguistics. Key concepts and developments in forensic linguistics are applied to the CorRUPT corpus, a corpus of (mainly) academic texts to identify and explain patterns and strategies used by university students to plagiarise. These can be used as an aid in the investigative process (e.g. by University disciplinary boards), and/or to provide forensic evidence of plagiarism. The strategies identified will be mentioned throughout this study, when relevant.

Firstly, this research investigates *how university students use information available (especially online) to plagiarise in their academic work*. This includes reuse of verbatim, word-for-word unacknowledged text, as well as cases where text is rephrased or paraphrased without an indication of the source. Particular attention is given to the use of translation without an indication of the source as a plagiarism strategy, as will be discussed in chapter 5.

In parallel with this study, an investigation is conducted into *how existing plagiarism detection software performs* (as described in chapter 6) and whether *software tools can be designed to improve the detection of the different instances of plagiarism*, as will be argued in chapter 7. Plagiarism detection software performs well in detecting some types of plagiarism, but not all, owing to the fact that relevant text elements are known to escape the computational processing tasks, hence the traditional view that computational linguistics has a limited effectiveness in the field of forensic linguistics. We shall see whether approaching plagiarism detection software design differently, and using common natural language processing (NLP) tools, can improve the detection results, not only in terms of the number of the instances detected, but also – and more importantly – of the quality of the detection procedure.

Thirdly, considering the debate on whether plagiarism consists of any unacknowledged borrowing or, on the contrary, whether it should be punished only when intentional, the research focuses on discussing, in chapter 4, *whether linguistic clues can*

be found in the text to demonstrate the plagiarist's intention to plagiarise, and *whether a certain instance of plagiarism is inadvertent or intentional*.

Finally, an investigation will be conducted into the applications of forensic linguistics to plagiarism detection. In particular, this research delves into how forensic linguistic analysis can contribute, on the one hand, to the *investigative process of instances of plagiarism* and, on the other hand, *whether and how that linguistic analysis can be used as evidence of plagiarism*. These two questions will underlie the whole thesis.

1.1.3 Methods for Operationalising the Research Questions

This section describes how the different theoretical perspectives approached in this study are operationalised. An investigation into plagiarism detection raises two problems that need to be addressed liberally. Firstly, plagiarism can and is studied across different disciplines, as discussed in section 1.3, including literary criticism, composition studies and education, intercultural studies, psychology and management, as well as computer sciences and natural language processing, law and linguistics, among others. However, each of these disciplines tends to investigate one particular aspect of plagiarism, leaving aside all the others. Whereas the educator is more interested in finding reasons for student plagiarism, and act upon the issues to correct them, and intercultural studies try to explain that plagiarism cannot be studied independently of the cultural context, the law is more focused on judging whether a particular norm has been breached. Whereas the computer scientist is mostly interested in programming the machine to identify instances of *any* textual reuse, including source code, the psychologist is more attracted to understanding what goes on in the plagiarist‘s mind when plagiarising. Whereas the manager is mostly worried about the loss of profit underlying counterfeit activities, the linguist works to identify and explain why a certain text extract represents plagiarism, leaving the interpretation of these results to the agent responsible for making decisions. Surprisingly, as the review of the literature demonstrates, cases of multidisciplinary research on plagiarism, and particularly on plagiarism detection are very rare, if they exist.

Forensic linguistics is in a privileged position to embrace this multidisciplinarity. Broadly defined as the interface between language and the law, forensic linguistics necessarily needs to consider other disciplines to go beyond describing *what* happens in the text, to understand and explain the meaning of those findings in a particular context. Research on plagiarism detection in forensic linguistics starts with a linguistic

analysis of textual material in forensic contexts (as opposed to other contexts, e.g. literature) to demonstrate, via appropriate description, why a certain text or text extract, which has not been attributed, has not been produced independently, and to provide explanations (often extra-textual) to sustain their claims.

Secondly, research on plagiarism detection imposes additional challenges on the methodology and data used. If the aim is to conduct an in-depth investigation of a particular instance of plagiarism that has been identified, a traditional analysis using a set of conventionalised methods will suffice. Conversely, if the aim is to *improve* plagiarism detection, two inter-related problems in particular are envisaged. On the one hand, an analysis of a particular set of data will allow the detection of the strategies used in that set, but not others. For instance, a thorough analysis of text extracts that were found to have copied literally from another source allows the linguist, and even the computer scientist, to draw conclusions on verbatim lifting, and correspondingly explain and programme the computer to detect those instances, but not instances of plagiarism based on lexical replacement – regardless of the volume of the data available. Scientific ethics determines that it is wrong to make assumptions for which one lacks evidence, so the linguist needs to adhere strictly to what is found in the text, which in the case of plagiarism detection may mean waiting for someone to intuitively detect an instance of plagiarism to conduct a proper linguistic (or computational) analysis. On the other hand, not following these scientific ethics guidelines strictly (or not as strictly as ideally expected), not only allows a comprehensive, *working* taxonomy of strategies to be proposed (and possibly later verified), but also frees the research into plagiarism detection from the boundaries imposed by the features contained in the data. The latter, focused on *possibilities*, rather than on *actual* instances, is not risk-free. The first main danger is that, by building upon the creative skills of the researchers to conceive of plagiarism possibilities, the research misses the focus on particular aspects. Secondly, that inability to test these possibilities over real data, and consequently the need to resort to fabricated data, may bias the results. This raises a 'chicken or the egg' type of causality dilemma, the most viable solution to which seems to build on operating a circular process of constant detection over fabricated data, testing over real data, refinement over real data and re-testing over real data.

The current research project is therefore operationalised as follows. The first research question, on the investigation of *how university students use information available (especially online) to plagiarise in their academic work*, is addressed via a quali-

tative linguistic analysis of the data. In particular, methods developed used with other languages, mostly English, are applied in this study of Portuguese. Therefore, verbatim, word-for-word plagiarism is addressed via simple pattern comparison, either manually or machine-assisted (e.g. using the 'compare documents' function of *Microsoft Word*, or specific plagiarism detection software). Cases involving reordering of the textual elements are detected via an analysis of the lexical items. More sophisticated forms of textual reuse, such as rephrasing and paraphrasing, require other methods and tools, such as thesauri. However, even using such tools, this is very difficult to handle computationally. Additionally, since string matching techniques tend to break down when new textual elements are used to interrupt a sequence of words, an analysis of lexical items is applied in these cases. Finally, translation of text from another language is approached using freely available machine-translation tools, as these are the ones that are at the disposal of the ordinary student user.

Existing plagiarism detection methods and tools will also be used whenever possible to test the examples available in the CorRUPT corpus. This implies assessing the performance of existing plagiarism detection software tools, and critically evaluating them, in order to investigate whether re-designing particular aspects could help *improve the detection of the different instances of plagiarism*. This improvement, albeit not representing a direct increase in the number of cases detected, may improve the quality of the detection method and indirectly increase that number, by (a) identifying cases that are otherwise missed; and (b) reducing the number of false positives, which can represent an increase in the amount of time available for investigating true positives, besides impacting the level of confidence in the method. Subsequently, the manual effort is devoted to the interpretation and analysis of the results, so as to demonstrate or otherwise disregard the evidential value of the linguistic analysis.

The detection method, however, cannot ultimately be considered independently of a precise definition of plagiarism, or even on the grounds that plagiarism is a consensual concept, understood on the same terms by everyone. In this context, it is important to define what is understood by plagiarism, not only in the literature, but more importantly by lecturers/tutors and students in general. For this purpose, a survey was conducted with lecturers/tutors and students in Portugal and in the United Kingdom. This survey includes 12 fictional scenarios (vignettes) that the participants are asked to read and subsequently rate on a scale as to whether they consider that the situations described represent a case of plagiarism or otherwise. Additionally, participants

have the opportunity to comment on each case individually, thereby justifying their choices. In particular, this allows the issue of intentionality to be raised, i.e. whether any case of unacknowledged textual reuse should be considered an instance of plagiarism or, on the contrary, whether an instance of plagiarism presupposes a greater or lesser degree of intention to mislead. This will allow the research to address the question of *whether a certain instance of plagiarism necessarily needs to be deemed as inadvertent or intentional*.

Although the main focus of this research is on the linguistic analysis of instances of plagiarism, with a view to applying it in forensic contexts, no one chapter is dedicated specifically to linguistic analysis. This analysis is, on the contrary, presented transversally, throughout the research, to illustrate and demonstrate the relevance of each approach discussed.

Finally, some terminological clarifications are required. The word *plagiarism* can be used as a countable or an uncountable noun, either to describe the overall act of reusing someone else's words without acknowledgement, or to describe an instance of that act. In this study, plagiarism is only used as an uncountable noun; the role filled by the countable noun is left to the phrases 'instance of plagiarism' or 'case of plagiarism'. On the other hand, it is not uncommon to find reference in the literature to texts that borrow without acknowledgement as being 'plagiarised texts'. Although the use of this term can be appropriate to describe an instance of derivative text, where the text is passively produced, without any originality, by lifting from another text, on the other hand it can raise issues of accuracy from a perspective of agency. In other words, in a case of plagiarism, there are typically two texts: one that acts upon the other actively, 'stealing' from it; and another text that is acted upon, from where text is stolen, in a passive role. In this study, I will call the derivative text that actively borrows from another, the *plagiarising text*. Conversely, I shall use *plagiarised text* to refer to the original from which the derivative text borrows.

1.2 The Structure of this Study

This research is structured as follows. The current chapter 1 introduces the research topic, presents the rationale behind this research and explains the research questions that it will attempt to address. In the following section 1.3, the core literature is reviewed. Plagiarism is presented in relation to authorship to discuss whether the violation of someone else's ownership of (copyrighted) material is the central prob-

lem of plagiarism. Since this study is focused on academic plagiarism, the topic is then discussed in relation to education and pedagogical approaches, as well as in relation to intercultural studies. In particular, it is necessary to understand whether plagiarism is understood equally and universally on the same defining terms or, on the contrary, whether different people have different conceptions of plagiarism, and, in this case, whether these different conceptions are mutually understandable or, conversely, whether they are mutually exclusive. In parallel, plagiarism is discussed in relation to institutional approaches. Definitions of plagiarism used by universities worldwide, as well as by international organisations, are discussed also in relation to this diversity or uniformity of concepts. The topic of automatic plagiarism detection is introduced to discuss the type(s) and nature of the instances of plagiarism that computers aim to detect. Subsequently, linguistic approaches to plagiarism detection in forensic contexts are discussed, and a relationship is established with the legal study of plagiarism. A comparison is made between the differences, clear in the literature, between linguistic and computational approaches. The section concludes with a definition of plagiarism that will be used transversally in this research.

In chapter 2, the methodology used in this research is briefly described. It introduces the methods used in the analysis presented in chapters 3, 4, 5, 6 and 7. This chapter is followed by chapter 3, which presents the analysis and the results of the survey conducted with different groups of participants (lecturers/tutors and students), from two different countries (UK and Portugal) to assess perceptions of plagiarism. Chapter 4 discusses theories of intention and intentionality, to propose a taxonomy of intention in cases of (academic) plagiarism. Chapter 5 introduces translation as a plagiarism strategy. In particular, different cases of plagiarism are discussed in which translation has been used to borrow without acknowledgement from other texts. Chapters 6 and 7 present existing plagiarism detection software. Their features and skills are presented, as well as their drawbacks, using practical, real examples of plagiarism. A set of improvements is proposed to overcome some of the identified drawbacks. This research concludes by presenting some final overarching remarks, identifying gaps that remained unsolved in this study and presenting suggestions for further work.

1.3 Research Context: The Literature on Plagiarism Research

The concept of *plagiarism* is unsurprisingly often intertwined with the concept of intellectual property in a broad sense, considering on the one hand the historical evolution

of ownership of 'moral', non-material goods, and on the other hand the need to establish a legal framework to cope with the infringement of those rights. Plagiarism has, therefore, attracted the attention of copyright lawyers, teachers and forensic linguists (Coulthard and Johnson, 2007), as well as publishers, computer engineers/software developers and computational linguists alike. The problem of plagiarism, and the corrective measures adopted to resolve it, therefore tend to fall into two categories: *ethical* and *legal*, each with different consequences.

On the legal side, the implications of reusing someone else's intellectual property[1] are mainly economic, usually involving compensation for financial damage or loss of profit resulting from the illegitimate exploration of someone else's work by a third party. Naturally, the infringement of ethical and moral principles also underlie this violation of someone else's property; but as the analysis of plagiarism cases taken to the Portuguese courts in recent years demonstrates, the outcome tends to be more favourable to the victim when the infringement has economic implications.

On the ethical and moral side, the infringement of intellectual property tends to have social, educational or disciplinary implications, as some of the cases discussed above demonstrate. As a result of plagiarising, ministers were forced to resign, journalists were suspended or compelled to issue a public apology, and university lecturers were dismissed.

The line dividing the ethical and the legal sides, however, can only be an artificial one. If, on the one hand, ethical and moral principles pave the way for laws, rules and regulations, on the other hand the adoption of ethical and moral sanctions, including disciplinary measures and requests for public apologies, cannot be applied with disregard for the law. In this sense, it is doubtful whether a student committing plagiarism in their academic work can/should be treated in the same way as a publisher infringing someone else's copyright. Therefore, since different contexts tend to dictate the application of different degrees of definition and, correspondingly, sanctions of a different nature, the focus will now be on academic plagiarism.

1.3.1 Academic Plagiarism

Academic writing imposes particular expectations upon the writer, specifically a perfect balance between the sources, the amount of original work, due acknowledgement and avoidance of over-citation. It is therefore the form that students must use when writing

[1] Intellectual property is used here in its broadest sense, to include words, works and ideas in general.

their texts in order to be able to meet certain expectations regarding their learning process. Academic writing conventions, when mastered appropriately, allow a clear distinction between the writer's own ideas, work and words, and the ideas, work and words of other authors while they elaborate on a topic building upon what has already been said and written by experts and authorities in the field (Carroll and Appleton, 2001). When writing academically, students are therefore required to build upon what they read and hear using their own words, or clearly signal the words of others, and contribute with their own thoughts and ideas in a new and original way, while they critique the opinions of their sources. In this sense, students are ultimately expected to produce a unique piece of academic writing (Anderson, 1998). However, as Howard (1995) demonstrated, *proper* academic writing requires very subtle, yet sophisticated mechanisms of authorship attribution, that span beyond quoting text, and that can and are, on the one hand, difficult to master by university students, and that, on the other, may lead one to mistake plagiarism for poor academic writing. Additionally, not all students are aware of the academic writing conventions and on the avoidance of plagiarism (Angèlil-Carter, 2000; Pecorari, 2008).

It has been argued, even if differently, by Howard (1995) and Scollon (1994, 1995), and later reinforced by Angèlil-Carter (2000) and Bloch (2012), that the concept of plagiarism cannot be studied on the grounds that one definition is clearly understandable by everyone, that it cannot be studied independently of pedagogical approaches, on the one hand, and intercultural studies, on the other, and that this determines how the 'crime' is punished.

Firstly, as Howard (1995) argued and Howard and Robillard (2008) later reinforced, any definition of plagiarism must allow room for educational and pedagogical practices, so as to protect individual authorship and, at the same time, allow students to learn from the words of others. If, on the one hand, it is a convention of the academy that, when writing academically, the student should be able to make a clear distinction between the words, works and ideas of others, and the words, works and ideas that they have contributed by following a set of more or less sophisticated pre-established rules, on the other hand, the academy is also part of the student's learning process. Adopting a pedagogical stance and using her own experience as an example, Howard (1995) claims that the type of plagiarism that her students committed (deleting some words from the source, altering grammatical structures, replacing words with synonyms), despite being considered a violation of the established rules, is nothing but 'patchwriting'

(Howard, 1995, 1999), i.e. an attempt at learning from others, rather than an act of cheating. She therefore sustains that the juridical definition of plagiarism cannot be applied in academic contexts independently of the plagiarist's intentionality. Hence the need to distinguish between *patchwriting* and *plagiarism*, the latter being

> [...] the representation of a source's words or ideas as one's own. Plagiarism occurs when a writer fails to supply quotation marks for exact quotations; fails to cite the sources of his or her ideas; or adopts the phrasing of his or her sources, with changes in grammar or word choice. (Howard, 1995: 799)

She elaborates on this definition to distinguish between three different forms of plagiarism: *cheating*, which consists of 'borrowing, purchasing, or otherwise obtaining work composed by someone else and submitting it under one's own name'; *non-attribution*, which consists of 'writing one's own paper but including passages copied exactly from the work of another' – which, if intentional, must be subject to penalties; and *patchwriting*, which means 'writing passages that are not copied exactly but that have nevertheless been borrowed from another source, with some changes' (Howard, 1995: 796).

Seconding Howard's claims, Pecorari (2008) later demonstrated that differing degrees of penalties are required in academic contexts where instances of plagiarism can indicate a lack of academic writing skills, rather than an intention to deceive. Therefore, textual reuse in the academy should be punished when the student attempts to deceive the reader, but not when the instances of plagiarism resemble a *patchwork* resulting from an unsuccessful attempt of the student at writing academically – especially considering that such 'patchwriting' is part of the educational experience of the students (Howard, 1995). From a pedagogical perspective, Howard claims, mimicking the language and concepts of others represents a form of expanding 'the lexical, stylistic and conceptual repertoires' (Howard, 1999: xviii), that when criminalised, as it tends to be in universities, can be counterproductive. In these terms, patchwriting should be punished when intentional, but should not be considered an offence if resulting from the plagiarist's unfamiliarity with the text or topic; on the contrary, in this case it should be considered part of the learning process, and hence a form of 'positive plagiarism' (Howard, 1995: 796). Partly in response to Howard's claims, Woolls (2003) argues that 'patchwriting' may be tolerated at a certain stage, but only under a 'controlled environment', so that it is identified early enough in the student's career, in order to avoid its 'natural' integration in the student's own style of

academic writing. The relevance of Howard's definition derives mostly from its including concepts of 'cheating', 'non-attribution of sources' (intentional or unintentional), and 'patchwriting', which are indispensable to the plagiarism equation.

Plagiarism however is also seen as a problem of student learning, and not simply a problem of academic writing. As Angèlil-Carter (2000) argued, it is not irrelevant that the incompetence of higher education students to properly write academically owes, in no small part, to the existence of prior, conflicting social discourses that may interfere with the academic discourse. Criticising the view that defines plagiarism as the 'deliberate theft of another person's intellectual property' (Angèlil-Carter, 2000: 2), she contends that, more than 'deliberate theft' by default, which results from a student's intention to deceive, plagiarism reveals a problem of academic literacy which results from the lack of a clear definition of the concept of plagiarism, together with the absence of clear policies and, correspondingly, adequate pedagogical approaches. On the contrary, with the disregard for the complexity of academic writing arising from the relationship between 'the educational environment, the nature of academic discourse and the nature of language' (Angèlil-Carter, 2000: 2), the concept of plagiarism is taken for granted and naturalised. She thus claims that 'plagiarism is an elusive concept, difficult to define, meaning different things in different contexts and for different textual genres' (Angèlil-Carter, 2000: 3), and therefore argues that the dialectic demanded by academic writing to avoid plagiarism by referencing, acknowledgement and attribution, on the one hand, and over-referencing, on the other, can represent a challenge to students. In these cases, especially when the student lacks linguistic resources, learning by imitation can therefore be a legitimate form of learning from sources (Angèlil-Carter, 2000). Although she admits that there are cases of intentional plagiarism among higher education students, and considers the principle that the intention to deceive is the most serious form of plagiarism, she argues for a distinction between *intentional* and *inadvertent* plagiarism, whereby the former is 'deceitful', as opposed to 'sloppy scholarship'. Even in cases where the instance of plagiarism is intentional, she claims, it is necessary to consider the student's motivation to plagiarise and weigh their educational context, which includes the student's familiarity with the discourse of the discipline, whether they have been taught the requested academic writing conventions, and what is specifically expected of the student to avoid instances of plagiarism (Angèlil-Carter, 2000).

Howard and Robillard (2008) later reiterated that each act of transgressive textual

or intertextual event cannot be attached to a 'moral frame', and Robillard (2008) concurs that plagiarism is a problem of discourses. When discussing her own pedagogical approach to authorship, that is focused on 'coinvestigation' as a means to promote mutual learning in the classroom from peers, she reports what, in her view, is the importance of a rhetorical discussion in approaching plagiarism in a social context that values individual creative work, and overvalues ('is obsessed with' (Robillard, 2008: 38)) originality and authenticity. Unlike other pedagogical approaches, however, hers sees the task of understanding the reasons why students plagiarise as 'a futile undertaking' (Robillard, 2008: 30). What she values is an approach that promotes learning by discussing the act of plagiarism, rather than the process of plagiarising. Building on the assumption that plagiarism is a form of writing, albeit socially condemnable (Howard and Robillard, 2008) (or, as Love (2002: 40) called it, 'discreditable'), her proposal sets out to instruct students on how to write by situating plagiarism as a form of authorship, among other forms of authorship (Robillard, 2008: 32) such as 'book reviewing', 'the creation of book clubs', 'literacy sponsorship', 'political polemic' and 'literary forgery'. Arguing that plagiarism is a matter of 'relationships', rather than a theft and appropriation 'of someone else's words', and hence a 'political' issue, with metonymic associations with immoral actions, she concludes, definitions of plagiarism cannot be limited to the textual elements (Robillard, 2008: 39).

Anson (2008) agrees that plagiarism is a pedagogical problem, whose responsibility has to be shared by lecturers/tutors. He adds that approaches to plagiarism that are motivated by a focus on plagiarism policies, detection methods and procedures and corresponding punishment measures divert the attention from the real problem: 'the goals of engagement, support, response, and intellectual scaffolding that are crucial for the development of students' advanced literacies' (Anson, 2008: 140).

The solution that Anson proposes to reduce the number of instances of academic plagiarism among students is to adopt a more engaging, learner-centred approach to learning and teaching, that consists of matching specific goals with 'goal-directed assignments', which means operating a shift from 'learning objectives' to 'learning outcomes' (D'Andrea, 1999). This proposal takes into account the approaches to 'surface' and 'deep' learning (Entwistle and Ramsden, 1983), and the five-level classification of learning proposed by Biggs (1987), to value understanding as a form of deep learning to the detriment of memorising, which is a form of surface learning. Copying and pasting from materials, which contribute to increasing plagiarism (Pears and Shields,

2008), can in Anson's opinion be discouraged by shifting the focus of assignments from assessment to learning.

In other words, promoting good teaching practices that consider redesigning assignments as learning tools, as opposed to 'products designed purely for assessment', and thus encouraging students to be more original while giving 'them ownership of their own work' (Anson, 2008: 141), can contribute (either by discouragement or impossibility) to preventing copying and pasting from materials to coursework – an act that, as Pears and Shields (2008) argue, contributes to increasing plagiarism. The assumptions in pedagogy that the methods of assessment, like learning and teaching methods, have a significant impact on the students performance have long been accepted in pedagogy. Howard had previously challenged the indispensable requirement that student assignments be authentic, by questioning: '[w]e expect authentic writing from our students, yet we do not write authentic assignments for them' (Howard, 2001). As argued by Wakeford (1999) a few years before Anson, assessment not only has the potential to guide learners in their learning process, but also the student perceptions of what elements are rewarded, or otherwise can have a significant impact on their learning behaviour, as well as on the learning outcomes. D'Andrea (1999) suggested that assessment methods are one of the elements that should be focused on the outcomes to be achieved, and their careful selection helps the students achieve the learning outcomes.

Plagiarism as a problem of social discourses also reflects on research on the East/West dichotomy. As a product of the implementation of international trade treaties (such as the *Berne Convention*) throughout the world, provisions that have been established by copyright agreements initially comprising mostly western countries were transferred to other contexts – including the academic context – thereby imposing a social order that clashed with the local culture. In the academic context in particular, this hegemony of the West (Bloch, 2012; Scollon, 1994, 1995) reflected in the adoption of (academic) policies against plagiarism and corresponding penalties by Eastern universities, with disregard for the local context, such as the 'cultural respect' for the Masters, unacceptable in Western terms, that materialised in the unsuccessful, yet unintentional attempt to build upon prior authoritative authors (Howard, 1995). Intercultural approaches to plagiarism therefore tend to challenge the principle that plagiarism is a universal concept, equally understandable by everyone. As the literature demonstrates, plagiarism is interpreted differently by different people (Angèlil-Carter,

2000; Bloch, 2012; Howard, 1995, 1999; Scollon, 1994, 1995). For example, Scollon (1994, 1995) and later Bloch (2012) discussed the problems of authorship attribution in non-native writing in English to claim that the presentation of facts in academic writing is inseparable from the individual author. Scollon (1995), specifically, used Goffman's *Frame Analysis* (Goffman, 1975) to argue that academic writing conventions presume a knowledge and command of the different degrees of 'lamination' to convey the different degrees of responsibility within the text, correctly attribute differing degrees of authorship, and avoid incorrect referencing. Therefore, Scollon claims, the problem for non-native speakers of English (in this case, Chinese-Cantonese students) is understanding the concept of *authorial stance* that takes responsibility for what is said, more than the lack of unawareness of the linguistic mechanisms of acknowledgement of authorship. This is due to the fact that, although these students can acknowledge facts, they see the authorial voice as a collective phenomenon, rather than an individual and autonomous one, and this impacts their use of referencing conventions. The individual author expected in academic writing in English is, therefore, an ideological construct that is in conflict with the cultural background of non-native speakers of English. And so is the Western definition of plagiarism; as people from different cultures are subject to conflicting discourses, the oversimplified perspective of plagiarism as a dualistic mode of 'wickedness' or 'incompetence' (Scollon, 1995: 4) represents the 'ideological arrogance' of the West (Scollon, 1994: 45).

In parallel with this lack of a geographically universal definition of plagiarism, Jameson (1993) describes a lack of a universal concept across different genres. She claims that the definition of plagiarism, which in its broadest sense consists of 'the misappropriation of materials (ideas, facts, words, structures) that were created, originated, or discovered by someone else' (Jameson, 1993: 19), is largely genre-related, so that the general conventions and constraints imposed on writing documentation tend to vary across genres. On the other hand, as Angèlil-Carter (2000) later claimed, not all students will be required – even expected – to acknowledge their sources when doing a job, after they leave the University. Partly as a result of this lack of homogeneity, Jameson too argues that a distinction should be made between plagiarism as the failure to use and follow documentation conventions and plagiarism as fraud, between intentional and inadvertent plagiarism. Although these arguments are convincing, her explanation of what is intentional plagiarism (stealing ideas from other authors, involving 'a whole work or extended parts of it'), when compared to inadver-

tent plagiarism ('a pastiche of passages interspersed with original prose') (Jameson, 1993: 22) is hardly unanimous. By giving precedence to quantity over the quality, this distinction makes the dangerous assumption that plagiarism is intentional when reusing a whole work – or a significant part of it – literally, but inadvertent if strings of text are alternated with pieces of own text. This classification would include in the inadvertent cases of plagiarism, in Howard's terms, both patchwriting and unintentional non-attribution of sources, but could misclassify cases of intentional plagiarism as unintentional. Despite this classification problem, her argument that plagiarism is largely genre-related remains valid for the most part, if we consider that the academy and the field of research are more demanding in terms of observing the rules of academic writing, referencing and acknowledgement.

Taken together, these theories challenge three commonly held assumptions: (a) that plagiarism can be assessed and punished on equal terms in academic and non-academic contexts, and moreover with disregard for the student's learning process; (b) that one identical definition of plagiarism can be used universally, and that such a definition is understood identically by everyone, regardless of the context (cultural and others); and (c) that even within academic contexts, plagiarising students should be treated equally. On the contrary, they suggest that keeping the students informed of what plagiarism is and how it can be prevented, rather than providing confusing definitions or presuming that the definition of plagiarism itself is 'common knowledge', is mandatory to assess and penalise instances of plagiarism fairly. Practice, however, is far from theory, and an analysis of plagiarism definitions used by universities demonstrates that they tend to vary. The following section presents some of the definitions used by some universities worldwide.

1.3.2 Institutional Approaches to Plagiarism

The simple dictionary definition of plagiarism as 'when someone uses another person's words, ideas, or work and pretends they are their own'[2], is fairly stable across at least several languages. But although this definition serves the aims of a dictionary definition, which is to provide a general, overall idea of the defined term or concept, it does not suit the aims of a university's definition, which is expected to describe in some detail what plagiarism is. In this sense, definitions of plagiarism used by higher edu-

[2]Longman Dictionary of Contemporary English Online – `http://www.ldoceonline.com/dictionary/plagiarism`

cation institutions worldwide tend to vary. The number of universities worldwide is too vast to allow for an in-depth analysis of a more comprehensive list of institutions. Given this impossibility, we selected a small number of universities from different geographical areas for a more detailed consideration. Although, owing to this small number, we have to be cautious about making generalisations, as I shall argue this discussion will contribute to understanding the different approaches to plagiarism worldwide, where these exist.

Anglo-American Universities, when compared to their Iberian counterparts, for example, usually have more clearly-defined definitions of plagiarism, and respective guidelines to avoid it. As Carroll and Appleton (2001) argued, institutional policies play an important role in informing students on plagiarism and how to avoid it, and most institutions have a handbook containing this information. In the UK, for example, vast research has been conducted, within the scope of the JISC or the Higher Education Academy (HEA), that demonstrated the importance of student-centred policies and methods to avoid plagiarism. Specifically, Carroll and Appleton (2001) found that, besides providing clear policies that are easily understandable by students, conducting induction courses that train students into avoiding plagiarism and redesigning assessments, creating individualised tasks and reconsidering the learning outcomes contributes to reducing the number of plagiarism cases. But, equally important, they found that encouraging the students' involvement and interest is crucial to reducing the number of plagiarism cases. More recently, de Nahlik (2011) demonstrated that it is a common belief that ethical values, and their teaching, impact the values and behaviours of students.

If Aston University is taken as an example, plagiarism is mentioned in at least four different official documents. The *Aston University Regulations on Student Discipline*, for example, claim that plagiarism is 'where a student uses, without acknowledgement, the work of other people and presents it as their own which may give an unfair advantage over others'. It is only slightly different from the one used in the *Aston Terms and Conditions, online enrollment 2008/2009*, which defines plagiarism as 'a form of cheating in which a student uses, without acknowledgement, the intellectual work of other people and presents it as his or her own'. A slightly different definition of plagiarism is presented by the Aston University School of Languages and Social Sciences, describing it as the 'deliberate unacknowledged borrowing' (including 'lifting chunks from another author, with just a few words or phrases changed to attempt to disguise

the plagiarism'), and considers inadvertent 'borrowing without acknowledgement ... bad academic practice'. The association of plagiarism to an intentional act reflects on the university's *Tariff of Penalties for Academic Offenses 2008-9*, where plagiarism is described as an academic offence, subsequently associating each 'type of offence' to a 'severity level', and correspondingly each severity level to a particular type of punishment[3]. Reflecting to a certain extent the pedagogical approaches that Howard, among others, considered to be desirable in academic contexts, the University considers that a distinction should be made between judgements of *plagiarism* and *cheating*, on the one hand, and *poor academic practice*, on the other. From a disciplinary perspective, this permits considering extenuating circumstances that may affect the product of the student's work, although, by focusing on the dichotomy between plagiarism and performance, it leaves aside the possibility that these faults may be part of the student's learning.

A different, more detailed definition is provided by the University of Manchester Faculty of Humanities[4] (ironically, in April 2009, this definition was still used by the Universitat Autònoma de Barcelona, with acknowledgement, in their website; the page is however no longer available):

> Plagiarism is presenting the ideas, work or words of other people without proper, clear and unambiguous acknowledgement. It also includes 'self-plagiarism' (which occurs where, for example, you submit work that you have presented for assessment on a previous occasion), and the submission of material from 'essay banks' (even if the authors of such material appear to be giving you permission to use it in this way). Obviously, the most blatant example of plagiarism would be to copy another student's work. Hence it is essential to make clear in your assignments the distinction between: the ideas and work of other people that you may have quite legitimately exploited and developed; and the ideas or material that you have personally contributed.

This definition is consistent with some of the guidelines provided by Carroll and Appleton (2001). In a short paragraph, a description is presented of the definition of plagiarism ('plagiarism is presenting the ideas, work or words of other people without proper, clear and unambiguous acknowledgement'). This definition is elaborated with examples, and an explanation of those examples ('it also includes 'self-plagiarism'

[3]Refer to `http://www.aston.ac.uk` for Aston University documents mentioned, and to `http://www.aston.ac.uk/lss` for the School of Languages and Social Sciences documents and website.

[4]University of Manchester – `http://www.humanities.manchester.ac.uk/studyskills/essentials/writing/avoiding_plagiarism.html`

(which occurs where, for example, you submit work that you have presented for assessment on a previous occasion), and the submission of material from 'essay banks' (even if the authors of such material appear to be giving you permission to use it in this way)'; 'obviously, the most blatant example of plagiarism would be to copy another student's work'). A solution to the problem is presented ('hence it is essential to make clear in your assignments the distinction between: the ideas and work of other people that you may have quite legitimately exploited and developed; and the ideas or material that you have personally contributed'). Additionally, the definition is complemented by a section of *do's and don'ts* of academic writing, and conversely what to avoid. Interestingly, although the definition is very inclusive (e.g. it discusses purchase of academic assignments from essay banks as a form of plagiarism), it describes collusion separately, as a different form of cheating, although the penalties are similar to those for plagiarism.

Unlike Aston University's definition, however, the University of Manchester definition does not include any information on the intentionality of the plagiarist, which suggests that all cases are equally treated, with the individual assessment of each case possibly considering extenuating circumstances and degrees of severity. However, that information is not provided.

Other universities adopt a more legalistic definition of plagiarism, as is the case of the Université de Sorbonne, which defined plagiarism as

> the act of copying a design, and omit the acknowledgement of the source. In the literary domain, it consists in copying, in full or in part, a work, and make it pass as their own. From a legal point of view, plagiarism is considered a serious infringement to copyright, and may be tried in equal terms to contra faction. The limit between imitation, inspiration and plagiarism is, sometimes, difficult to determine.[5]

What is interesting about this definition is the degree of intertextuality to the copyright laws, which in most countries are based on the Berne convention. The similarity, including the use of the terms *copyright* – and its infringement thereof –, *contra faction* and *design*, is striking, using technical, legal jargon, contrary to what is expected from an educational institution.

[5]This definition is my own translation from French, and was available online in April 2009 at the Université de Sorbonne's website – `http://www.paris-sorbonne.fr/fr/spip.php?article4752`. The page is currently unavailable, and no reference was found to plagiarism recently on the University's website.

In the same vein, the Université de Montréal (in Canada) establishes a comparison between plagiarism and fraud: 'plagiarism consists of passing off someone else's texts or ideas as our own', whereas fraud is 'an act of deceit intended to gain a personal advantage, often to the disadvantage of others'[6]. The distinction between fraud in general and plagiarism in particular is not clear, as the two terms tend to co-occur in the text; on the other hand, the examples given, which would be expected to provide pedagogical examples, albeit illustrative of fraudulent activities, are not useful to understanding what plagiarism is.

In Germany, the University of Applied Sciences of Berlin, in its Portal Plagiarism[7], claims that it 'is very difficult to define [plagiarism] exactly', and pedagogically asks: 'is it just a one-to-one copy of sentences? How much has to be copied in order to be considered a plagiarism? Is it still plagiarism if you edit the text? If you only use the order of arguments and do not copy the words?' Curiously enough, they do not provide their own definition of plagiarism; they cite the definitions of Gerhard Fröhlich, University of Linz, Austria, and of the German author Paul Englisch; the former defines plagiarism as 'the unauthorized use of the intellectual property of other persons, also known as 'theft"; the latter argues that plagiarism

> is the free decision of an author or an artist to take a not unsubstantial portion of the intellectual content of another person and use this in his own work with the intention of blurring this enforced loan [to] through appropriate means so that the pretense of it being one's own work is encouraged in the reader or viewer.

In Asia, despite the argument for the cultural differences that interfere with the writing expectations imposed by the 'Western tradition', Universities are adopting policies and guidelines identical to those of the 'West'. In the description of the English module contents, the University of Tokyo includes a note advertising 'ZERO tolerance for plagiarism. Students who plagiarize will not only fail the course but will forfeit all the credits for this semester. There will be absolutely no room for negotiation on this matter'[8]. This is however not exclusive to the English module. The University has available a document on referencing and plagiarism[9], which describes what plagiarism is and

[6] Université de Montréal – `http://www.integrite.umontreal.ca/definitions/fraude.html` (own translation).

[7] Portal Plagiarism – `http://plagiat.htw-berlin.de/start-en/`

[8] University of Tokyo – alternative, shortened url – `http://tinyurl.com/c2vko5v`.

[9] `http://lecture.ecc.u-tokyo.ac.jp/~cwpgally/references/200610_Quoting_and_plagiarism_English.pdf`

instructs how to reference, while admitting the possibility of 'accidental plagiarism'. Interestingly, these guidelines include a section on how to manage translation from the sources, indicating that these need to be quoted just like an original.

1.3.3 Plagiarism and the Law

Since the publication of 'Stolen Words: The Classic Book on Plagiarism' (Mallon, 2001), plagiarism has been often connoted with metaphors of criminal acts, including *theft* (Angèlil-Carter, 2000), *misappropriation* (Jameson, 1993) or *illegal appropriation of ideas* (Turell, 2008), *violation*, *kidnapping* (Johnson, 1997) and *crime* and *wordnapping* (Scanlan, 2004). Some of these authors (Angèlil-Carter, 2000; Jameson, 1993) concur that this association is not independent of the attribution of property rights to intangible goods, in the late eighteenth century and mid-nineteenth century, which granted the right to individual property and required a legal framework to cope with the infringement of those rights. Additionally, if we take Garner's definition of fraud as knowingly misrepresenting the truth or concealing 'a material fact to induce another to act to his or her detriment' (Garner, 2009: 731), then we have to agree that plagiarism is a fraudulent behaviour. However, plagiarism is not always considered on the grounds of the oversimplified dichotomist relation between the moral – or ethical – and legal grounds, and perspectives of how plagiarism cases should be handled, and whether and how they should be subject to punishment, have not been unanimous over the centuries. The metaphorical association between the violation of intangible, non-material property and the violation of material goods, which attaches to plagiarism a certain 'illegal' nature, is therefore undeniable. Plagiarism may consequently be considered a violation of both moral and property rights (Leitão, 2011).

Etymologically, the word plagiarism is derived from the Latin *plagiarius* ('kidnapper') and from the Greek *plagion* ('a kidnapping')[10]. The concept evolved over time with the evolution of the concept of authorship, and started being used more frequently after the invention of printing (15th Century), the rise of writing as trade and the decline of patronage (16th Century), the acceptance of the legal claim to texts (early 18th Century), the emergency of the originality in writing (Neoclassicism), and the development of the belief of ownership of non-material goods (mid-18th Century) (Jameson, 1993). This view of the evolution of plagiarism as being directly related to the evolution of

[10]Compact Oxford English Dictionary (online) – `http://www.askoxford.com/dictionaries/compact\_oed/?view=uk`

authorship as a result of the invention of the press is plausible, but not consensual. In his discussion of what an author is, Foucault (1979) convincingly argues that originally signing a text as the author was more an issue of taking responsibility for what was said than taking credit for it.

Nowadays, most legislation on copyright, including plagiarism, is a result of the proliferation of general principles of copyright law, owing for the most part to the almost universal adoption of the Berne and Paris Conventions, especially in terms of originality and requirements of being fixed in material form. These conventions have influenced intellectual property law in general, and copyright law in particular in several countries (Bently and Sherman, 2009). In simple terms, these conventions guarantee both economic and moral rights. The *moral rights* grant an author the right, among others, to 'claim authorship' and 'to object to certain modifications and other derogatory actions' (article 6bis of the Berne Convention[11]), thereby preserving the 'paternity' of the work (Pereira, 2003). The *economic rights*, on the other hand, grant the owner of the work the exclusive right to profit from that work (Pereira, 2003).

Over time, however, the concept of plagiarism has oscillated between a discussion of whether it is immoral or illegal. Lindey (1952), cited in Jameson (1993), established a parallel between plagiarism and copyright infringement to claim that the former is immoral, whereas the latter is illegal:

> Plagiarism covers a wider field; infringement involves more serious consequences. For purposes of plagiarism, the material stolen need not be in copyright; for infringement, it must be. There can be no plagiarism without the thief's posing as originator; infringement may occur even though proper authorship credit is given.

This distinction is later echoed in other works. Anderson (1998), for example, builds upon similar assumptions to claim that plagiarism is mostly considered an *ethical* problem, whereas copyright infringement is a *legal* issue. Later, Goldstein (2003) reaffirmed this belief by considering that plagiarism is neither a legal doctrine nor a legal offence, which makes it liable to be enforced by academic authorities, rather than by courts. Garner (2009) concurred that *generally* plagiarism is immoral but not illegal.

This debate is not surprising, considering that in most legal systems (including the Portuguese and the English) the concept of 'illegal' in the construction of the social norm is commonly based on the principle of that which is 'immoral', so that what is

[11] Berne Convention for the Protection of Literary and Artistic Work – `http://www.wipo.int/treaties/en/ip/berne/trtdocs_wo001.html#P123_20726`

immoral often becomes illegal, and as a consequence the view of plagiarism as a purely ethical issue has been refuted (Eiras and Fortes, 2010; Finnis, 1991). In fact, this dualistic view of plagiarism as having both moral *and* legal implications, involving moral and property rights respectively, was considered by Merton (1968), several decades ago, to explain that the violation of rules of acknowledgement imply 'moral and sometimes legal sanctions visited upon those judged to have violated the norm through the kinds of grand and petty intellectual larceny which we know as plagiarism'. This is the perspective that tends to be reflected in most of the policies adopted by organisations worldwide, that tend to consider that plagiarism is in practice an act of misconduct, conducted with the intention of misleading others (White, 2004), and hence the result of a fraudulent activity. In the UK, the UKRIO (UK Research Integrity Office)[12] describes plagiarism, alongside other activities such as fabrication, falsification and misrepresentation of data, as a type of misconduct in research, with serious disciplinary implications. In the USA, an identical approach has been adopted by the Public Health Service, the National Science Foundation and the National Academy of Sciences, that include plagiarism in their definition of scientific misconduct (Hansen and Hansen, 1995).

In Portugal, plagiarism is equated with theft and contrafaction, being legally framed by copyright and related rights law, whose most updated version (2008) reflects the dispositions of the 2004 EU Directive. By law, all intellectual creations of literary, scientific and artistic nature, which have been made public via any medium, regardless of their genre, media, merit, means or objective, are entitled to protection, but not the original ideas. That protection encompasses the two types of rights foreseen in the Berne Convention, protecting both the author's financial and personal (also referred to as 'moral') rights. The Portuguese law establishes that using a work without the authorisation of the author represents a crime of misappropriation (article 195), whereas reusing someone else's work, in whole or in part, as one's own, to the extent that there is no marked individuality between the two, represents the crime of counter-faction (article 196). In this case, the following four criteria must be met for that conduct to be considered an infringement: it must (i) be deceitful; (ii) use other people's work and pass it off as one's own; (iii) be a mere reproduction of someone else's work; and it must (iv) reproduce someone else's work in such an identical way that it has no individual value.

[12]UKRIO – `http://www.ukrio.org/`

The violation of these property rights involves criminal sanctions, including imprisonment and fines (article 197). Ascensão (1992: 65) justified these requirements on the grounds that 'plagiarism is not a mere copy; it is more insidious, because it appropriates the creative essence of the work assuming a different form or shape'. Identical legal sanctions are foreseen in the law for the violation of moral rights, by intentionally taking the authorship of a work, or harming its integrity or the reputation of its author (article 198). Plagiarism thus represents both a moral and a financial loss for authors, the latter deriving from the inability to financially benefit from their own rights, as they wish. In cases of plagiarism, and particularly academic plagiarism, it is mostly the guarantee of the moral rights of the original author of the work that needs to be considered. The fact that the works entitled to protection are immaterial and ubiquitous allows their simultaneous use by different people, which compromises the original author's ability to control the use of his/her own work (Pereira, 2003: 20). In this sense, article 27 of the copyright law claims that the author is the intellectual creator of the work, whose name is referenced as such in the work. Use of summaries or excerpts is allowed, e.g. for educational or scientific purposes, but those references should always bear the identifying elements of the intellectual work (article 76). Taken together, these articles guarantee the legal protection of the author, not only against copyright infringement, but also against plagiarism.

The application of these legal premises also reflects on the policies of international organisations. The IEEE (originally called *Institute of Electrical and Electronics Engineers, Inc.*), for example, defines plagiarism as 'the use of someone else's prior ideas, processes, results, or words without explicitly acknowledging the original author and source', and considers that 'plagiarism in any form is unacceptable and is considered a serious breach of professional conduct, with potentially severe ethical and legal consequences' (IEEE, 2006: 57).

Once more, plagiarism is represented as an act of misconduct that is subject to both ethical and legal implications. However, two particular assumptions of this approach to plagiarism can be challenged; firstly, that plagiarism can be determined by chronological factors; and secondly, that its severity can be determined by calculating the amount of copying between the first (the supposed original) and the most recent (the supposed instance of plagiarism) documents, according to one of five levels (*uncredited verbatim copying of a full paper*, *uncredited verbatim copying of a large portion (up to 50%) from a paper*, *uncredited verbatim copying of individual elements*

(paragraph(s), sentence(s), illustration(s), etc.), *uncredited improper paraphrasing of pages or paragraphs*; and *credited Verbatim Copying of a Major Portion of a Paper without Clear Delineation*) (IEEE, 2006: 57). These two aspects, which have been discussed at reasonable length in the field of forensic linguistics, will be discussed in the following section, 1.3.4.

1.3.4 Linguistic Approaches to Plagiarism

Traditionally, plagiarism has been considered mainly a problem of *unlawful borrowing of words* and linguists may be inclined to agree, not the least because as Coulthard and Johnson (2007) clearly put it, the type of plagiarism that they are able to demonstrate is mainly linguistic plagiarism. Although plagiarism can take place in any 'text', in the social semiotic sense of any type of unit where meaning is realised and produced (Kress, 2001: 544), or 'text genre', as understood by Saville-Troike (1982) and later expanded by Swales to comprise 'a class of communicative events, the members of which share some set of communicative purposes' (Swales, 1990: 58), detecting plagiarism of a song, a film or a photograph is beyond the scope of the linguist's skills. However, 'linguistic plagiarism' can have different meanings; it can mean either *text that has been reused from other works*, as is the case of text that is lifted verbatim from another source; or it can mean *any instance of plagiarism that gains form in linguistic text*. This is an important distinction, considering that the former does not contemplate lifting of works and ideas, but the latter does. The latter would permit the analysis of cases using translation, but the former would not, although the object – the need to find evidence of plagiarism – is shared by both.

Finding evidence of plagiarism and demonstrating that it represents a fraudulent activity, whose deceptive nature results e.g. from lying (Eiras and Fortes, 2010), presupposes the ability to detect the instance(s) of plagiarism and fraud. The latter, in particular, can be a challenging task, especially if we consider that linguistic analysis has limitations in detecting deception (Eggington, 2008). But this is the area to which linguistic evidence can contribute. This type of evidence, provided via a forensic linguistic analysis of the texts, has been increasingly used in cases of fraud, where moral and/or economic rights are suspected to have been violated, both in academic and non-academic contexts.

Turell (2008) contends that, in countries within the Civil Law tradition (forensic) linguists are seldom called upon as expert witnesses in cases of plagiarism, either

because cases of plagiarism are rarely taken to court or because linguistic evidence is not traditionally accepted in court of law. However, forensic linguistics, which consists of using linguistic methods and analyses in forensic contexts, has been used effectively to investigate and provide evidence of fraud, e.g. in the detection and investigation of text reuse, including plagiarism and collusion. Although linguists have been increasingly asked to detect, investigate and/or confirm – or refute – instances of student plagiarism in the academy, as reported by Coulthard and Johnson (2007), the potential of linguistic analysis to proving textual reuse in non-academic contexts is also demonstrated. Citing the example of the document 'Iraq: Its Infrastructure of Concealment, Deception and Intimidation' (also known as the 'Dodgy Dossier'), which the British government presented to the United Nations in 2003 to justify their invasion of Iraq, the authors discuss the extraordinarily high textual identity between this report and a prior academic article. A linguistic analysis of the two texts demonstrated that the official governmental document had been substantially plagiarised from this article, with changes in spelling only (from American English to British English). Similarly, Turell (2004, 2008) discusses a case that was taken recently to the Spanish courts involving copyright disputes. Upon suspicion that a translation of Shakespeare's *Julius Caesar* into Spanish derived from another previously published translation, rather than having been produced independently, the two translations (the suspect translation and the original) were compared against each other and against two other translations.

The focus on translated works is particularly interesting methodologically, from a forensic linguistic perspective, since a translated text is in this sense considered a derivative text, in that it is a rewriting of an existing text in a new language. This by itself determines some of the translator's choices; the contents and the form in general, as well as its gist and the idea, are expected to be as close as possible to the original (i.e. the text it derives from), although the wording is necessarily new and hence provides a new original. The challenge for the forensic linguistic analysis is that a higher volume of overlap among translations is to be expected, and consequently this could impact the statistical significance of the results. The linguistic analysis, however, demonstrated a statistically significant amount of *overlapping vocabulary*, *shared once-only words*, *words unique to each file*, *shared once-only phrases*, and *identical or similar phrasing* between the suspect translation and the original translation. She thus showed that the use of linguistic evidence can help determine a type of plagiarism that is often decided in court of law: translation of literary works.

Investigations of instances of plagiarism are usually based on a comparison of the suspect text(s) against possible originals to find linguistic data to investigate whether the suspect text is a derivative text, or otherwise an original. In its simplest form, this analysis could imply matching the suspect text against the source and highlighting the identical strings, which are usually sequences of a few or several words that are copied verbatim, without any (major) changes to the original from another source and used without acknowledgement. In cases of verbatim plagiarism like this, where the original text is lifted literally, word-for-word, a simple comparison of the suspect and the original is sufficient to identify the overlapping, identical phrases, sentences or even paragraphs.

However, the investigation is more complex when the derivative text is frequently edited, whether to disguise the authorship and make it pass as their own, or in a failed attempt to write properly – as is often the case in academic writing. Such changes can consist, for example, in altering the word order or reformulating the sentence structures, in paraphrasing the original text or changing the cohesion and the coherence of the original. These strategies, which involve simple to sophisticated alterations in grammar, punctuation, syntax, and semantics or even in vocabulary and discourse, make the detection procedure more difficult. Changes in grammar, punctuation, syntax and word order usually imply text re-ordering in a sense that the sequences of identical words are interrupted, building an apparently distinct original – albeit containing the same, non-original ideas, and possibly reusing some vocabulary. As a result of the reuse of a number of identical words in a different order, the verbatim detection procedure fails to identify sequences of identical words of a length that is significant enough to be considered plagiarism.

In order to overcome problems of this type, more sophisticated methods are required, such as the one originally used by Johnson (1997): discarding strings and sequences or chains of words (the methods traditionally used to detect verbatim plagiarism) and grammatical items, which are closed sets of words (and hence smaller in number, and likely to be shared anyway), she concentrated on the analysis of *shared lexical items*. After calculating the percentage of lexical types in the set of three suspect documents against that of a set of three non-suspect documents, she found that the percentage of overlapping types (i.e. the number of types of lexical vocabulary occurring in the text) in the latter set amounted to 20%, compared to 72% in the former. This demonstrates that an analysis of lexical overlap is robust even with changes

in syntax and word reordering, allowing the detection of instances of plagiarism that are usually missed by search of identical strings. The theoretical explanation provided by Coulthard (2004) and Coulthard and Johnson (2007) for these empirical results is grounded on the principle of linguistic uniqueness, i.e. that even the same person writing on the same topic on different occasions would be expected to word the text differently; in the case of two or more texts each authored by different people, lexical overlap would therefore indicate either that one is derivative from the other(s) or that they have been produced collaboratively.

This type of lexical analysis may, however, be of limited applicability in the analysis of the coherence and cohesion of instances of plagiarism, which can be disguised by edits introduced to the derivative texts usually involving changes in the words of the original to reflect a textual or extra-textual, 'real-world' reality that does not necessarily match that of the source text. In other words, the plagiarist may or may not retain the strings and sequences/chains of words of the original, as well as lexical items related to the topic of the text, but elements of coherence and cohesion can be adapted to convey a coherent link to the plagiarist's reality. Conversely, those instances can be detected more effectively by identifying inconsistencies usually revealed in *referential style* (e.g. inconsistent use of imperative or infinitive verb tenses in forms of address), *decontextualisation* (e.g. by omitting parts of text that otherwise contribute to setting the context of the text reused), and *inversion of structural elements* that result in conceptual inconsistencies (Turell, 2008). Although extremely useful when the derivative text has flaws at these levels, these linguistic markers tend to break in cases where alterations are successfully made to retain the coherence and the cohesion of the plagiarised text; in this case the linguistic analysis of the text needs to concentrate on the analysis of differences surrounding identical textual elements, more than on the analysis of similarities and inconsistencies.

Prior authorship and volume of lifting are two other criteria used to identify instances of plagiarism, and Turell (2008) makes a good case for both of them. Prior authorship, which is usually determined by the date of publication, helps resolve issues of directionality, by chronologically determining which text is the original and which one is the derivative. As a method, it raises problems in cases where the dates of publication are very close, or when the two texts are contemporary (Turell, 2008). On the other hand, considerations of volume, which are based on the assumption that the higher the percentage of overlapping text, the more likely it is that two (or more) texts have not

been produced independently, are relevant both in academic and non-academic contexts. In academic contexts, universities seem to often base their definitions of plagiarism on the assumption of 'substantial' borrowing (Coulthard and Johnson, 2007). In non-academic contexts, consideration of different levels of volume of lifting have been used to determine different levels (or degrees) of plagiarism, as was described in section 1.3.3. Empirical evidence demonstrates that using quantitative measures such as similarity in overlapping vocabulary, shared once-only words, unique vocabulary and shared once-only phrases (Johnson, 1997; Turell, 2004; Woolls and Coulthard, 1998; Woolls, 2003) can effectively contribute to the start of the analysis, but it is also admitted that '[t]aken in isolation, it is possible that all these measurements do not discriminate sufficiently' (Turell, 2008: 288). Woolls (2010) even points out the need to question judgements based on attaching a percentage to the degree of similarity, while on the other drawing attention to the dangers of approaching word and even lexical overlap blindly, with illustrative examples of texts that have a high word overlap, but do not contain a semantic relationship. Alternatively, biased judgements of plagiarism based solely on quantitative criteria (plagiarism thresholds) can be avoided by using a combination of quantitative and qualitative analyses. These help demonstrate, based on the principle of idiolect and linguistic uniqueness (Coulthard, 2004; Coulthard and Johnson, 2007), that it is very unlikely that two different people at different occasions produce identical text, and the amount of identical text across the documents from the same sets could indicate that they were either (a) produced by the same person(s), (b) produced by different person(s), with or without the knowledge of the other(s), or (c) be both based on a third text.

1.3.5 Computational Forensic Linguistics and Plagiarism Detection

Computational (forensic) linguistics approaches the task of plagiarism detection from two different angles: plagiarism detection by comparison, whereby a suspect text is compared against a closed set of tests, or alternatively against an open set, such as the Internet (in cases where the possible sources are not known); plagiarism detection by stylistic analysis, which consists of determining whether a suspect text has been plagiarised by identifying linguistic inconsistencies in the text. The former, which is a procedure of 'plagiarism identification' by comparison, has come to be known (especially among computational linguists) as 'external plagiarism detection'; the latter, which consists of applying stylometric analyses to authorship, is known among them

as ‘intrinsic plagiarism detection’, since the text is only compared with itself.

Finding evidence of plagiarism involves many terms of an equation, especially in an era of massive textual production, and although the comparison of a ‘closed’ set of document, such as a document pair comparison, can and does provide good results, comparing a suspect text against an ‘open’ set of unknown texts (or systematically searching for plagiarism in non-suspect texts) through ‘manual’ search and comparison is a challenge, if not a human impossibility. Creating and refining plagiarism detection software is therefore crucial to provide such evidence, and this is an area particularly suitable for computational forensic linguistics.

Computational forensic linguistics, which Woolls describes as ‘a branch of computational linguistics’ (Woolls, 2010: 576), is the development and application of computational tools in forensic contexts. The field has attracted the interest of computational linguists over the last decades, especially in the areas of authorship analysis, with a focus on stylistic analysis and stylometry (Abbasi and Chen, 2008; Argamon and Juola, 2011; Argamon *et al.*, 2007, 2009; Ayala *et al.*, 2011; Hirst and Feiguina, 2007; Koppel and Schler, 2004; Luyckx, 2011; Mikros and Perifanos, 2011; Smith, 1994; Solorio *et al.*, 2011; Stamatatos, 2006). Although some of these works focus on the application of computational methods to traditional analysis of the authorship in non-forensic contexts, such as the one described by Hänlein (1998), the most recent ones in particular (e.g. Argamon and Juola (2011); Ayala *et al.* (2011); Stamatatos (2006); Luyckx (2011)) present research that follows what happens in forensic contexts. Although they may not use real forensic data, at least they use new media as corpora (e.g. news editorials, emails, tweets) that resemble the natural language that would be expected from a corpus of texts of forensic nature.

Prolific research in the field of computational forensic linguistics has also been conducted for applications related to (semi)automatic plagiarism detection. Early work includes that of Brin *et al.* (1995), who investigated copy detection mechanisms. Bernstein and Zobel (2004) later presented an approach oriented to the identification of co-derivative documents, which is useful in any forensic context, involving not only cases of plagiarism, but also other applications where the focus is on determining whether two or more texts have been produced independently or otherwise. In the same vein, Iyer and Singh (2005) and Stein and Eissen (2006) presented computational approaches to analyse the similarity between documents, with Cha (2007) later focusing specifically on similarity measures. In parallel with these, a vast range of

studies concentrated specifically on plagiarism detection (Barrón-Cedeño and Rosso, 2009; Hoad and Zobel, 2003; Kestemont *et al.*, 2011; Meyer Zu Eissen and Stein, 2006; Oberreuter *et al.*, 2011; Rao *et al.*, 2011; Stamatatos, 2009a; Stein *et al.*, 2007; Uzuner *et al.*, 2005), each with different aims or focusing on different aspects of language processing. Hoad and Zobel (2003) and Stein *et al.* (2007) conducted research on plagiarism detection in general, by analysing versionised and plagiarised documents and discussing computational strategies to identify plagiarised documents. Uzuner *et al.* (2005), on the other hand, based their plagiarism detection methods specifically on the use of syntactic information, and later Meyer Zu Eissen and Stein (2006) focused on 'intrinsic plagiarism detection', i.e. on the analysis of the writer's style, or their stylistic inconsistencies, to detect plagiarism. Approaches to intrinsic plagiarism detection have also been adopted by Stamatatos (2009a) and Kestemont *et al.* (2011). Their methods varied only slightly, in that whereas the latter focused specifically on calculating distance scores by analysing trigrams, the former focused on determining and analysing the *n*-gram profiles of the writers. Earlier approaches to plagiarism detection based on the analysis of *n*-grams had been used by Barrón-Cedeño and Rosso (2009), who concentrated on detecting plagiarism by comparing a suspect document against a corpus of potential originals. Research on plagiarism detection, both 'intrinsic' and 'external', has more recently been conducted by Oberreuter *et al.* (2011) and Rao *et al.* (2011), the latter basing their approach on discourse markers. Although the focus of these methods is not on forensic contexts, and although most or all of them use fabricated *corpora*, commonly containing invented examples, their application to forensic contexts is not impossible.

A significant use of computational linguistics tools in forensic contexts was reported by Johnson (1997), who used computer tools to detect collusion in student essays. Using the method described in section 1.3.4, she found that a set of three suspect student essays had not been produced independently. Her work was followed by Woolls and Coulthard (1998), who described a set of 'tools for the trade', that linguists could use in forensic contexts. Woolls (2003) later followed up this work, exploring the development of the tools. In particular, he describes the improvements made to CopyCatch, including the ability to handle multiple document files, of different formats, the output of sentence matches, besides shared phrases and vocabulary, and the presentation of file pairs with the shared vocabulary, and how it is used in the text. The *Gold* version of the software also allowed a comparison of files or sets of files against other sets. Reusing

some of the work described in these works, Woolls (2010) showed the complexity of detecting plagiarism, by using examples of English language. He demonstrated that even the apparently simple tasks of determining where the similarities appear in the text, and assessing their significance, are hampered by the complexities raised by the concept of similarity. Arguing that, when comparing human and computer detection of similarity between strings of text and documents, computers have the ability to handle large volumes of data, quickly and consistently, without the drawbacks of 'mental fatigue', he concludes that computer software 'can only be an approximation' (Woolls, 2010) to the capabilities of the human reader. More general aspects of plagiarism and plagiarism detection, within the scope of computational forensic linguistics, had already been approached in his previous research (Woolls, 2006), and continued in his subsequent work (Woolls, 2012).

1.3.6 Plagiarism and Authorship

The question that most forensic linguists – or even lecturers/tutors – ask when faced with a text of suspected plagiarism is whether the person who signed the text is the actual author, or on the contrary whether the text is a composition resulting from the work(s) of other people. Plagiarism builds upon two commonly held assumptions: (a) that the text was not written, in whole or in part, by the person who claims to have written it; (b) that the text reused from another source does not acknowledge the original author. Particularly in the academy, it is an expectation that the works are both original and authentic (unique (Anderson, 1998)). Taken together, these two conditions usually represent a violation of the moral, if not financial, rights of another text, which are usually resolved with reference to theories of authorship and authorship attribution.

The issue of authorship has, however, been questioned for several years, with the study of the individual uniqueness at the centre of attribution studies. Love (2002) claimed that it is the object of study of authorship attribution to focus on how the individuality of each person reflects in their writing. This individuality is explained in detail by Coulthard who, building upon the principle of 'uniqueness of utterance' and the theory of *idiolect*, argues that, even in academic writing, the expectation is that two different authors can and may 'choose any overlapping, but by no means identical, set of lexico-grammatical items' (Coulthard, 2004: 434) to express similar meanings. The principle of the individual enactment of each person in writing, however, does not pass

unchallenged. One of the problems is determining this individuality, considering that the process of socialisation, built upon sharing the language among the members of a society, can and often does 'overwrite' (Love, 2002: 4) some of the unique characteristics of the individual. Previous findings reported by Johnson (1997), specifically in the context of plagiarism, empirically demonstrated the existence of the uniqueness of each writer, even in academic contexts, and despite this socialisation and language sharing. She compared two sets of student texts on the same topic, to conclude that a certain amount of textual overlap is to be expected from two or more people writing on the same topic, but that choice is by no means identical. In part due to these findings, Coulthard (2004) concluded that the use of identical lexico-grammatical items by two different people, or by the same person at two different times, may therefore indicate that it derives from another text, rather than being produced independently.

In parallel with the discussion surrounding the 'uniqueness of utterance', runs another problem: determining to what extent a text can be original. In the 1960s and 1970s, structuralists and post-structuralists started challenging the Romantic and post-Romantic concept of authorship, arguing for a shift from the importance of the biographical subject, to focus on the text. Rather radically, Barthes defended the 'death of the author' to claim that the author-centred understanding of authority, and consequently originality, was wrong, and that the focus had to be shifted on to the language: 'it is language which speaks, not the author' (Barthes, 1977: 143). In a call for intertextuality, he argued that the text was not the message of an 'Author-God', but instead 'a multi-dimensional space', where a combination of non-original, previous writings and derivations from real-life situations 'blend and clash', to build a 'tissue of quotations drawn from the innumerable centres of culture' (Barthes, 1977: 146). In other words, writing becomes a neutral activity, where the language replaces the author, and the author's task is only to mix and counter writing, 'in such a way as never to rest on any one of them' (Barthes, 1977: 146), to then reach its true 'place' of destination – the reader. The assumption that the reading becomes 'the true place of the writing' marks the 'birth of the reader', at the expense of 'the death of the Author' (Barthes, 1977: 147–148), so that finally a text is freed from the boundaries of its authorship. This approach to intertextuality, which argues that there are no original texts, but rather that all texts are intertextual, raise particular questions to plagiarism: if there are no original texts, all derivative ideas consequently come to be understood, as (Love, 2002) argues, as 'influences'.

Foucault, not so much arguing for the 'death of the author', but discussing instead the 'disappearance of that author', questions the concept of the author. In Foucault's terms, this disappearance of the author enables the identification of the 'author-function'. In his discussion of the relationship holding between the author and the text, Foucault (1979) portrays the author as a function of 'discourse', to determine the circumstances under which an individual can fulfil the function of author. Unlike Barthes, he is not categorical in determining the 'death of the author', and does not deny the existence of the author. On the contrary, he argues that the author's name is functional in the sense that it works as a system of classification, that contributes, among others, to separating one text from the others. In this sense, works have become objects of appropriation, this property being originally motivated, not by financial aspects, but by the inherent responsibility (and corresponding risks) of their authors as transgressive discourses; in other words, by the legal and institutional context of the texts.

A second feature of authorship is that the 'author-function' is neither constant, nor universal, i.e. it is not uniform across all discourses, at any time, in all cultures. The authenticity of the literary texts was once determined, not by the identity of the author, but by their age. Conversely, scientific texts were treated differently in the Middle Ages and in the 17th and 18th centuries; the name of the author, which in the Middle Ages was a requirement to guarantee the truthfulness of the text, came to be accessory when methods of verification were developed: 'the role of the author disappeared as an index of truthfulness' (Foucault, 1979: 20). Interestingly, at this time the name of the author was required to determine the acceptance of literary texts.

The third feature of the 'author-function' is related to the construction of the author. Foucault claims that different genres[13], or the same genre at different historical periods, have been constructed differently. He argues that the 'author-function' involves procedures that are more complex than the 'spontaneous attribution of a text to its creator' (Foucault, 1979: 23). Citing Saint Jerome, he draws attention to the similarities between the Christian exegesis in authorship authentication and modern authorship attribution, according to which attribution to the same author required: (a) 'a standard level of quality'; (b) show a 'conceptual or theoretical coherence'; (c) be stylistically uniform; and (d) a chronological element, that writings ulterior to the death of the author would be excluded as their work (Foucault, 1979: 22).

[13]Foucault calls it 'forms of discourse' (Foucault, 1979: 21).

A fourth feature is that the 'author-function' often refers to a plurality of egos, rather than to one individual. Foucault gives the example of mathematical treatises, where the 'I' composing the treatise and the 'I' demonstrating it do not necessarily refer to one and the same subject.

A discussion of the complex nature of the author is offered by Goffman (1981) in his analysis of the 'production format' of an utterance. In his identification of the diverse roles of the speaker, Goffman (1981: 144) argues that, depending on the particular circumstances, the speaker can play one of the following roles:

- *Animator*: 'an individual active in the role of utterance production', i.e. the person speaking;
- *Author*: 'someone who has selected the sentiments that are being expressed and the words in which they are encoded', i.e. the person responsible for formulating their own text;
- *Principal*: 'someone whose position is established by the words that are spoken, someone whose beliefs have been told, someone who is committed to what the words say', i.e. the speaker who is publicly claiming their own position, therefore playing a social identity or a social role.

Goffman saves the role of 'speaker' for someone who simultaneously animates the text that they have formulated, while establishing their own position.

However, as McCawley (1999) recognises, Goffman's three roles are not absolute, but instead dependent on frames. This classification of the role of the 'speaker' in academic contexts can pose specific complexities, not the least in establishing the role that the student should play when writing academically. The student can be the speaker, if s/he composes the utterance, produces it graphically – if we apply Goffman's production of sounds to writing – and commits himself/herself to their text. However, academic writing usually implies integrating the voice of others in one's own writing, and as a consequence the roles played by the student-writer and the original writer are shared, and even tend to overlap at certain points. For example, in the case where a student quotes Goffman in their assignment, the student would undoubtedly be the animator, whereas Goffman would be the author; but who would be the principal? Goffman would definitely be the principal of the quoted text, the one who commits himself to the quoted utterance. However, the student also commits to the circumstances and the text surrounding the quotation, and in a sense, also by committing

himself/herself to the utterance, s/he can also be the principal. If the student, instead of quoting, paraphrased Goffman, the roles should not alter much. In this sense, the violation of the role of the author seems to be the one contributing most for judgement of an instance as being plagiarism. But there are also expectations that the student commits himself/herself to the text, even to the extracts that he animated, but did not author, and this expectation, which requires that a clear distinction is made between the student's commitment and the original author's commitment, is the one that may be more difficult to meet.

The complexity of authorship is also discussed by Love, who considers that 'a model of authorship requires a repertoire of practices, techniques and functions' (Love, 2002: 33). He provides a framework for written texts based on four different phases of authorship, which are determined by the nature of the individual agency involved in the authorship of the text: *precursory*, *executive*, *declarative* and *revisionary* authorship. Although his multifunctional model addressed mainly historic and literary attribution cases, the complexity and diversity of functions of authorship in forensic contexts was demonstrated by Grant (2008), who argued that questions of a forensic nature can reflect that complexity in many different ways.

Precursory authorship describes the incorporation of earlier work into one's own, new text. Operating as 'source' or 'influence' of the new text, the function of precursory author should not be taken to challenge the individuality of the new text, but instead s/he is seen as a collaborator that participates in writing that text. In non-academic writing, precursory authorship is so explicit that it is commonly not acknowledged. Love gives parodies as an example, and Coulthard and Johnson (2007: 184) point to the intentional intertextual references of T. S. Elliot. But precursory authorship also reflects in academic writing. In instances where the contribution of the precursory author is acknowledged, this participation is indicated by quotation or, when paraphrasing, by citation (i.e. indicating the author's name, the name of the work, year and page numbers), if applicable. Conversely, in cases of unacknowledged borrowing, use of precursory authorship is seen as plagiarism.

Executive authorship describes the function of the author who formulates the text, making decisions regarding the selection and coselection of textual elements, and subsequently for compiling the text. In cases of academic writing, the role of executive authorship is the one that the student plays when writing his/her essay; s/he chooses the textual elements that they will use to commit to the ideas that they aim to express.

Executive authorship, in this sense, is single. Executive authorship can also be collaborative, for instance in cases of group assignments, where the acknowledged work of different executive authors is put together. In cases of individual assignments where students have worked collaboratively, that collaborative work is not acknowledged, so the students are responsible for an act of collusion; from a perspective of authorship, however, the assignment is still one of collaborative executive authorship.

Declarative authorship describes the function of the author who validates the text, and comes to possess the words expressed in the text. In cases of plagiarism, like ghost-writing, Love's model allows for both the plagiarist and the ghost-writer to play the role of author, even if not that of executive author; but unlike plagiarism, ghost-writing tends to be morally acceptable, in cases where the declarative author is a famous person, e.g. writing their autobiography. Love insists that declarative authors still play a genuine role in authorship: although they might not have contributed to the creative process, they still participate in the process of authorship, not the least by giving their name, taking the risks and the credits for the work. Declarative authorship applies both to cases of traditional plagiarism, where a student borrows from another text without acknowledgement, as well as to cases of students purchasing essays from essay banks. Morally, both are condemnable, although from a legal perspective the financial rights of the executive author, in the case of essay purchase, are guaranteed.

Revisionary authorship describes the function of the person who revises the text, polishing and correcting it. This is the function performed by the executive author when revising their own text, but also by other reviewers, including editors before a text is published. From the perspective of authorship attribution, Love claims that revisionary authorship is often difficult to separate from primary composition. In this case, the distinction may be not one of telling between two different authors, but instead one of discriminating two different phases.

Love's framework of the functions of authorship demonstrates that the analysis of instances of plagiarism can often be a grey area. Academic expectations regarding student writing is that they present an 'original' and 'authentic' piece of writing, but, as Love demonstrates, an author can be 'original' in one of several instances. It seems clear that what is expected from the students is that they formulate their own text, thereby engaging in the ideas that they express, thus playing the role of original *executive author*. Precursory authors can contribute to their assignment, as long as these contributions are clearly indicated. This is a requirement that can impact the act

of academic writing, in that the deference and loyalty to the precursory authors that Love (2002) claims was so common in the Middle Ages and Renaissance, still seems to permeate some current cultural writing practices. On the other hand, revisionary authorship of the type offered to a student by his/her supervisor, whereby the latter suggests – and the former accepts – a set of revisions, is not reproachable in the academy. However, from an ethical perspective, this influence on the final result of the executive authorship can be challenged.

1.3.7 Punitive Trends

An extensive body of literature has demonstrated that plagiarism is metaphorically and metonymically associated with criminal discourse: 'crime', 'theft', 'lifting', 'violation', 'kidnapping', 'wordnapping'. Howard and Robillard (2008) argue that approaches and research to academic plagiarism in recent years came to mirror the media representations of plagiarism as being universal, based on generally-accepted definitions, with universal policies and generally-applicable preventive measures and punishments. This dual role played by the media had been earlier described by Cohen (1972) as 'carrier' and 'producer' of *moral panics*.

Moral panics are examples of collective behaviour that arise when 'folk devils' (the actors of deviant actions) practice certain social events that are viewed as a threat to social values, and hence raise collective 'concern, anxiety, indignation or panic' (Cohen, 1972: 10). Social events that are labelled as 'moral panics' usually result from a particularly illustrative case, whose significance and extent are subsequently exaggerated in themselves and in comparison with other actually more problematic social events, and are constructed upon a few essential elements, among others (Cohen, 1972: xii): a *suitable enemy*, usually an easy target that is given little power, whose deviant, vilified acts against a *suitable victim* – an ordinary, honest person, with whom the public can easily identify – contribute to building a general, false *consensus* of vulnerability, and that those acts, rather than insulated events, represent a general threat. This false consensus around the panic is fuelled by the media, who vilify the 'folk devils'. The panicky reactions, when combined with the perceived need for the protection of certain values, are preconditions for defining new social problems, and subsequently create new rules and normative concerns. New definitions of 'right' and 'wrong' are consequently created, and deviance to act according to the social norm draws attention, attracting punitive measures (Becker, 1963).

The media attention that plagiarism attracted in recent years suggests that the topic may be reaching a stage of moral panic, if the five essential elements of moral panics (Cohen, 1972: xxvi) are considered; (i) there is a *concern* about plagiarism being a potential threat, with comments that plagiarism is on the rise not always statistically grounded; (ii) this threat results in *hostility* towards the plagiarists, the 'folk devils', the actors of the deviant actions – students are failed or subject to severe disciplinary action, ministers are forced to resign and face social condemnation, etc.; (iii) plagiarism has reached a stage of general *consensus* that this threat exists and needs to be acted upon and solved, not only by systematically screening academic work for plagiarism, but especially by punishing the offenders; (iv) the mismatch between the number of cases, or their real significance, and the actual seriousness involved if action is not taken is often exaggerated and *disproportionate* ; (v) the fact that plagiarism is an old problem that has attracted a lot of attention recently demonstrates that it is *volatile*, and as a consequence it can remain a social concern for some time, but it can also dissipate suddenly. As a result of this panicky perspective, a culture of disbelief is promoted that naturalises academic plagiarism, whereby the students are seen, not only as potential, but more importantly as likely plagiarists. That deviance then 'emerges and is stabilized as an artefact of social control' (Cohen, 1972: 8).

This portrayal of students as 'folk devils' suits well Anson's claims that higher education institutions increasingly judge student plagiarism 'through a lens of criminality' (Anson, 2008: 140), adopting complex documents and procedures, such as policies, with the double aim of legitimising the punishment imposed on students, while simultaneously legally protecting themselves. This punitive context, where teachers are situated as 'plagiarism police' (Robillard, 2008), and where a dominant culture promotes control strategies, describes what Garland (2001) calls a *culture of control*.

Over the last three decades of the twentieth century, the UK and the USA in particular have gone through deep social and cultural changes, which Garland (2001: viii) calls 'the coming of late modernity', that impacted the perspectives of criminal justice, law and social order, and consequently the social responses to crime, first in the USA, and later in Britain. These changes led to a trivialisation of crime, which views it as normal and commonplace, committed by rational, normal, fully responsible individuals. Albeit admitting that the media can have a shaping effect in these late modern view of criminal justice, Garland (2001) is not in agreement with Cohen (1972). On the contrary, he claims that those responses stem from the collective experience of

crime in these societies. Higher rates of crime, which came to be seen as a normal social fact, demanded a collective adaptation – a new *crime complex* (Garland, 2001: xi) – and it was this adaptation that gave salience to crime. The cultural assumptions underlying criminal justice were the ones that changed more significantly, rather than the structures of criminal justice themselves. These changes are, he argues, mainly due to two reasons: the social organisation of late modernity and the conservative, free market politics that dominated the UK and the USA towards the end of the 20th Century. These responses to crime control brought along stricter measures and more frequent punishment, with the private sector working in policing and incarceration activities, electronic monitoring. A strong belief in restorative, rather than rehabilitative justice, ultimately gave rise to a much higher rate of executions and incarcerations in the USA, for example, compared to identical nations.

This new approach to crime control came to legitimise a policy that builds upon punitive, neo-liberal sentiments – even a punitive populism – rather than what Garland calls a 'penal welfarism', and give rise to an institutionalised culture of control, that is supported by particular configurations in cultural, political and economic segments of society. As a result, new practices of controlling behaviour and doing justice were put in place. Among some of the most important indices of change in that coming of late modernity, Garland identifies: the decline of the rehabilitative ideal, in favour of retributive penal measures; the re-emergence of punitive sanctions and expressive justice, which legitimated a retributive discourse that promotes a sense of vengeance, in detriment to rehabilitative sanctions, and which contributed to a language of condemnation and punishment entering official discourse; changes in the emotional tone of crime policy, which shifted from a view of the offender as the compassionately less fortunate and disadvantaged, and from a socially just solution, to a collective entitlement to anger, in which the fear of crime is foregrounded – which criminal policies seek to address; the return of the victim, which came to occupy the centre of the criminal justice systems, as a collective and symbolic, rather than individual character, and whose interests came to serve as justification for the punitive measures[14]; the protection of the public, which came to dominate the discussions of penal policy; the politicisation of the new populism, which led to a general consensus, with no polarisation, in all political parties, around criminal law and on the need for tight crime control measures, and which values public opinion, more than expert evidence; the reinvention of the

[14]Garland cites as an interesting example the fact that many laws passed nowadays are named after the victims.

prison, not as a rehabilitative place, but rather as an effective place of retributive punishment; the transformation of criminological thought from a perspective of justice that valued corrective and welfarist measures, to new control theories that situate the offender as anti-social and self-centred, needing discipline and effective controls that inhibit them; the expanding infrastructure of crime prevention and community safety, which promotes the involvement of the community, including self-policing, and the dissemination of ideas and practices related to crime prevention; the commercialisation of crime control, which involves the civil society in crime reduction activities, and transfers the values of the private sector to public sector agencies (e.g. prisons), often in favour of the dominant commercial interests; new management styles and working practices, determined by performance measures, which are supported by an ideal of cost-effectiveness; and a perpetual sense of crisis, which raises doubts regarding the coherence and effectiveness of current crime control measures, and consequently to the discredit of the sector professionals and the loss of faith in the criminal justice system.

The rationale for these new crime control strategies do not stem, in Garland's terms, from their ability to resolve the problems that they are supposed to address, but rather from the fact that they 'characterize problems and identify solutions' (Garland, 2001: 26) that are consistent with the dominant culture. Undeniably, much of this, and particularly the preference for retributive, instead of rehabilitative measures, reflects upon current punitive approaches to plagiarism.

My primary interest is understanding the divergence between punitive approaches to plagiarism, which are especially marked in Anglian countries, as opposed to the rejection of punishment in cases of plagiarism. As discussed in section 1.3.2, countries within the Anglian tradition have traditionally focused on designing academic integrity policies and detailed definitions that explain what plagiarism is, how to avoid it and what the consequences are if those rules are violated. More recently, as a result of findings that these policies were not sufficiently clear, and that students were faced with thousands of documents at the start of their degree programme (Carroll and Appleton, 2001), the focus shifted to new methods. Specifically, they found that redesigning student assessments and creating individualised tasks, based on the students' learning outcomes, represented an important contribution. In addition to this, students are commonly invited to induction programmes that provide them guidance on how to write academically and avoid plagiarism, even if they may miss these sessions.

Taken together, these efforts to inform and, more importantly, to train the students on avoiding plagiarism legitimise a set of rules that the students are expected to observe.

1.4 Defining Plagiarism

Plagiarism is intentionally or knowingly reusing someone else's words, works or ideas without a proper, clear and unambiguous acknowledgement. In the academy, the most common form of plagiarism is when a student copies a text written by someone else, and uses it as his/her own work, intentionally or knowingly. A student commits plagiarism if s/he, reusing someone else's text: (a) omits the citation (i.e. reference) to the real sources; (b) cites his/her sources, but not sufficiently clearly to prevent the reader from immediately understanding if the text belongs to the student or to the source; or (c) cites his/her sources, but that citation is ambiguous, resulting in the reader understanding the text as belonging to the student and to the person whose work is acknowledged at the same time. The most serious form of plagiarism is when it is intentional. In this case, the student acts with the aim of deliberately deceiving the readers, by making them believe that s/he is the original author of the text. However, it is also plagiarism if the student does not act with the aim of passing the text off as his/her own, knowing that s/he can deceive the readers, and does not reformulate the text to prevent this act of deception.

It is also plagiarism if a student, despite not using the exact same words of another source, copies someone else's work or ideas. One example of this type of plagiarism is translation, which happens when a student copies a text from a source in another language, translates it, or has it translated to a second language, and then reuses it in his/her assignments without acknowledging the sources. This is not a common case of plagiarism of words, because the words that are used are not identical to the original ones; but it is plagiarism of ideas and/or work. Students are allowed to use translations of texts from other languages in their assignments, but they should always indicate this clearly and give due credit to the sources. Another example where the work and/or ideas are retained, although the words are different, is paraphrasing, which consists of expressing the same ideas using different words and text structures. Plagiarism of works and ideas also applies, but not exclusively, to cases where the student reuses (without acknowledging it) tables, drawings, designs, processes and images, as well as to situations where the student makes changes to the original text, for example by

reordering the words or replacing some of the words with equivalent meaning (such as synonyms).

In academic contexts, plagiarism is a type of academic dishonesty and fraud. An academic fraudulent activity that is identical to plagiarism is collusion, i.e. when two or more students work collaboratively on the same assignment and subsequently submit it as their own individual work. Other instances of academic dishonesty include: cheating in exams; fabricating data in research projects, including inventing data that does not exist; ghost-writing, whereby the student signs and submits a work that s/he has not done (includes buying essays from other people, for example from essay banks); self-plagiarism, whereby an author reuses his/her own previous work as new, without acknowledging it.

In research contexts, both honorary authorship and 'salami slicing' are also considered dishonest. The first consists of including someone's name as one of the authors of a paper when s/he has not contributed to the work (either with the aim of honouring a more senior researcher, or of using someone's reputation to increase the popularity of the research). The latter consists of producing multiple research articles, by fragmenting ('slicing') the same set of data.

1.5 Chapter Summary

Although there has been a considerable amount of previous research into plagiarism and plagiarism detection, some aspects of plagiarism detection have attracted little research attention. The extensive research conducted in the field of linguistics, in general, and forensic linguistics in particular, contributed to the identification of a set of linguistic devices that can be used successfully, to the detriment of others, to analyse plagiarism instances across the same language. Measures such as those of lexical overlap and *hapax legomena* proposed by Johnson (1997) were later thoroughly tested and proven by Turell (2004, 2007, 2008), including in languages other than English, and Coulthard (2004), and used successfully by Woolls and Coulthard (1998) and Woolls (2003) to develop plagiarism detection software.

However, the study of translation as a plagiarism strategy, whereby someone translates a foreign language text and uses it as their own, although a problem in academic, as well as non-academic contexts – as Kaplan and Torbati (2007) demonstrated – remains under-researched, as will be discussed in chapter 5. The increasing mobility of students across the world not only implies that students have to write academically

in a different cultural context, often following distinct academic writing rules, but also that there is a strong likelihood that, if they lift texts from other languages and translate them into the target language, they will pass undetected. In non-academic contexts, translation as a plagiarism strategy has implications that span beyond the traditional tenets of academic plagiarism, to the realm of copyright. This is a complex problem, raising particular issues, especially (a) whether it is possible to detect plagiarism across different languages, considering that different 'linguistic signs' are formulated, thus hampering any comparison between the two; and (b) whether, theoretically, this is even considered plagiarism, owing to the fact that it is the ideas or the works that are reused without acknowledgement, and not the *exact* words.

In parallel, computational linguistics research into plagiarism has attracted the attention of computational linguists from all over the world, and consequently made a considerable contribution to plagiarism detection. One of the problems, however, as Woolls (2010) explained, is that computational linguistic analyses are too often conducted by people with a deep knowledge of computer science, but not of linguistics. As a consequence, the detection procedure is usually based on the development of particular algorithms that revolve around detecting shorter or longer strings of identical or near-identical text, disregarding the fact that, as Coulthard *et al.* (2010) argue, most intentional plagiarism use strategies other than this one.

One of the key tenets of this study is that plagiarism is a moral and an ethical issue. On the moral side, plagiarism brings social implications, with the power to ruin the reputation of the plagiarist; on the legal side, it implies the infringement of moral rights, and often financial rights, both of which are punishable by law. Even in the academy, where cases of student plagiarism in general do not involve financial rights, the punishment measures involved can be severe. Under these circumstances, an approach to plagiarism detection informed by forensic linguistics, and by its ethical implications, and that is able to investigate and provide thoroughly explained and theoretically-grounded evidence is preferred, to the detriment of other approaches.

However, as the literature on pedagogical approaches to plagiarism demonstrates, instances of academic plagiarism cannot be universally evaluated. On the one hand, definitions of plagiarism tend to vary across different contexts, and the assumption that the same definition is understandable by everyone has long been discarded; on the contrary, the belief that student plagiarism can be the result of an attempt to write academically (Howard, 1995), or even a form of authorship (Love, 2002), need to be

considered. In parallel with this multiplicity of definitions, and considering that plagiarism reflects a problem of authorship, two additional problems need to be entered into the equation; firstly, the concept of plagiarism has varied significantly over time, and even in recent decades philosophical approaches to authorship have demonstrated that each writer can play distinct, though not always mutually-exclusive, roles as authors; secondly, although plagiarism reflects a problem of authorship – or of its due acknowledgement – the expectation in the academy is that students are capable of integrating in their own writing the multiple voices of their sources. These factors contribute, on the one hand, to questioning the type(s) of authorship that is(are) expected from student writers, while rejecting the possibility of single-voiced writings, where only one writing style is to be found. Admitting that fraud is difficult to determine linguistically (Eggington, 2008), as is intention, and that this may raise ethical issues to the linguistics expert (Finegan, 1993, 2009), a taxonomy of degrees of intention, graded by severity, will be proposed in chapter 4.

This chapter introduced and presented the rationale for this research on plagiarism detection. The most relevant literature on the topic was then reviewed, establishing a relationship between the different disciplines that a study of this nature involves, and explaining why a combination of those disciplines is relevant. The next point to be discussed is the overall methodology adopted to conduct this study, and which will be used in the following chapters of this research.

Chapter 2

Methodology

2.1 Introduction

Chapter 1 reviewed the literature on plagiarism and approaches to plagiarism detection, and presented the rationale for this research. Since plagiarism, as discussed, has social, as well as moral and legal implications, it is a case for forensic linguistics. But despite the breadth and dept of the research conducted in recent years, there are some gaps that still need to be bridged. Consequently, it was argued that a suitable approach to plagiarism detection in the field of Forensic Linguistics should consider, not only the linguistic analysis, but also the social and legal context in which instances of plagiarism have been assessed. It is therefore hypothesised that a research of this type has implications that span beyond the boundaries of linguistics and computer sciences.

As discussed in the previous chapter, plagiarism has traditionally been approached on the assumption that plagiarism consists of textual duplication. However, this represents an oversimplification that does not allow plagiarism detection to be properly addressed. Among others, the following issues in particular are raised.

Firstly, it is questioned whether plagiarism is understood universally in the same terms, i.e. whether different groups of people have identical or otherwise distinct perceptions of what plagiarism is, and of how it should be handled. This question is addressed in chapter 3, which presents the results of a survey conducted with students and lecturers/tutors, in Portugal and in the UK, to determine whether plagiarism is understood identically across different countries or, on the contrary, whether different people tend to have different concepts of plagiarism. This allows for a comparison of the perceptions of plagiarism within groups, as well as between groups of participants.

Secondly, a debate is raised as to whether realistically different degrees of severity are attached to different instances of plagiarism, and whether these degrees of severity are dependent on the assessment of the degree of intention of the plagiarist, or whether they are intention-agnostic. This issue has been discussed in the literature for several decades, and assumptions have been made based on this discussion, but the extent to which group perceptions reflect or otherwise compete against these assumptions has not been given much, if any, research attention. This issue is approached in chapter 3.

The perceptions of the survey participants are analysed quantitatively, as well as qualitatively, in order to assess whether intention is relevant or otherwise to determining the severity of the instances of plagiarism, and the corresponding penalties. If the degree of intention is to be considered in the judgement of those instances, it is then relevant to ask what is the role of Forensic Linguistics on this matter. If the role of the forensic linguist is limited to detecting and identifying the strings that are lifted from other texts, while leaving the assessment of the intention to a decision-making body, such as a disciplinary board, then a linguistic comparison between the suspect text and the alleged original suffices. If, however, the role of the forensic linguist is taken to include all his/her expertise, as well as the forensic linguist's critical input so as to allow decision-making bodies to make more informed decisions, then contributing all information available is crucial, including clues that may indicate the intention to deceive. This issue is discussed in chapter 4, in relation to the findings reported in chapter 3. In particular, the discussion of intention is approached from a moral and legal perspective, and a replication of legal theories of intention is proposed for the linguistic analysis of instances of plagiarism. Moreover, a link is established between different linguistic strategies and the degrees of intention behind those strategies.

In parallel with this, most research has been conducted using English language data, to the detriment of other languages (with the exception of Spanish, owing to the contribution of Turell (2004, 2007, 2008)). Some of these findings, of which those of Johnson (1997) are a good example, can be replicated with success in other languages. Others, however, are language specific and need further research attention. Perhaps owing to this focus on English texts, some areas are under-researched, such as the use of translation to plagiarise from texts written in a foreign language. If textual plagiarism is the focus of the forensic linguist's task, then research on detection of plagiarism involving translation is unnecessary; if, on the contrary, it is believed that

translating a text from another language without acknowledgement and passing the text off as our own, despite using different linguistic signs, is also plagiarism – then further research in the field is necessary, particularly in languages other than English. This issue is addressed in chapter 5, which discusses several instances of texts that have been plagiarised from other languages. A linguistic analysis of the data is conducted and presented to investigate whether textual reuse from foreign language texts can be detected empirically, and whether the empirical findings can be explained via current translation and linguistic theories.

Considering the increasing flow of information, one can hypothesise that the detection procedure cannot run independently of both a linguistic analysis and a computational approach. However, the interaction between linguistics and computer sciences for purposes of plagiarism detection has attracted little research attention, although plagiarism detection has attracted the interest of both disciplines individually. With the exception of the work reported by Woolls and Coulthard (1998), which included linguistic input in the programming, linguists have traditionally focused on the manual linguistic analysis of instances of plagiarism, whereas computer scientists have focused mostly on the tasks of finding duplicates or near-duplicates. One can therefore hypothesise that there may be room for improvement in plagiarism detection software. This issue is addressed in chapters 6 and 7. First, theoretical approaches to authorship and computational plagiarism detection are presented to investigate whether detection software and procedures have been empirically– or theoretically–grounded, or both. This issue is addressed via the exploration of research literature on the topic. An evaluation of existing software detection tools is subsequently presented, in order to identify the functionalities currently offered by plagiarism detection systems. A description of the most popular and common plagiarism detection tools is presented in chapter 6. These tools are assessed, based on their performance using examples from the CorRUPT corpus, my own corpus of reused and plagiarised texts, or, when access to these tools is restricted, on the review of the literature. Based on the description of the software features, and, where possible, on the system performance using real data, gaps are identified and explanations are provided for those performance issues. One can hypothesise that most of the existing gaps of plagiarism detection software can be addressed by already existing tools. The main performance gaps are identified to suggest detection system improvements, in chapter 7. Where applicable, examples the CorRUPT corpus are used to illustrate the relevance of those tools.

This research therefore adopts a forensic linguistic approach to plagiarism detection that aims to identify the linguistic patterns and strategies used by forensic linguists to plagiarise, and ultimately contribute to improving plagiarism detection methods and procedures in general, and computational approaches to plagiarism detection in particular.

These research questions and hypotheses are based on mainly two types of data: (a) the survey data; and (b) the textual data. The survey data provides information on the perception that students and lecturers/tutors have of plagiarism; the textual data includes instances of plagiarism that allow for an identification of the linguistic strategies used to plagiarise. The sections that follow include a description of these data, as well the research design, including the data collection procedures and the groups of participants involved in each set.

Since the involvement of human participants is a sensitive topic, the ethical considerations adopted are also discussed. The analytical framework is subsequently briefly presented. As each of the topics listed previously will be discussed in each chapter individually (except for computational plagiarism detection, which is divided into two chapters), a more detailed description of the methodology adopted to address each of the research questions is presented in each of the corresponding chapters.

The next section presents the nature of the two sets of data collected, as well as the data collection procedures in each case.

2.2 Nature of the Data and Data Collection

This research is based on the analysis of two types of data: a survey that aims to understand the perceptions that students and lecturers/tutors have of plagiarism, and texts containing instances of plagiarism.

2.2.1 Survey Data

The survey data, which is described in more detail in chapter 3, consists of the responses of students and lecturers/tutors from two universities in two different countries (Aston University, UK, and University of Porto, Portugal) to a survey on perceptions of plagiarism. The survey, in which students and lecturers/tutors from all areas of study in those two universities were invited to participate, consisted of reading a set of 12 fictional scenarios (vignettes) describing cases of plagiarism, and say whether they

agreed, or disagreed with the claims that the student in each scenario had plagiarised. For each vignette, the respondents were invited to choose one of seven options available: '1' being 'Strongly disagree' with the claims of plagiarism and '7' being 'Strongly agree' with those claims.

Prior to the final survey, a pilot was conducted. This pilot aimed, on the one hand, to test the survey, so as to ensure that everything was in order for the final survey. Although this pilot was conducted in the two countries, it was advertised only in institutions other than those expected for the final survey. In Portugal, all participants in the pilot survey were from the University of Minho; in the UK, participants were from several universities (including the University of Birmingham and the University of Leeds, as well as from the Buckinghamshire New University). Aston University did not participate in this phase. At this stage, some technical problems were resolved (for example, some optional fields were available initially only as 'compulsory'). Secondly, although the number of vignettes aimed for the final survey was 12, initially 36 were made available. This allowed the most discriminatory scenarios to be selected, following the procedure described in chapter 3. The results of the final survey were submitted to a quantitative statistical analysis, a 2 x 2 ANOVA (analysis of variance), so as to determine whether significant differences existed between (a) participant roles (students and lecturers/tutors), (b) two institutions (Aston and Porto), and (c) if there was any interaction, i.e. whereby each group of participants (students and lecturers/tutors) behaved differently in the other institute (Aston and Porto).

A more detailed description of the method, as well as the presentation of the analysis and the discussion of the results are provided in chapter 3.

2.2.2 Textual Data

A set of textual data is used in this research. Considering that many cases of plagiarism still pass unnoticed, especially in countries where systematic screening for plagiarism is uncommon (as is the case of Portugal), the volume of plagiarism data available are scarce. Therefore, this set of data includes texts of two different genres: academic and non-academic texts.

Non-academic texts include two cases of news plagiarism in Portugal that are public, and therefore do not have access restrictions. The first is the case of a journalist who was accused by the Portuguese quality newspaper *Público* of having plagiarised mainly from the *New Scientist* and *Wikipedia*. The issue was first raised by a reader of

the newspaper, who upon reading the text intuitively realised that he had read some of the information somewhere else. Upon further investigation, this reader realised that at least part of the text derived from *Wikipedia*. The newspaper investigated the case further and concluded that the journalist had indeed plagiarised from other sources written in English. Initially, she denied the accusations but eventually admitted the lifting and was invited to submit a public apology. The second case is that of another journalist who had plagiarised a news piece on the Oscar Academy Awards, more specifically on Ellen DeGeneres. Like the other case, suspicion on this one was raised by another reader. However, unlike her colleague, the journalist immediately admitted the lifting from a Reuters news wire, submitted the public apology and the case was dismissed.

Academic texts include six cases of academic plagiarism that were considered by the lecturers/tutors to have been plagiarised. Although the number of assignments is small, they are of considerable length, and involve different strategies. This allowed an analysis of cases that go beyond cases of literal, verbatim textual reuse.

The full analysis of the textual data is not provided as a whole in the main body of this thesis for two reasons. Firstly, considering the nature of this study, providing individual examples to illustrate the findings would be more appropriate than providing the full analysis. The strategies used by the plagiarists are often repeated, so using individual examples avoids redundancy. On the other hand, since the academic data are sensitive, including only specific, anonymised examples guarantees an additional preservation of the confidentiality of the data, and consequently of the participants' identity. For reasons of consistency, the same method is applied to the non-academic data.

The next section 2.3 describes the participants involved in both cases.

2.3 Participants

Plagiarism in general and academic plagiarism in particular is a sensitive topic, with special ethical implications, involving special risks, both for students and their lecturers/tutors or supervisors, not to mention the negative institutional publicity that it attracts. This research involved these three types of participants.

2.3.1 Student Participants

The student group is the main, obvious group of participants for which ethical concerns can be raised. The main risk for students is that as participants, they may be identified, notwithstanding all efforts to anonymise their writing. A thorough investigation of the texts may disclose the author's identity and, if incriminating behaviour is revealed, lead to punishment or disciplinary action, with serious academic and professional implications.

Therefore, this research is based on the analysis of:

1. The responses to the survey;
2. The essays and the dissertations written by students.

As plagiarism is a socially sensitive topic, ethical issues may arise, starting with the *survey*. Although this issue should be overcome by using vignettes, as compared to other methods used to survey the students' perceptions (e.g. interviews to identify whether students have plagiarised), the risk is that the participants find they may have acted as described in one or more vignettes. Participants can therefore be exposed to some mental stress, embarrassment, shame, find that their privacy is invaded and/or be faced with the mental dilemma that they might have inadvertently plagiarised.

In order to mitigate this risk of mental stress, participants were informed that the scenarios depicted in the vignettes are fictional. Secondly, the fact that they were able to take the survey online avoids any confrontation or interaction with any human subjects when making their judgements. Thirdly, no personal, identifying data collected from the participants was or will be disclosed (especially the personal data that they have optionally provided to enter into the prize draw). Additionally, the identifying data collected will not be stored after the research project is finished.

Prior to taking the survey, participants were asked to confirm that they agreed with a consent form, which provided them with information on:

- The purpose of the project;
- Why they have been asked to take part;
- What they will be asked to do;
- Possible problems; and
- The confidentiality of information.

As the survey was conducted online, accepting the consent form by clicking the 'Continue' button to proceed was explicitly agreed to correspond to signing the form.

As far as the corpus of data of instances of plagiarism collected are concerned, the main risk is that the participants may find themselves in a compromising situation, which might lead under extreme situations to being accused of misconduct and improper behaviour. It was therefore crucial to ensure their right to privacy, in accordance with the legal procedures in force.

Given the nature of this research, it was impracticable to ask students to sign consent forms to use their dissertations in this research. Although this might seem to constitute a breach of ethical principles, there are several reasons for this. First, signing such forms is not required, neither by the Portuguese law, nor by the University regulations. Secondly, the cases were provided by the plagiarists' lecturers/tutors after they have been considered to have plagiarised, and consequently had already been agreed or discussed with the students accused of plagiarism.

However, to ensure their privacy as much as possible, all identifying elements were anonymised. The participants' names were replaced by a combination of letters and numbers (the key to which will be accessible to the researcher and his supervisors only); and other identifying elements were deleted (e.g. address, e-mail, reference to funding agencies). Moreover, all research materials, both electronic and hard copy, were stored securely at all times, and no participant was identified in any way in all outputs of this research (e.g. reports, presentations, papers, articles, etc). All research materials relating to the project that contain identifying information will be destroyed (e.g. papers shredded, electronic files deleted) following completion of the project, plus one year. Since the main focus of this research consists of a textual analysis, all strings of text used in the findings will also be anonymised. This aims mainly to mitigate the risk that the general public might find relevant identifying elements potentially leading to professional penalties, and not so much that the lecturers/tutors and their institutions might use the findings of the research to screen the authors – the latter not being a real risk, since the lecturers/tutors and their institutions may use the methodology described to screen their student essays for plagiarism.

2.3.2 Lecturer/tutor Participants

Since the number of disciplines involved is finite, guaranteeing the privacy of the students also implies guaranteeing the privacy of their lecturers/tutors and departments,

as well as avoiding a direct relationship between the students, the topic, and the lecturers/tutors. If this relationship can be discovered in any way, the students' identity might be revealed in their writing, and result in risks also to their lecturers/tutors' careers, who might have not in the first instance detected the plagiarism – or failed to report it. Additionally, the participating departments are not named, and the disciplines are only identified broadly, in general terms.

As far as the survey is concerned, the approach adopted with teaching staff is the same as the one adopted with students. However, unlike students, lecturers/tutors may be exposed to additional stress if they find that they are responsible (wholly or in part) for their students – or even themselves – acting as described in the situations depicted in the vignettes. To mitigate this risk, the survey was conducted online, thus avoiding any confrontation or interaction with any human subjects when making their judgements. Besides, lecturers/tutors, like students, are informed that the scenarios depicted in the vignettes are fictional. As proposed for students, no personal, identifying data was collected when conducting the survey. Before completing the survey, lecturers/tutors were invited to enter their name and e-mail address, with the only intent that no participant takes the survey twice, as well as to enter the prize draw. Additionally, no identifying data will be stored, including the IP addresses of the computers used to respond to the survey.

Prior to taking the survey, lecturer/tutor participants were also asked to accept a consent form, which provides them with information on:

- The purpose of the project;
- Why they have been asked to take part;
- What they will be asked to do;
- Possible problems; and
- The confidentiality of information.

As is the case of student participants, accepting the online consent form by clicking the 'Continue' button explicitly corresponds to signing a paper version of the form.

2.3.3 Institutional Participants

Prior to any data collection, the permission of the University of Porto was obtained to collect academic assignments written by students of the University. This includes a

letter from the Rector/Vice-Chancellor and from the Faculty/School Directors and/or Department Directors and/or lecturers/tutors. However, this permission is not necessary if the lecturers/tutors individually grant access to the student assignments. This is the case of some of the assignments, whose access was granted by lecturers/tutors of another Portuguese University – the University of Minho – and no other permission was sought.

The main risk for the universities is that the discussion of these cases of plagiarism among the university students may expose the Institution to negative publicity, e.g. in the media. Therefore, in order to mitigate this risk, confidentiality has been granted. All identifying elements were anonymised; identifying information regarding human participants and departments involved in possible instances of plagiarism will not be publicly revealed. The institutions will be given the choice to decide whether acknowledgement is made of its contribution in the research outputs. Besides, permission obtained from lecturers/tutors for the data collected can only be used for the purposes for which it was obtained; i.e., this research project.

2.3.4 Ethical Considerations

The research design and the data collection method, storage and analysis was approved by the Ethics Committee of the School of Languages and Social Sciences at Aston University.

From an ethical perspective, the present research does not raise any issues related to the researcher's physical integrity. However, although the measures stated above may ensure that ethical principles are in place to guarantee a proper conduct, some moral issues may arise. In particular, the research may impose moral dilemmas as to whether or not to report instances of plagiarism to the university authorities.

This research involves investigation of deception among university students. Whether plagiarism is considered to be 'only' unethical or also illegal, I find that as a researcher I have a moral duty to share any conclusions regarding the instances of misconduct or improper behaviour to the appropriate authorities. Additionally, this moral duty also applies to authors whose works may have been lifted, and to other readers who are deceived. This requirement often supersedes any confidentiality or anonymity agreements made, explicitly or implicitly, with the participants.

In any case, as this is a novel research in the Portuguese context, adequate measures regarding confidentiality will be undertaken not to breach the students' and lec-

turers/tutors' anonymity.

As far as value judgements are concerned, as has been claimed, 'control over what people do obviously has a moral dimension' (Robson, 2002: 66). As researcher, it is my aim to seek an understanding of plagiarism and report it correctly and reliably. Therefore, the range of vignettes provided aimed to cover a wide spectrum of scenarios and instances of plagiarism, in order to avoid any unconscious imposition of the researcher's own values and expectations. Moreover, the survey was conducted in English and in Portuguese, so as to minimise any impact from lack of linguistic skills of the respondents.

On balance, this research admittedly involves some amount of risk, both for the participants (individual and institutional) and for the researcher (namely, in terms of moral dilemmas). However, considering the nature of the research, which aims to investigate plagiarism by considering the involvement of ethical issues bordering the illegal, measures were in place throughout the research project to identify and mitigate the risks as far as possible. Namely, knowledge and/or consent of the participants was obtained (where relevant), information about the true nature of the research was disclosed, and measures were deployed to minimise the exposure of individual participants to stress. Finally, measures were adopted to guarantee that the participants were treated fairly, with due respect and consideration, without violation of their rights. Overall, the potential gains obtained from this research therefore outweigh the expected disadvantages.

2.4 Analytical Framework

Considering the two types of data collected in this study, two different analytical approaches are used.

2.4.1 Survey Data

The considerable amount of survey data collected is analysed statistically to investigate whether any differences are reported between the two participant roles involved (students and lecturers/tutors), from the two institutions (Aston and Porto), whether any interaction occurred, whereby each 'role' behaved differently in each condition of the variable 'institute', and whether those differences are statistically significant. The cases where a statistically significant difference – or an interaction – was found are

then analysed in more detail. This includes a qualitative analysis of the feedback provided by the participants, in order to find possible explanations for their scores.

2.4.2 Textual Data

The textual data collected, which included the instances of plagiarism and the corresponding originals, are part of the CorRUPT corpus, my own Corpus of Reused and Plagiarised Texts. The data were manually compared in order to identify, via descriptive linguistics, the strategies used by the plagiarists to borrow, without acknowledgement, from their corresponding originals. The cases of plagiarism that were found were submitted to a lexico-grammatical analysis, in order to identify the instances of plagiarism, explain the linguistic operations performed by the plagiarists, and justify why those cases represent instances of plagiarism. Although not all possible linguistic strategies are present in the CorRUPT corpus, since other linguistic strategies can be potentially used to plagiarise, only the examples found and explained were used to illustrate the points raised and discussed.

Therefore, although this study does not include a specific chapter where all the cases are analysed systematically, the linguistic analysis permeates all the chapters where linguistic data are analysed, in particular chapters 4, 5, 6 and 7.

2.5 Chapter Summary

This chapter started by presenting the nature of the data collected. Firstly, a description of the *survey data* was provided, including the nature of the data and the collection procedure followed. An explanation was provided for the rationale behind the survey, and the procedure adopted to conduct the pilot survey, select the final vignettes and conduct the final survey was described. The nature of the real *textual data* and the procedures adopted to collect it were then described. A justification was provided for the amount of data used, and the data collection procedure was detailed. An explanation was provided for the two types of data used (academic and non-academic), together with the collection method adopted in an attempt to avoid stressful situations among the participants.

The following section described the participants and both types of data used, and explained that all the survey participants contributed voluntarily to the research. Both groups of participants (students and lecturers/tutors) from the two countries (Portugal

and the UK) were informed of the design, aims and objectives of the survey. They were also aware that they accepted the terms advertised, by clicking the 'Continue' button; and they were informed of the procedure adopted to respond to the survey questions.

A description of the contributors (participants) to the textual data was provided. It was clear that the non-academic data was collected for analysis without obtaining any particular permission, considering that the data used are publicly available; on the contrary, access to the academic data was granted by the lecturers/tutors responsible for the corresponding modules. When access was granted to the data, the outcome of the cases had already been agreed or discussed with the students involved.

The third group of participants – Institutions – was then described. Although these institutions did not participate directly, they could be indirectly involved in the cases of plagiarism, e.g. due to the negative publicity that the plagiarism cases might attract. A description was provided of the activities implemented to avoid such publicity and any possible negative impact on the institutions.

The measures adopted to observe ethical issues were explained, and the possible ethical implications were discussed. This was followed by a description of how these issues were addressed. The potential participants were explained that these measures followed the ethics policy submitted to and approved by the School of Languages and Social Sciences at Aston University.

Finally a brief description was presented of the analytical framework adopted. The participants were explained that, on the one hand, a statistical analysis of the survey data was conducted, which was followed by a qualitative analysis of the optional feedback provided by the participants, whereas the textual data, on the other hand, were analysed qualitatively, using methods of descriptive linguistics.

The next chapter 3 presents a detailed description of the survey conducted, as well as the results obtained, and discusses the analyses of these results, suggesting possible interpretations for these. The linguistic analysis will be used in the remaining chapters, in particular to illustrate the points presented.

Chapter 3

Understanding Plagiarism: A Survey to Measure Attitudes towards Plagiarism

3.1 Introduction

The discussion of plagiarism presented in chapter 1 demonstrated that part of the problem in addressing plagiarism in general, and plagiarism detection in particular, is due to the lack of a consensual, universal agreement on what plagiarism is. Given the multiplicity of definitions available and their respective degree of detail, not only across different cultures, but also across different countries, regions and even disciplines and textual genres, it is very unlikely that different people, at different times and in different contexts have an identical, even similar concept when they refer to plagiarism. In the academy alone, the lack of a consensual definition of plagiarism, equally understandable and understood by everyone raises problems of several kinds, the answers to which may seem obvious to professional academics, but not necessarily so to students. The first main problem for students is understanding what is expected of them when they think and write academically, and under which circumstances can and should they refer to other authors. Which and why certain authors' details are required, but others are not? What types of materials are subject to acknowledgement? Should works and ideas also be acknowledged, or just the words? If works and ideas also require acknowledgement, how should they be marked in the text? Additionally, when acknowledging words, do all exact words need to be inserted in between quote marks? And do they have to be always followed by the author's name, date of publication and page numbers – even when this makes the writing awkward? When (re)using text from other languages, should the text be quoted in the original language, or translated

freely? If the latter, should quote marks be used?

The list of doubts is endless, and followed by the fact that academic writing standards require an extreme precision and level of detail which even skilled students may have difficulty in grasping. If we take reporting verbs as an example, understanding why different verbs (e.g. *state*, *utter* and *declare*) mean and imply different things in different contexts may be problematic. If the basic purpose of referencing is to make a clear distinction in one's own text between the works, words and ideas of others, and the works, words and ideas that one has contributed, why is it relevant that one or other referencing style in particular be used? Why does one institution favour, even demand, one style to the detriment of another? A just and fair assessment of instances of plagiarism that considers the particular circumstances behind it not only is desirable, but also crucial, especially considering that very often plagiarism involves a range of serious academic and even legal implications. In order to verify this hypothesis of diverse, even conflicting, conceptions of plagiarism, the perceptions of lecturers/tutors and students were collected to investigate how they define plagiarism and how they see these and other questions.

This chapter investigates whether participants of different groups (e.g. lecturers/tutors and students) from different cultures and countries have identical or similar perceptions of plagiarism or whether they tend to see and judge instances of plagiarism differently. In order to assess these perceptions, a survey was conducted involving two groups of participants (lecturers/tutors and students) from two different countries (UK and Portugal).

3.2 Operationalising the Research

3.2.1 Collection Instruments and Procedure

Since this primary research involves human participants, ethical permission was sought and obtained from the Ethics Committee of the School of Languages and Social Sciences at Aston University. To collect the perceptions of students and lecturers/tutors, a survey was conducted using attitude scales. The respondents were provided with a set of vignettes, i.e. story like, fictional, 'short narratives' (Dõrnyei, 2007: 255) or scenarios that describe real life, typical or representative events or experiences. Each respondent was asked to grade each scenario in a continuum, using a summated rating (Likert) scale (Robson, 2002), from '1' to '7' – '1' being 'Strongly disagree' with the

claims of plagiarism and '7' being 'Strongly agree' with those claims. According to this scale, '4' is the neutral point, 'Neither disagree, nor agree'. The use of vignettes was chosen for their proven usefulness in research on sensitive topics, where participants may feel threatened or uncomfortable about giving their perceptions, attitudes, beliefs and judgements (Hughes, 1998). As they depict real-life situations of other people's lives, these fictional narratives allow the participants to assess the scenarios in a more 'distanced', less threatening way.

3.2.2 Vignettes Design

These vignettes (initially 36) were written based on real cases of plagiarism in different settings, collected mainly from personal communication with lecturers, tutors, researchers and students. All these vignettes asked the same question: whether the main character (typically a student or lecturer/tutor), upon having adopted a particular action, should be charged with plagiarism, and consequently be subject to the corresponding sanctions. The following is an example of these vignettes (Vignette 8):

> Maria has just submitted her Master's dissertation and was surprised when her school office called her to visit her supervisor. When Maria asked why, she was told her supervisor had a very serious issue to discuss with her. When she met some friends later that day, she told them about the call, but said she was intrigued about the meeting, as was confident that her dissertation on computer engineering was excellent. She had read thoroughly on the topic and built upon the theories of four different famous theoreticians to then justify her approach, which she thinks is innovative. When her supervisor discussed it with her, he claimed she had plagiarised. In an attempt to disprove the accusations of plagiarism she is now facing, Maria, who was well aware of her work, explained that she had always intended to acknowledge the other theories, and claimed that such an acknowledgment was in fact vital to reinforce her argument. Besides, she was well aware that she needed to provide authors' names, the year of publication and page numbers whenever she cited anyone – besides including it in the reference list. She told her supervisor she was quite sure that was what she had done. Her supervisor however counter-argued that although the reader might understand which part is her own work and which part is the work of others, and who that work belongs to, the acknowledgment is not clear in the sense that it does not follow the School conventions; i.e. she had not followed the referencing style adopted by the school.

All individual names, module names and specific settings mentioned in the text of the vignettes are fictional. Since two countries with two different languages are

involved, the language used required particular attention. Given that most Portuguese university students speak English as a second language, the survey could have been conducted only in English. However, in order to avoid any possible impact of the lack of language skills on the assessment of the scenarios, these were provided in Portuguese to Portuguese participants and in English to British participants. A version of the texts was produced in Portuguese, rather than a simple translation of the original in English. Awareness of the wording of the scenarios may inadvertently impact the participants' responses, and the same risk holds true for conducting the survey in different languages. However, the risks implied by this procedure are minor when compared to the impact of administering the survey only in English, even to speakers of English as a second language.

The survey was advertised by email and was conducted online, using the Bristol Online Surveys (BOS) service[1]. The BOS service, subscribed by Aston University, is free for Aston University users, and was preferred to other unpaid survey services for its flexibility and ease of use, as well as for its security and guarantee of confidentiality, and above all for the unlimited number of questions and responses permitted, unlike other services. Moreover, BOS also allows the collected data to be exported to Excel format, as well as to SPSS for further analysis.

3.2.3 Access, Selection of Participants, Sampling and Analysis

The survey was conducted with 673 respondents, distributed as follows:

Pilot Survey				Final Survey			
UK Universities		Minho		Aston		Porto	
Lecturer	Student	Lecturer	Student	Lecturer	Student	Lecturer	Student
5	11	15	14	20	192	52	364

Table 3.1: Number of participants in the survey.

The total number of respondents (n=673) exceeds the minimum number initially expected (160 participants overall, of which 140 respondents in the main study and 20 respondents in the pilot study).

The final survey was conducted in one University in the UK (Aston University) and another in Portugal (University of Porto). Prior to this survey, participants of each group (students and lecturers/tutors) in the two countries (UK and Portugal) were invited to

[1] *Bristol Online Surveys* – `http://www.survey.bris.ac.uk/`

respond to a pilot study in order to test it and then allow for the selection of the most discriminating vignettes out of the 36 original ones. Although the text of the vignettes was the same across the four groups, four different URLs were provided, one per group per country so as to collect only data that is relevant for the particular group, thereby avoiding drop-outs due to the requirement to complete irrelevant fields. E-mails advertising the survey were sent to the student and lecturer/tutor e-mail addresses, inviting them to participate. In order to encourage responses, given the estimated relatively long time necessary to take the survey (1.5 to 2 hours in the pilot study, and up to one hour in the main study), each participant was offered the chance to enter a prize draw. Additionally, to avoid any possible contamination among participants, the pilot survey was advertised only with potential participants in universities other than those chosen for the final study.

Once the minimum number of responses for each individual group was collected, the results were analysed to select the 12 most discriminating of these 36 vignettes. Discriminative power (DP) analysis (Robson, 2002) was the method chosen to select these vignettes, because it allows all the respondents whose scores tend to fall towards the middle of the scale (hence being more neutral) to be discarded, retaining only the responses of the highest and the lowest scorers (which are not necessarily the highest and lowest scores for every single vignette) to calculate their discriminative power. As these represent the less unanimous cases, they are naturally more prone to debate and clashing, mutually-exclusive opinions. Additionally, and more importantly, discriminative power analysis is extremely rigorous, enabling the sum of the scores obtained from the perceptions for an analysis on the overall scale while avoiding any biases. It has therefore been described as particularly effective to investigate differences observed between groups of participants, such as perceptions, considering the specific circumstances involved in the cases (Burns and Burns, 2009). Following this method, the averages for lecturer/tutor scores on the one hand, and for student scores on the other, in the UK and in Portugal, were calculated, and one list was prepared per group of participants, with vignettes sorted in descending order. Unsurprisingly, each of the two lists included vignettes that were not among the highest scores of the other list, but five out of the 36 scenarios were common to the two lists. Since it was a requirement of the final study that the vignettes used to measure the attitudes of lecturers/tutors and students be the same, in order to allow for the comparability of the results, a shortlist of 12 vignettes with the most discriminative power (i.e. the

top scoring scenarios) was prepared, which combined the top scorers of each of the two lecturer/tutor and student lists. The lowest scoring shortlisted vignette, and consequently in the final study, was vignette 34, with an index of discriminative power (DP) of 3.33. All other vignettes in the final survey scored higher, up to DP = 5.50.

The final survey was subsequently prepared including these 12 most discriminative scenarios, and advertised with potential participants of the two groups in the two universities expected (Aston University, UK and University of Porto, Portugal).

As in the pilot survey, the aims and objectives of the study, and what was required of them, were explained to the participants, and they were asked to click 'Continue' if they accepted. They were then asked to give their opinion on whether they thought each case described was an instance of plagiarism – and plagiarism only – or otherwise, whether they thought that the person accused should be dismissed. As in the pilot, a Likert scale was used, from 1 – 7, where 1 = 'Strongly disagree' and 7 = 'Strongly agree', and 4 = neutral, meaning 'Neither agree nor disagree'. Unlike the pilot, however, the final survey allowed participants to input their comments, e.g. reasons for their choice, in a free text, 'comments box'. Although providing comments was not compulsory, these comments provided a clearer understanding of the choices made by the respondents using the attitude scale, as is discussed in section 3.3.

The results of the final study were finally analysed considering the responses of all the participants of the two *roles: students (n = 556) and teachers (n = 72)*, from the two different *institutions: Aston (n = 212) and Porto (n = 416)*. This analysis was followed by a 2 by 2 between-groups Univariate Analysis of Variance (ANOVA), using SPSS on the seven mean rating scores ('Strongly disagree', 'Partly disagree', 'Disagree', 'Neither agree nor disagree', 'Agree', 'Partly agree', 'Strongly agree'), in order to test the null hypothesis (H^0) that all the population/population with different roles (students and lecturers), from the two institutions have identical or differing perceptions of plagiarism[2]. Firstly, a statistical analysis of the overall sum of scores was conducted to analyse, considering the complete set of vignettes, the perceptions of each group overall, for each of the variables 'institute' and 'role', as well as for the 'interaction', and whether the statistical differences observed were significant. Subsequently, since the results obtained from the analysis of the overall sum of scores included the sum of each individual vignette, even those that can be statistically non-significant, each vignette was analysed individually in order to identify those vignettes where a statistical

[2]In the description of the vignettes in this analysis, 'institute' refers to the 'institution', and 'teacher' refers to 'lecturer/tutor'. 'Institute' and 'teacher' were used for ease of reference in SPSS.

significant difference was observed, as opposed to those where the difference is not significant. Those cases where the difference was statistically significant were then analysed qualitatively, in order to allow for a discussion of the particular circumstances involved in the scenario. In particular, the comments made by the respondents were screened in order to find plausible explanations for the participants' scores in those particular cases, and better understand those perceptions. The results of the analysis are those described in the following section.

3.3 Perceptions of Plagiarism

The first stage of the analysis of the results was to compare the groups of participants, in order to determine whether they as a group have identical or distinct perceptions of plagiarism. The perceptions of the four groups are presented in figure 3.1, which includes the results of the overall sum of scores of the 12 vignettes discussed, and which is labeled as dependent variable 'total'. The table 3.3 below presents the overall results of the sum of scores of each individual vignette, considering the two main effects ('institute' and 'role'), as well as the interaction ('institute * role'). The statistical analysis of the main effect 'institute' (Aston and Porto) for each individual case is presented in table 3.4, whereas the observed results of the main effect 'role' (students and lecturers/tutors) and of the interaction are presented in tables 3.5 and 3.6, respectively. The cases where a statistical significance has been found are highlighted in **bold typeface**, with asterisks indicating their significance, as shown in table 3.2 below.

Symbol	*p* value
*	$\boldsymbol{p < .05}$
**	$\boldsymbol{p < .01}$
***	$\boldsymbol{p < .001}$

Table 3.2: Meaning of asterisks in the tables.

The analysis of the sum of scores of the 12 vignettes presented in figure 3.1 indicates that a significant difference was found, both in the main effect 'role' ($F_{1,624}$ = 12.95, $p < .001$, partial η^2 = 0.020), and at the level of the main effect 'institute' ($F_{1,624}$ = 11.73, p = .001, partial η^2 = 0.018). Overall, these results suggest that, considering the four groups of participants surveyed – Aston students and Aston lecturers/tutors,

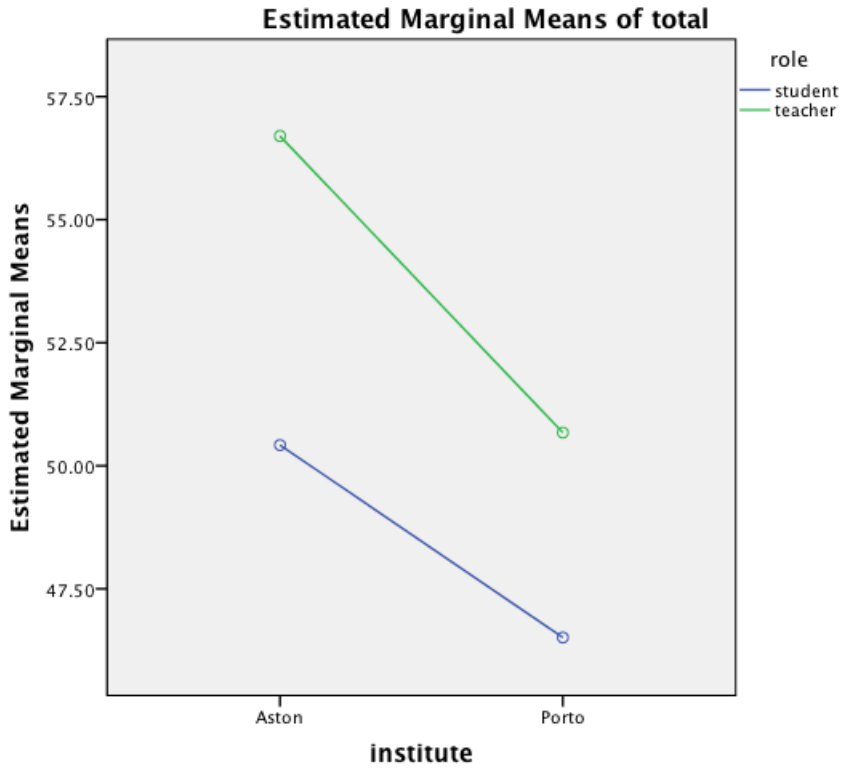

Figure 3.1: Overall results profile plot.

Source	df	F	*p*	Partial η^2
Institute	**1**	**11.73**	**.001****	**0.018**
Role	**1**	**12.95**	**< .001*****	**0.020**
Interaction	1	.53	.466	0.001

Table 3.3: Analysis of the main effects and interaction of DV 'Total'.

on the one hand, and Porto students and Porto lecturers/tutors, on the other – participants from Aston in general scored consistently higher than those participants from Porto. Participants from Porto in general tended to score lower than their Aston colleagues, indicating that they hesitated more than Aston respondents in agreeing with the charges of plagiarism. This is corroborated by the summary of the results of the tests of between-subjects effects, presented in table 3.3, where *statistically significant* differences are indicated in bold. Interestingly, Porto lecturers/tutors scored only very slightly higher than Aston students; the *t*-test shows that the difference between Aston students and Porto lecturers/tutors has an associated value of $t(242) = -.159$, $p = .874$, thus indicating that the difference between the perceptions of Aston students and Porto lecturers/tutors is not statistically significant. However, when considering the

survey results of each vignette individually, a statistically significant difference was observed in only four of the 12 cases in the main effect 'institute' and in only three of the 12 vignettes in the case of the main effect 'role'. This suggests that in all other cases, especially when observing the profile plots for the group of all scenarios, despite the apparent different perceptions observed between the participants, those differences are not significant statistically. Since they do not allow drawing assumptions, they are therefore not worth discussing.

However, where statistically significant differences are found, a few questions need to be addressed, in particular:

a) When those cases where participants from Aston and from Porto, as well as those cases where student and lecturer/tutor participants score differently, do these scores indicate diametrically opposed perceptions or on the contrary, despite being different, are they within the same plane in a dichotomic *agree–do not agree* system?

b) Are there any explanations provided by the participants that justify their choice?

c) If so, do these explanations demonstrate taking the particular circumstances of the scenarios into account, or on the contrary do they tend to judge the subject of the vignette simply as 'guilty' or 'not guilty'?

d) Do the results and the explanations provided suggest that the *degree of intentionality* behind each case influenced the judgement?

e) Do these scorings suggest that only – or mainly – textual material is subject to plagiarism, or is borrowing of works and ideas also perceived as plagiarism?

The following subsections will attempt to answer at least some of these questions.

3.3.1 Analysing the Perceptions of Plagiarism in Aston and Porto

These survey results data demonstrate that the difference between the two institutions is not only considerable, but also *statistically significant* ($F_{1,624}$ = 11.73, p = .001, partial η^2 = 0.018), indicating that, by scoring higher, Aston participants tend to agree more than those participants from Porto with the claims that the situations described represent instances of plagiarism.

Nonetheless, the analysis of each vignette individually demonstrates that the results observed for each scenario are not always statistically significant, as illustrated in the table 3.4.

Vignette	df	F	*p*	Partial η^2
1	**1**	**4.28**	**.039***	**0.007**
2	1	0.29	.590	0.000
3	1	0.19	.660	0.000
4	1	0.13	.715	0.000
5	1	0.10	.747	0.000
6	1	0.97	.326	0.002
7	**1**	**12.51**	**$<$ .001*****	**0.020**
8	**1**	**31.99**	**$<$ .001*****	**0.049**
9	1	2.25	.134	0.004
10	1	0.95	.331	0.002
11	1	2.13	.145	0.003
12	**1**	**6.42**	**.012***	**0.010**

Table 3.4: Analysis of the main effect 'institute'.

A statistically significant difference, indicated in bold in the tables, was observed in only four of the 12 cases in the main effect 'institute', which indicates that in all other cases, despite the apparent different perceptions observed between the participants, those differences are not significant statistically, and as a consequence not worth discussing.

The analysis of *mean values* of each of the four dependent variables where results are statistically significant for the main effect 'institute' indicates that Aston and Porto participants agree that the subject of the vignettes has plagiarised in two cases (vignettes 7 and 12), but do not agree that the student has plagiarised in the other two (vignettes 1 and 8). In these scenarios, Aston participants score consistently higher in all cases, and consider that the subject of the vignettes plagiarised. Porto participants score consistently lower in all the four vignettes, and in the case of the vignettes 1 and 8, these participants even consider the students as not having plagiarised, and therefore *do not agree* with the accusations of plagiarism. In order to find possible reasons for these different scores, the two sets of vignettes are now discussed to search for ex-

planations to account for these disparate perceptions. In the analysis of the vignettes, the following codes will be used for participants:

- *UKS* - UK students
- *UKT* - UK lecturers/tutors
- *PTS* - Portuguese students
- *PTT* - Portuguese lecturers/tutors

Vignette 7 describes the case of a student of architecture abroad, who, having worked hard on an original project, discovered at a late stage that his idea did not work. Under the pressure to submit his final project, he borrowed a few features from two other previous projects, without acknowledgement. However, what made the assessment of the instance of plagiarism difficult was that, despite borrowing a few ideas from other projects, he used them to create an original work, based on a new idea. When asked if he did not consider plagiarism to be a problem, he argued that it 'is only a problem if and when you get caught'.

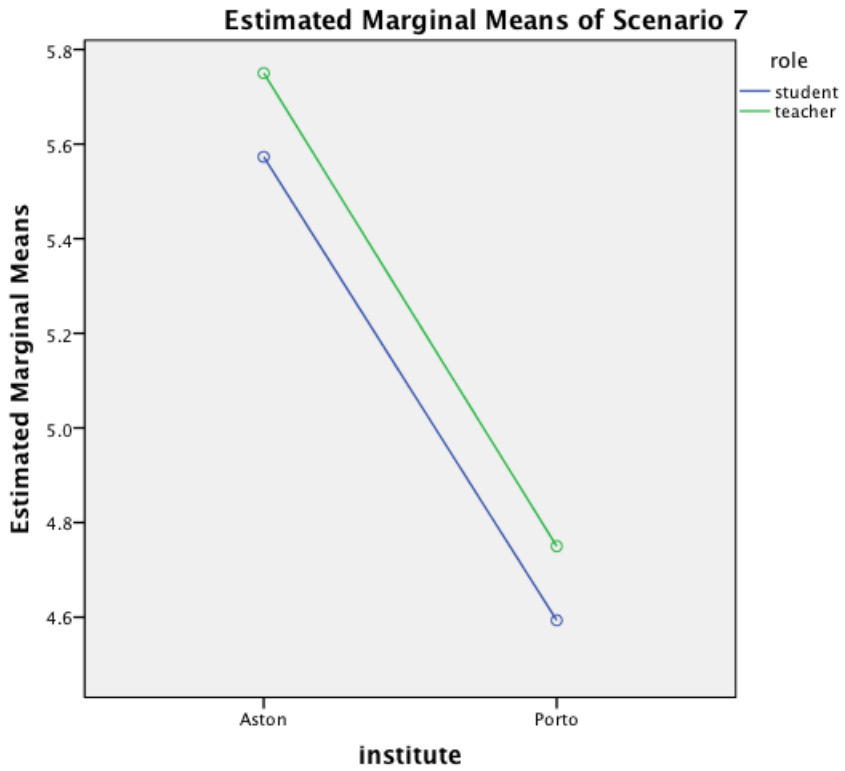

Figure 3.2: Profile plots for vignette 7.

The profile plots for this vignette, shown in figure 3.2, indicate that the difference between the scores in both institutions is significant, with Aston participants scoring

much higher than Porto participants. Additionally, both institutions score much higher, although the scores at Porto fall abruptly when compared to the scores at Aston (as would be expected from the analysis of the mean and 95% Confidence Intervals, shown in figure 3.3).

4. institute * role

Dependent Variable: Scenario 7

institute	role	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Aston	student	5.573	.145	5.287	5.858
	teacher	5.750	.450	4.865	6.635
Porto	student	4.593	.106	4.386	4.801
	teacher	4.750	.279	4.201	5.299

Figure 3.3: Mean and 95% confidence intervals for vignette 7.

The analysis of the difference between the two institutions reveals a statistically significant result ($F_{1,624}$ = 12.51, $p < .001$, partial η^2 = 0.020), indicating that the perceptions of the participants from both institutions are very different. Whereas Porto participants score much lower, apparently less reluctant to accept the fact that the student had created something new, Aston participants are less dismissive, considering that the student undoubtedly plagiarised. One can therefore ask which factors contribute to these different scores.

On the one hand, at a lower level, participants discuss the fact that what is involved is an exercise of academic writing, so that academic writing rules and conventions need to be taken into consideration, and consequently the outcome should be decided with this in mind. It is clear that the sources should be acknowledged, since from an academic writing perspective '[t]here is nothing called borrowing from previous research with no reference' (UKS-117). On the other hand, it is common practice and a known rule of the academy that academic work needs to build upon previous work (Carroll and Appleton, 2001) to be sufficiently grounded, and 'that work must be acknowledged' (UKS-43) because the new ideas resulted from the knowledge that the old ones existed (UKS-53). By admittedly failing to acknowledge his 'sources', the student in this scenario violated the key academic rule of referencing, consequently passing someone else's work off as his own.

For Porto participants, the student's failure to observe an expectation of academic writing, i.e. his failure to reference (PTS-86) when both a reference and a critique of previous works was due (PTS-105), means also a lack of academic achievement on his part; academically, achievements also depend on knowing, recognising and

relying upon previous work (PTS-48), and this would add value to, rather than devalue, his own work. Creating something new, which is one of the student's arguments, 'can only be legitimate if previous works are acknowledged' (PTS-180), and under no circumstance can that justify the lack of acknowledgment (PTS-329). Although the degree of severity attached to the student's actions is different, as the significantly different scores indicate, both Aston and Porto participants therefore agree that failure to acknowledge previous works is plagiarism (PTS-109, UKS-21) – and a 'blatant' one (UKS-100) – and represents a type of behaviour that should be punished (PTS-109, UKS-66). It is therefore relevant to investigate this matter further, in order to find possible explanations for these different scores. Three hypotheses remain that may explain this difference; the first is the fact that the materials borrowed are non-linguistic materials; the second is that the student was from abroad; the third is that the student's behaviour may indicate that he acted intentionally, and therefore in violation of general moral or ethical principles.

The hypothesis that this difference may be due to the non-linguistic nature of the borrowed materials does not stand in this case, considering that this issue is raised by the participants, but not in any particularly relevant way; on the contrary, acknowledgment of works (UKS-43) and ideas (UKS-53) is due. Accordingly, the fact that he is a foreign student is dismissed by the claims that he purposefully used other people's work without acknowledgement for his own benefit, so that the fact that he is a foreign student abroad is irrelevant (UKS-189). On the contrary, the participants' comments reveal that they attach a significant degree of importance to the fact that the student acted intentionally. If, on the one hand, the student's claim is irrelevant (PTT-52, UKT-12) to decide on the outcome of the punishment, considering that the violation of the rule is sufficient to justify the case for plagiarism, on the other hand it contributes to the general agreement that the student undoubtedly plagiarised.

According to the participants, the student's final comment means that he 'did not take plagiarism seriously' (PTS-153), which brings along with it a violation of several moral principles. Firstly, he dismissed the principles of academic writing altogether, in particular the knowledge and recognition of the sources (PTT-34), thus failing to submit an original piece of work (UKS-77), which would give him credit anyway, albeit of a different nature (UKS-144). Then, by not providing a reference to his sources, the student prevents readers possibly interested in having access to those sources from finding them (UKS-85), which reveals a lack of professionalism (PTS-58). To this

moral principle participants add the student's moral obligation to respect the right to ownership of the legitimate authors of the original ideas, and to recognise the effort required to produce them (UKS-136), as well as the respect owed to his colleagues (PTS-58). Acting otherwise, in violation of any of these principles alone, shows that he 'clearly has no concept of right and wrong and why plagiarism is a problem' (UKS-147).

However, and maybe more importantly, he was also 'consciously opportunistic' (PTT-14), in violation of an important moral principle. The fact that he has no concept of right and wrong, despite potentially limiting his sense of personal liability to avoiding being caught, demonstrates that he not only 'knew [that] he had plagiarised' (UKS-24), but more importantly that he acted cunningly, knowingly and intentionally (UKS-79, UKS-171), with deliberate intent to deceive his examiners (UKS-173) and the university (PTS-150), with no sense of guilt (UKS-118), rather than out of ignorance or accidentally (UKS-29, UKT-5). Even more seriously, he is 'perfectly happy' with it (UKS-113). Finally, the moral wrong of the student also reflects on his claims that the only problem of plagiarism is being caught, and this is a perspective that the respondents, acting as jurors in a court of law, are unwilling to tolerate. Like murderers, judgements of plagiarism are not, the participants argue, in the eye of the beholder; metaphorically speaking, 'a murderer does not need to be caught to be a murderer' (PTS-249). Judgements of plagiarism therefore run independently of the actor's opinion (PTS-58), and of whether the plagiarist is caught or not (UKS-50, UKS-178). This clear intentionality on the part of the student (PTS-155), his 'lack of good intentions' (PTS-328), therefore annuls any consideration of extenuating circumstances.

The consistently high scores in this scenario reflect the participants' perception that the student not only violated the rules of academic writing, but he also violated several moral principles cumulatively. The violation of the latter, built upon the student's guilt (UKS-145, UKS-111, UKS-144, UKS-179), is that which the participants seem to be less willing to tolerate, with one mitigating circumstance – the creation of a work that is new – being accepted by the Portuguese participants more than by the UK participants. However, it is not dismissed by either group.

Participants in general also score higher in vignette 12, as is demonstrated by the profile plots presented in figure 3.4, suggesting that they agree with the accusations of plagiarism.

This vignette describes the instance of a student who borrowed works from other authors without acknowledging them, because she felt that citing too many references

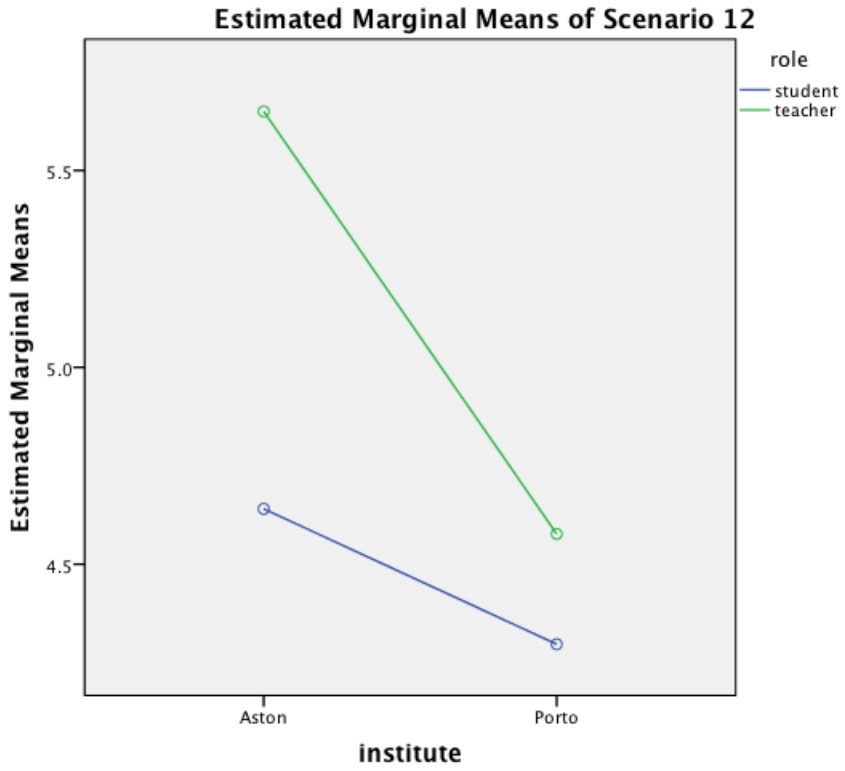

Figure 3.4: Profile plots for vignette 12.

would weaken her arguments. The student wrote the other authors' works and ideas using her own words, in order to avoid being challenged and in an attempt to make the work pass as her own, and then prepared herself for her viva, reading her dissertation thoroughly to be familiar with the work and answer any questions that she might be asked. The suspicions of the examiners were raised because they thought that the work was very well articulated, while there was no evidence that she could have read so extensively on the topic, or done the research herself. However, they were unable to prove that she had not authored the dissertation herself, since she was able to answer all the questions correctly.

Although both groups score above the neutral point, indicating that they agree with the charges, Aston participants score much higher than Porto participants, which accounts for a statistically significant difference between the two institutions ($F_{1,624}$ = 6.42, p = .012, partial η^2 = 0.010). This indicates that, although Porto scores quite highly, Aston participants assess this situation more readily and more severely as plagiarism than Porto participants.

This scenario raises several complex issues that need to be taken into account when considering cases of plagiarism, and that in particular reflect the pervasiveness

of the definition of plagiarism (PTT-27). The first is that by rephrasing other people's works and ideas, the student used her own words – a strategy that is often encouraged in the academy. However, using one's own words to rephrase other people's works and ideas is only legitimate in academic writing when the latter are acknowledged. By failing to credit the original authors, the student violated one of the main rules of academic writing – i.e. the rule of referencing – but the violation of this rule can represent different things, in different contexts. On the one hand, it can mean that the student, being unaware of the academic writing rules, did not know that she had to reference previous work or, alternatively, that she knew that acknowledgement of her sources was due, but lacked the competence to acknowledge them properly. In either case, this could reflect the student's lack of academic writing skills and a lack of proper training. On the other hand, it can result from the student's reluctance to observe the established procedures, rather than lack of competence and training. In this case, despite being aware of the applicable referencing conventions, the student may decide to knowingly, even consciously, attempt to obfuscate the unacknowledged borrowing. The act of plagiarism is intentional, and therefore represents a higher level violation also of ethical and moral principles.

The respondents' perceptions reflect these issues. Firstly, the question is raised of whether the type of materials or processes that the student borrowed are subject to plagiarism. In particular, suggesting that accusations of plagiarism depend on the evidence provided, such as the analysis of the excerpts taken from other works, cases involving types of plagiarism other than the textual, 'classic' plagiarism (UKT-12) may be undetectable, and hence impossible to prove unless the plagiarist confesses to having voluntarily hidden their sources (PTT-32). However, reinforcing the fact that the judgement of an instance as plagiarism is independent of the challenges that the detection task needs to face, and consequently that plagiarism can and does exist regardless of whether it is detected, participants argue that all types of work can be plagiarised (e.g. UKT-2, UKT-5, UKT-13, UKS-32, UKS-36, UKS-112, UKS-113, UKS-127, UKS-146), even if very subtly (PTT-5), including arguments and ideas, as well as core concepts, thoughts and even reflections (PTS-232). Although they recognise that the strategy used by the student is similar to paraphrasing, the participants are adamant in arguing that, as with paraphrasing, acknowledgement is due (UKS-26, UKS-49).

Secondly, by omitting the sources in her bibliography (UKS-31), the student failed

to provide proof of research, and missed the opportunity to reinforce and strengthen her ideas (UKS-72) by building upon the literature. As a result of failing to apply these academic writing rules, the student shows a lack of achievement that is questioned in the light of whether she had proper training (UKT-19); if this is a result of lack of competence, then the lecturer/tutor and the university should also take responsibility for her lack of achievement. Additionally, the context in which the academic assignment is set, as well as identifying the aims and objectives for the dissertation, and making these clear to the students is also crucial; if precedence is given to the element of originality, then the examiners should be able to understand that the work presented by the student was a sum of several different works; however, if the element that is mostly valued in the assignment is the student's ability to competently present a certain topic, using appropriate metalanguage, then punishing the student for plagiarism is not the most appropriate action (PTT-27). These aspects are crucial not only to avoid instances of plagiarism, but also to determine whether an instance represents a problem of academic achievement or a case of academic malpractice.

These issues are countered by the fact that, by knowingly and consciously obfuscating her sources and revising thoroughly to be familiar with the work (UKS-34, UKS-54) and subsequently address any challenges, the student was conscious of the wrongdoing; in itself, this represents a violation of moral and ethical principles, and this violation is so markedly intentional that the violation of the pragmatic rule embodied in the lack of acknowledgement of her sources, although not passing unnoticed, is backgrounded. Conversely, the violation of the moral principle of wrongdoing is foregrounded by the fact that the student actually intended to disguise the authorship of her dissertation and deceive her readers, making the text pass as her own (UKT-18, UKS-54, UKS-192, PTS-153). This intention contributed to disguising the text, transforming it into a coherent and cohesive whole that preserved it from the suspicion of her readers. The problem then is not one of simply failing to properly acknowledge the work. Instead, knowingly taking unfair advantage of the work of others (UKS-99, UKS-110, UKS-123) is a serious wrong, which increases the seriousness of the violation. On balance, the respondents' perceptions therefore rate this case as an obvious instance of plagiarism; by using 'unattributed work' (UKT-2), the student failed to observe the rules of academic writing, but it is the intention behind her acts that make this a clear case of plagiarism.

In both these vignettes, despite the significant statistical difference between As-

ton and Porto participants, they are overall unanimous that the students should be penalised for plagiarism in both instances. The tendency to disagree with the accusations of plagiarism across the two institutions increases in parallel with the apparent tension between the violation of the rule and the violation of the principle, as is the case of vignettes 1 and 8, which are judged diametrically by the participants of the two institutes as plagiarism and not plagiarism, respectively. A closer analysis of these particular vignettes and a discussion of the results in more detail to investigate possible reasons for this significant discrepancy should help understand the circumstances underlying these apparently disparate perceptions.

Vignette 1 describes an instance of a student who was accused of plagiarism because she failed to follow the referencing conventions adopted by the School. Although she had referenced all the works, words and ideas that she borrowed from her sources by indicating for every instance the name of the original author, she only included the year of the works cited in a few cases, and failed to provide page numbers in several instances.

The profile plots for this vignette, illustrated in figure 3.5, suggest that there is a considerable difference between the participants from both institutions, with those participants from Porto scoring significantly lower than those from Aston.

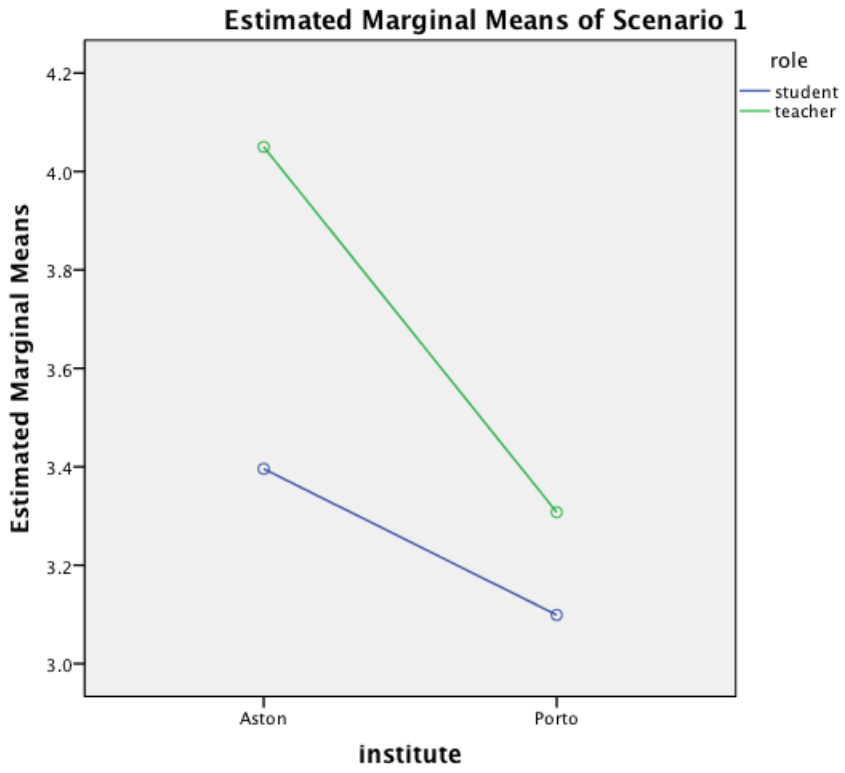

Figure 3.5: Profile plots for vignette 1.

The lines indicate that the scores of those participants from Porto fall considerably more (even abruptly, in the case of lecturers/tutors) than those of the Aston participants, and the analysis of the difference between the two institutions reveals a statistically significant result ($F_{1,624}$ = 4.28, p = .039, partial η^2 = 0.007), thus showing an actually relevant difference between the two institutions. This difference reflects the final assessment of the severity of plagiarism in this case which, maybe owing to the particular circumstances of the scenario, is not unanimous. The responses from Aston indicate that the British participants are overall less inclined to tolerate the non-compliance with the established rules than Portuguese participants. However, the level of agreement that the student should be penalised varies in degree of severity, even within the same institution – although the overall statistics are very rigorous. In one end of the scale, the main argument for this lack of tolerance is that the institution is entitled to setting the standards, and subsequently by making the students aware of these standards at an early stage e.g. by signing 'some kind of student contract which covers plagiarism' (UKT-12), and which states that improper referencing is a serious offence, the institution has the right to demand that the students abide by those rules and not others; this implies clearly marking the verbatim sentences using quotes along with the other elements of citation required, not only to signpost one's own work and opinion and the work and opinion of others, but also to allow the references to be 'confirmed *easily*' (UKS-4) by the readers. In this particular case, the seriousness is increased by the fact that the student included the year of publication only in some cases, while omitting it in others. The simple existence of this contract means that the students know the rules, or not the least that they do not have a reason not to know them, in which case students involved in cases like this cannot prove their innocence. Any violation of these 'very simple' (UKS-135) rules and conventions results in a misleading text, and implies a violation of the guidelines that had been made known to the students at an early stage in university, and of which they have acknowledged to be aware. As a result, 'paying the consequences' (UKS-135), with 'no arguments' (UKS-121) appears to be a fair outcome in these Aston participants' terms, and reflects a type of reasoning that provides a justification to consider that the student has plagiarised. Moreover, to counter arguments that some students may not have read the regulations and the contract properly, participants claim that punishing for plagiarism in these cases would be 'a lesson to read things properly' (UKS-147). In this sense, the university would have a case against the student, and whether the student acted

intentionally or not is irrelevant for the participants, to whom not providing enough information or 'not giving it close enough to the source' (UKT-12) can represent a case of inadvertent plagiarism, but technically is still plagiarism.

However, some participants hesitate in assessing the case as plagiarism outright, and consider that this could either be a case of 'plagiarism, ignorance or laziness' (UKS-153). Despite eventually scoring the case as plagiarism, they critically admit that the 'system is very harsh' (UKS-152), and argue that the student, despite not having cited properly owing to her failure to observe the conventions associated with and expected from academic writing, included a reference to authors, which shows that she 'understood the spirit of plagiarism' (UKT-18). These perceptions therefore suggest that a distinction should be made between flouting the rules and violating the principle, and consequently between *intentional* and *inadvertent* plagiarism. By providing the names of the original authors, the student demonstrates that she did not pretend to have authored those words and ideas, and that she did not misreference on purpose. Her actions were not intentional, and should as a consequence be considered less of a serious offence. This is the main argument presented by those participants that scored the lowest, i.e. plagiarism should be punished by disciplinary action when intentional, but not when it is inadvertent. When inadvertent, and especially when the student is at the outset of his/her degree, this violation of the conventions can indicate a lack of competence on the part of the student, and these circumstances should be taken to be mitigating. The student may have not been properly taught at school, or possibly is unfamiliar 'with the specific university system' (UKS-56), as in the case of foreign students. Or these inaccuracies may be the result of inappropriate management of the higher number of sources that are expected from longer essays. Since this is a matter of inconsistent referencing, there is some agreement in these circumstances that the student should be marked down, receive some support, and subsequently be given a second chance; but not be punished, academically and administratively, for plagiarism, since assuming that the wrong referencing was done on purpose is to a certain extent unfair.

Participants from Porto show a more relaxed attitude when assessing the student's actions, and less confident in passing judgements of plagiarism. In general, they justify their option by arguing that, although the conventions require that page numbers and year of publication be used, and that the student should be aware of the applicable conventions, penalising the student for plagiarism is an excessive measure (PTT-25)

– especially if the student is at an early stage of university. The power of existing regulations is partly dismissed on the grounds that such regulations are not always easily accessible, or not widely distributed – or, when available, they are often unclear (PTS-271). Unlike their Aston colleagues, Porto participants are less unanimous in considering that a student has the obligation to know all the regulations; some exculpate the student by admitting that she is entitled to not knowing them all in detail, whereas others argue that not knowing the regulations is the student's fault. Nonetheless, Porto participants tend to agree that the student's misreferencing is a case of methodological inaccuracy (PTT-2), frequently owing to the lack of proper previous training or clashing instructions from tutors, with an impact on the formal presentation of the text – more than a matter of barefaced plagiarism, which reflects more on the contents (PTT-38). Assessing this instance as an example of the former, more than representing it as the latter, Porto participants agree with some of their Aston colleagues that the student should be penalised for a lesser offence, be taught how to reference correctly (as part of her university education) (PTT-25) and invited to resubmit observing the university regulations, rather than being punished for plagiarism. The role of the student's intention is foregrounded by Porto participants, apparently more than by their British counterparts, by quoting intention as one of the main defining terms of plagiarism; if there is not an intention to steal someone else's work, then it cannot be considered plagiarism (PTT-27). Even if the lack of accuracy makes it harder to verify the sources, participants claim that the student cannot be penalised for giving the tutor more trouble in confirming the sources (PTS-122). On balance, they suggest that there should be a compromise between good sense and a sound practical judgment, and a blindfolded, outright enforcement of the rules (PTS-39). By this token, they contend that by citing her sources, especially when doing so systematically – even if incorrectly – the student demonstrated that she did not omit the details on purpose, and did not intend to deceive. Although incomplete, those references had some information, which makes a difference when comparing to cases with no reference at all (PTS-207). Therefore, since not following all the conventions does not mean that the student has plagiarised, i.e. that plagiarising is different from simply breaking a set of academic writing rules, most participants agree on penalising the student, but academically, for *what* she actually did (poor academic writing), rather than disciplinarily for plagiarism.

Vignette 8 presents another debate between the rule and the principle. It describes the case of a student who, having worked hard on her Master's dissertation to present

an original piece of work, was nonetheless accused of plagiarism. Despite having referenced all her sources, and included authors' names, years and page numbers (when applicable), besides providing a reference list containing all the sources she used, her supervisor argued that the charges of plagiarism are due to the fact that she used a referencing style and conventions that are different from the ones adopted and expected by her school; in other words, although she had made it clear in her dissertation which ones were her own works, words and ideas and which ones belonged to other authors, and provided the dates of publication and page numbers where those ideas can be found, she used referencing conventions that would be acceptable, or even recommended by other institutions, but which do not match those required by her school. The profile plots in figure 3.6 below indicate that Aston participants score considerably higher than Porto participants:

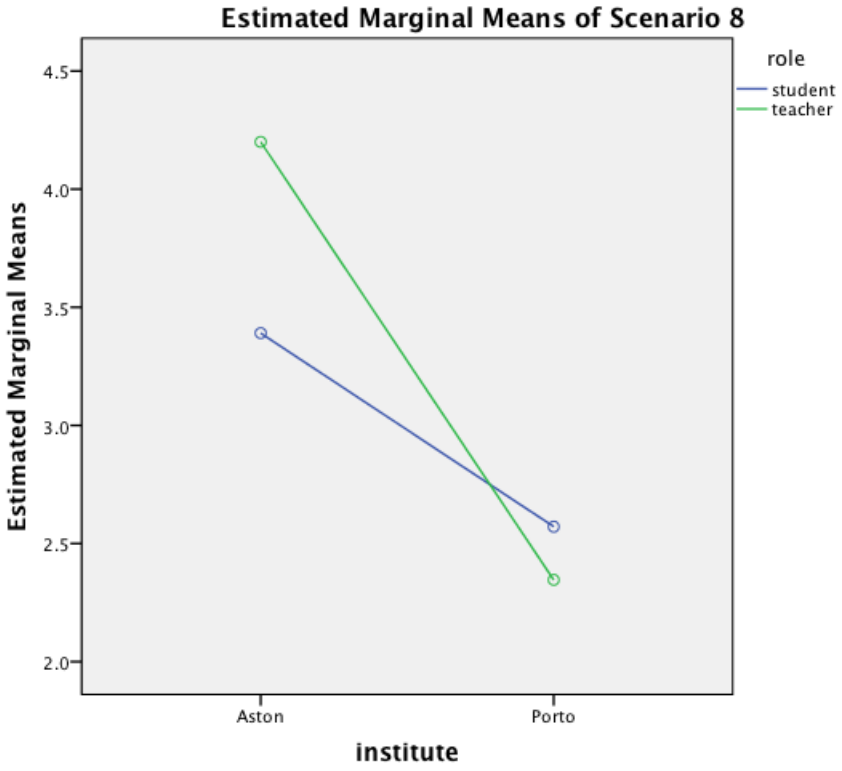

Figure 3.6: Profile plots for vignette 8.

The analysis of the results observed, presented in table 3.4, shows a statistically significant difference between the two institutions ($F_{1,624}$ = 31.99, $p < .001$, partial η^2 = 0.049), indicating that plagiarism is perceived significantly differently by the participants of the two universities; the participants from Aston (especially lecturers/tutors) score the highest, with a mean score of 3.80, indicating that they agree with the claims

that the student has plagiarised; on the contrary, the participants from Porto score much lower, with a mean score of 2.46, and therefore rejecting the claims that the student has plagiarised.

Once again, several questions are raised. By using a referencing style other than the one adopted by her university, the student failed to observe an established rule of academic writing. Although the student argued that she still acknowledged her sources properly, her supervisor claimed that this represented a problem, firstly because it indicated an academic writing issue, and secondly because it could make it harder for the reader to trace the sources. If the former, this instance would represent mainly a lack of academic writing skills that could be resolved by providing further training; however, by fully referencing her sources, the student challenges the claims that she tried to pass someone else's work as her own; on the contrary, the fact that she provided the authors' names, years of publication and page numbers confirms her claims that she wanted to build upon other people's work to reinforce her original combination of ideas. In this case, the violation of the established rules exempts the student from any moral wrongdoing; the violation of the rule occurred at a pragmatic level, and the moral principle remained intact. Nevertheless, she failed to follow the conventions adopted by her institution, and by so doing she potentially made it more difficult for her examiners to monitor her sources; this could indicate an intentional strategy used by the student to distract the reader from the sources, and consequently at a higher level represent a case of violation of the moral and ethical principles of trying to pass off other people's work as her own.

Although Aston participants admit that the inaccurate referencing provided by the student might be a problem of academic writing skills (UKS-54) even if at Master's level (UKS-39), they flag this scenario as plagiarism. Considering, in particular, that the student made a clear distinction between her work and the work of others (UKS-34, UKS-124) by referencing in fact her sources (UKS-67, UKS-111, UKS-179), this would represent an instance of 'bad referencing' (UKS-25), in which case the student should be invited to 'resubmit' (UKS-175), redoing 'the references in the correct style' (UKS-21) and observing the university conventions (UKS-36). For Porto participants, the distinction between an academic writing problem and plagiarism is crucial, and so should be the resulting penalties (PTT-49); lack of academic writing skills should reflect in the mark (PTT-52, PTS-174, PTS-301), but not in the judgement as plagiarism, which involves disciplinary action. The student referenced, and the failure to reference

in accordance with the set regulations is just a matter of formatting (PTT-1), an administrative (PTT-5) or methodological issue, or at most a technical detail that reflects in the lack of precision in the referencing system used (PTT-35). This is an important, yet formal aspect that aims to ensure that the sources used are easily identified and thereby confirm that authorship has not been violated; but it is not the only way of guaranteeing appropriate attribution (PTT-38), and any deviations can be corrected (PTS-180). Considering that the work is duly referenced, that due acknowledgement is thorough and well done (PTT-8, PTT-48), that the sources can be verified (PTS-328), and in particular that it does neither harm, nor challenge the original authorship and copyright (PTS-252, PTS-264), the problem should therefore be considered for what it is: a mistake (PTT-8, PTT-48) or disregard for the conventions, and not plagiarism (PTT-2, PTT-23, PTS-62, PTS-116, PTS-153, PTS-362).

Aston participants, on the one hand, tend to agree with these arguments and with the claims that following the established referencing conventions 'is a technicality' (UKS-95). Some of them therefore consider that adhering 'to one form of referencing over another is an arbitrary waste of time in academia' (UKS-188), since what is relevant is that the reader knows 'who wrote what, and where the original material can be found' (UKS-86), and that the student did not pass or try to pass someone else's work as their own (UKS-92) – especially when a particular style is used consistently (UKS-188). In this sense, 'improper referencing is not an act of plagiarism' (UKS-108); on the contrary, the fact that a large number of resources using different referencing styles are available (UKS-191) reinforces the argument for at least a certain degree of freedom. For these reasons, *not following* a particular referencing style is different from *not referencing* at all (UKS-96), and failing to observe the established conventions, despite accurately referencing the sources (UKS-79), is completely different from *intentionally plagiarising* someone else's works (UKS-33). On the other hand, in general they counter-argue that any circumvention of the established conventions should be discussed with the supervisor first and a proper justification should be provided (UKT-1), otherwise the student is 'guilty of plagiarism by carelessness' (UKT-11). Claiming that 'referencing guidelines are there for a reason' (UKT-1) and must 'be followed to the letter' (UKT-12), they give precedence to the violation of the convention, regardless of the existence of false representation (UKS-50), whose severity could be determined by analysing 'the depth of the deviation from the conventional reference model adopted by the school' (UKS-167).

Despite the apparent resignation to the established conventions, participants point out the need to adopt corrective measures that are proportionate to the student's wrong. By that token, although not following a particular referencing style, not referencing at all, failing to observe the established conventions, borrowing without acknowledgement, carelessly and intentionally plagiarising someone else's works are all instances of plagiarism, and each brings along different degrees of severity. At one end of the scale, the severity of the accusations can be extenuated by the fact that the responsibility for the student's fault is also shared by the supervisor (who could have provided help to cite correctly) and the university (by providing clearer guidelines) (UKS-32, PTS-186, PTS-86, PTT-142). Considering that the violation of the referencing rule resulted from the lack of understanding of the conventions and that the referencing system adopted by the institution is the one to be used, the supervisor should have drawn the student's attention to this before submission (PTT-45, PTT-10, PTS-226), and consequently avoided the violation by discussing those guidelines with the student to help her know 'where she was' (UKT-19) and train her (UKS-136). Failure to act may indicate a lack of supervision during the whole process (PTS-242), which is reinforced by the fact that the student actually cited the works used. These perceptions, shared by participants of both institutions, challenge the principle of why an institution should have their own conventions, and not permit other conventions (PTT-9, PTT-5). If there was a universally accepted referencing style, then the school referencing conventions should be reviewed (PTS-149) accordingly, and not following a particular institution's conventions would not equate with a wrong referencing (PTS-73). Otherwise, the lack of universal principles means that not all readers will follow the same referencing conventions, and therefore the definition and the concept of plagiarism should not be based on the familiarity of the reader with a certain system (PTT-34).

At the other end of the scale, the most serious instances, and correspondingly the most serious penalties, involve intention. In this sense, plagiarism is perceived mostly, but not only, by Portuguese participants as a moral issue that is violated by unscrupulous people when they take unfair advantage of someone else's work (PTT-27). The universal principle upon which plagiarism relies, and the one that needs to be absolutely respected, consists of making clear the work that is original and the work that is based on the sources, so what is relevant is that a distinction is made between issues of referencing conventions and the inappropriate borrowing of someone else's work,

between the ethical and the pragmatic, between the violation of the rules and matter of form (which must be considered academically) and the violation of the principle, which is substantial matter (and which must be considered disciplinarily). The issue of intention therefore plays a crucial role in assessing the most serious instances of plagiarism.

The results observed for the perceptions of plagiarism indicate that, in this tension between the rule and the principle, the Aston participants consider that both the rule, and possibly the principle have been violated, whereas the Porto participants, conversely, seem to believe that the principle remained intact, and that breaking the rule is less important, considering that the student made it clear which works she borrowed, and from where. This tension reflects even more clearly in the interaction that is observed between the two main effects 'institute' and 'role', and will be discussed further in subsection 3.3.3.

3.3.2 Analysing the Perceptions of Plagiarism among Students and Lecturers/tutors

The analysis of the main effect 'role' also demonstrates a statistically significant difference between students and lecturers/tutors ($F_{1,624}$ = 12.95, $p < .001$, partial η^2 = 0.020). The latter score consistently higher than the former, indicating that students have a less well-defined, or more permissive view of plagiarism than lecturers/tutors when faced with the cases of plagiarism described in the vignettes, whereas the latter, by agreeing with the charges of plagiarism described and with the corresponding sanctions more consistently than the student participants, show a tendency to be less reluctant to punish.

Once again, the analysis of each vignette individually demonstrates that the results observed for each scenario are not always statistically significant, as illustrated in the table 3.5.

A statistically significant difference, indicated in bold in the tables, was observed in only three of the 12 cases in the main effect 'role', which indicates that in all other cases, despite the apparent different perceptions observed between students and lecturers/tutors, those differences are not significant statistically.

The analysis of the *mean values* of each of the three dependent variables of the main effect *role* where results are statistically significant indicates that student and lecturer/tutor participants agree that the subject of the vignettes has plagiarised in the

Vignette	df	F	*p*	Partial η^2
1	1	2.95	.086	0.005
2	1	0.59	.442	0.001
3	1	1.37	.243	0.002
4	1	0.71	.401	0.001
5	1	3.60	.058	0.006
6	1	1.04	.308	0.002
7	1	0.36	.551	0.001
8	1	1.53	.217	0.002
9	**1**	**31.83**	**$<$ .001*** **	**0.049**
10	**1**	**7.83**	**.005** **	**0.012**
11	1	0.00	.983	0.000
12	**1**	**5.32**	**.021* **	**0.008**

Table 3.5: Analysis of the main effect 'role'.

three cases (vignettes 9, 10 and 12), but the intensity of agreement varies. In all these three scenarios, lecturers/tutors score consistently higher in all cases. Students score consistently lower across the three vignettes, and in the case of the vignette 9 their level of agreement is just above the middle of the scale. Although the perceptions of the two groups of participants are relatively unanimous, and no research is required to investigate disagreement between the two groups, it is, however, important to investigate the reasons that are behind the statistically significant differences between the two groups. In order to find possible reasons for these differences, these vignettes are now discussed.

The first case describes the instance of a student who developed a new model with a friend, and both used it; the student used the model in his dissertation, and his friend used it in a journal article that he published, i.e. the student worked collaboratively with a friend, and both agreed that they should use it as their own. The examiners however argued that collaborative work was subject to acknowledgement, so the student should have acknowledged that the model was developed with a friend.

Unlike other previous cases, this instance does not raise copyright issues, but it may raise the question of collusion, since the work was created by both authors, and used individually by each of them without acknowledging the other. The case becomes

even more complex if we consider that, owing to the fact that neither the student nor probably his friend attempted to use the collaborative work and make it pass solely as their own, the principle of authorship was not violated; but the rule of submission of unacknowledged joint work was violated: by submitting the work resulting from a collaborative effort, the student adopted a procedure that is not allowed in the academy.

The profile plots resulting from the statistical analysis of the perceptions of students and lecturers/tutors presented in 3.7 indicate that the difference between the two roles is considerable.

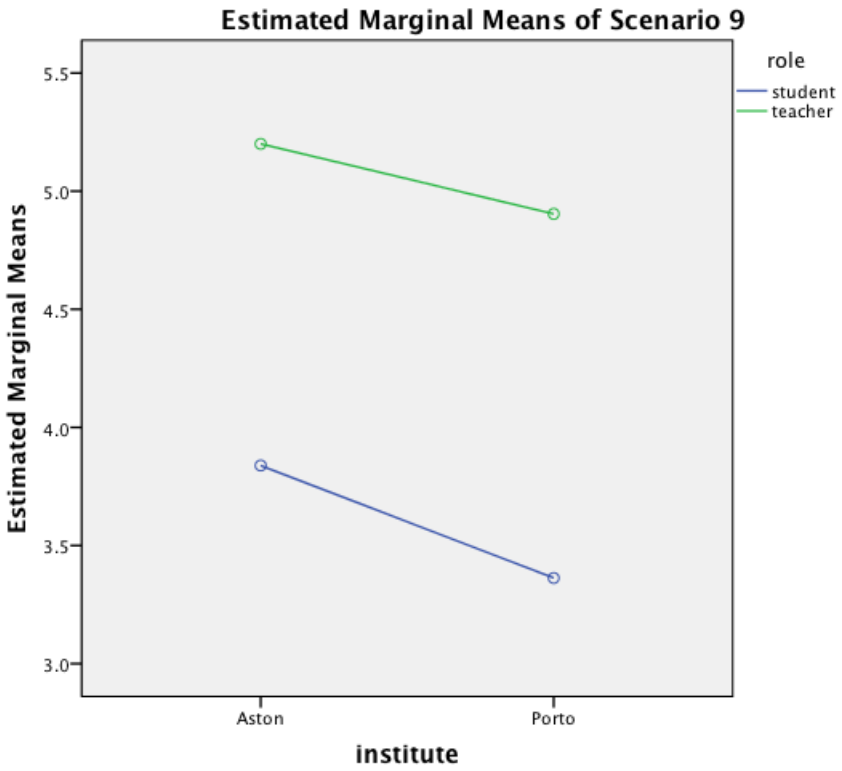

Figure 3.7: Profile plots for vignette 9.

Students score their perception at 3.60, compared to the score of 5.05 of lecturers/tutors. Interestingly, the blue line (corresponding to the students) drops more abruptly, as a result of Porto students scoring below 3.50, which suggests that they do not consider that the student has plagiarised. The statistical analysis of the data, shown in the table 3.5, indicates a statistically significant difference between the roles of student and lecturer/tutor ($F_{1,624}$ = 31.83, $p < .001$, partial η^2 = 0.049), suggesting that lecturers/tutors are quite confident that, by failing to acknowledge the common authorship, the student had plagiarised, whereas the students are more reluctant to assess the case as plagiarism. Although, in general, they tend to agree with the

charges that the student has plagiarised, they score the scenario considerably lower to suggest that, possibly due to mitigating circumstances, the student is hardly responsible for plagiarism. The analysis of the mean and 95% confidence intervals, shown in figure 3.8, indicates that the Porto students tend to consider that the student had not plagiarised.

4. institute * role

Dependent Variable: Scenario 9

institute	role	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
Aston	student	3.839	.134	3.576	4.101
	teacher	5.200	.414	4.387	6.013
Porto	student	3.363	.097	3.172	3.553
	teacher	4.904	.257	4.400	5.408

Figure 3.8: Mean and 95% confidence intervals for vignette 9.

This scenario raises some complex questions. By reusing the work developed previously with a friend in his own assignment, the student may have doubly infringed the rules of submission of academic work. Firstly, he submitted work that he authored jointly with a friend, although single authorship is expected and, secondly, he reused his own previous work without acknowledgement. These issues could, however, be subject to two diametrically opposed interpretations; on the one hand, the student's acts could be considered a result of his lack of awareness of the rules, and hence an example of academic incompetence; or they could be taken as an act of deception, and hence a violation of a moral conduct whose seriousness would depend on whether the student acted intentionally of negligently. Alternatively, they could be judged more moderately, anywhere between these two ends of the scale.

The analysis of the responses indicate that, in general, the lecturer/tutor participants are the ones that hesitate less in passing judgements of plagiarism. The rules of academic writing are clear, and any deviation from those rules means that the student, to a greater or lesser extent, has plagiarised. The student participants also tend to agree that, in this scenario, the student should have referenced and acknowledged the sources (UKS-20, UKS-141, UKS-188, PTS-23), since using work that has been published under someone else's name without acknowledgement indicates to the reader that it originally belongs to the student, and this act of misleading represents plagiarism (PTS-23, PTS-212, UKS-123). However, the analysis of the feedback collected from the respondents reveals that this perception is not always unanimous, and that there are two complementary, overlapping (Sunderland, 2004) but sometimes

competing, clashing discourses that interfere with the respondents' perceptions: the legal and the academic.

In legal contexts, the most important aspect is that the intellectual property, the ownership of work, words and ideas, are not violated, and this violation can be prevented by obtaining permission from the property rights owner to (re)use the materials. In cases where more than one publication use the same materials, the one published first usually has more credibility in taking the authorship (UKS-107), and although prior authorship is not the only criterium used to determine the ownership of intellectual property, it is a crucial one in settling disputes. From this perspective, by making an agreement with his friend, who first published the collaborative work, the student obtained permission to use the work, but although this permission may be sufficient to resolve legal, intellectual property rights issues (PTS-57), it is not sufficient to meet the academic requirements. The fact that the article was published first thus reinforces the importance of proper referencing in all subsequent works (PTS-362), including the student's dissertation (UKS-110, PTS-48).

Academic procedures, although framed by a legal context, run differently from and tend to be even more demanding than simple legal procedures, in order to prevent practices that albeit legal from the point of view of ownership (e.g. 'commercial transactions' (PTS-249)), can represent academic malpractice. In this context, not only do the legal procedures have to be observed, but the expectations regarding the rules of academic writing behaviour have to be strict. Even if the official permission of the original author is not required to reuse parts of published work in one's own research, all instances of other people's work and joint authorship have to be formally acknowledged by referencing. Consequently, if the student is not officially one of the authors of the published article, then reference is due in his dissertation to an article that was written by someone else (PTS-207).

Secondly, it is an academic requirement that work submitted for assessment is done individually (UKS-19), although acknowledged collaborative work can be used in cases of co-authoring (PTS-79, PTS-364), with permission of the other co-author(s) (PTS-189, PTS-153). It is a rule that the distinction between one's own work and the work of others is clear, with the true authorship of each section properly acknowledged (UKT-1, UKT-9, UKT-12), otherwise the use of collaborative work is labelled inappropriate, and considered to be collusion, usually implying punishment for all those involved in the collaborative process. In this particular case, a reference was due both to himself

and to his friend (UKS-43, PTS-158, PTS-174, PTS-184, PTS-329), not the least to avoid claiming sole authorship (PTT-32) or refute any attempt at doing so (UKS-163).

Conversely, if those requirements are met, then the student in this scenario could be found not guilty of plagiarism (UKS-189). Additionally, even one's own previous work needs to be clearly marked, so that the reader knows *exactly* which work one has contributed, as well as the stage at which that contribution was made. In order to avoid accusations of self-plagiarism (UKS-65, UKS-100), if authorship had been granted in the article, then the student should have cited his own article in his dissertation (PTS-105, PTS-226). This is a requirement of proper academic writing that might, under some circumstances, clash with definitions of plagiarism. If plagiarism is defined as the unacknowledged borrowing of someone else's works, in *strictu sensu* the fact that the student reused work that he developed in collaboration with a friend, and that is also partly his own work, then the student cannot be said to have violated another author's original work altogether.

The perspective of plagiarism as the violation of someone else's ownership of an original work, with moral implications reflected in the successful or unsuccessful attempt to claim authorship for someone else's work, rather than as the violation of a technical (UKS-156) or purely bureaucratic issue that is irrelevant when the student has contributed to the work (PTS-123), raises the question of whether when judging cases of plagiarism, the observation of the rule takes precedence over the moral principle, or otherwise. This failure to comply with the rules can result from negligence, lack of awareness or even unknowing violation of the procedures, or from an intention to deceive. The latter, representing the violation of the moral principle of truthfulness, results in deception – or not the least in an attempt at deceiving. Even when perceiving this case as technically representing an instance of plagiarism (UKS-113), which is based on the violation of a rule, students consider that the rules should be applied with good common sense (PTS-39), and that some moral principles need to be taken into account when judging more severely. A carefully presented dissertation should include all the authors but, even so, arguments for failing the student are not strong enough.

On the one hand, the fact the the student used work to which he contributed legitimises at least in part his reuse (UKS-98, UKS-99, UKS-170). On the other hand, seen as a technical issue, this lack of acknowledgement can be resolved by providing the required references (UKS-99), and hence disclosing the authorship (UKS-156), which

would be sufficient to withdraw the accusations against the student. Under these circumstances the lack of citation is only a formal mistake that may reflect on the student's mark (PTS-261), but not in the final outcome as severe plagiarism, because this is a single case throughout the whole dissertation (PTS-360), and, more importantly, this action was possibly unintentional (UKS-143), or else a result of unclear guidelines (UKS-72) or misunderstanding (UKT-18). Consequently, the fact that the student did not reference correctly might not be entirely his fault, even though he is still responsible for not having referenced (UKS-72). These perceptions suggest that the students are more inclined to judge plagiarism less strictly when there is no indication of intention (UKS-29, UKS-113, PTS-338), and value the principle in detriment to the rule.

If moral principles are to take precedence over the rules, then disciplinary considerations need to challenge the fact that the student's name as a co-author was omitted from the journal article, for which co-authorship was due (PTT-45, PTT-52) as a result of both having contributed their ideas. Under these extenuating circumstances (PTT-37), the lack of intention on his part, although not dismissing the charges of plagiarism altogether, is not irrelevant, and even deserves some lenience (UKT-10). However, if the student is responsible for plagiarism, then equal corrective measures should be proportionally applied to his friend (PTT-45, PTS-281), since his name should have been included in the article in the first place (PTS-122).

The identical obligation to acknowledgment in academic and published work raises a pertinent, but hardly mentioned discussion as to whether student plagiarism is assessed on the same terms as plagiarism in other published works. Interestingly, the perception that this is not exclusively the student's problem, indicating that the rules and the corresponding corrective measures apply both to academic work and to research (and consequently also to his friend) is more salient among student respondents than lecturers/tutors.

Taken together, these perceptions indicate that there is a tension between the principle, the rule and the practice that requires further investigation into the particular circumstances of the case (UKS-163). Although Aston respondents tend to support the rule that the lack of citation is sufficient to judge a text as an instance of plagiarism (UKS-54, UKS-136), the violation of a rule of academic writing is less important than the intention behind the offence of the moral principle, and should not be judged independently of the particular circumstances involved. The degree of intentionality, the awareness of the rules, the clarity of the university guidelines, or the student's

academic writing skills are all factors contributing to the judgement of plagiarism and, although a student at Master's level is expected to know how to act in a situation like the one described (UKT-12), they all contribute to assessing an instance of plagiarism as more or less severe. This reluctance to judge the student's acts as a severe case of plagiarism is not independent of the fact that the student has not apparently hidden joint authorship intentionally. Alternatively, the scores may be due to the fact that the students either (a) consider that violation of copyright is one of the main components (if not the main component) of plagiarism, in which case the student would not be responsible for any breach, or (b) they do not consider that collusion is plagiarism (it should be noted that the participants were asked to give their opinion as to whether the subject of the vignettes should be charged with *plagiarism* only).

The weight of the moral assessment of instances of plagiarism also reflects in vignette 10, which describes the instance of a student who borrowed a text from an original written in another language variant, without acknowledgement, by introducing changes in spelling, morpho-syntax and vocabulary, so as to adapt it to her language variant. These changes in wording contributed to producing a different text, despite reusing some phrases of the original. However, the works and ideas are the same as the original, therefore violating both the rule and the principle; the former is violated as a result of the student failing to acknowledge the original; the latter is violated because the student knowingly reused the text without acknowledgement, and furthermore changed the text to make it read coherently 'like an original', i.e. like a text written in her own language variant. This case raises once again the question of whether the student's infringement of the rules of academic writing are a result of an intention to deceive or, on the contrary, of her lack of academic writing skills.

The profile plots for this case, presented in figure 3.9, suggest that all participants score quite high: the analysis of the means and 95% confidence intervals shows a lower bound of 4.58 among the participants scoring the lowest (Aston students).

The statistical analysis of the data presented in table 3.5 reveals a statistically significant difference between the two groups of participants in the main effect 'role' ($F_{1,624}$ = 7.83, p = .005, partial η^2 = 0.012), meaning that only the different scores between lecturers/tutors and students are statistically significant, whereas the apparently considerable difference between the scores of the two groups of lecturers/tutors, on the one hand, and the two groups of students, on the other, is not. An analysis of the two groups (students and lecturers/tutors) individually reveals that, unsurpris-

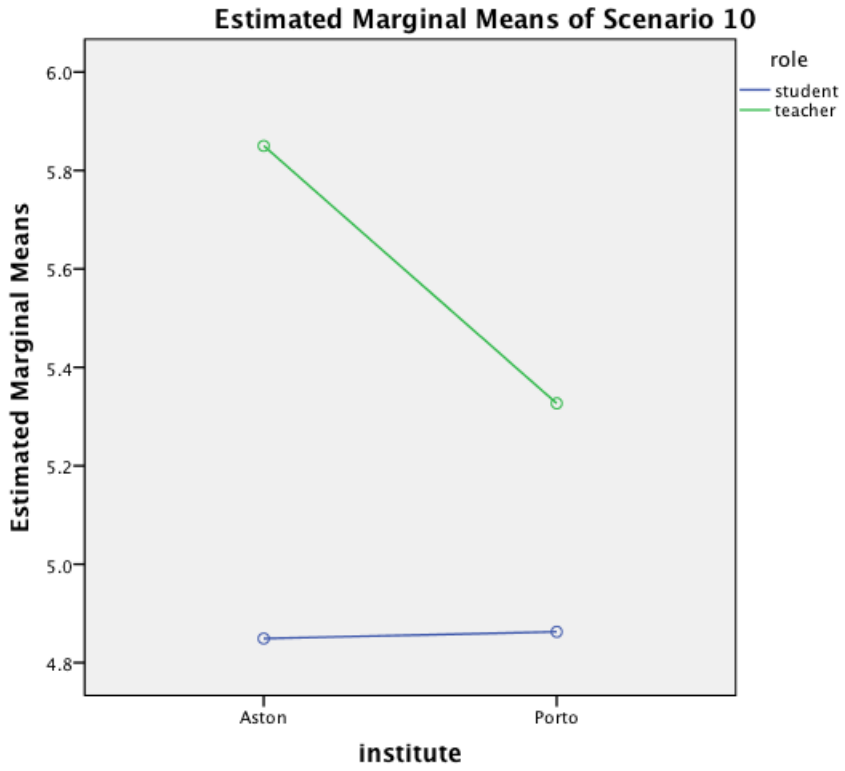

Figure 3.9: Profile plots for vignette 10.

ingly lecturers/tutors assess this type of plagiarism more severely than students and, as educators, are less tolerant to what they believe is academic dishonesty, although students, despite scoring lower than lecturers/tutors, agree with the accusations of plagiarism, and are equally unwilling to exculpate it – either because it violates the established rules, or because on principle they refuse to accept their colleagues gaining an unfair advantage over themselves.

The judgements of the participants therefore oscillate between censoring the violation of the rules, reflected on the student's acts, or punishing the infringement of the principles, which reflects on the student's intentions. Two questions in particular can be asked to justify the lack of acknowledgment of the student, who borrowed works and ideas inappropriately, namely (a) whether this lack of acknowledgment was due to the lack of academic writing skills and proper training in academic writing, or (b) whether it is a result of the student's knowing or (even more serious) intentional attempt to deceive. The perceptions of the participants suggest that the former does not seem to stand, especially owing to the fact that the thesis lacks originality, and fails to present any innovative material, as is expected from a student of this level. Perceptions that the student has clearly plagiarised are grounded on the fact that her thesis

is based on a previous one (UKT-16, PTT-14, PTT-52), which was object of 'translation' and adaptation from a source in another language variant (UKS-36, PTT-5), and lacking any innovative material (PTT-5).

In this respect, the participants' perceptions are clear; if the writing is based on someone else's previous ideas or text, then reference is due (PTS-228, PTS-242). Secondly, the situation is described as being identical to paraphrasing (UKS-26, PTS-135) or translation (PTS-62, PTS-207), so should be handled in similar terms; as in cases of translation or paraphrasing, the student is entitled to borrow the text, but credit is still owed to the original author, and consequently this should be clearly indicated in the text (e.g. by using italics). In this case, the work remains someone else's work (UKS-35), and reference is due (UKS-107, UKS-178, PTS-34, PTS-47, PTS-116, PTS-271) regardless of the language used (UKS-189), or of how many changes have been made (UKS-47) to reassemble a different text (UKS-32). The original idea, the fundamental concepts from where the knowledge originated (UKS-79), still belongs to another author. By failing to do so, the student borrowed the ideas inappropriately (PTS-34, PTS-62, UKS-36). In this case, unacknowledged use (UKS-31, UKS-77, UKS-110, UKS-121, UKS-142, UKS-144) represents plagiarism of ideas (UKS-96), and consequently a breach of originality (UKS-153). As a consequence, although the student changed words and sentence structures, those changes are not sufficient to make the text her own (PTS-70, PTS-73); given the importance attributed to content, rather than form, linguistic alterations are not sufficient to avoid the plagiarism of the ideas (UKS-83, UKS-113).

From the perspective of the violation of the rules, the clear lack of acknowledgement suffices to penalise her for plagiarism but, even if the lack of acknowledgment was remotely due to the fact that the student lacked academic writing competence, her work (maybe owing to the fact that she edited the text, rather than rephrasing it in her own words and acknowledging it (UKS-136, UKS-156, PTS-286)) resembled the original enough to be traced and challenged by the supervisor – which suggests that only minor changes and edits have been introduced to the text.

At the same time, the fact that those linguistic changes at the level of spelling, morpho-syntax and vocabulary were made coherently to read like an original, together with the lack of acknowledgment, indicate that the student had an intention to pass off someone else's work as her own, and hence intended to deceive, and it is this violation of the moral principle that the participants tolerate less. The results observed demon-

strate that all participants consider that the student plagiarised, suggesting that she intended to deceive by deliberately taking someone else's work and passing it off as her own. For the lecturers/tutors, by omitting the reference to the source, and more importantly by not including even a minor, incomplete reference to the original work, the student demonstrates that she reused the structure and the ideas with intent to plagiarise (UKT-2, UKT-18, PTT-27). Even if considering that the student did not understand that taking someone else's words, altering them and passing them off as one's own is plagiarism (UKT-12), the fact that the issue has been brought to her attention and that she was offered the opportunity to resubmit demonstrates the good will of her supervisor, and should be enough from an academic perspective (UKT-5, UKS-56). On the contrary, her refusal to resubmit signals her ill-faith, and reinforces her degree of intentionality for using someone else's work and passing (or attempting to pass) it off as her own (PTT-32, UKS-49). Consequently, the linguistic alterations, which reflect on the changes of sentence structures or replacement with synonyms, when combined with the lack of acknowledgment, are indicators of intention, and hence not only are they insufficient to reject the accusations of plagiarism, they contribute to proving the accusations and reinforcing the arguments for the student's plagiarism (PTS-62, PTS-109, UKS-54).

Taken together, these circumstances contribute to the general agreement that the student acted knowingly (PTS-301), and more importantly acted with clear (PTS-364), 'evident and indisputable' (PTS-180) intention, and therefore, having deceived her colleagues and the examiners by gaining or trying to gain an unfair advantage over her colleagues, she is justly subject to accusations and punishment for plagiarism (PTS-153, PTS-261, UKS-65, UKS-125).

The other instance where participants in general agree with the accusations of plagiarism, by scoring quite high, is vignette 12, as discussed in section 3.3.1 above.

This vignette, presented above, describes the instance of a student who borrowed from previous works without acknowledgement, and used her own words, in an attempt to avoid too much referencing. The student subsequently prepared herself very well for her viva, in order to be familiar with the ideas, the terminology and the text, and hence be able to address any challenges made by the examiners.

In addition to participants from Aston scoring higher then those from Porto, as discussed above, lecturers/tutors also score higher than students in both institutions, thus following a pattern that is similar to the one in the overall, total sum of scores.

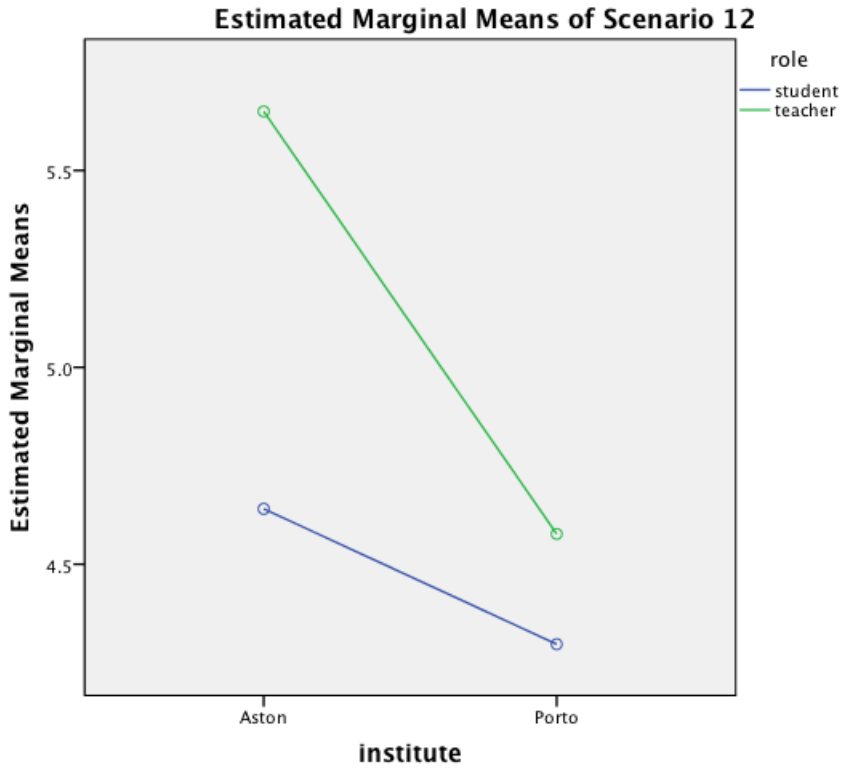

Figure 3.10: Profile plots for vignette 12.

The profile plots illustrated in figure 3.10 indicate a considerable difference between the scores of students and lecturers/tutors, with the latter scoring consistently higher than the former. This particular case differs however from the sum of scores in that the lecturer/tutor scores tend to fall more abruptly, meaning that the difference between the two role conditions (students and lecturers/tutors) in Aston is much greater than the difference between the same role conditions in Porto, maybe owing to the fact that Aston lecturers/tutors score particularly high in this case. The difference between the two role conditions overall – students and lecturers/tutors – is statistically significant ($F_{1,624}$ = 5.32, p = .021, partial η^2 = 0.008), suggesting that although both students and lecturers/tutors consider that the student has plagiarised (the lower bound is 4.29, therefore above the neutral point of 3.50, and the mean value is 4.47), lecturer/tutors are more demanding in assessing the case (with a lower bound value of 4.59, and a mean value of 5.11). Both groups of participants therefore consider that this scenario presents a case of blatant plagiarism, perceiving that the student intended to disguise the authorship and deceive her readers, in particular her examiners.

This vignette raises several aspects that reflect the pervasiveness of the definition of plagiarism (PTT-27). The first main issue is that by omitting the acknowledgment

of the sources, the student infringed a rule – the lack of acknowledgement – but not the required referencing conventions, since the sources that she cited were done correctly. Additionally, two other problems are raised. On the one hand, this is not a case of 'classic' plagiarism, in the sense that it is not the words that are borrowed verbatim from the source. On the contrary, the student borrowed the works and ideas, but used different words, and therefore the issue is once again raised of whether works and ideas, besides words, can be improperly borrowed from other sources. Additionally, the problem arises of determining whether plagiarism can be detected by reference to linguistic elements used to express works and ideas, or whether only 'textual plagiarism' (Pecorari, 2008) can de detected. On the other hand, the judgement of the case also depends on determining whether the student broke the rule as a result of lack of academic writing skills, or alternatively whether she was expected to show that she has read on the topic and demonstrate that she gained the required knowledge, in which case citing the sources would be backgrounded. The latter option is discarded altogether, both because at Master's level the student is supposed to be familiar with the required academic writing conventions, and because her lack of referencing was not consistent; although she failed to reference in some cases, she acknowledged her sources in other cases. This makes breaking the rules inadmissible, and hence demonstrates that not only was she aware of her actions, knowing of what she was doing, but also that she acted consciously and deliberately. It therefore indicates that her actions represent an act of deception, which is unfair to colleagues, examiners and the system – and hence, by taking advantage of others, is immoral.

On balance, students consider that this form of borrowing, owing to the fact that the student used 'unattributed work' (UKT-2) when proper acknowledgement of the other authors' works and ideas was due (UKT-12), is barefaced plagiarism (PTS-180). Moreover, these considerations are reinforced by the argument that all types of work are subject to charges of plagiarism (e.g. UKS-32, UKS-36, UKS-112, UKS-113, UKS-127, UKS-146), including arguments and ideas, as well as core concepts, thoughts and even reflections (PTS-232). For these participants, it is clear that even if the student knew the ideas, they were not originally hers, and since knowledge is not sufficient to discard the need for referencing, acknowledgement is due (UKS-53, PTS-39). Such cases are therefore clearly perceived as plagiarism (e.g. UKS-20), even if involving linguistic strategies that are more sophisticated than word for word borrowing, such as paraphrasing (UKS-26, UKS-49); and like paraphrasing, referencing is also

due in these cases.

For lecturers/tutors this instance does not represent a case of classic plagiarism (UKT-12), in which the student borrows – and reuses – the text from other sources. This therefore makes it much more difficult for the examiners to challenge the student, and subsequently to prove that the student has actually plagiarised – unless the plagiarist confesses to it (PTT-32). This is a practical issue that is of particular concern to them, because they are responsible for marking their students' assignments and feel the pressure to detect plagiarism of works and ideas, besides the more 'straightforward' form of textual plagiarism, and their detection task may be hindered by the inability to prove the instance of borrowing. Despite the extra work required to detect these instances of plagiarism, lecturers/tutors concur with the students' claims that borrowing works and ideas is plagiarism, albeit one of a very subtle form (UKT-2, UKT-5, UKT-13, PTT-5). In parallel, they also recognise that simply changing the original words to one's own is not sufficient – and is even irrelevant (UKS-47) – to avoid the accusations of plagiarism.

Both lecturers/tutors and students agree that the behaviour adopted by the subject of the vignette is very serious, even inadmissible (UKS-77) in academic contexts. But even more serious than breaking the rules of academic writing is the intention behind that violation, which represents a violation of the principle of improper borrowing. Her clear intention to appropriate someone else's work, by deliberately passing it off as her own (UKT-18, UKS-54, UKS-192, PTS-153), correlates with her attempt to revise thoroughly and avoid any problems with addressing the examiners' challenges (UKS-34, UKS-54). These actions not only reveal her intentions, but more importantly demonstrate that she was conscious of the wrongdoing, and knowingly chose to act as she did (UKS-99, UKS-110, UKS-123). The intention behind these actions clearly contributes to the students' responses: they consider that actions of this type deserve a severe punishment, and consequently it is unfair that students who consciously chose to act like this and deceive the system may get away with it (UKS-145, UKS-99, PTS-48, PTS-364), as a result of being clever (UKS-176), or a stroke of luck (UKS-99). In this sense, despite admitting that the student might have misreferenced and failed to provide proof of research (UKS-31), thus missing an opportunity to reinforce her own ideas (UKS-72), as a consequence of her lack of proper training in academic writing (UKT-19), which is crucial to avoid instances of plagiarism, the participants tend not to dismiss her behaviour of malpractice in favour of her lack of academic achievement.

This vignette hence represents a case of double violation; violation of the rules of academic writing and violation of the principle of academic honesty, according to which plagiarism is an ethical and moral wrong. The former reflects on the lack of in-text citation, as well as in the omission of references in the bibliography section; the latter, on the other hand, reflects on the student's choice to omit a reference to her sources, and even more importantly on her deliberate attempt to disguise the origin of the works and ideas reused in her dissertation, with the intent to ultimately deceive the examiners. Taken together, these circumstances indicate that the violation of the principle is markedly intentional, and this has the potential to explain, at least in part, the unanimous assessment of this scenario as plagiarism.

3.3.3 Analysing the Interaction between the Institution and the Role

The two previous subsections discussed the vignettes where significant differences were found between the two conditions of each main effect 'institute' ('Aston' and 'Porto') and 'role' ('lecturers/tutors' and 'students'). This section discusses the third element, the interaction between the 'institute' and the 'role'.

Unlike the previous main effects, the tests of between-subject effects when considering the total of the sum of scores reveals that the 'interaction' between the two main effects 'institute' and 'role' is not statistically significant ($F_{1,624}$ = 0.53, p = .466, partial η^2 = 0.001), indicating that, overall, each variable shows a different behaviour in each condition of the other variable. Consequently considering that these results are not statistically significant, they are not worth discussing in general. However, the analysis of each vignette individually, as illustrated in the table 3.6 below, demonstrates that a statistically significant interaction can exist between the two conditions of each main effect, whereby one variable (e.g. 'role') shows a different behaviour in each condition ('Aston' and 'Porto') of the other variable ('institute'), in particular circumstances:

Specifically, the results observed for each scenario demonstrate that this interaction is *statistically significant* ($F_{1,624}$ = 4.79, p = .029, partial η^2 = 0.008) in one case – vignette 8) –, and a near-significant result was found in vignette 4 ($F_{1,624}$ = 3.43, p = .064, partial η^2 = 0.005). In the other cases where an interaction exists (vignettes 6 and 11), this is absolutely non-significant. Although these two cases are not sufficiently strong to influence the results observed in the overall sum of scores, a further investigation into the two instances can provide some explanations for the significant and near-significant interactions observed.

Vignette	df	F	*p*	Partial η^2
1	1	0.79	.376	0.001
2	1	0.37	.545	0.001
3	1	0.12	.728	0.000
4	1	3.43	.064	0.005
5	1	0.24	.625	0.000
6	1	1.28	.259	0.002
7	1	0.00	.971	0.000
8	**1**	**4.79**	**.029***	**0.008**
9	1	0.12	.727	0.000
10	1	1.05	.306	0.002
11	1	0.36	.550	0.001
12	1	1.70	.193	0.003

Table 3.6: Analysis of the interaction 'institute * role'.

Vignette 4 describes the instance of a student who, arguing that common knowledge is not subject to acknowledgement, reused the same structure of another work without any referencing, while reusing some of the wording and editing other parts of the text via changes, additions and deletions. It is usually of common agreement that expressing common knowledge in a different wording is not plagiarism; however, borrowing works and ideas that are not common knowledge, but other people's original work is, particularly if the wording is, for the most part and despite minor edits, also borrowed. The borrowing becomes more serious if it is considered that the student knowingly reused other people's original works and ideas as if they were his own, upon claiming that he did not use the text.

Although all participants score the case as being clearly an instance of plagiarism, Porto students are those scoring the lowest with a mean value of 4.29, when compared to 4.86 in the case of their Aston colleagues. Porto lecturers/tutors are those scoring the highest, with a mean of 4.98, when compared to 4.60 of their Aston colleagues. Interestingly, Aston lecturers/tutors score lower than their students, contributing the most to this near-interaction. The profile plots of vignette 4, presented in figure 3.11 below, illustrate this interaction.

The perceptions of the participants partly reflect this pattern, as well as the dilem-

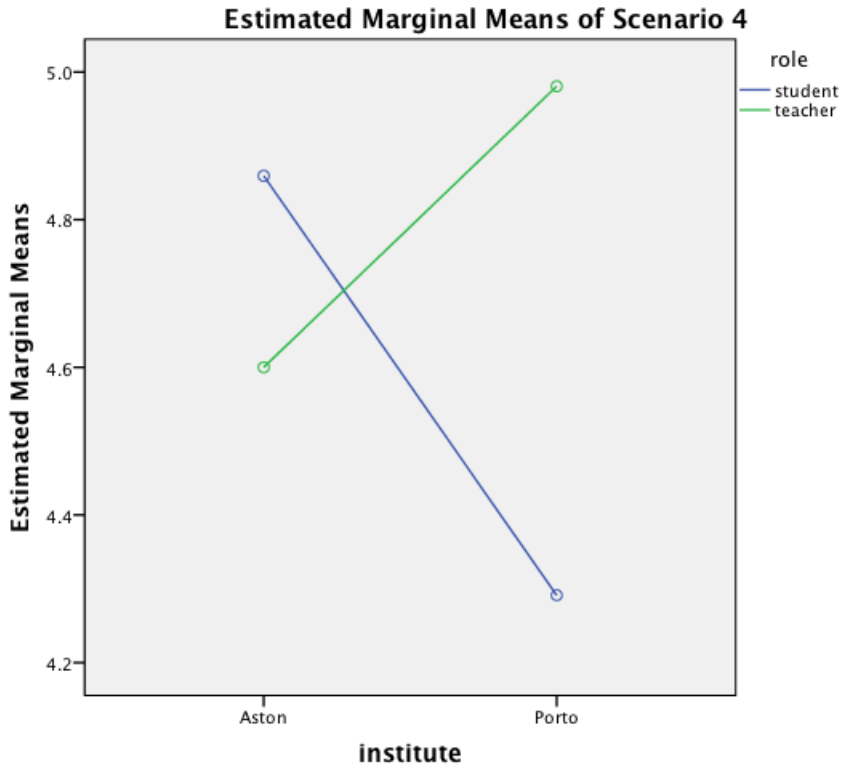

Figure 3.11: Profile plots for vignette 4.

mas involved in this scenario, although they do not always account indisputably for the overall scores. On the contrary, despite scoring differently, participants tend to share some justifications for their perceptions. The first main dilemma is establishing to what extent an idea is general, therefore not entitled to scientific protection and not requiring acknowledgement. If this perspective is the one to take precedence over others, then the student explanation suffices and he is not responsible for plagiarising. Although this possibility is admitted by Porto lecturers/tutors (PTT-49), it is favoured by their Aston colleagues, who like Porto students consider that this can represent more a case of 'poor referencing' (UKT-10) or lack of academic writing training (PTS-228), which led the student into assuming that no acknowledgement was required. In this sense, this instance of plagiarism would more appropriately represent a case of lack of academic skills than misconduct, the judgement of which can also be dependent on the amount of borrowing (UKT-2).

If, on the contrary, it is considered that works and ideas, besides words, are subject to acknowledgement, then the student is responsible for using someone else's work, and hence for plagiarism. Although this is a possibility that Aston lecturers/tutors admit (UKT-6), it is the perception favoured by Porto lecturers/tutors, who tend to consider

that the structure created by someone else was retained (PTT-5, PTT-11) without acknowledgement, and that by so doing the student 'abused' (PTT-38) someone else's work. Porto students also agree that ideas are subject to acknowledgement (PTS-8, PTS-34, PTS-125, PTS-153, PTS-364), in particular to avoid inappropriate appropriation of someone else's intellectual property (PTS-116).

In addition, moral principles are raised in respect of academic writing; the nature of academic work requires that its author/writer acknowledges work upon which they have built to produce their own work, and this production demands some effort on the part of the student. Doing the work the 'easy way', without making that effort, by reusing the work of others without acknowledgement reveals a moral wrong on the part of the student, and in a sense an attempt at deceiving their audience. It is therefore demonstrative of the student's intention. Interestingly, on the punitive side Porto lecturers/tutors (PTT-27), like Aston students, are more unwilling to dismiss this lack of effort, 'a case of the lazies' (UKS-99) that, Porto students admit, is identical to simply revising (PTS-158), whereas Aston lecturers/tutors tend to value in particular the fact the the student was aware of his acts, and knowingly manipulated his source (UKT-12); by arguing general knowledge when he did not even have that knowledge before reading his source, the student was 'desingenuous' (UKT-17). Aston and Porto students also considered the student morally wrong, and argue that the student's acts are identical to 'stealing' (UKS-113), especially when deliberately attempting to pass it off as one's own (UKS-31, UKS-187), and hence 'not right' (PTS-286).

These perceptions demonstrate that, despite scoring a considerable difference, all groups of participants in general consider that the student has plagiarised, either because he failed to acknowledge previous work and ideas, because he failed to observe the academic writing rules of referencing, which should always be observed, or because his misbehaviour is a result of his lack of effort or deliberate attempt to deceive his readers. The latter, which can be seen most clearly in the scores of the Porto lecturers/tutors, demonstrates that the violation of the principle is the one that Porto lecturers/tutors are less willing to tolerate. Aston students, scoring second and only slightly lower than Porto lecturers/tutors, also tend to be willing to punish the student, both for breaking the rules of academic writing, and particularly for his immoral and deliberate attempt to deceive others. Aston lecturers/tutors, despite considering that the student may have acted with intention, prefer, in this case, to weigh the amount of borrowing, and consider the fact that the violation of the rules is a result of lack of

skills. Porto students, scoring the lowest of the four groups, clearly argue that taking someone else's works and ideas is plagiarism, and that by acting as he did the student acted immorally, but consider that extenuating circumstances need to be taken into account, including understanding what is general knowledge and lack of training.

These results therefore suggest that the existence of these moral dilemmas can account, at least in part, for the interaction, but drawing this conclusion from the analysis of one single case, and more importantly considering that this instance is not even statistically significant, is pure speculation. In order to investigate this further, ideally further interactions should be analysed. The analysis of vignette 8 can therefore contribute to understanding this phenomenon, especially considering that the differences observed are significant.

The profile plots of vignette 8, presented in in figure 3.12 below, illustrate this interaction.

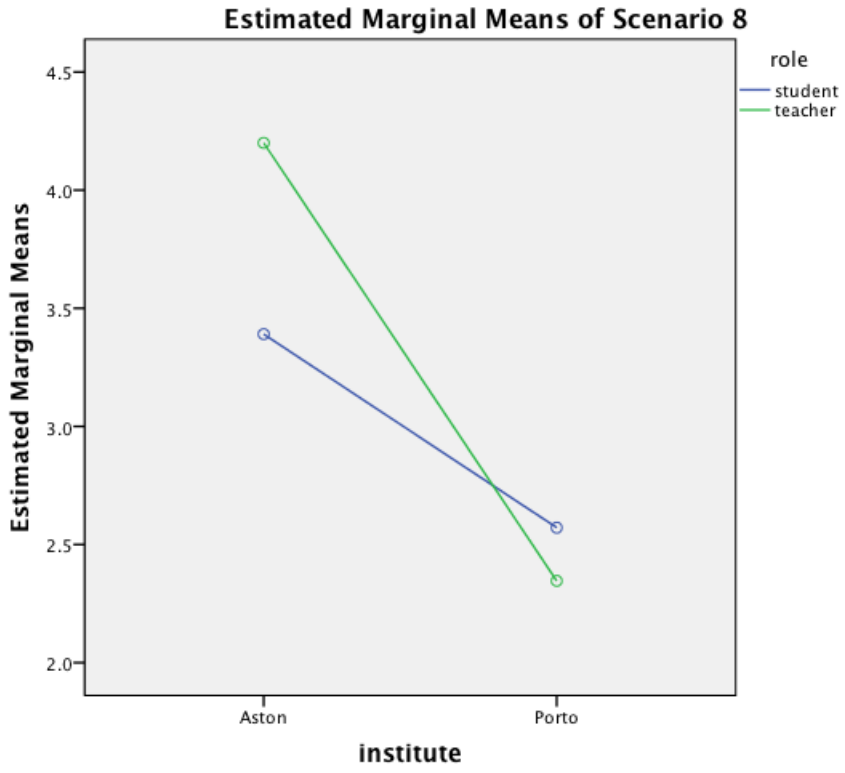

Figure 3.12: Profile plots for vignette 8.

As this figure suggests, the perceptions of each group of participants (lecturers/tutors and students) that the character in the vignette has plagiarised are different in each of the institutes 'Aston' and 'Porto': the blue line, indicating the student scores, falls slightly towards the Porto students area, and is crossed by the green line indicating the

lecturer/tutor scores, which fall abruptly. The statistically significant results related to the interaction 'institute * role' suggest therefore a change of effect resulting from the 'role' variable behaving differently in each condition of the variable 'institute', indicating that Porto lecturers/tutors are the ones that disagree the most with the accusations of plagiarism, and can be found to behave in an opposite way to their Aston colleagues. Therefore, whereas Aston lecturers/tutors think (even if the score is not extremely high) that the student should be penalised for plagiarism, their Porto colleagues have the opposite perception.

This scenario presents the case of the student who, having thoroughly referenced all her sources and provided a complete reference list, was accused by her supervisor of having plagiarised, on the grounds that she had not followed the school referencing style and conventions. Student participants score higher in Aston than in Porto, but although the upper bound of Aston students is slightly above the 'neither agree, nor disagree', neutral 3.50, the means of 3.39 and 2.57 in Aston and Porto, respectively, suggest that both groups of students tend to consider the instance as not representing plagiarism. The lecturers/tutors from Aston score the highest, with a mean of 4.20; conversely, the lecturers/tutors from Porto score the lowest, with a mean of 2.35 – even lower than the Porto students' score. Although the lower bound for Aston lecturers/tutors is 3.45 (i.e. below the neutral 3.50), the upper bound is very high (4.95) when compared even to the second highest score (Aston students, at 3.63), which influences the overall means for the group.

Since the lecturers/tutors are those scoring on opposite ends, it is relevant to find possible explanations for these scores in the justifications provided by the respondents, focusing in particular on these two groups. Aston educators tend to foreground the fact that students are provided with sufficient guidelines and instructions, so they should not only know how to write academically, but also know what plagiarism is and how to avoid it – and moreover, since those guidelines usually specify a particular referencing style 'for a reason' (UKT-1), they must 'be followed to the letter' (UKT-12). Other referencing styles and conventions can admittedly be used, but before adopting alternative referencing conventions, the student should discuss them with the supervisor first, and provide a proper justification for their choice (UKT-1); otherwise, they risk being 'guilty of plagiarism by carelessness' (UKT-11). At the same time, this perspective backgrounds the distinction between using different referencing conventions and not referencing at all. On the contrary, Porto educators not only make a distinction

between the two, but more importantly they foreground the fact that what is crucial is that the sources are referenced correctly and thoroughly (PTT-8, PTT-48, PTS-271), regardless of the referencing style and conventions used; choosing one or other set of referencing conventions is simply a formal matter of formatting (PTT-1), an administrative (PTT-5), methodological or technical issue that reflects, as a mistake (PTT-8, PTT-48) or lack of precision, in the non-observation of the referencing system used (PTT-35, PTT-37) – and consequently on the mark (PTT-52, PTS-174, PTS-301) – but not in judgements of plagiarism (PTT-2, PTT-23, PTS-62, PTS-116, PTS-153, PTS-362) involving disciplinary action. In particular, the argument for these claims is that following an alternative set of conventions neither harms, nor challenges the original authorship and copyright (PTS-252), while still allowing the easy verification of the sources (PTT-38, PTS-180); on the contrary, an instance of plagiarism consists of stealing someone else's ideas, usually as a result of obfuscating the original authorship or preventing their verification (PTS-328). In this sense, plagiarism represents the violation of a universal, moral principle of making clear the work of others and the efforts one has contributed, by unscrupulous people who attempt to take an unfair advantage of someone else's work (PTT-27), more than a violation of a set of individual rules enforced individually by a certain institution, and is therefore subject to different penalties (PTT-49).

This perception of plagiarism by Porto lecturers/tutors as dichotomic contrasts with their Aston colleagues', to whom plagiarism is perceived in a continuum of severity, giving them room to conceive of the existence of different degrees of plagiarism that consider the existence of extenuating circumstances. In this particular instance, two circumstances may have contributed to the student failing to observe the established conventions; on the one hand, she did clearly not plagiarise intentionally, but more as a result of 'a problem of understanding the conventions' (UKT-18); on the other hand, her failure to observe the conventions may also reflect her lack of understanding of the rules, for which the supervisor is also responsible, owing to the fact that he did not discuss the guidelines with her to let her 'know where she was' (UKT-19) – a perception that their Porto colleagues also share (PTT-45, PTT-10, PTS-226). The corrective measures suggested sporadically run in parallel with the different degrees of severity. Interestingly, in order to accommodate those extenuating circumstances the arguments that there 'should be a warning, followed by a re-submission' (UKT-5), rather than an outright submission of the student to disciplinary action, are more

consistent with a case of poor academic writing than plagiarism.

Although the perceptions of lecturers/tutors at Aston and Porto vary, with the latter considering that the student did not plagiarise at all, whereas the former, albeit considering that the student did not act intentionally, argues that she is still responsible for plagiarising, there is at least some agreement between both groups that she did not act with ill-faith (PTT-32). The differing perceptions raise once again the tension between the violation of the moral principle of acknowledgement of authorship, as opposed to the inobservance of purely technical referencing rules and conventions, that rather than being universal tend to vary according to the institution, discipline or text genre. The latter does not imply any moral wrongdoing on the part of the student, since unless the student intentionally tried to obfuscate her sources and take someone else's work as her own, the moral principle remained intact; the former does. For Aston lecturers/tutors, this implies charging the student with plagiarism in either case, even if accepting the extenuating circumstances surrounding the instance; for their Porto colleagues, on the contrary, those charges should be dismissed and the student should, at most, be asked to correct the referencing and resubmit, or – even more plausible – accept the referencing system as is, as long as it is used consistently throughout the dissertation.

3.4 Understanding the Reasons for the Differing Perceptions

The analysis of the two main effects 'institute' and 'role' above indicate that, in the cases where a statistically significant difference has been found, participants from Aston score consistently higher than participants from Porto, and lecturers/tutors score consistently higher than students. This indicates that Aston participants tend to agree with the accusations of plagiarism more than Porto participants, and students are less willing to penalise the students for plagiarism than lecturers/tutors. The analysis of the interaction between 'institute' and 'role' demonstrates that exception in two different cases, in vignettes 4 and 8, but only the latter is statistically significant. In this case, the variable 'role' shows a different behaviour in each condition ('Aston' and 'Porto') of the other variable, 'institute', so that exceptionally lecturers/tutors from Porto score lower than the students, hence disagreeing the most with the accusations of plagiarism. Interestingly, vignette 4, despite not being statistically significant, shows an identical behaviour, but on the 'Aston' condition of the variable 'institute', whereby lecturers/tutors score lower than the students from the same institution.

Several reasons can explain the different perceptions, both at the level of the main effects, as well as at the level of the interaction. One is that UK higher education institutions have been marketed internationality for much longer than Portugal. As the high numbers of overseas students demonstrate, UK universities attract people from all over the World, in no small part due to the quality of their education. UK universities therefore need to maintain their credibility, which includes strict controls over academic dishonesty.

At the same time, the UK higher education institutions have been systematically handling cases of plagiarism for several decades, whereas in Portugal institutions only rarely make reference in their documents, including their websites, to academic integrity in general, and to plagiarism in particular. UK institutions have carefully drafted policies, and clearly defined procedures to tackle those cases, and it is the lecturer/tutor's responsibility to identify and denounce instances of plagiarism, and pass the cases on to a 'Disciplinary Board' that will make a decision on behalf of the University. In Portugal, the lecturer/tutor is responsible for identifying the instances of plagiarism, interpreting and assessing their degree of severity, and deciding on the punitive measures to be adopted, usually without the explicit support of the institution. Undoubtedly, this approach puts more pressure on lecturers/tutors, when compared to the ordinary procedure in the UK, necessarily influencing their confidence and degree of certainty in passing final judgements, and as a consequence implying a second thought before making final decisions.

The weight of responsibility also reflects on the students. Students in the UK are compelled to know what plagiarism is, and how to avoid it, from an early stage (Coulthard and Johnson, 2007). On the one hand, upon institutional request, they are required to acknowledge and accept codes of practice discussing plagiarism and academic integrity. On the other hand, UK students will know from personal experience or by 'word-of-mouth' of cases where colleagues have been severely punished for borrowing other people's work without acknowledgement. Additionally, if Aston University is taken as an example, UK students are required to confirm upon submission of written work that they have observed the required regulations, including those related to plagiarism. Conversely, Portuguese students do not tend to be offered a systematic approach, and are rarely familiar with the topic. At most, they hear about it in the news, due to high profile cases that are publicly discussed in the media, or from (some of) their lecturers/tutors, who are usually the most conscious of the need to train ethical in-

dividuals. Moreover, some reports[3] state that pupils from a young age are encouraged by their teachers to copy and paste information from the Internet. By concentrating on the fact that this exercise helps the students find and compile online information, their teachers seem oblivious to the fact that they are promoting a writing practice that, albeit eventually considered intolerable at a later stage, may pass unnoticed up to University level. As they are faced with different standards of original academic writing, UK participants are potentially more conscious of the problem when compared to Portuguese participants, even more demanding with themselves and their colleagues, and consequently more unwilling to tolerate their peers' academic dishonesty and unfair advantage.

In most of the vignettes there is a tension between the academic wrong, i.e. the violation of the pragmatic rule of referencing in academic writing, and the moral wrong, i.e. the violation of the moral, universal principles of authorship, truthfulness and fairness towards others. But this reflects in particular in vignettes 4 and 8, and is a plausible explanation for the interaction observed. In the case of vignette 4, the student failed to acknowledge his sources on the grounds that what he borrowed was general knowledge; in the other case (vignette 8), the student acknowledged all sources properly, but used a referencing style that was different from the one adopted by her school. These two cases are especially illustrative of this struggle, in that, for Portuguese participants, an instance of plagiarism presupposes that the plagiarist reused materials without acknowledgment intentionally, or at least knowingly, whereas for Aston participants the violation of the pragmatic rules suffices for sustaining the accusations. As the reversed positions of the two roles in the Aston condition demonstrate, that is particularly true for lecturers/tutors, which can explain specifically why Aston participants tend to grade plagiarism more in a continuum of severity. The survey results demonstrate that UK participants perceive that the cases of unacknowledged borrowing, in general, should be investigated further. Depending on the degree of severity, a Disciplinary Board then determines the seriousness of the offence, considering aspects of volume, circumstances and even intention. The degree of intention is not, however, the primary focus of the offence. But the existence of an intention behind the challenged act leads to a more unanimous assessment of plagiarism among the different participants, of which vignettes 1 and 8 are illustrative examples. Unsurprisingly, students and lecturers/tutors tend to agree more in cases where the (bad) intentions are

[3]E.g. *Miúdos Seguros na Net* – `http://www.miudossegurosna.net/artigos/2008-05-16.html`

clear. Intention is the element that Porto participants seem to favour to the detriment of other circumstances, such as the volume of borrowing. As a consequence, they are more keen on assessing the plagiarists dichotomically as guilty or not guilty. In addition, in cases where the discussion revolves around malpractice vs. lack of academic achievement, as is the case of vignette 12, the participants do not tend to dismiss the former in favour of the latter.

Surprisingly, despite the observed differences being sometimes statistically significant, in most cases all participants are consistently in agreement, sharing identical ideas and comments. The difference is on the degree of severity that they attach to each instance. In this case, Porto participants, as discussed above, tend to be less punitive than their Aston colleagues; but their perceptions are not, in general, diametrically opposed, so that even in Portugal plagiarism involves disciplinary action in the mind of the respondents. It is then relevant to investigate other possible reasons for those differences. One possible explanation is the tension holding between the violation of the rule, as compared to the violation of the principle; the other is the theories around the punitive turn. These are discussed in the following subsections.

3.4.1 Violation of the Rule and Violation of the Principle

Research on plagiarism over time, as discussed in chapter 1, has demonstrated if anything that there is not a single agreement on what constitutes plagiarism, and more importantly on how potential cases of plagiarism should be handled. On the one hand, considerations that academic plagiarism may reveal a lack of academic writing skills, rather than represent academic misconduct (Howard, 1995; Pecorari, 2008) pointed out that students in particular do not necessarily see plagiarism in much the same way as their lecturer/tutors. Additionally, there can be a mismatch between the expectations that lecturer/tutors have of how students write academically, and in particular of how they cite their sources, in that the referencing style that students learnt to use may differ from the one that their lecturer/tutors are ready to accept. The discussion of the survey results presented in the previous section revealed that lecturers/tutors and students tend to see plagiarism differently, and that difference is often significant. In particular, students tend to score their agreement with the charges of plagiarism lower than lecturers/tutors, and their comments suggest that they are more willing to accept and understand the students' faults that lead them into plagiarising. However, when faced with hypothetical real cases, their opinion is not as diametrically opposed

to that of the lecturers/tutors, as Howard (1995) and Pecorari (2008) might sometimes suggest.

On the other hand, people around the world have been adopting different academic writing practices, involving differing referencing styles and, perhaps even more relevant, differing writing conventions. In this sense, intercultural studies (Scollon, 1994, 1995) have argued for the need to fight against the hegemonic referencing tradition of the West, frequently promoted by international agreements and regulations, by increasing the tolerance towards cultural difference. Nearly twenty years later, Bloch (2012) claimed that there is now clearly a concern with plagiarism in the East, although plagiarism, like originality, may be defined differently from the West. This oversimplified East/West dichotomy risks erring on the side of caution, and leaves aside important terms of the plagiarism equation. In particular, by discussing this problem dualistically, rather than admitting that, even in the 'West', there are multiple approaches to academic writing, originality and plagiarism that make proper acknowledgment easy to understand but difficult to meet, these perspectives dismiss the fact that subtle, non-obvious deviations to a norm can sometimes take more time and effort to understand, explain and justify than disparate systems. If one may raise some problems in understanding the 'Western' moral principle of ownership of works and ideas, the other presupposes understanding particular rules and conventions whose fine-grained detail is difficult to master, even by skilled students. As one of the survey respondents indicated, the EndNote bibliography management software alone includes over 5,000 referencing styles from which to chose – not to mention the customisation possibilities that it offers – to the total confusion of any inexperienced writer. The problem goes therefore farther than understanding the East/West dichotomy.

At the same time, individual institutions such as Universities defend their own right to establish and define beforehand, via statements, regulations and codes of practice, the rules of a game that prospective students agree to play. A good example of this strategy is described by Coulthard and Johnson (2007) in relation to the University of Birmingham, whose students were asked to sign a statement declaring that they knew what plagiarism was, and that they agreed not to plagiarise. Subsequent violations of such statements, regulations and codes of practice imply a violation of the rules, so that linguistic evidence that two – or more – texts have not been produced independently is the only proof required to charge a student with plagiarism, and subsequently apply the corresponding penalties (usually of differing degrees of severity, depending

on a quantitative assessment of the extent of borrowing). Determining the penalties based on the violation of the rules, despite straightforwardly just and legally supported by contractual premises is, however, bound to disregard the moral elements of principle (e.g. whether the instance of plagiarism was intentional, or whether the plagiarist simply lacked the skills to do better), and raises issues of fairness if mitigating, circumstantial elements are to be taken into account. There does not seem to be a general, universal agreement on this matter. The survey results above reveal that some perceptions would favour an interpretation of the principles underlying the instances of plagiarism, and assess their corresponding degree of severity based on the distinction between right and wrong, good and bad, intentional and unintentional, knowing or inadvertent, and so on; others, claiming that intentions can hardly be guessed, would avoid making judgements or speculating on the reasons and the morals behind those instances of plagiarism, and restrict their assessment to the compliance or otherwise with the established and agreed rules.

These paradoxical perceptions of the principle vs the rule reignite the discussion of whether plagiarism is inherently a moral/ethical issue, or on the contrary a legal matter, or both, and whether any one of these should be foregrounded at the expense of the other. In one case not complying with the rules and conventions that one has accepted, and by which one has formally and officially agreed to abide could represent several different things, some of which are more morally reproachable than others. At its simplest, a violation of this contract could mean that one has read and understood the rules, and having agreed to abide by them, then decided to violate them for one's own benefit. This would mean knowingly breaking the rule, hence constituting a violation of the principle behind it, as well as the violation of the rule itself. Violations of this type can be more or less evident, or otherwise completely imperceptible; in cases where the plagiarist intentionally or – to be less extreme – knowingly obfuscates the violation of the rule, that apparent non-violation would still encompass a violation of the principle. A more complex scenario is where a student has read and assumed to have understood the regulations, and therefore agreed to comply with them, but then fails to follow the agreed conventions due to lack of knowledge. In cases of this type, providing an account or a justification for that violation and demonstrating that it was not intentional would not make it less of a violation of the rule, but could imply that the moral principle behind the rule remained (almost) intact.

A violation of the rule does not, under these circumstances, imply a violation of the

principle – even if that lack of knowledge can bring along additional moral implications such as, for example, the inability to interpret a text; or a violation of the rules and conventions owing to the fact that the student signed and agreed to abide by a contract that s/he has *not* actually read. In this latter case, the violation of the rule would also mean a violation of a moral principle, but one of a different type, i.e. the principle of truth: not actually reading the contract and stating that one has read it, implies lying when signing a statement confirming that one was aware of the regulations. However, the violation of general moral principles like this span beyond the ethical principles underlying academic referencing, and therefore are outside the scope of this research. The fact that so much discussion is being held demonstrates, if anything, a lack of consensus on this matter, and indicates that the concept of plagiarism is far from being universal. Interestingly, the survey results suggest that Porto participants tend to define plagiarism in terms of the violation of the moral principle more than the Aston participants, for whom the definition of plagiarism depends simply on whether an established rule has been broken. Consequently, Porto participants in general are more willing to give the benefit of the doubt, and rely on the intentions behind the act to pass their judgement than Aston participants, who albeit not disregarding the relevance of the violation of the principle, hesitate less in concluding that breaking the established rules is sufficient to charge a student with plagiarism.

Attributing these differences to the fact that UK participants have a stronger and longer tradition of academic rules, regulations and writing conventions, which allows them to determine precisely how to (re)act when faced with certain circumstances is a possible justification for the statistically significant differences found. It is also plausible that those differences, and particularly the fact that the Portuguese participants apparently attach more importance to the intentions behind the act of improper borrowing, owe to the long-standing tradition of Catholic morals.

However, legal reasons can also account for these cross-cultural differences. Portugal and the UK have two different legal systems. The former, within the Civil Law tradition, builds upon legal normative codes that frame social activities. Conversely, the UK legal system is uncodified, lacking even a criminal code. As a result, rules and regulations set over time are those that determine the social norm. A comparative analysis of these legal systems can help understand these different views.

3.4.2 Common Law and Civil Law

The different legal systems of England (and Wales) and Portugal help explain the importance attached to the violation of the rule and the violation of the principle, respectively. Dainow (1966) provides a clear explanation of the differences between the two legal systems. In countries within the Common Law tradition, judicial decisions provide the basis of the law. Common Law is a generally uncodified legal system, which consists of applying rules that were generalised over time, as a result of judicial decisions that were previously made, and whose application to all subsequent identical cases is expected. By adopting a doctrine of 'precedent' – 'earlier decision' (Dainow, 1966: 425) –, Common Law is therefore based on particular rules set forth by judicial decisions resulting from particular situations involving particular details. Consequently, since the expectation that the same problems demand the same results, rules set by previous identical cases are applied. Previous decisions therefore determine the outcome of a similar case, whereas new rules are created when new problems raise new cases. However, the judge is offered alternatives to the precedent if the situation at stake is similar but not identical. In this case, they can apply the same solution; distinguish the situation, by applying it 'limited to the specific fact situation that it controlled' (Dainow, 1966: 425), or overrule the precedent, if they believe this precedent to be wrong. Dainow (1966) strongly argues that, by applying the rule and following a precedent, on the one hand, while allowing room for distinguishing or overruling that precedent, on the other, guarantees both the stability and the continuity of the law, as well as the flexibility and adjustment required by new situations.

Conversely, countries within the Civil Law tradition, like Portugal, have comprehensive, codified sets of laws (legal codes) that are prepared by legislative bodies, and that attempt to address as many legal situations as possible that can be brought before a court of law, while simultaneously describing the corresponding legal procedure and the respective punishment. Their comprehensiveness, as Dainow (1966: 424) argues, lies 'not in the details but in the principles', i.e. by focusing on the principles regulating social activity, rather than on the specifics of that activity, these codes allow an application of the law to new legal situations that may arise. Moreover, these general principles demand a liberal, comprehensive interpretation of the law by the judge, whose decisions on new, particular cases always start with a look on the broad, general, legal principles. In this case, judgements are not made based on general rules, and the rule of precedent – though sometimes considered, e.g. in cases of

jurisprudence – is not binding.

These differences between the two legal systems therefore seem to be reflected in the overall perceptions that Portuguese and British participants have of plagiarism, with the former prioritising the principle, and the latter foregrounding the rule. As Dainow concludes when comparing both systems, whereas 'the common law judge creates the rules of the law' (Dainow, 1966: 434), in the Civil Law, the judge seeks justification for his/her decisions in the general principles and in the written text of the law. This is not meant to indicate an absolute absence of legislation from common law systems, or of judicial decisions from Civil Law countries, rather the contrary. The statutes enforced by common law countries are an example of legislation, just like making jurisprudence after a considerable number of cases are decided in the same way is an example of how judges in Civil Law countries use judicial decisions. He therefore concludes that both legislative decisions and judicial decisions exist in both systems. However, a different relative importance is given in each system, with Civil Law favouring legislative decisions, and Common Law giving preference to judicial decisions. Ultimately, the preference given by Civil Law and Common Law to the principle and the rule, respectively, is unquestionable; legislation in common law countries is still formulated considering a set of rules, and conversely, even when making a judicial decision based on jurisprudence, judges in civil law countries still need to make principle-based decisions.

This comparative analysis of the two legal systems can help explain why the Portuguese and British participants give preference to the principle or to the rule, respectively, in detriment to the other. But it does not entirely explain why, even when agreeing on the outcome, British participants, scoring higher, are less reluctant to punish the suspect plagiarist. Research conducted in the field of criminology in recent years can help explain these survey results in terms of the punitive turn. The following section attempts to provide some explanations for the disparate perceptions of plagiarism between Aston and Porto participants, by discussing plagiarism in relation to theories of the *punitive turn*.

3.4.3 *The Punitive Turn*

The differences observed in the results between students and lecturers/tutors are not surprising. But it is surprising that participants from Aston overall are consistently less willing to tolerate referencing and plagiarism-related misbehaviours than participants

from Porto. These results therefore suggest that the different scores may be due to the contextual and cultural differences existing between the two countries, in particular the quantity and quality of information available on academic writing and avoidance of plagiarism and the influence of Catholic morals, in the case of Portugal, and which may underlie the predisposition to judge and punish more systematically or otherwise. As discussed in section 1.3.2, the UK has traditionally implemented detailed policies and definitions of plagiarism, together with student induction sessions and training on academic integrity and plagiarism. And although students may or may not attend these sessions, as Carroll and Appleton (2001) claim, both the lecturers/tutors and those students who do attend consider that, by failing to comply with their obligation, those students cannot use ignorance of the rules in their own defence. Portuguese participants, by the contrary, seem more willing to give the benefit of the doubt and accept the lack of skills, competence and knowledge of the others, and possibly accept the flaws of a system that is neither clear, nor systematic, in its approach to plagiarism detection. It can therefore be speculated that, owing to some or all of these circumstances, Aston participants hesitate less in passing judgements of plagiarism, whereas Portuguese participants tend to be more permissive.

On the contrary, research conducted in the field of the *punitive turn* throughout the world can help explain why Porto participants are more hesitant in passing judgements of plagiarism than Aston participants.

In his article, Muncie (2008) discusses the evolution of a 'resurgent authoritarianism', a punitive turn in Western Europe and the USA in recent years, and how this evolution has affected juvenile justice. Citing Wacquant (1999) and Garland (2001), he discusses how Western European societies, in order to attract international investment, came to value economic competitiveness at the expense of the social state, adopting 'similar economic, social *and* criminal justice policies' (Muncie, 2008: 117), consequently importing a 'punitive turn' from the USA, and implementing it as a means to control the economically excluded – often as an end in itself, and frequently in violation of (young) people's rights. In particular, he discusses how in the UK needs-based interventions came to be replaced with a tighter penal control, increased sanctioning, mass imprisonment, interventions based on risk assessment and zero tolerance, and how this came to encourage a punitive mentality that eventually impacted daily social relations – all in the name of a 'new correctionalism' (Muncie, 2008: 108). As a consequence of these changes, juvenile protection and support gave way to politically

legitimised principles of retribution, responsibility and accountability, and repressive measures came to be favoured to the detriment of educational actions.

In Western Europe, this is particularly marked in the UK, where the age of criminal responsibility is one of the lowest (8 in Scotland, 10 in England and Wales, when compared to 16 in Portugal), sentence lengths have been increased, non-criminal activities have been treated as anti-social behaviour, and with a record for imprisonment of 12 year-olds. Moreover, Muncie (2008) reports England and Wales as possessing the second highest rate of incarceration in the world, after the USA, and the highest in Western Europe – 148 per 100,000, compared to 121 per 100,000 in Portugal. Tonry (2001) argues that the best explanations for severe punishment – or tolerance – in Western (European) countries can be found at a national or cultural level, and Muncie (2008) critically adds that assuming that North-American trends in penal policy can be transferred internationally, as academics and policy makers throughout Anglophone countries have been doing, is wrong, as the 'children first' policy of Wales, *versus* the 'offender first' policy of England demonstrates. This can explain, at least in part, why survey respondents in the UK tend to be more repressive than those in Portugal.

The punitive mentality tends to be based upon the perceived increase in violence, with the corresponding increase in the number of incarcerated, which in turn contributes to increasing the control and as a consequence legitimising the cultures of control. This is a perspective that Estrada (2001) challenges, by arguing that this perceived increase is simply 'an ideological shift' (Estrada, 2001: 647) fuelled by the media attention attached to it, and that affects the sensitivity of the public and the propensity that the general public has to report crimes, and which reflects changes in perceptions and reactions to violence, including the ease with which technological developments allowed crimes to be reported almost immediately.

If we consider, as is often the case, that plagiarism is a crime, it is then relevant to ask whether plagiarism has actually increased in recent years, or whether the perceived increase in plagiarism as a criminal activity is simply a result of a shift from an informal and sporadic to a formal and systematic detection – in which case this perceived rise would be the result of a different, simplified approach to detection routines and procedures.

These theories that explain the (until now) more marked punitive turn in the UK, when compared to Portugal, contribute, if not to justify, at least to understand the significantly different and more lenient perceptions that Porto students and lecturers/tutors

have, when compared to their Aston counterparts. Additionally, the shift from a social to an individual state that is associated with that punitive turn helps understand the importance attached by Aston respondents to individual responsibility and accountability of the students, in contrast with the (social) responsibility that Porto respondents attach to the universities and the role of the supervisor as agents of education.

The different perceptions between Aston and Porto participants can also be a reflection of approaches to authorship. Countries that attach more weight to the individual than to the social state are also expected to attach more value to individual property, including authorship. In this respect, it seems paradoxical that, the USA, for example, did not ratify the Berne Convention for the Protection of Literary and Artistic Works until the late 1980s. This paradox can however be explained by the fact that the ratification of the Convention required changes to the US copyright law, especially in terms of moral rights and requirements for the registration of copyright (Pereira, 2003). It was feared that these changes would make the property of US citizens more vulnerable in the international market.

3.5 Chapter Summary

This chapter investigated the different perceptions that students and lecturers/tutors, in different countries, have of plagiarism. A survey was conducted in Portugal and the UK, with students and lecturers/tutors. The statistical analysis of the results demonstrated that the differences in the perceptions of plagiarism are noticeable. Participants in the UK tend to score higher than participants in Portugal, demonstrating that they agree more with the charges of plagiarism than their Portuguese colleagues. Additionally, students tend to score lower than lecturers/tutors, even if in general they also agree with the accusations of plagiarism. The interpretation of the results also suggested that Portuguese participants tend to base their judgement on the plagiarist's intention, more than their UK colleagues. However, the outcome of each case as plagiarism or otherwise is shared in most of the vignettes by all or most of the groups of participants.

A possible explanation for this apparently less lenient attitude of the UK participants was offered. The chapter concluded with a discussion of how the tendency to punish is intimately related to the perceptions of individual responsibility and accountability, as well as to the availability of preventive measures and procedures. The case for the forensic linguist is also demonstrated. Specifically, some participants point out

the need for a closer analysis and comparison of the challenged passages before determining an outcome that can have serious disciplinary implications. This will be discussed next, in chapter 4.

Chapter 4

Intention and Intentionality

4.1 Introduction

This chapter discusses the concepts of 'intention' and 'intentionality' in relation to plagiarism. It challenges two commonly held assumptions: on the one hand, the claims that plagiarism can be detected, including by linguists and computer scientists, without any concerns for the potential plagiarist's ultimate purposes or circumstances, i.e. their intentions; and, on the other hand, the assumption that plagiarism is only subject to punishment if and when the plagiarist is on balance proven to have borrowed from other sources with no mitigating circumstances.

Based on previous research findings into the study of plagiarism, as explained in section 1.3, three hypotheses are discussed. The first hypothesis is that plagiarism detection can be run independently of any concerns with the potential plagiarist's intention. This hypothesis is based on the research into plagiarism detection, such as the work of Coulthard and Johnson (2007) on forensic linguistics and plagiarism, in which the authors found that from a linguistic perspective plagiarism detection should take into account only the texts being analysed. Taking academic plagiarism as an example, this research found that having the students sign a statement acknowledging that they know what plagiarism is, that they accept to abide by the academic rules and regulations set forth by the Institution, and that they thereby commit not to plagiarise other works, allows the investigators (such as lecturers/tutors, as well as forensic linguists) to concentrate on the type of work that is competent to detect the borrowing: the linguistic analysis of the text, with disregard for the plagiarist's intention.

The second hypothesis is based mainly on research into teaching and learning and intercultural studies (e.g. Anderson (1998) and Scollon (1995)), and argues that as

plagiarism is a severe accusation, its judgements cannot be passed regardless of the corresponding circumstances. This research found that judgments of plagiarism so far have been based on a 'one-size-fits-all' approach, whereas approaches to authorship tend to vary across genres and cultures. Not taking into consideration cultural diversity is therefore a form of hegemonic power, an attempt of Western cultures and ideologies to impose on other, non-dominant cultures.

A third hypothesis investigates whether all cases of unacknowledged textual borrowing should be considered plagiarism, or whether there are any mitigating circumstances that make textual borrowing a lesser offence. Additionally, it asks whether it is possible to borrow text without acknowledgement, while making no attempt to deceive the reader(s), and whether plagiarism is necessarily intentional. This hypothesis capitalises on some of the strengths and advantages of both the previous two hypotheses. It is rooted on the assumption that plagiarism cannot be tried based simply on textual identity, neither can it be always discounted for all the possible mitigating circumstances. Additionally, while claiming that plagiarism equates with intention, it is argued that forensic linguistic analysis can provide some insight as to the plagiarist's intention. How this can be demonstrated by the linguistic analysis will be explored in this chapter.

4.2 Defining Intention and Intentionality

The assessment of instances of plagiarism has inevitably been related to considerations of intention, as is demonstrated by the claims that a distinction has to be made between cases where students borrow improperly, but inadvertently, from other texts in a attempt at writing academically, and cases where a student consciously borrows from other texts without acknowledgement in order to pass off the text as his or her own (Anderson, 1998; Howard, 1995; Pecorari, 2008). Similarly, as the findings reported in chapter 3 indicate, it is generally agreed that plagiarism should be punished when intentional, but not when unintentional – or at least, not equally severely.

The need for this distinction is reflected in most academic integrity policies, especially in Anglo-American countries, as was discussed in chapter 1. However, plagiarism detection methods and procedures have not tended to follow the same rationale. This is due, at least in part, to competence limitations; as Coulthard and Johnson (2007) argue, linguists, for instance, are only competent to deal with linguistic plagiarism. As a consequence, detecting the plagiarist's intentions is beyond their expertise,

and the same applies to computer scientists, when designing plagiarism detection software. In both cases, the aim is to detect instances of textual borrowing from other sources, leaving the investigation of intentionality to disciplinary boards. But as Finegan argues, 'it is for the trier of fact to render a verdict' (Finegan, 2009: 274), but it is for the linguistics expert to render an opinion that helps that 'trier of fact' to make a decision. The question is to what extent linguistic evidence can provide clues to intention.

Intention has long been studied across several genres, from linguistics to the law, including psychology. But as has been argued (Eggington, 2008; Shuy, 2001), entering a mind and discovering the exact intentions of that person is an impossible task, especially for linguists. However, although the actual intentions cannot realistically be identified, 'clues to such intentions' (Shuy, 2001: 444) can be determined by means of a linguistic analysis of the topics and response strategies of the speaker – or, it can be argued, of a writer. The fact that Shuy approaches *speakers*, rather than *writers*, cannot be accidental, as the setting he chooses to discuss intention and intentionality is that of a court of law. In this setting, defendants or suspects are interrogated in an attempt to determine their intentions. Nevertheless, considering that it is in the defendants' best interests to protect themselves, the likelihood that the court will have a truthful answer is very small, not to mention that sometimes these intentions are even alien to them (Shuy, 2001). It is therefore necessary to rely on other clues to intentions in an attempt to infer the actual intentions. One of the main clues is the topics introduced by the speaker (e.g. by agreeing or disagreeing with the other speaker's topic, by elaborating on the topic to indicate acceptance/unacceptance, agreement/disagreement of the other person's stance, or by changing the subject), as well as the speaker's responses to the topics of others.

In some respects, the procedure adopted by universities to handle cases of academic plagiarism is not very different. Universities currently tend to investigate cases of plagiarism by analysing any written evidence that they may have, and then interviewing the potential plagiarist by asking them questions. The problem, however, is very similar to the one described by Shuy for crime-related interrogations in courts of law, i.e. 'it is very difficult, if not impossible, to prove that such responses indicate agreement to participate in a crime' (Shuy, 2001: 445). Yet, clues to intention can be found in written, as in spoken discourse. Linguistics and philosophy of language have

been trying to demonstrate, at least since the 1960s[1], that language is explicitly used with a goal in mind, so the speakers have an intention as they speak or write, and use motive to produce the text. These are utterance-level intentions, and have attracted the attention of other genres, including the computer sciences (e.g. Grosz and Sidner (1986)). *Intention*, in this case, equates with meaning, i.e. that which is meant by a person when they make an utterance (Grice, 1969). In other words, this type of intention consists of the *communicative intention* that is realised when the message gets across, and depends on drawing conclusions by the participants.

These are the kind of intentions underlying pragmatics and the Speech Act Theory (Austin, 1962), which considers that an utterer's intention is to induce a response in an audience; 'doing things with words', discourse participants try to accomplish *x* by saying *y*. In many of Grice's examples, however, the utterer is usually seen as wanting or wishing their interlocutor to know what they mean, i.e. to recognise their intentions. As he claims, 'the success of intentions of the kind involved in communication requires those to whom communications or near-communications are addressed to be capable in the circumstances of having certain thoughts and drawing certain conclusions' (Grice, 1969: 158).

Intentions, however, cannot all be assessed on the same terms. Utterance-level intentions – or communicative intentions – are 'public' intentions, which are meant to be recognised by the other discourse participants. And even when such intentions are more or less hidden, discourse analysts (e.g. Fairclough and Wodak (1997); Johnstone (2008)) have shown how to reveal some of these intentions, for example by analysing genre (such as adverts, job applications, informational texts, political speeches or research texts, among others) in terms of their communicative purposes – and judge them on the grounds of whether they fulfil those purposes. On the contrary, the kind of intentions referred to by Shuy, and which define the type of intentions that are most relevant to the present study, are of a 'private', rather than public nature, in that they are meant *not* to be recognised. It is therefore daunting to assume that intentions and the hidden agendas such as the ones unveiled by discourse analysts can be identified in the same way. Firstly, if hidden agendas and the hegemonic workings of power and ideology such as the ones studied by Critical Discourse Analysis (CDA) can be operated unconsciously, or even subconsciously, intentions presuppose, at least in part, a degree of control, and acting consciously. *Acting intentionally* and

[1] See for instance Grice (1969) and Austin (1962).

acting for a reason are not, in this respect, necessarily synonymous.

Despite admitting, therefore, that public and private intentions operate differently, the procedure used by critical discourse analysts to unveil hidden agendas can potentially be used to identify clues to intentions, especially because, like hidden agendas, clues to intentions can be identified by analysing the discourse of the speaker or writer, even when the reasons behind those private intentions are less conscious, in approximately the same way as the workings of power and ideology are hegemonic, in CDA. Alvarez (2009) argues that all things done for a reason are intentional, but that the opposite is not necessarily true, i.e. that not all things done intentionally are done for a reason. 'Doing something for no reason' is doing something intentionally, but with no end or purpose in it. This includes doing things 'because we want to or feel like doing them' (Alvarez, 2009: 300). Doing something intentionally – and for a reason – therefore requires two features: 'some *awareness* of what one is doing, and some degree of *control* over it' (Alvarez, 2009: 294) – i.e., while pursuing a goal. 'Reason', in Alvarez's words, is believing and wanting something.

Intentionality is also a major line of research in Psychology, with a focus on understanding the possible intentions of actors in social interaction and as a means to create impressions. Hart and Albarracín (2010) conducted research oriented to investigating whether formal linguistic features could be used, not to determine another person's intentions, but rather whether they could influence other people's impressions of intentionality. In this particular case, the authors focused on the verbal aspect of such descriptions, namely the imperfective aspect (*was doing*) in comparison with the perfective aspect (*did*). Their hypotheses derive from understanding the imperfective aspect to describe a dynamic, unfolding sequence of events, meant to include more concrete actions. Conversely, they argue, the perfective aspect is used as a complete whole, and hence as being finished. The imperfective, describing the action components in more detail, makes more details available to infer someone else's intention(s). This is also used in studies of narratives, e.g. police interviewing.

Despite all the debate about intentionality across the different disciplines, intentional plagiarism cannot be discussed without considering intention in the light of the law in general, and the criminal law in particular. Criminal law is usually rooted in two basic elements: the *actus reus* and the *mens rea*. The *actus reus* is considered to be the 'guilty act', the 'wrongful deed that comprises the physical components of a crime' (Garner, 2009: 41), thus connoting a physical result. Hence, for example, wishing

someone's death is not a criminal act. In order to be criminalisable, the deed often needs – although not always – to be accompanied by a mental element, the *mens rea*, the 'guilty mind'. Having a 'guilty mind' or a 'criminal mind' is the second element often used in law by the prosecution to 'secure a conviction' (Garner, 2009: 1075).

However, proving the mental element is not always necessary. As Duff (2007) demonstrates, a certain action need not be deemed intentional to be considered punishable in the light of criminal action. Some crimes, called 'basic crimes' (Kugler, 2002), are punishable whether the outcome results from the actor's intention or not, and hence do not necessarily require a 'criminal mind' (i.e. determining the actor's intention) to be punishable. In such instances, proving the actor's 'criminal mind' is only necessary to determine the degree of the actor's moral blameworthiness, and to categorise an offence as an *aggravated* or a *lesser* act. Although the result in itself (the *actus reus*) is punishable, the mental element (the *mens rea*) makes it morally more or less serious, so that an offence is considered to be *aggravated* if intentional, but is a *lesser* one if done knowingly, yet not in order to achieve the result. In this respect, a driver need not hit a pedestrian intentionally to be liable to criminal law; it suffices to demonstrate that the driver was reckless, e.g. by driving carelessly or failing to do the required car maintenance. A parallel can be established between driving and plagiarism. The plagiarist who knowingly and wilfully plagiarises with the *intention* of achieving a certain result, usually a good grade, is responsible for an aggravated offence, whereas the student who knowingly plagiarises, but who does not have an intention to do so, is responsible for a lesser offence. The discussion then shifts from intentionality to responsibility and liability. In the example above, the driver's or the plagiarist's intentionality need not be proven for them to be considered guilty of the offence in any case. The seriousness of the crime, which determines that the resulting offence is an *aggravated* or a *lesser* one, however, depends on the distinction between two additional concepts, those of responsibility and liability. As Duff argues, although liability is grounded on responsibility, 'responsibility does not entail liability' (Duff, 2007: 16). In other words, although a person can be liable to punishment only if they are responsible for it, they can accept responsibility while rejecting liability by offering an 'exculpatory answer' (Duff, 2007: 16).

The concept of intention is not, therefore, dichotomic, since most criminal law systems presuppose a ranking of intention in terms of its seriousness, i.e. whether a person does something with intention, whether that intention is simply oblique, or

whether that person's acts lack intention. In the continuum of seriousness, *intention* is attributed the highest ranking, followed by *oblique intention* and, finally, the *lack of intention*, which is attributed to unintentional actions. This model of intention, which applies to English criminal law (Finnis, 1991) and, with some minor differences, to Portuguese criminal law (Almeida, 2010; Pereira and Lafayette, 2008), describes intention (or doing something with intent) as doing something for which one does not have a lawful justification or excuse that one knows will have those results; acting intentionally thus means acting for a reason. This definition is based on the principle that, while doing something, the doer aims to obtain a certain result. The student who plagiarises in order to obtain a higher grade, for example, does so intentionally, even if such intention is not an end in itself, but rather a means to an end (i.e. the student does not plagiarise to commit plagiarism plain and simple, but as a means to obtain a good grade). Oblique intention, on the other hand, is applied to actions for which intent is only indirect, i.e. actions that are not unintentional, but are not intended either. This type of intention, which is often described as 'culpable recklessness', presupposes that the results derive, as a necessary consequence, from one's actions. That is the case where a student, despite knowing that what he or she is doing may bring along punishment for their actions, do not attempt to prevent that result. Last, unintentional actions are those happening by accident, mistake or lack of foresight. These actions presuppose that the doer did not foresee the consequences of his or her actions. Instances of unintentional plagiarism would therefore include cases where a student is unaware of the referencing conventions, or believes that what he or she is doing is actually right.

Finnis (1991) claims that the concept of intention is often confused in the philosophical, legal and academic arenas, mainly because: (a) the distinction between free-choice (i.e. a deliberate, rationally motivated alternative) and spontaneity (i.e., not acting for a reason, not rationally motivated action), and between rational and subrational motivation often remains undefined; and (b) there is a need for a theory of ethics that clearly describes the norms of human conduct regarding what one intends and does and what one foresees as a side-effect (and therefore is able to control) of what one intends and does. He thus considers that intention is instantiated on 'choice', in the sense of a plan or proposal (even if one conceived at that moment, rather than a premeditated one) considered rationally, i.e. 'by practical reasoning and deliberation' (Finnis, 1991: 36), in face of other alternatives available and equally considered.

He adds that what one intends is part of what is chosen as one's plan, either as an end or as a means. The main difference between means and ends is that means are instrumental purposes that one chooses to reach an end; ends, on the contrary to the means, consist of the practical reasoning behind the action. Therefore, whereas means are immediate purposes, ends are ultimate purposes.

A sound theory of intention nevertheless needs to consider other concepts, which Finnis (1991) clarifies, are those of *desire*, *side-effect* and *foresight*. Intention does not in this sense necessarily presuppose desire; a doer can intend without desiring. In this sense, 'desire' equates with 'emotional desire', in contrast with intention, which consists of 'volitional desire'. As Finnis demonstrates, means do not necessarily derive from the appeal to feeling; the doer may find the means repugnant in the sense of appeal to feelings and emotion, and still feel motivated to proceed with the plan or proposal in order to reach an end. Following from this distinction, it is then necessary to distinguish between the effects that one wants (i.e. intended effects) and side-effects. In the light of the English judicial doctrine, Finnis (1991) argues, side-effects are all effects which are not intended, neither as ends, nor as means. Side-effects do not feature in the plan or proposal adopted. The practical reasoning behind the plan or proposal determines which states of affair are side-effects and which states of affair are means. This distinction is particularly relevant to the study of plagiarism. One's acceptance of side-effects can only be justified as long as it is proportionate (Finnis, 1991: 56); i.e., actions must be shown to contribute to a legitimate aim; the effect should be significantly outweighed by the importance and benefits of this action; and there should be no reasonable (and less offensive) alternative. When the plan or proposal does not include in the practical reasoning the effect of one's actions (even if one may still welcome the effects of those actions), the action is considered unintentional.

The third concept that needs clarification is that of foresight. The concept of foresight is seen differently by academics and by the judicial system; whereas academics consider that what is foreseen as certain is intended, the legal system relies on a tenuous distinction between that which is 'highly probable' (i.e. not necessarily intended) and 'virtually certain' (i.e. 'practically certain' (Kugler, 2002), necessarily intended). By definition, if one foresees the result of one's actions as possible, then considering the principle of practical reasoning, the plan or proposal should consider the effects of one's actions. However, as Finnis (1991: 33) claims, the fact that one foresees a result as likely or even certain does not necessarily mean that one intends that result.

In criminal law, intending means wanting something to exist or occur, being aware that it exists, or being almost certain that it exists or that it will exist or occur. Acting intentionally against any instantiation of a basic human good – either by destroying, damaging, or impeding it – is, therefore, choosing to act in accordance with a prior state of mind that is contrary to what is established by human conduct as moral. 'Choosing', in this sense, implies reasoning to act purposefully and obtain an intelligible good. Therefore, acting intentionally, as Finnis (1991: 62) argues, does not consist so much of an 'internal feeling or impression', but rather implies that one chooses and sets oneself to execute the plan and obtain the intelligible benefits that make an end appealing to reason.

In face of this discussion, it can be argued that 'intentional plagiarism', in the sense of an act as an end in itself, can hardly exist, since it is very unlikely that someone would plagiarise just for the sake of borrowing without acknowledgement. Rather, plagiarising, in this case, is but a means to reach an end: obtaining a higher grade. Taking academic plagiarism as an example, a student would most certainly not want plagiarism to exist or occur if they could obtain the same effect otherwise. But plagiarism cannot be considered a side-effect, either, since side-effects are not intended, as ends or as means. In this case, intention lies with the fact that the student knowingly uses the means to reach an end. Additionally, whether the intention is shaped by the means or by the end is irrelevant if both means and ends are legally forbidden. In the student paradigm, it would be easier to demonstrate that the student included intention in their plan or proposal – if not as end, at least as means – as long as they knew the conditions of acceptance, awareness and causation. Conversely, it is easy to argue for the existence of 'unintentional plagiarism', i.e. that which happens by accident, mistake or lack of foresight.

Plagiarism can therefore only be considered as intended action when the plagiarist knows the accepted norms (i.e., not the ones that they accept to comply with, but the ones they are expected to conform to), is aware of them and knows the causes implied if they choose to execute their plan. In other words, plagiarism is intentional when the plagiarist knowingly chooses to plagiarise, when they foresee and want the result. This choice to plagiarise is free when the plagiarist, given alternative, even if incompatible choices, is rationally motivated to select one rather than the other. Plagiarism would be considered intentional if derived from 'recklessness' or 'negligence', although with differing degrees of moral blameworthiness, as mentioned in section 4.2.

There is no common agreement, however, on the concepts of 'recklessness' and 'negligence'. Garner, for example, proposes a distinction whereby 'recklessness' involves culpably taking a risk, while being actually aware of it, and 'negligence' involves 'blameworthy inadvertence' (Garner, 2009: 882). However, in practice, legal systems, including Portuguese criminal law (Eiras and Fortes, 2010) claim that negligence means acting with insufficient attention, in accordance with the circumstances and the doer's capabilities. Negligence presupposes that one does not adopt the necessary measures to avoid the 'immoral', foreseeable results; the doer is either indifferent to the effects or, not foreseeing the effects, was obliged to foresee them. If negligence is conscious or unconscious is then irrelevant when the doer is expected to have foreseen the outcome.

In cases of intentional plagiarism, the plagiarist's intention presupposes an intention to deceive; the text is knowingly, wilfully or recklessly (re)produced with the intention that others (the readers) rely on it as being original, and thus the plagiarist obtains some gain. In criminal cases, Azuelos-Atias (2007) claims, intention is determined by the 'defendant's responsibility for the offence, and by the ability to establish a causal relation between the 'defendant's state of mind' (the *mens rea*) and the 'injuries' inflicted (the *actus reus*). However, contrary to what she claims, intention cannot be simply demonstrated by the immediacy of cause and effect. In other words, an effect does not immediately follow a cause, so finding evidence of intentional plagiarism by demonstrating the proximity of the action to the event through the analysis of a sequence of events is not sufficient. On the contrary, linguistic elements can provide evidence of intention, not the least from 'implied intent', in that 'a person's state of mind that can be inferred from speech or conduct, or from language used in an instrument to which the person is a party' (Garner, 2009: 882).

If plagiarism among higher education students is often a problem of literacy, with prior, conflicting social discourses that may interfere with academic discourse (Angèlil-Carter, 2000), we then have to aver that a distinction should be made between intentional and inadvertent plagiarism. When it is part of the learning process and results from the plagiarist's unfamiliarity with the text or topic, it should be considered 'positive plagiarism' (Howard, 1995: 796) and hence not an offence; plagiarism should be prosecuted when intentional. The intention behind the instances of (academic) plagiarism therefore determines the nature of the disciplinary action adopted. In parallel with this, a discussion ensues of whether plagiarism is 'immoral' or 'illegal'. It has been argued

(Garner, 2009; Goldstein, 2003) that plagiarism is immoral but not illegal. However, most legal systems (including the English and the Portuguese) base their definition of 'illegal' on the principle of 'immoral', so that what is immoral often becomes illegal (Finnis, 1991; Eiras and Fortes, 2010); on the other hand, offence of moral rights often implies corrective measures that are identical to those involved in infringement of financial rights. Plagiarism then is subject to trial in courts of law, as has been demonstrated (Turell, 2008), and is not just subject to the social moral sanctions that are expected from immoral actions. Taken together, these studies indicate that judging plagiarism is neither a problem of debate on the ethical vs. illegal principles, nor a problem of lack of agreement on whether intentional plagiarism should be penalised, but rather a problem of determining textual evidence to demonstrate intention.

Since the principle of plagiarism relies on the fact that the text is, in whole or in part, based on someone else's work, words or ideas, I argue that the more a text is manipulated to disguise the original authorship and suggest that it was written by oneself – and, correspondingly, the more difficult it is to trace it back to its original sources –, the higher the degree of intentionality, and hence the degree of seriousness.

4.3 A Taxonomy of Intention

Approaches to academic plagiarism have demonstrated, as discussed in section 1.3, that the assessment of instances of plagiarism cannot be considered independently of the plagiarist's intention, especially considering the fact that plagiarism is so often associated with metaphors of crime. On the other hand, as the perceptions of plagiarism discussed in chapter 3 illustrate, it has been more or less generally agreed that accusations of plagiarism and their corresponding penalties need to take into account the plagiarist's intention to deceive and, moreover, that different degrees of intention should be associated with different levels of punishment. It is therefore relevant to investigate how problems of intention and oblique intention in plagiarism should be dealt with. Additionally, if plagiarism is to be graded according to whether it is intentional, indirectly (oblique) intentional or unintentional – and, if oblique intentional, whether it is done knowingly, recklessly or negligently – how can a linguistic analysis of the borrowing reflect these different classes and subclasses?

A linguistic analysis of a suspect text, and its comparison with potential originals, has the potential to identify the suspect strings of text. These can be borrowed verbatim from the known texts, altered, with more or less effort e.g. to paraphrase the

original text, or even paraphrased from a foreign language and reused without acknowledgement. From a linguistic perspective, several linguistic devices and strategies are or can be used to plagiarise text. Identifying these linguistic devices and strategies is therefore a crucial investigative tool, since it has the potential to provide monitoring and control methods of plagiarism detection. Additionally, these devices also work as a preventive measure in that they identify potential items of suspected plagiarism, and hence operate as an educational measure. Another advantage of using these linguistic devices effectively for plagiarism detection is that they avoid the 'inquisitorial nature' of interviews, which usually bring along unnecessary stress for the accused, especially in cases where the borrowing is inadvertent, and especially when the accused is not a plagiarist until this has been proven.

Different linguistic devices assist differently in identifying the instances of plagiarised text and in providing evidence of intention. For example, it seems obvious that the more sophisticated the borrowing, and the stronger the effort and the more complex the workaround to change the text, the more intentional the action – the *actus reus* – seems to be, and the more it seems to show the mental state – the actor's *mens rea*. Nevertheless, proposing a taxonomy of intentional plagiarism actions, including a description of a comprehensive set of possibilities is a very difficult task, if not nearly to impossible. Yet, instances of plagiarism can realistically be assigned to three main, broad categories, i.e. *intentional plagiarism*, *oblique intentional plagiarism* and *unintentional plagiarism* and, within these, assigned four different mental states, from *purposeful* to *knowing*, *reckless* and *negligent* actions.

4.3.1 Intentional Plagiarism

Intentional plagiarism, which aims at achieving undeserved credit, in this sense means purposely using the works, words and ideas of another person as one's own. One example of this type of plagiarism can be paraphrasing, which plays a crucial role in providing evidence of intention; a case where the original phrasing is kept, changing as many as just a few or all the meaningful (i.e. the lexical) words and omitting the sources, for example, is an obvious strategy of deception, as it consists of changing the wording in order to obfuscate the original authorship of the borrowed works and ideas. It presupposes that the plagiarist is aware of their actions, and moreover an involvement of the control condition, i.e. that the plagiarist, given a set of alternatives, and being capable of doing otherwise, chooses to plagiarise. This latter element, the

'moral responsibility' (Alvarez, 2009: 295), contributes to demonstrating the offender's mental element.

Plagiarism of this type includes also instances where the plagiarist uses citation in a way that aims to mislead authorship – including ambiguous, unclear and confusing referencing that makes it impossible to trace back the sources. Citing some (but not all) the authors quoted in a text, using pronouns confusingly in a way that makes it difficult for the reader to understand whom the writer is attributing the ideas to, and using quotation marks with some, but not all, the quotations are some examples of intentional plagiarism. Other cases include providing wrong page numbers to distract authorship, using references fuzzily, so that it becomes very difficult to identify what belongs to whom, and using text in a way that it becomes difficult for the reader to determine whether it was the writer who wrote it or the author (i.e. the source) to whom it is attributed. Likewise, this involves cases where the plagiarist not only attempts to pass someone else's work as their own, but also instances where they intentionally try to pass misattributions that are possibly identified problems with academic writing.

A similar case is where a student copies someone else's work, words, or ideas – including data and analysis –, and acknowledges them only partially, for instance by citing them only once regarding a minor aspect, and finally consider the text as their own. The same applies to paraphrasing the original text without acknowledging the source, in such a way that relevant lexical items are replaced, on the paradigmatic axis, with words from the same semantic field, including synonyms, hyponyms, etc. It is also intentional to copy the text presented in the original, without attribution (i.e. without quotes and the author), while changing the order to disguise the source. Frequently, plagiarists borrow from multiple sources, and the inconsistency of the referencing styles is often indicative of that borrowing. Some other times, they borrow the text verbatim from an original that cites other sources. In this case, the plagiarist cites direct speech, including the quotes and the citations, but fails to cite reported speech. This represents a problem situation, in that the reader is led into believing that none of the text is borrowed, since sources are acknowledged, and does not hesitate to attribute the authorship of the reported speech to the plagiarist. Moreover, the plagiarist can, in this situation, be given credit for academic writing skills that they lack.

The same principles of intentionality also apply to translating the source text from a foreign language into a target language and using it as one's own, without acknowledgement, as will be discussed in chapter 5. Translating texts from other languages

and using them in their assignments can be a legitimate or illegitimate strategy used by students; the unskilled writer might fail to cite correctly, while not missing the author's name next to the borrowed text. The use of translation in this case is therefore legitimate, even if it does not result in successful academic writing. Conversely, an intentional plagiarist would attempt to disguise the text, which implies an illegitimate use of translation; they would translate the text freely or automatically, edit it to make it pass as their own and omit the source. The argument for the plagiarist's ignorance of the applicable academic writing conventions in this case is partly convincing, but not fully. Although attributing some of these situations to the lack of writing skills, rather than to the intention to plagiarise, seems realistic, it is more likely that the unskilled writer would 'over-cite', more than 'under-cite' or cite misleadingly.

Intentional plagiarism implies both responsibility and liability on the part of the plagiarist, who lacks an 'exculpatory answer' (Duff, 2007) for the offence, and hence is answerable for it.

4.3.2 Oblique Intentional Plagiarism

Oblique intention in plagiarism, on the other hand, is demonstrated by instances where the plagiarist is or should be aware of the conventions, but fails to meet the necessary requirements as a result of recklessness. Instances of oblique intentional plagiarism include cases where the plagiarist borrows the text by verbatim copying it sparingly, with no attribution at all, in a way that indicates that the author of the text lacks academic competence to reference correctly. Another case of oblique intention is where a plagiarist uses 'patchwriting' (see section 1.3) to attribute sources clumsily, whilst failing to cite all the due instances. As Howard (1995) argued, patchwriting in the academy often results from cases where the student, lacking sufficient writing skills, fails to paraphrase correctly, and may result in confusing writing styles or using the same wording or syntax as the original, where alternative options are available. This may also include copying the text in the order presented in the original, without attribution (i.e. without quotes and the author). One of the features that makes this an oblique intention, rather than an intentional action, is that in both these cases the instances of plagiarism are easily detectable; in the former, this could be identified intuitively or even automatically, based on the confusing or inconsistent writing styles; the latter case could be identified automatically, by using any software detection package.

An additional example of oblique intention consists of using sources whose author-

ship is more volatile (e.g. the internet) to avert attribution. Under these circumstances, the intentional plagiarist would disguise the text to make it pass as their own, whereas the oblique intentional plagiarist, on the other hand, would miss the fine-grained conventions, while providing some more or less evident clues that the text is the original work of someone else. The case of the student who legitimately, but unsuccessfully translates a text from a foreign language, as presented in section 4.3.1, would in this sense represent a case of oblique intentional plagiarism.

Oblique intentional plagiarists tend to have excuses for their actions. Hence, they are responsible for the instances of plagiarism they produce, but not liable to them on the grounds that they have an 'exculpatory answer' (Duff, 2007) that intentional plagiarists lack.

4.3.3 Unintentional Plagiarism

Contrary to intentional plagiarism, unintentional plagiarism includes instances of citation that, although not conforming to the established conventions, provide hints to the sources and make it clear that the works and ideas are not one's own. This instance of plagiarism is usually due to lack of writing skills, or ascribed to diverse cultural conventions. In the scope of such cases, it is usually clear who the original author is – or at least that the potential plagiarist is not the original author – but the reader feels that the text is not properly attributed to the sources.

Along the same lines, unintentional plagiarism also includes verbatim copying the words of the original, without attribution, when such words are not lexical primes. This usually results from the fact that the plagiarist, being aware that the words of the original are 'commonly used', does not think that they need to cite them/quote them.

Unintentional plagiarism, in the academy, may be penalised for poor academic writing, but should be exempt from plagiarism, in that it negates responsibility – and hence, lacking responsibility, it is not liable either.

4.4 Assessing the Severity of Intention

Research into academic plagiarism detection (e.g. Pecorari (2008)) has traditionally used textual analysis to identify the 'borrowed' text, followed by an analysis of the potential plagiarist's intention by interviewing the suspects, before determining whether the unacknowledged borrowing was intentional or, on the contrary, deriving from lack of

academic writing skills (i.e. the student's inability to conform to specific academic writing conventions). This approach, however, poses some problems. Firstly, interviewing students is not always feasible. In some countries (e.g. in Portugal) the investigation of whether a student has plagiarised is left to the investigative skills of the lecturers or tutors (as has been confirmed by personal communication with several people), and very rarely are Institutions called upon to act on the issue. Conversely, in countries such as the UK, where plagiarism is (at least in principle) taken more seriously, the issue is usually dealt with by a disciplinary board, which will hardly be able to cope with the matter if the number of cases of plagiarism increases. Additionally, the 'interviewing approach' does not always necessarily demonstrate that the plagiarist has plagiarised intentionally or unintentionally, as its success largely depends on the 'lying skills' of the interviewee. As a consequence, an incompetent student may, in an act of shyness, be more reluctant to admit that he or she lacks academic writing skills, and hence pass off as a plagiarist; conversely, an 'intentional plagiarist' may be sufficiently assertive to deny his or her intention, and demonstrate that his or her acts were due more to lack of academic writing skills than to intentional plagiarism – in which case their instances of plagiarism would be considered unintentional.

Another approach, described by Coulthard and Johnson (2007), is that students are asked to sign a statement confirming that they know both what plagiarism and the corresponding referencing conventions are. If, on the one hand, this has the advantage of allowing linguists to concentrate uniquely on the kind of plagiarism they are most comfortable in detecting (i.e. textual overlap), on the other hand it disregards educational concerns that a candidate might *think* he or she knows – but not actually knows – about what plagiarism and the applicable academic writing conventions are. Moreover, this principle seems to oversee the fact that candidates, as demonstrated by Carroll and Appleton (2001), are given 3,000 documents at the start of a 3-year undergraduate degree. Additionally, as they suggest with reference to the recommendations of the National Union of Students (NUS), these documents need to be drafted in a friendly language, that the students are able to understand. Therefore, including plagiarism in a student handbook does not necessarily mean that they have understood it.

Linguistic evidence, on the contrary, has the potential to provide clues to indicate that a certain text has borrowed, intentionally or unintentionally, from a previous text, and assist lecturers/tutors, supervisors and educational institutions that may hesitate

to pass judgment on the reuse of the original text by students – thus helping them make a decision without necessarily having to consider contextual information (Pecorari, 2008). This assertion is based on the simple premise that, the more a text is manipulated to disguise the original authorship and suggest that it was written by oneself, the higher the degree of intentionality – and the higher the degree of severity. An explanation of the most commonly used strategies and textual devices used to borrow text is provided in the section 4.3 above. However, owing to the fact that the nature of the several strategies and devices varies greatly, and so does the degree of severity attached to them, they are bound to perform differently in demonstrating the writer's intention. In the next section 4.5, some of these linguistic devices are discussed, as well as their corresponding degree of severity.

4.5 Linguistic Evidence of Intention

Detecting plagiarism and determining the degree of blameworthiness requires considering different categories, like in criminal action: *intention* (deliberately acting for a reason), *oblique intention* (actions that are neither intentional, nor unintentional, usually involving foresight of the outcome as certain) and *lack of intention* (those happening by accident, mistake or lack of foresight). However, useful as these categories may be, they can hardly account for all the required degrees of moral blameworthiness behind each instance of plagiarism. For instance, determining that a certain instance of plagiarism is either intentional or oblique intentional does not account for all the instances where the actor's *mens rea* is *purposeful*, *knowledgeable*, *reckless* or *negligent*. From a moral perspective, it seems intuitively appropriate to consider that acting *purposely* (i.e. acting with the intention of bringing about that result) is morally more blameworthy than acting with *knowledge* (i.e. with the knowledge that the result is practically certain) that this result might ensue. Likewise, acting *knowingly* seems to be more blameworthy than acting *recklessly*, and acting with *recklessness* is more serious than acting with *negligence*. It therefore seems morally appropriate to consider that one acts with *direct intention* when the *actus reus* is performed in order to achieve a certain result; otherwise, intention can only be considered to be *oblique*.

The discussion presented in this section is based, not only on the distinction between intention, oblique intention and lack of intention, but also on the four different degrees of culpability discussed previously: *purpose*, *knowledge*, *recklessness* and *negligence*. Different linguistic devices are attributed to each of these classes of in-

tention and to each degree of culpability, and their degree of moral blameworthiness is based on the potential of each of these devices to disguise the original authorship. The hypothesis proposed is based on the principle that, the stronger the effort to disguise an unattributed source, the higher the degree of moral blameworthiness. This description aims to be a representative, rather than a comprehensive one. The linguistic devices presented in this section therefore illustrate the type of devices that can be used, but a comprehensive description should include many more devices. Additionally, these devices are language-dependent, and analysing some of them seems more relevant in one language than another. The approach adopted here is, therefore, able to work with English, but the main focus is on Portuguese – the reason for this being that Portuguese is more syntactically, lexically and morphologically diverse, hence allowing for more combinations than English.

Several different linguistic devices are analysed, starting with *paraphrasing*. Paraphrasing can be based on a complex process whereby the text is recast altogether, using neither the syntax, not the lexis of the original. At its simplest, however, it takes place when words in a string are replaced with semantically-related words. It occurs when the plagiarist replaces the lexical items in a phrase, sentence or paragraph with words from the same semantic field, with related meanings (including synonyms, antonyms, superordinates, co-hyponyms, etc). Contrary to verbatim plagiarism, where the changes are made on the syntagmatic axis, paraphrasing acts on the paradigmatic axis. This is one of the main plagiarism strategies, and one of the most difficult to detect, either because the text is unable to show an exact match, or because the lexical items have been changed by the plagiarist. This issue is even more relevant in languages with complex syntactic, morphological and lexical structures (such as Portuguese), where a simple variation to the lemmatisation of a word may be sufficient to rephrase text and mislead automatic, string-matching plagiarism detection systems. Paraphrasing of this type is intentional in that the writer is required to consciously make significant changes to the original, whilst hiding any reference to it. When used in academic writing, this is a particularly complex device to analyse from a forensic linguistic perspective, in that students are often encouraged to use their own words effectively, either to describe what other authors claimed, or to take their own notes. Instances like these are quite different from a writer using their own words to refer to another author's works, words and ideas in that 'note-taking in one's own words' is usually accompanied by a reference to the original. A simple, illustrative example of this feature can be

found in the text of S1.4: although the string ‘a forma e o conteúdo da publicação’ is maintained, this collocation is so common that it can hardly be proved as an instance of plagiarism. A proper analysis of this instance therefore needs to also take into account the other words of the sentence, i.e. ‘Esta opção irá condicionar’; although this would at first sight seem very different from the original (‘opção que condicionará’), in reality this is only apparent: the first change consists of introducing a demonstrative pronoun, ‘esta’, that is missing from the original. However, the main change is operated on the verb: ‘condicionará’ is changed to ‘irá condicionar’, and the future tense is transferred from the main verb (‘condicionará’) to the auxiliary verb (‘irá’). The main verb then changes from future tense to infinitive, as required by Portuguese grammar when the tense is applied to the auxiliary verb.

However, it is not just the verb tenses that tend to be changed. Like verbs, nouns are subject to change, especially when the terminology of one language variant differs from the one in which the plagiarised text is created. An example of this change can be found in S4.7, where the original ‘superposição’ is replaced with the corresponding variant term, ‘sobreposição’. Changes in terminology, however, are not exclusive to text reuse from other variants: in S5.3, for example, the original European Portuguese (EP) ‘fotodocumentalismo’ is replaced with (also EP) ‘fotodocumentário’.

Derivative texts are also created by deleting elements that give the text specification. This consists of making the text more general, and hence possibly more difficult to trace to its original sources. In S1.1, for example, the student deletes three words from the original (‘um ou dois’) to make the derivative text less specific. The same applies to S1.2, where the examples of issues related to the school community, in brackets, and the examples of publication, in dashes, are deleted from the derivative text. Other examples are shown in S3.16 and S3.24, where detailed information in brackets is deleted, or in S3.20 where the additional information after the column (‘Su padre o su maestro’) is not reused.

What makes deleting specification a case of oblique intention, rather than a case of intention, is that deleting elements from a text is often an editing process, an attempt to improve the text, rather than an intention to deceive. Yet, an analysis of this linguistic device is able to provide evidence of the *directionality* of the borrowing. In such instances, the plagiarising text borrows from the original text and deletes all unnecessary elements, by adding none or little new information to the derivative text.

Conversely, it is not uncommon to find cases where specification is added to the

original to create the derivative text. This includes cases where the added words contribute little to making a new text, usually providing synonyms to existing words – and hence repeating existing phrases. An example of this strategy can be found in S5.7, where 'na fotografia noticiosa' is added (unnecessarily, as it is a synonym of the phrase that precedes it, 'fotografia de notícias'). Contrary to deleting specification as a result of minor text editing, adding specification implies knowingly inserting text (often to obfuscate the lifting), and hence is more intentional than deleting specification.

Automatic detection systems can easily be misled by introducing in the verbatim-copied text new elements that do not exist in the original, such as spelling mistakes or other text. Adding words to reused verbatim strings differs from adding specification in that, whereas the latter is an intentional attempt to improve the text, the former is more likely to be an intention to mislead. One example of this is illustrated by S2.1, where 'no corrente ano de 2006 todos estes números já estão ligados ao passado...' is newly introduced. The adjustments made to punctuation ('...' is maintained) to keep consistency with the original are particularly noticeable. A similar strategy is used by S4 in S4.10, where 'Com vermelhos' is replaced with 'Com cores vermelhas' i.e., a new noun ('cores') is introduced, and the noun ('vermelhos') is transformed into an adjective and its gender changed accordingly ('vermelhas').

As the textual analysis of plagiarised text shows, reusing the original text and introducing some changes to make it coherent is not uncommon. This is, for example, the case where the original is similar to the topic, but not exactly on it. Omitting and replacing textual elements to make the derivative text coherent, for example, demonstrates intention in that it requires a conscious stance from the writer to adapt the text to their aims. Examples of this feature can be found in S1.8: the sentence 'Há receituários de desenho gráfico que podem ser consultados para procurar inspiração', as it makes reference to design strategies used by newspapers, makes sense in the context of the original (the article was aimed to instruct schools on how to create a school newspaper). However, it is irrelevant in the context of the derivative text, as this essay was supposedly written by design students. Hence, the string was deleted from the plagiarised text.

Plagiarism is often disguised by changes in spelling and morphology, and is in that sense intentional. Portuguese students are often said to plagiarise Brazilian texts, which then requires them to change the spelling in accordance with the European Portuguese conventions: for example, where European Portuguese uses 'cç', 'ct' and

'pt', Brazilian Portuguese uses only 'ç', 't' and 't', respectively. In parallel, in many instances where the British English spelling uses 's', the American English spelling uses the 'z', as has been demonstrated in the analysis of the *Iraq dossier*. Although this is a straightforward and obvious change in spelling, it is often sufficient to mislead automatic detection systems. Some examples of this can be found in S4.3 and S4.13 ('objetos' and 'objeto' is rewritten as 'objectos' and 'objecto', respectively), S4.9 ('Adotando' is repurposed as 'Adoptando'), S4.11 ('fator' is altered to 'factor'), and S4.14 ('ação' is rewritten as 'acção' in the plagiarising text).

Morphological changes also operate between variants, for example at the level of pronouns and adverbs. For instance, where Brazilian Portuguese tends to use a pronoun only, European Portuguese requires a determiner followed by a pronoun, as in S4.1 ('seus princípios' is reworded as 'os seus princípios' in the derivative version) and S4.17 ('sua influência' is altered to 'a sua influência'). Likewise, adverbs are used differently, as is demonstrated by the example in S4.21; where one variant uses 'junto', the other uses 'juntamente'.

Another linguistic device often used to plagiarise consists of changing the word order. Although it can be argued that simply changing the word order is not paraphrasing in the strict sense that the original words are not replaced with other words, the syntactic operations underlying it are identical to those involved in paraphrasing, in the sense that a new combination, containing the same words, is created. This change passes undetected by most plagiarism detection systems, and prevents the source from being identified in search engines like *Google*, when searching for the exact strings. Changing the word order is, in this sense, intentional in that the meaning provided, as well as the individual words, are reused consciously without acknowledgement. Examples of this feature can be found in S4.16, where 'foram assim chamados' is changed to 'assim foram chamados', and in S5.2, where 'No sentido lato, entendemos por fotojornalismo' was replaced with 'por fotojornalismo no sentido lato, entendemos'.

Plagiarism in general, and academic plagiarism in particular, is often revealed by the use of reporting verbs. Most instances of reporting verbs used in instances of academic plagiarism demonstrate a verbatim borrowing from the original, without further changes, just like the quoted text preceding or following the reporting verb. Examples of this can be found in S3.25 (the original 'defendía así' is translated as 'defendía assim'), S4.5 (the verb 'afirmara' is retained) and S4.8 (the phrase 'com diz', despite grammatically wrong, is reused in the plagiarising version). However, this type of verb

is often reworked to distance it from the sources, e.g. by replacing a verb by a preposition where language so permits. S5.8 is an illustrative example: the reporting verb of the original, 'sustentada', is replaced with a preposition in the derivative text 'Para'. The text follows without quotation marks, but is (almost) verbatim copied from the original. Therefore, whereas verbatim copying is more a case of recklessness, reworking reporting verbs involves a conscious attempt to change the original, and is thus intentional. In many instances of plagiarism, the reused/adjusted reporting verbs, and the quotations that precede or follow them, are never based on the direct references, but instead copied from other texts where the original references had been cited.

On balance, it might be thought that using paraphrasing and borrowing reporting verbs without acknowledgement can demonstrate intention or oblique intention. Instances of intention include replacing words on the paradigmatic axis with semantically-related words, adding specification, omitting and replacing textual elements to make the derivative text coherent, adding words to reused verbatim strings, changing the spelling and morphology, changing the word order and introducing changes to reporting verbs. Oblique intention includes deleting specification and verbatim using reporting verbs.

If text manipulation is taken to equate with intention, it is not uncommon for the plagiarist to knowingly borrow unacknowledged text as a means to reach an end – i.e., obtain the desired grades. In section 4.2, it was strongly argued that the judgment of intention relies on the fact that plagiarising students knowingly chose to plagiarise, or foresaw the result and yet decided to implement their plan. However, the degree of manipulation and reworking required by the types of alterations presented in this section demonstrates that these cases cannot be considered accidental or inadvertent. On the contrary, this type of manipulation requires an awareness of the respective conventions, and hence knowledge of the causes implied by their choice to execute their plan. Ultimately, the plagiarist's intention to deceive is demonstrated by the intention that the readers/lecturers/tutors rely on it as being original, and mark them as such, and only under these circumstances should the plagiarist be deemed liable for their actions.

4.6 Chapter Summary

In this chapter, it was argued that the paradigms of criminal law can be used with a significant degree of success for plagiarism detection. The most serious degree of

intention is intentional plagiarism, and consists of instances where the text is manipulated to disguise the sources and obtain undeserved credit. Replacing words on the paradigmatic axis with semantically-related words is one of the main strategies used by students to plagiarise, and one that is difficult to detect by plagiarism detection software; however, many other strategies can be involved. Cases reported often show that the borrowing is mostly verbatim from the original. In these cases, the plagiarists copy both the verbs used to report the borrowed text and the quotations (in-text citations) from the original. The reused reporting verbs, and the quotations that accompany them, are mainly copied from other texts where the original sources had been cited. It has been argued that this does not demonstrate the lack of intention to deceive, but on the contrary that it should fall within the same degree of intention as verbatim plagiarism.

A more serious degree of intention to plagiarise is indicated by cases where the reporting verb is manipulated or replaced with a preposition, in an attempt to distance the borrowed text from the original sources. On the other end of the continuum is unintentional 'plagiarism', which has been argued not to consist of plagiarism. This includes cases of 'patchwriting' (Howard, 1995) and poor academic writing skills, which are a problem of the students' education and training, rather than an 'intention to deceive'. In the middle of the scale is oblique intentional plagiarism, which includes cases of culpable recklessness; this includes instances where those responsible for the plagiarism intend the effects only indirectly: they foresee the effects of their actions but do not impede their actions.

One of the main challenges lies however in determining cases of plagiarism resulting from an attempt at being creative, i.e. where an effort is made to change the original that may either be perceived as an effort to write academically, or conversely as an intention to deceive. This is a borderline case, in that it can easily shift between the scope of 'highly probable' and the 'virtually certain'.

It has therefore been argued that assessing instances of plagiarism and determining the respective punishment cannot be performed independently of judgements of intentionality. In this respect, a forensic linguistic analysis of improper textual borrowing has the potential to identify and provide evidence of intention in plagiarism. Additionally, it was demonstrated, by establishing a parallel with criminal law, that principles of responsibility and liability are important to consider blameworthiness in textual borrowing.

The chapter concluded by claiming that not all apparent situations of unacknowledged textual borrowing should be considered plagiarism, and that there may be mitigating circumstances that can exempt a suspect from accusations of plagiarism.

Chapter 5

Translingual Plagiarism

5.1 Introduction

This chapter discusses the concept of plagiarism in relation to translation, whether a text in one language can be plagiarised from a text in another language, and whether and how such plagiarism can be detected, described and demonstrated.

Firstly, it discusses the concept of translation over time, considering different approaches to translation, from more traditional perspectives of human translation where a transfer of meaning from one language into another is implied, to more recent theories where translation is seen as meaning negotiation between two different languages and cultures, and to approaches that propose a combination of different theories. Based on this discussion, it is assumed that any text (in a semiotic, and not just linguistic sense (Kress, 2001)) is subject to the principles of originality. It investigates, in accordance with the long-standing linguistic canons of originality, whether only particular words and word combinations are entitled to legal *protection* across different countries, or whether such *protection* is also extended to the meanings conveyed by those words and word combinations. This conceptualisation contributes to explaining and accounting for linguistic plagiarism that goes beyond monolingual text reuse, to include also translingual plagiarism, where one text derives from an original in another language.

In the subsequent section, the discussion focuses on the investigation of grammatical and discursive negotiations that operate in translation, as well as on the identification of linguistic and interlingual cues that may lead to considering that a text has been plagiarised. The first step focuses on the detection task, and consists of identifying markers of interlanguage and mistranslation arising from the failure to meet language

standards, usually as a result of lack of translation skills, as well as from interlingual aspects. Empirical evidence is used to demonstrate that, although an assessment of the quality of a translation can be useful in intuitively determining that a text has been translated (i.e. it *reads like a translation*), more information is required to cope with cases of sophisticated translation. Therefore, relying on 'sentence-based grammars' is not sufficient; instead, considering other elements such as discourse grammar is also required, as these impact the informational packaging of the text, as well as its use of cohesion and coherence.

A discussion of the theory of interlanguage is consequently necessary, as well as a debate on how it influences both second-language acquisition and performance in the native language. It is tentatively suggested, with the assistance of research on bilingualism and second language acquisition, that performance in the use of the native language is minimally or maximally impacted by the speaker or writer's interlanguage, and that the latter impact reflects not only on the second language, but also on the native language of the speaker.

Thirdly, a methodology is proposed to detect plagiarism across translated texts, illustrated with different language pair combinations (e.g. Portuguese/English, English/Portuguese, Spanish/Portuguese). Using a simple method based on machine translation (MT), it is firstly hypothesised that any text can be investigated for 'translated plagiarism' with a minimum effort. Subsequently, a categorisation is proposed to test and explain how and why one text derives from a text in another language, and hence potentially prove or disprove the preliminary results of the detection phase.

Finally, the chapter discusses whether the detection method can be relatively independent of the linguistic analysis, or, on the contrary, whether it should rely upon it, to conclude that the linguistic analysis has a great potential in assisting with the detection method, and may be crucial to investigate, analyse and explain textual reuse.

5.2 The Case for Translingual Plagiarism

Plagiarism has been thoroughly investigated across different genres (e.g. academic, literary, etc.), disciplines (e.g. engineering, business, linguistics, etc.), and contexts (e.g. Western/Eastern cultures, industry/education, etc.) (Anderson, 1998; Angèlil-Carter, 2000; Bennett, 2005; Howard, 1995; Jameson, 1993; Johnson, 1997; Kaplan and Torbati, 2007; Pecorari, 2008; Robillard and Howard, 2008; Scollon, 1995), as discussed in section 1.3. The use of translation as a plagiarism strategy is a known

issue, in academic as well as non-academic contexts. In non-academic contexts, two reporters of the Portuguese quality newspaper *Público* were recently found to have plagiarised from other news pieces, and in the academy students have likewise been reported to have translated from other languages and passed off the text as their own. In academic contexts, as mentioned in section 1.3.2, translation has been a concern of institutions worldwide (especially in non-English speaking countries), which include a reference to translation in their plagiarism definitions – and how to cite it properly – in order to avoid student plagiarism. However, none of these studies has approached translation as a plagiarism strategy, and research into this area has been very limited, or has demonstrated disappointing results. Jones (2009), arguing that students have the creativity to devise new methods to plagiarise on a regular basis, reports a method that used translation to detect plagiarism. But his research interest focused less on the use of translation to plagiarise, than on an obfuscation technique. Specifically, he reported that students used translation to convert a text into a foreign language, and then back-translate it into their own language in order to obtain a different wording. His detection method consisted, therefore, in reversing this procedure to detect same-language plagiarism. He concludes that detecting plagiarism when this strategy is used is very difficult, if not impossible, advising lecturers/tutors to devise assessment tasks that avoid this type of strategy.

On the other hand, where detection of plagiarism across texts in different languages has been researched, the focus is on pure computational, rather than linguistically grounded approaches. Ceska *et al.* (2008) proposed an approach that consists of pre-processing the two texts, in order to transform them into a language-independent form and subsequently compare the two. The performance of this method depends, however, on the availability of a parallel thesaurus of the two languages involved, as well as on the size of that thesaurus, which limits the number of words that the system is able to successfully index. Additionally, since the suspect and potential original texts need to be pre-processed before making the comparison, the suspect original needs to be known in advance. Similarly, Corezola Pereira *et al.* (2010) offered a method that consists of a 5-phase complex procedure that includes language normalization, retrieval of candidate documents, classifier training, plagiarism analysis, and post-processing. The results reported are poorer than those achieved over the monolingual plagiarism detection procedure, even using an artificial plagiarism corpus. Limited results have also been reported by Barrón-Cedeño *et al.* (2010). They compared a new

detection method consisting of a combination of machine translation and monolingual similarity analysis against two previous methods devised by Potthast *et al.* (2011a). Although they reported good results with their two methods, with one performing better with syntactically identical language pairs, and the other showing a good performance with 'exact' translations, Barrón-Cedeño *et al.* (2010) concluded that both methods presented poor results with their corpus. Their analysis confirmed that these methods were largely dependent on the syntactic identity between the suspect text and the original, on the size of the resources available and on the computational capacity (with better results demanding extremely high processing capabilities). On the other hand, their machine translation and monolingual similarity comparison method was demonstrated to perform better than those offered by Potthast *et al.* (2011a), but requires previous translation of all documents, which may become expensive and unrealistic.

More recently, Pataki (2012) offered a translated plagiarism detection method that is based on a distant function search, in order to search for 'all possible translations'. The method, whose results are admittedly poor, is based on sentence chunking, so that the comparison between the suspect text and the possible translations is made on a sentence-by-sentence basis.

Several reasons may explain why this area remains relatively under-researched. One may be the fact that the concept of plagiarism adopted may consider that translating from another language is not plagiarism in the sense that the text (re)used is neither textually identical, nor similar to the original one. Firstly, an investigation of texts suspected to have plagiarised an original in another language challenges linguistic concepts such as the ones proposed by Johnson (1997). Since the original and the suspect instance of plagiarism are in different languages, the identity and/or similarity between words, strings and grammar becomes significantly more difficult to demonstrate. Secondly, plagiarism of this type, where 'a language A text written by author A is translated into language B and the translator appears as the author of the original translated text' (Turell, 2004: 7) – and which Turell (2008: 271) calls 'plagiarism in translation' – does not reuse the 'linguistic text' from the original, thus being often categorised as *plagiarism of ideas*, in contrast with *linguistic plagiarism*. Therefore, the text that is lifted from the original could be roughly considered to be a *semiotic text*, more than simply a linguistic text. Additionally, the detection task is challenging, and the methods proposed, despite their computational sophistication, can hardly obtain satisfactory results.

Unsurprisingly, therefore, most investigations into linguistic plagiarism until this date have been limited to monolingual plagiarism, where a text borrows from another original in the same language, and relatively little research attention has been paid to plagiarism across different languages, by means of translation. Where research has considered translation and plagiarism (e.g. Turell (2004, 2007)), the analysis has focused mainly on comparing the translations of the same original to find instances of plagiarism among the translated texts. This task, which consists of comparing the translated texts, rather than a translated text and an original in another language, is therefore strikingly similar to monolingual plagiarism detection; in other words, the suspect translation is not checked for identity and similarity to the other language original, but against another translation of the same original in the same language. This may imply using the original for reference, but the main task consists of analysing same-language texts.

Detecting plagiarism of this type is hampered by the limitations imposed on the linguistic analyses, since the type of plagiarism that linguists are mostly competent to deal with, as was demonstrated by Coulthard and Johnson (2007), is *linguistic plagiarism*. Therefore, a very strong effort is required to detect surreptitious theft from other languages, owing to the fact that although the ideas are the same, the wording is necessarily different. On the other hand, given the distinct wording, finding duplicates and near-duplicates, which is an ordinary procedure in computational plagiarism detection, as discussed in chapter 6, is limited by the (technical) ability to cope with it. It adds to this that most research on plagiarism is *English-centred*, i.e. it focuses on texts written in English and is mostly conducted in English-speaking countries; for example, if we take the Internet in general as an example, a large percentage of texts are nowadays written in English[1], so the demand for texts in other languages is comparatively much smaller.

The absence of research motivation in this area, which seems to have until recently dominated 'English-centered' research, is very likely to change in the future, especially as even English universities are growingly acknowledging the need for approaches to plagiarism detection across texts in different languages. A recent post to the Plagiarism mailing list[2] raised the question of whether a bibliography of foreign-language

[1] According to the Wikipedia (`http://en.wikipedia.org/wiki/Internet`), in June 2011 it was estimated that the English language accounted for 27% of the Internet, followed by Chinese (23%) and Spanish (with only 8%).

[2] PlagiarismAdvice.Org – `PLAGIARISM@JISCMAIL.AC.UK`, 15/11/2011.

texts should be accepted, on the grounds that the lecturer would be unable to 'translate the material', and that this might make it impossible 'to verify the source'. Text reuse from originals in another language is, therefore, admittedly a problem. And this applies both in and outside the academy. In 2007, for instance, a journalist of the Portuguese quality newspaper *Público* was accused of having plagiarised the *Wikipedia* and the *NewScientist* in her piece on sunscreens, published in the newspaper's Sunday supplement. The case received considerable attention. A webpage was dedicated to it, including the original and the derived texts[3], and the newspaper used it as an example of malpractice. In the academia, cases of (this type of) plagiarism rarely acquire a high profile and make their way into the newspapers, but several educators[4] demonstrated that the phenomenon is either on the rise, or more attention is being paid to it – or both.

Owing to the technological developments of the last decades, as was argued by Coulthard and Johnson (2007), the Internet eased the access to more information more readily, making that information particularly susceptible to plagiarism. But the problem is not limited to monolingual texts, and it has been claimed that these technological developments made it easier to access 'global and multilingual contents' (Maurer *et al.*, 2006: 1079). Actually, the latter authors argued that current machine detection systems, even those that work well with monolingual textual material, tend to break down '[w]hen plagiarism crosses language boundaries' (Maurer *et al.*, 2006: 1080). This, they anticipated, will remain a challenge for many years.

Recent research has however demonstrated that plagiarism of ideas, similarly to what happens with linguistic plagiarism, can be investigated, described, explained and proved from the perspective of linguistic analysis. This has been shown in particular by Turell (2004), who compared four translations of Shakespeare's *Julius Caesar* that were published in Spain to show that the use of quantitative linguistic evidence can help determine plagiarism between translations. Interestingly, Turell's study did not compare one text in one given target language (TL) to an original in another, source language (SL), but instead made a comparison of the four TL texts translated from the same SL text, to determine the extent of plagiarism, by explaining how much overlap could be expected among the TL texts. It therefore focused on the analysis to conclude, not that the ideas that were lifted were not those of the original, but on the contrary those of another translation of the same source. This is an illustrative exam-

[3]The documents, last accessed on 27 June 2011, can be found at `http://tinyurl.com/6fkzngv`.

[4]Personal communication.

ple of plagiarism analysis that used a textual comparison to demonstrate an instance of plagiarism of ideas, where some ideas were translated using the same TL words, even when such wording is unexpected.

Conversely, the textual comparison of two texts in different languages presents additional problems; firstly, the text is not duplicated or nearly duplicated from another translation in the same language, but it is an 'interpretation' of the original in another language. This challenges the textual analysis since, as was argued by Johnson (1997), a case of plagiarism cannot be proved unless 'clear lexical parallels' and 'identical lexical strings' can be found between the texts. The limitations of this detection and investigation therefore consist mainly of not being able to use textual analysis at this stage to help confirm or discard the hypothesis that there has been plagiarism of ideas, and are reinforced by the fact that the possibilities of translation of one textual string are almost endless. The relationship between one SL and another TL string is one-to-many, so that one string in one source language text may have numerous translation possibilities into another, target language text.

However, if plagiarism is to be considered any surreptitious theft of words, works and ideas, this type of plagiarism cannot be paradoxically put aside as a lesser plagiarism strategy. This will be termed 'translingual plagiarism'.

Translingual plagiarism has been found to be more precise terminologically than other alternatives to refer to this type of plagiarism, such as 'cross-lingual plagiarism' and 'translated plagiarism'. 'Translated plagiarism' is imprecise, since it may be used to refer both to cases where one TL text was translated and lifted without acknowledgement from another source language text, as well as to cases where one translated text lifts from another translation, in the same target language. 'Cross-lingual plagiarism', on the other hand, is closer to the analysis of texts that are derivative from an original in another language, but as the prefix 'cross' indicates, it focuses on the particular *intersection*, rather than on the *transection* of the texts. Hence, it may include all languages considered, avoiding any concepts of directionality. 'Translingual plagiarism' therefore suggests that the interaction between the two texts is transversal, crossing each other, one being active and the other one being passive, rather than the two texts being equal, contrary to what the prefix 'cross' might suggest. In terms of directionality, it suggests the analysis occurs from one language to another.

A method is presented below (see section 5.5) to detect and investigate *translingual plagiarism*. However, first it is important to consider translation historically and

conceptually, in order to define translation and understand how this can impact translingual plagiarism description and detection. The impact of interlanguage on both human and machine translation is discussed subsequently, in section 5.4.

5.3 Defining Translation

How translation is envisaged, how it is briefed, and its purpose largely determine the extent to which a translated text may or may not be considered plagiarism. Firstly, the brief helps contextualise the translation task; it allows the translator to learn more about the project, to understand her/his customer requirements, and to perform the job. Secondly, a translation is also impacted by the translator's and the customer's purposes, as much as it is by its own purpose. In this sense, a translation done by a professional translator for a company is bound to be different from an amateur translation, done for one's own personal use. In parallel, a translation is dependent on a translator's and/or customer's agenda, i.e. the matters to which the translation seeks to attend. These might be selling a company's goods or services, in the case of a professional translator working on a marketing document for a company, or translating the text from another language and passing it off as one's own, in the case of a student plagiarising an essay. In the former, the translator would be expected to provide a text that targets the product or service, and that is aimed at selling; in the latter, the student could be expected to produce a translation that is literal, word-for-word, but not necessarily so. These aspects, as translation studies have demonstrated over the last decades, not only influence how translation is approached, but how they are also influenced by it.

5.3.1 Traditional Translation Theories and Human Translation

Over time, translation has been studied from several different, often contradictory perspectives, most of which are based on the concepts of transfer and equivalence from one language to another – whether that transfer operates at the level of semantics, (surface) structure or other elements of the source language text. It is therefore not uncommon, for instance, to find metaphors associated with those concepts, which view the translator as operating a relationship between language and travel, with translation mediating the links between language and culture (Cronin, 2000); even more frequently, the translator is seen as a traitor who is not faithful to the original text

(Bassnett, 2002).

Faithfulness lies at the heart of the *Horace* model of translation, which is historically and chronologically the first major translation model, dating a few centuries BC. The model is based on the simple premise of *fidus interpres*, a phrase commonly attributed to Horace, the Roman poet after whom the model is named. However, as Lefevere and Bassnett (1998) claim, *fidus interpres* did not mean being faithful to the text, but rather to the customers, for the satisfaction of both parties involved in an act of interpreting. The Horatian model presupposed that although the translation task involved a negotiation between two languages that were of unequal status, where one – Latin – was privileged over the others, the concept of equivalence relied on elements such as the function, the design and even the target audience of the text, which are still recognised by contemporary theories of translation. For centuries, however, the Horatian model was overshadowed by the subsequent model of translation, the *Jerome* model (Lefevere and Bassnett, 1998).

The *Jerome* model of translation, which dominated the West from around the 4th until the 18th Century, is named after Saint Jerome, a Christian church father who translated the Bible into Latin. Saint Jerome believed that translating such a sacred text, which embodied the word of God, demanded being faithful to the source language text, with as little interference as possible. Ideally, the text should be transposed linearly and mechanically into the target language, i.e. by matching each word in the original with the corresponding word in the target text, in such a way that anyone with access to a dictionary or word list would be able to perform it. Although syntactically this strategy could cause serious problems that rendered a text unintelligible, as argued by Lefevere and Bassnett (1998), it remained the ideal model of translation (including of texts other than biblical) until recent centuries – and in practice is still one of the strategies used by underskilled translators. The linguistic text thus occupied the central place, which owing to its sacred nature was unchangeable, demanding absolute faithfulness. The influence of the Bible then ceased to be as powerful as it had once been, so the debate over faithfulness in translation moved on to a perspective where equivalence no longer operated as an imposition, but rather as a strategy freely adopted by translators to 'ensure ... that a given text is received by the target audience in optimal conditions' (Lefevere and Bassnett, 1998: 3).

Schleiermacher, however, was worried that having a translation read as a natural text in the target language would lead to a loss of the translated text. He thus argued

for a 'qualitative distinction between a 'true' and a 'mechanical' translation' (Gentzler, 2011: 62), which provided translators with the necessary ability to translate art. The principle behind the *Schleiermacher* model is that translation should be performed in such a way that the reader is able to grasp the language behind the original text; ultimately, a translation should read like a translation in order not to 'trick' the reader into believing that s/he is reading an original text. In other words, the reader of a German translation of a French text should get a feeling of French, just like the reader of a German translation of an English original should feel the English of the original. Schleiermacher illustrated the relevance of preserving the 'otherness' of the source text with examples of translation of artistic texts to contend that the German reader should be entitled to experience the French culture, even if that reader was unable to read French.

Although the worth of Schleiermacher's model has been appraised by theories of translation studies, especially in a post-modern view of literary translation, it can be challenged by technical translation and localisation professionals, whose aim is not to bring the reader to the author, as he claimed, but rather to bring the text to the reader. In many regards, contrary to literary translation, technical translation in general, but especially localisation, is intended to adapt the local conventions of the source text to those of the target text region. The two types of translation also differ in terms of authorship and authorial status; whereas by translating a literary text, the translator is entitled to a status (in the translation) that is equivalent to that of the author of the original[5], the translator of a technical text is usually not entitled to identical rights, with the copyright, in this case, usually belonging to the company commissioning the translation.

The argument for the combination of different translation strategies and theories gained strength, to a great extent as a result of Derrida's post-structuralist theories, and paved the way to more recent, 'post-colonial' research on translation (Bassnett, 2002). The work of Bassnett and Lefevere on cultural interaction is a good example of this post-structuralist approach. Rather than arguing for the 'faithfulness' to the original, these theories approach translation as a process of making meaning in a new language, and consequently not as 'the transfer of texts from one language into

[5]This status is granted by copyright law in many countries, including the Copyright Law of the USA, the UK Copyright Designs and Patent Act and the Portuguese copyright law, and is consubstantiated by international conventions, such as the *Bern Convention* and *The Universal Copyright Convention* (UCC), that provide protection in most countries, subject nonetheless to national laws.

another' (Bassnett, 2002: 6), but on the contrary as meaning negotiation between two different languages.

Nevertheless, (linguistic) texts indisputably are made of language material – grammar, semantics, pragmatics, discourse – and the translator's work consists of disassembling and unpacking this material and reassembling the signs of the original to compose a new text (rather than copying the original), in a new language. The task is then inevitably similar to one of creative writing of the target text, even if translation has not always been able to free itself from linguistic boundaries, contrary to what is at times desired by translation studies. In a comparative analysis of translation in the West and in China, Lefevere (1998a) suggests that the oral, collective tradition of translation of the latter contrasts with the written, single-author tradition of translation in the West, where the 'original' lies behind the translated text. As a consequence, the translated text remains bizarrely tied to the source text, with features of two different syntactic structures.

Such bizarreness is not, according to Toury (1995), necessarily indicative of an attempt to disguise the text to make it seem an original. On the contrary, he suggests, the awareness of using non-standard, even bizarre textual and linguistic features, indicating that a certain text has been translated, often results in consciously producing an original as if it were a translation. These fictitious translated texts, which pass as a translation even though they lack a corresponding source text in another language, are called *pseudotranslations* (Toury, 1995). Toury offers some explanations for pseudotranslations, and for what may lead authors into adopting it as a strategy. One of the main possible explanations is that the author might want to introduce elements of a foreign language in a new one, with the intention of introducing changes in the writing of one language, for instance. Conversely, it may reflect an act of subordination to another language that might be considered more prestigious, important or dominant, or an attempt to manipulate how a text is received by a certain target audience, by means of creating otherness, so as to avoid censorship against the author. Or, simply, it can be used as an act of 'culture planning' (Toury, 1995: 43). According to Bassnett (1998: 34), Toury attributes those 'indices of foreignness' that arise from the use of bizarre and non-standard forms of the language to the fact that a certain culture has a notion of translation, and most importantly that certain functions that are recognised and acknowledged by the members of that culture, are associated with that notion and its realisations.

The relevance of Toury's theory to research into translation and plagiarism lies, not so much with the reasons why authors choose pseudotranslations (that analysis is of more relevance to cultural studies than it is to plagiarism research), but rather with its potential to help understand and explain the linguistic devices – i.e. the non-standard linguistic forms at the basis of any (written) (linguistic) text – that can be expected from a translation. These non-standard forms are often conceived as being identifiable with the help of 'textual grids', a concept first used by Bourdieu to mean 'the collection of acceptable literary forms and genres in which texts can be expressed' (Gentzler, 1998: xiii), in which the definition of translation as a kind of language transfer is embedded. Those grids are a non-mechanistic process that gives rise to patterns of expectations in the target audiences, and determines the strategies adopted by the translators to cope with the translation. Lefevere (1998b) claims that the submission to these textual and conceptual grids, which determine the role of texts in a certain culture and the corresponding role in the target culture, take precedence over the translator's knowledge or even the mastery of languages. However, it is (also) in the linguistic forms that non-standard language is reflected.

The three distinct models of translation discussed above, together with Toury's theory of pseudotranslation, have in common the same element: the *translator's agenda*, even if this agenda may be more or less evident, or more or less hidden, depending on the circumstances. The Jerome model, for instance, is based on spreading the sacred word of God, ultimately attempting to evangelise the Latin-speaking world. It does so by translating the source text literally, while dodging metaphoric and figurative language that could possibly lead to a 'free' interpretation of the sacred translated text. Conversely, the Horatian model of translation attempts to please the target audience, by doing a work that meets their expectations, and that is in accordance with the satisfaction of both the initiator of the translation and their target audience. In Horace's time, the status of good translator was achieved by 'who could be trusted, who got the job done on time and to the satisfaction of both parties' (Lefevere and Bassnett, 1998: 3). The translator's agenda behind Schleiermacher's model, on the other hand, differs in that it aims to free the translators from the traditional canons of translation, while providing them with (an artistic-like) permission for rewriting, and hence a higher status. Given a status that is identical to a writer, the translator is no longer seen predominantly as the person who picks words from word lists to translate the text mechanically, but rather as the person who has control over the text. This raises

questions of authority and power, as the translation becomes the original in the eyes of the reader, when the reader does not speak the language in which that original was written (Lefevere, 1998a), or when the translation is sufficiently obfuscated to disguise its authorship.

Although these agendas are distinct from the plagiarist's agenda when translating a text from another language, they all have in common the fact that the translator is the one who holds the power, who knows whom the original belongs to, where it came from, and, ultimately, the one who chooses the strategy (most probably, in a sub-conscious manner) to obfuscate it. The issues of power and the translator's agenda therefore hold for amateur translators working for personal use, as they do for professional translators, especially because the diversity of translation strategies, and their influence on someone's writing, prevents someone from being accused of plagiarism simply because the text is written in a certain way that 'reads like a translation'.

5.3.2 Translation Theories and Machine Translation

These theories have been applied, more or less indirectly, to machine translation (MT) systems. These systems, which are intended to translate text to optimal quality with the minimum possible human intervention, have evolved over the last decades, now achieving results that are immeasurably better from those of the ideal systems of the 1960s, when machines were expected to perform translation in a way that is similar to that performed by humans. The need for human pre- and post-processing of machine translated texts has been acknowledged in the meantime (Slocum, 1984), but the main principle of MT remains the same: being able to perform a complete translation independently, with no human intervention, using specific software, grammars, and sets of rules (Seneff, 1992).

Feeding a set of rules into the system to 'teach' it how to translate is the basis of rule-based approaches. However, after several years of experiments it was demonstrated that using just rule-based methods proved to be inflexible, and raised problems of reusability. Once the system was used in domains other than those for which it was trained, results were of poor quality and rules had to be re-written (Macherey *et al.*, 2001). Research into the field of rule-based translation was not dropped altogether, but the focus shifted to data-driven methods, in particular statistical MT approaches (Koehn, 2010). Contrary to rule-based systems, statistical MT approaches consist of integrating in one system one or several methods in combination, usually aligning

naturally-occurring text by matching source and target texts. These methods may vary, depending on the linguistic units that they aim to process, such as words, phrases or sentences; they can also build upon 'language models', which consider for example n-grams, or even lexical models, that consider lexical translation and, when more sophisticated, can also take into account alterations (e.g. deletions, additions, duplications). This alignment information is then retrieved based on probabilistic models, which consider the relative, rather than absolute position of the words. In other words, these MT systems tend to consider, not only the words, but also the differences in the word positions (Vogel *et al.*, 1996), but can also take into account standard distributions to model the system, or even collect statistical information e.g. on the word co-selection and calculate the probability distributions accordingly (Koehn, 2010).

There have been different approaches to statistical machine translation over the last decades. The system proposed by Vogel *et al.* (1996), for example, consisted of aligning the source and target texts on a word-to-word basis. The model was found to be limited in that it established a correspondence of each word in the source to one word in the target text, failing to identify and process groups of words (Och *et al.*, 1999). Additionally, this posed the risk that words be taken as a grammatical unit in translation, with one single function, and hence invariable. In an attempt to address this problem, later research proposed methods to align source and target texts above word level; in this respect, Och *et al.* (1999) demonstrated that the text could be aligned at phrase level, first, and then at word level. Statistical machine translation systems then evolved to working on a many-to-many basis (Macherey *et al.*, 2001), where many word combinations in one source could possibly correspond to many word combinations in the target. This provided results considered significantly better than the ones of rule-based systems.

One example of the application of statistical machine translation is *Google Translate*[6]. The tool was developed to translate words, sentences and pages instantly – and for free – from and to 58 different languages. It uses previous human translations to align the source and the target, and thus collect a set of patterns[7]. *Google*'s translation tool gained a popularity[8] that other systems missed. *Google*, however,

[6] *Google Translate* — `http://translate.google.com`

[7] See `http://translate.google.com/about/index.html`.

[8] According to the 'Google Translate Usage Statistics Website' (`http://trends.builtwith.com/widgets/Google-Translate-Widget`), on 27/06/2011, 99,089 websites used *Google Translate Widget* alone, including more than 4,000 of the top million visited websites in the World.

recognises a limitation that not all users of their translation system seem to be aware of; that the translation is generated by machines, and hence is not perfect. These machines use *web crawling* to crawl webpages in order to find source and target texts, translated by humans. Consequently, the higher the number of human translated texts available, the higher the likelihood that the translation has been evaluated by humans, and consequently that it is good – or, at least, closer to human translation. *Google*'s solution to this quality assurance lies with post-processing and assessment of the quality of the translation that, it can be speculated, is hardly ever performed by skilled professionals. Conversely, using the method of web crawling increases the risk that machine-translated, non-processed texts are fed into the system, thus increasing the amount of 'noise', and replicating it as the usage of the system increases.

At the same time, the assessment of the quality of a translation has traditionally considered only judgements of the results, leaving aside any considerations of the quality of the performance of the translation task. It is, therefore, common to judge a translation as literal or word-for-word based on those non-standard linguistic forms, while neglecting the possibility that such forms might result from other constraints. Toury's theory of pseudotranslation above can contribute to explaining the translator's performance by analysing the indices of foreignness contained in the text, so that a text first intuitively thought to have lifted from a source in another language might be disproved as such by attributing those indices to other factors, such as subconscious, non-standard linguistic forms. These forms occupy a 'middle ground' between the source and the target language, which Selinker (1972) calls *interlanguage* – a concept that is discussed in the following section.

5.4 Interlanguage: The Go-between

Studies on bilingualism and second language acquisition have tried, for several decades, to describe the use of non-standard forms by non-native speakers of a language. Traditionally, such non-standard forms have been attributed to a phenomenon that Weinreich (1953) called 'interference', which he studied in relation to bilingualism. Interference is a deviation from a standard language norm, and reflects the introduction of foreign elements at various linguistic levels, such as morphology, syntax and vocabulary. It results in the rearrangement of patterns which are attributed to the fact that a bilingual speaker is the one who is familiar with more than one language, and can alternately use two languages, with such an alternation giving rise to a language

contact situation. Additionally, Weinreich pointed out that interlingual identifications resulted also from either 'formal similarity' or 'similarity in preexisting functions' (Weinreich, 1953: 39) between phonemes, grammatical relationships and semantic features in the two languages.

Selinker (1972), partly in response to the belief that Weinreich's theory left some unanswered questions, but based on Weinreich's practical assumption of *interlingual identifications*, argued that second language learning involves three linguistic systems; the system of the mother tongue; the system of a target language; and the competence of a speaker in a second language – the 'interlanguage' (IL). Interlanguage was a term coined by Selinker to mean 'a separate linguistic system based on the observable output which results from a learner's attempted production of a TL norm' (Selinker, 1972: 214). He argued that speaking a second language involves an attempt to achieve a 'meaningful performance' in the system of the target language, i.e. 'an attempt to express meanings which he may already have, in a language which he is in a process of learning' (Selinker, 1972: 210). Selinker's argument that those 'interlingual identifications' exist within a latent psychological structure that is 'activated when one attempts to learn a second language' (Selinker, 1972: 211) are at the basis of his interlanguage theory.

The *Interlanguage* theory is based on the principle that the same meanings are not expressed identically, i.e. using identical sets of utterances, by a native speaker and a learner of that language. To demonstrate and explain his point, Selinker analysed three sets of utterances or 'behavioural events' that he considered to be relevant to interlingual identifications:

- utterances in the learner's native language (NL) produced by the learner;
- interlanguage (IL) utterances produced by the learner; and
- target language (TL) utterances produced by the native speakers of that TL.

Based on this analysis, Selinker compared the competent native-speaker, who acquires the language and its principles of organisation without being explicitly taught, to the incompetent learner, who focuses on one norm of the language s/he is attempting to perform, and overgeneralises TL linguistic material, to claim that speakers of a particular NL tend to keep, in their IL relative to a particular TL, 'fossilizable linguistic phenomena' (Selinker, 1972: 215), i.e. rules, subsystems and linguistic items of that NL. As Finegan (2012: 522) explains, 'fossilization underlies the nonnative speech

characteristics of someone who may have spoken the target language for some time but has stopped the process of learning'. An 'interlingual situation' arises when a particular combination of NL, TL and IL elements is obtained, often resulting from a speaker or writer's conscious or subconscious realisation that they lack 'linguistic competence with regard to some aspect of the TL' (Selinker, 1972: 219). This lack of linguistic competence, as Finegan (2012: 522) argues, reflects more on certain linguistic aspects than others: once the learning process stops, interlanguage stabilises, and additional language acquisition ceases, except for vocabulary. Selinker concluded, by experimental, empirical demonstration, that second-language learners tend to 'backslide' from a TL norm toward an IL norm, and not actually toward the speaker's NL or randomly.

In his earlier studies, in particular, Selinker considers that the influence of interlanguage operates unidirectionally, rather than bidirectionally, so that consequently the impact of interlanguage is studied only on the second language, and not on the native language of the speaker or writer; hence, his focus on investigating second-language acquisition, which has typically considered the native-language interference in second language learning. However, the hypothesis that IL can also influence the native language has been considered by studies of bilingualism, especially in childhood (e.g. Weinreich), to talk about 'two coexistent systems', rather than a 'merged single system' (Weinreich, 1953: 9). Selinker himself does not reject this hypothesis, at least not explicitly, despite his criticism of Weinreich's theory. Several authors have strongly argued for the fact that language systems are not subject to change[9] once 'mature', under the assumption that the weaker linguistic system is attached to that of the stronger language and is dominated by it (Weinreich, 1953); yet, other authors have proved otherwise. A specific example is that of Major (1992), who documented, in his study with adults, that native language loss can occur even in cases where there is language contact, or where a subject is learning a second language. Likewise, Pavlenko (2000) later argued that adults' 'matured', native language systems are unstable and permeable. Selinker's argument of 'interlingual identifications uniting the three linguistic systems (NL, IL, and TL) psychologically' (Selinker, 1972: 221), even if applied to a context of an individual producing utterances in a TL, leaves room to study the influence of the IL on the NL of the speaker or writer. In fact, Selinker's terminology allows for an ambiguous interpretation; on the one hand, his theory considers NL

[9]For a good discussion on this topic, see Pavlenko (2000).

(native language), IL (interlanguage) and TL (target language), where TL is taken to mean a second language, while on the other hand not explicitly claiming that TL and second language are synonyms. Therefore, this interlanguage theory allows for an interpretation of the process of IL influence on another language as a bidirectional one, which is a consideration that he was later shown to be willing to accept (Selinker and Lakshmanan, 1992).

Therefore, if IL is a bidirectional process, it seems plausible to admit that the success of language learning and use, and not just second-language acquisition, depends on a speaker or writer's ability to reorganise the linguistic material from an IL to attempt an identification with a particular language, not just the NL. Accordingly, the 'surface structures of IL sentences' (Selinker, 1972: 229) would apply to language learning in general, and not just second-language learning. Even more importantly perhaps, Selinker's claim that fossilisation (the linguistic phenomena that speakers of a certain language will tend to keep in their IL when speaking or writing another language) underlies the surface structures of a speaker's IL (Selinker, 1972), and that these phenomena tend to reappear regularly in IL performance, applies also to bilingualism and translation. However, his focus on the psychological latent structures that are reactivated in the speaker's mind, rather than acquired by instruction, to explain the fact that some speakers manage to achieve a native-speaker 'competence' (i.e., a performance that is indistinguishable from native speakers) might represent an oversimplified problematisation. Firstly, because it tends to see competence in the Chomskian sense of a native speaker's competence that is monolingual, instead of conceiving of the possibility that native speakers of nearly any language are exposed to existing multilingual contexts (Major, 1992). Moreover, if it is a matter of competence, than it also applies to the native language as an effect of the second language. As Fillmore (1991) demonstrated based on the results of a study conducted with children who were native speakers of languages other than English, living in the United States and that were exposed to environments where English was the dominant language, such massive exposure can affect both the process of learning English and the retention and use of their primary languages. A similar conclusion was reached later by Pavlenko (2000), who demonstrated that, not only can the interference (i.e. involuntary influence) of the native language (L1) be noticeable in a second language (L2), but also L2 is taken to influence L1. She found that L2 can impact the competence, performance, processing and conceptual representations based on L1, including aspects of

phonology, morphosyntax, lexis, semantics, pragmatics, and rhetoric, thus challenging Selinker's argument that the lack of 'native-speaker'-like competence reflects mainly in syntax. Additionally, she suggested that such influence is greater as a result of a longer exposure to L2 or of a high level of L2 proficiency, although other extralinguistic factors have also been found by previous studies (e.g. Seliger and Vago (1991)) to account for that influence – including language prestige, social status and the desire to integrate and assimilate the L2 environment.

As these latter approaches to interlanguage – which consider language loss and language transfer – indicate, IL may be taken to impact the native language, as much as it does influence the second language. In the light of these approaches, it can be argued that the concept of the bilingual translator, who has an impermeable native-language competence in at least two languages, is a rare and sometimes 'idealised' version of reality. To a certain extent, every speaker is exposed to one language more prominently than to another, and this prominence, if we consider previous research on interlanguage, tends to overtake the other (Fillmore, 1991; Pavlenko, 2000). The 'overtaking language' would then play the role of dominant language (i.e. the language to which a speaker or writer is more prominently exposed at a certain time, and not necessarily his/her native language, or the language they are more comfortable with), and be expected to operate at different levels, depending on the speaker. Hypothetically, however, it could be speculated that it would be greater if the speaker or writer is not bilingual, or a professional, trained translator, owing to the fact that the individual speaker or writer would be focused mainly on trying to reproduce the source language text.

These theories are also inevitably bound to influence the translation process, and even more so when this translation is not performed by a skilled, trained translator, who is conscious about the preventive measures to adopt in order to produce a version of the original document that is – or attempts to be – free from non-standard linguistic forms. Regardless of whether translation is seen as transfer or as equivalence between two different cultural and linguistic systems, it necessarily implies an interaction between two languages, one being influenced by the other, either owing to the existence of one intermediary language (the interlanguage), or owing to language transfer influencing both native languages (in the case of bilingual speakers/writers), or the native and second language.

Based on this assumption, in the following section the concepts of translation and

interlanguage will be applied to propose a set of linguistic markers and discursive devices that can describe and explain the linguistic processes behind interlingual plagiarism, and then propose a method to detect translingual plagiarism.

5.5 Translingual Plagiarism: An Analytical Framework

Translingual plagiarism can involve mainly two scenarios. The first is where a plagiarist takes an original work published in another language, translates it into his/her language and publishes it as their own. A second scenario is where the plagiarist takes an original work that has been published in one language, translates it into another language (his/her own, or the language in which the derivative text is expected to be written) and passes it off as their own. Contrary to the first scenario, which is aimed for wider audiences, the second one describes a case where plagiarism is for personal use, and is usually intended to be read by smaller audiences. Consequently, the first scenario usually has legal implications, involving publishing issues and copyright violation, whereas the second scenario is often judged morally. The translation job is also expected to be performed differently in the two scenarios; in the first scenario, a professional, careful translation would be expected; conversely, in the second scenario a less proficient translation is more likely to be the result. Considering the reasons why someone (e.g. a student) plagiarises, as discussed in section 1.3, whether that reason is lack of time, pressure, mental fatigue or even lack of academic writing skills, or simply laziness, when plagiarising a translation, the plagiarist could hardly be expected to produce a high quality translation. Under these circumstances, it is very likely that the student would make a minimal effort, and spend the least possible amount of time on the task; speculatively, his/her main concern would be translating the meanings, by doing their own interpretation of the text and rewriting it in another language. A second way to perform this task quickly – and for free – would be to use one publicly available machine translation system to input the original text and instantaneously obtain the translated text. The latter suits the plagiarist intentions well, since it permits translating the main ideas in a short space of time (Slocum, 1984), while allowing the reader to have a gist of text even when the translation produced is rough (Koehn, 2010). More careful plagiarism could then be completed by editing and revising the text for the surface structure.

However, as discussed in section 5.3, translating consists of more than simply transferring text from one language into another; it involves a complex process of

meaning negotiation, which accordingly requires negotiating lexico-grammar, both at the sentence- and discourse-level. Moreover, given the mutual language transfer and influence of two linguistic systems, as discussed in 5.4, it is not always easy, even for trained translators, to (re-)write a text in another language that conveys the meaning(s) of the original without 'compromising' the translation, by avoiding hints to the fact that the text originated in a foreign language. As a consequence, translated texts are often permeated with linguistic forms of the original, source language text, and hence are said to 'read like a translation'; in this case, the translated text shows non-standard linguistic elements and 'indices of foreignness' – or a combination of such elements. One such example is the case of passivisation in Portuguese and French, on the one hand, and in English on the other hand; Portuguese and French use passive structures differently from English because they have an 'impersonal' grammatical strategy to convey information that can only be conveyed in English using passive structures; likewise, the traditional SVO order in English, albeit allowing a slight degree of flexibility, can be reordered diversely in Portuguese, and accomplish more indirect realisations than English. The mismatches resulting from uncommon reorganisations of the words, especially when in large doses, can lead a native speaker of the target language into intuitively evaluating the text as a 'poor translation' – i.e. unnatural, prone to errors, and often indicative of the original source language. In this sense, a poor translation can be easily identified as not being an original text, and raise suspicion that the text may have been plagiarised.

Within cases of plagiarism involving texts translated from other languages, this is the type that is easiest to detect, despite the skills required to distinguish a poorly translated text from a poorly written one. On the contrary, a good translation (one that 'reads like an original', and seems to have been produced independently of any other source text) is less prone to raise that type of suspicion, unless, of course, the reader has previously read the original text. Good translations tend to mediate meanings impeccably between a source and a target language, in such a way that the original text tends to lose its original surface structure, which is then replaced with that of the target language. Therefore, a poor translation could mean violating the target language standards at all levels of language, including morphology, morphosyntax, semantics, pragmatics and discourse, for example by combining short sentences into longer ones, making a different use of personal pronouns, using passives when impersonal structures are more common, translating multi-word units, phrases and idioms

literally when their use is not common in the target language, etc.

5.5.1 Interlingual Transfer and Influence

In order to describe and explain the changes operated by means of the interlingual transfer and influence of a source language on a target language, a systematic approach is necessary. Among the scarce research conducted in the field – and usually as part of studies on bilingualism and second language acquisition (SLA) – the work of Pavlenko (2000) in particular was found to be highly relevant, since it proposes a 5-class – yet, admittedly open – classificatory framework, which she named 'unitary classificatory framework' (Pavlenko, 2000), to describe the linguistic processes that operate at the level of language transfer and language influence, namely L2 influence on L1 phonology, morphosyntax, lexicon and semantics, concepts, pragmatics and rhetoric.

The framework used in this chapter is an adaptation of Pavlenko's framework, since not all categories and linguistic processes used by her original model are applicable to the analysis of linguistic plagiarism. Firstly, this framework will be used for a purpose other than the one used by bilingualism, multilingualism or SLA theories: to describe the linguistic systems used, and not to assess the linguistic competence of speakers and writers. Secondly, in order to avoid the identification of L1 with native language and L2 (and possibly L3, L4, L5, ...) with target language(s), L1 will be renamed TL (target language) and L2 will be renamed SL (source language). This shall also make the description of the directionality of the reuse easier. Additionally, the focus shifts from the speaker to the text, so that rather than concentrating on the L1 (native language) speaker, the analysis can focus on a SL text and on one or more TL texts. Considerations of which language is the speaker's native language and which one is the speaker's second language are not as relevant at this stage as those involved in the identification of the source and target language texts. Conceptually, this also permits categorising cases where a plagiarist uses texts from more than one source language. Moreover, the linguistic devices required by translingual plagiarism detection do not exactly match the ones used by the unitary classificatory framework to analyse language transfer and influence in bilingualism theories. Influence on morphosyntax, lexicon and semantics, concepts and pragmatics – and, possibly, rhetoric – remains relevant, whereas influence on phonology can be discarded. Concurrently, new devices need to be considered, such as punctuation, spelling, and discourse grammar.

Influence on morphosyntax will consider, for example: sentence structure (word order rules that are missed or restructured); extension of SL rules for agreement, prepositions, adverbs, adjectives, pronouns, pre-/post-positioning; paradigmatic TL conjugations; and verb usage. SL influence on lexicon and semantics can operate at the levels of the lexicon, semantic networks or lexical processing, and will therefore consider especially forms of lexical lifting, such as loanwords (these are lexical borrowings *per se*, i.e. lexical items from one language adapted phonologically or morphologically for use in another); loan blends (hybrid forms which combine elements of both languages); loan shifts (which are often referred to as 'semantic extension', i.e. TL words which acquire the SL meaning); and loan translations (also known as calques, i.e. literal translations of SL words, phrases, or expressions). The influence of SL on TL concepts will focus on the linguistic, rather than the psycholinguistic aspect, and includes cases where concepts of the source language are transposed linguistically (but not necessarily coherently) to the target language. Other instances that need consideration are punctuation and spelling (e.g. to describe cases where the source language punctuation conventions and particular spelling are brought to the target language text), as well as discourse grammar and informational organisation and packaging, to enable issues of coherence and cohesion, and aspects of theme and rheme, to be addressed.

The following are the five categories proposed, adapted from Pavlenko's 'unitary classificatory framework':

- *borrowing transfer*: consists of adding new SL elements to the TL text;
- *convergence*: consists of creating a unitary system that includes elements that belong neither to the SL, nor to the TL;
- *shift*: departs from structures or values of the TL, by approximating to those of the SL;
- *restructuring transfer*: consists of incorporating SL elements into the TL, originating changes, substitutions, or partial shifts;
- *attrition*: loss of some TL elements as a result of the SL influence.

This classificatory framework is mostly relevant to linguistically explain the strategies and the moves used by the plagiarist when lifting from an original in another

language. However, as the next section will demonstrate, cases of translingual plagiarism can also be demonstrated by simple comparison, after the source is identified. A method aimed at this identification is discussed in the next section.

5.5.2 Detecting Translingual Plagiarism

The distinction in section 5.2 between *linguistic plagiarism* and *plagiarism of ideas* – or copying ideas – is key to investigating cases of plagiarism in translation, since, although linguistic plagiarism and plagiarism of ideas can be investigated independently, in fact the detection of linguistic plagiarism can lead to the detection of the plagiarism of ideas. This section addresses the detection of translingual plagiarism – which in a sense is plagiarism of ideas – using linguistic analysis customary in detecting linguistic plagiarism.

The borrowing of ideas is often considered in cases of copyright infringement, which applies for instance when someone gets copyrighted material translated into another language and published without the consent of the copyright holder, i.e. the person(s) holding 'the exclusive right to reproduce, adapt, distribute, perform and display the work' (Garner, 2009: 386), that is usually the author or the publisher[10]. A copyrighted work can therefore be plagiarised if someone else uses it and passes it off as their own – in which case the copyright is also infringed. However, given a hypothetical case where the original author gives someone else permission to use their own work without acknowledgement, this may not represent copyright infringement, but is still plagiarism.

This section leaves aside considerations of copyright to focus instead on translation-based plagiarism. Jones (2009) reports that, among native speakers of English, it is becoming increasingly common to back-translate a text originally written in English into another language, and then translate it again into English (using a MT tool) to change the wording of the original, effortlessly. And as discussed in section 5.5, cases of translingual plagiarism are among those that are most complex to detect. Nevertheless, taking Selinker's concept of interlanguage, its influence on the surface structure of a text, and the principle that interlanguage may influence both the native and the second language, together with Pavlenko's claims that such influence reflects on both

[10] In Portugal, and possibly in other countries, 'copyright' is referred to in law as 'author's rights', but it has been argued (Ascensão, 1992) that the law should change to protect more the non-transferable moral rights of the authors, and less the financial rights of the publishers.

the second and the native language (whenever they act as the target language), it is hypothesised that instances of plagiarism can be detected by means of an analysis that reverses the translation process. This process works by back-translating a suspect text into the suspected original language, or using a software package to 'guess' the most probable language of the original. The linguistic markers and discursive devices described in section 5.5 can then be used to find empirical evidence that contributes to a theory aimed at describing translingual plagiarism, using different language pair combinations (e.g. Portuguese/English, English/Portuguese, Portuguese/Spanish). These criteria are used to provide investigative, descriptive and analytical clues to plagiarism. Subsequently, a computational approach is proposed to detect this type of plagiarism that suits forensic linguistic research into plagiarism well.

Translingual plagiarism can either result from (a) a professional translation of an original in another language; (b) an 'amateur' translation performed by an untrained person; or (c) a work performed using machine translation (possibly more common in the current IT era). Contrary to professional translation, the results of 'amateur' translation and MT are often considered to be poor in that they tend to be (too) literal, word-for-word, and reflect the surface structure of the original, which often impact syntax, semantics and lexico-grammar. For instance, previous research[11] demonstrated that news texts written in one language (*Language A*), based on newspaper articles/newswires originally written in another language (*Language B*), tend to retain a structure that is more similar to texts written in the language of the original source (i.e. *Language B*) than to other texts (regardless of their genre), written in the target language in which they are published (i.e. *Language A*). Hence, a news report on national politics would be written in a style that is different from another report on international politics, in spite of the possibility that both might be signed by native speakers of the same language, and the (more or less) tight editorial decisions made nowadays by newsrooms[12].

This empirical study supports, for the most part, the arguments for language transfer and language influence of the source language on the target language, confirming the assumption that interlanguage acts mainly at the level of the surface structure of the text (Selinker, 1972). As professional translators will know from experience, trans-

[11] Paper presented at the 9th Conference of the International Association of Forensic Linguists.

[12] For an illustrative example, refer to the *Reuters Handbook of Journalism* (`http://handbook.reuters.com/index.php/Main_Page`) or the *Livro de Estilo do Público* (`http://static.publico.pt/nos/livro_estilo/nova/index.html`).

lating a text into another language and making it read 'naturally' is not as easy and simple as it might at first suggest. It is consequently easy to fail to meet language standards, for instance at the level of syntax, lexico-grammar, morphology and informational organisation, as well as of cohesion (i.e. at the level of 'the non-structural resources for discourse' (Halliday, 1994: 309)) and coherence (i.e. at the level of 'the realization of semiotic orders 'above' the language' (Halliday, 1994: 339)). These, together with considerations of misinterpretation of the source text, terminology and named entities and punctuation, are some of the linguistic markers and discursive strategies that can hypothetically contribute to the detection of translingual plagiarism, while at the same time providing clues of directionality, in order to determine which text is the original and which one is the derivative. Determination of directionality is particularly relevant in cases where text production is proven to be contemporary (Turell, 2008), or where the likelihood that one person had access to the other person's text is not evident.

Consequently, by translating the suspect texts back to the expected language of the original, linguists may trigger potential cases of plagiarism for investigation, as this procedure allows one first to find the putative original, and then to compare the two texts for lexical parallels and identical lexical strings that are said to prove the instances of plagiarism. A few empirical experiments have been conducted, first with one sentence taken from *The Guardian* newspaper to see how the procedure works, and then using real data from the CorRUPT corpus. The results of these experiments indicate that free machine translation, e.g. *Google Translate*[13], performs well in finding the potential source. For the purposes of these experiments, a machine translation system is preferred to human translation because it returns the results instantaneously, and avoids any possible human translators' bias. In practice, it involves:

1. Taking a (suspect) text written in *Language A* (in this case, Portuguese);
2. Translating it into *Language B* (in this case, English) using *Google Translate*;
3. Checking the translated text for non-standard linguistic forms;
4. 'Googling' the translated strings.

The third step, checking the translated text for non-standard linguistic forms, is

[13] *Google Translate* (`http://translate.google.com`) is used for its degree of popularity (see section 5.3), and for the continuing, daily updates to the system.

not absolutely necessary, but allows narrowing down the search strings in more fuzzy cases, and hence increases search efficiency.

As a result of the procedure above, the more fluent the machine-translated text (i.e. the less it violates the surface structure, as well as the discourse structure of the target language, *Language B*), the greater the likelihood that the text has been produced using a discursive strategy that is close to the one of the target language. This can be explained by the fact that machine translation has now reached a stage where lexical items and even terminology and named entities are frequently replaced with a considerable degree of correctness, but the surface structures are often missed or transferred incorrectly. Conversely, if the machine-translated text is *odd*, then the more likely it is that it has been originally produced in a language other than the target language, and consequently may indicate that the text is original. The following example of a headline posted on *The Guardian* website on 11th June 2009 illustrates this point:

`Three guilty of Ben Kinsella murder.`

Google Translate returned the following translation to Portuguese:

`Três acusados de homicídio Ben Kinsella.`

The headline in Portuguese is grammatically correct and could be found in a Portuguese newspaper, but any newspaper reader would easily notice that the headline shows indices of foreignness; if the news had been originally written in Portuguese, a different headline would be expected. To confirm this assumption, a media professional was asked to rewrite the headline in Portuguese based on the information conveyed by the English headline. The proposal she suggested was:

`Três assassinos de Ben Kinsella condenados.`

Syntactically, this sentence sounds more 'natural' in Portuguese than the translated version above; semantically, the original headline provides information that the translated headline misses, but the rewritten headline emphasises: the adjective 'guilty' indicates that the murderers were convicted, but the machine translated headline only indicates that they were 'accused'; the rewritten headline ('Three convicted for murder of Ben Kinsela') is semantically similar to the original headline, and conveys the core information. But machine-translating the sentence

```
Três acusados de homicídio Ben Kinsella.
```

back to English returned the headline

```
Three charged with Ben Kinsella murder.
```

The following examples, in tables 5.1 to 5.5, also illustrate this method. The sentences used are from a case of newspaper plagiarism, involving the Portuguese quality newspaper *Público*, and are part of the CorRUPT corpus. For each table, the alleged instance of plagiarism is indicated in the first row ('PT'), followed by the corresponding machine-translated version ('PT-EN') in the second line; the original version, in English ('EN'), is provided in the third line for comparison. The identical, overlapping text is highlighted in **bold** typeface. Only the exact matches are highlighted in bold; consequently, the amount of overlapping text would increase if synonyms and words with the same lemma were also considered. These examples demonstrate that most instances of the machine-translated text (row 'PT-EN') are identical to the English original (in row 'EN'). Therefore, although the translation sometimes shows some errors (which derive mainly from the edits introduced in the suspect translated text), they could easily be read and understood by a native speaker of English.

These results demonstrate a high degree of lexical overlap, besides the identity in surface structure, indicating that machine translation is able to handle lexical items efficiently.

PT	Faz com que a melanina se combine com o oxigénio, o que produz o escurecimento da pele.
PT-EN	**Causes the melanin to combine with oxygen**, **which** causes darkening of **the skin**.
EN	**causes the melanin to combine with oxygen** (oxidize), **which** creates the actual tan color in **the skin**.

Table 5.1: Example of translingual plagiarism.

Figure 5.1 proposes a method for describing a procedure of translingual plagiarism detection.

Section 'A' of the diagram describes a suspect text written in Portuguese that was machine-translated into English. Although the text shows some (mainly) grammatical errors, it is sufficiently fluent to be easily understood by a native speaker of English.

PT	Pode ser quase completamente bloqueada pelos protectores solares.
PT-EN	It can be **almost completely blocked by sunscreen**
EN	is **almost completely blocked by** virtually all **sunscreen**s

Table 5.2: Example of translingual plagiarism.

PT	A chave deste novo autobronzeador está num extracto de plantas chamado forskolina que, nas experiências da equipa, protegeu ratinhos sem pêlo de radiação ultravioleta e permitiu-lhes desenvolver um bronzeado natural, estimulando os seus melanócitos.
PT-EN	**The key** to this new self-tanning is **a plant extract called forskolin** that the experience of the team, **protected** hairless **mice** to ultraviolet radiation **and allowed them to develop a natural tan by stimulating** their **melanocytes**.
EN	**The key** chemical, **a plant extract called forskolin**, **protected mice** against UV rays **and allowed them to develop a natural tan by stimulating** pigment-producing cells called **melanocytes**.

Table 5.3: Example of translingual plagiarism.

However, an Internet search using the machine-translated string confirms that the text is almost identical to a previous text, with a lexical overlap of (say) 80%[14].

Section 'B' of the diagram shows a case where the same procedure described in 'A' is applied, but where the automatic translation returns a text with a higher index of foreignness. This text is prone to be less fluent than the one described in 'A', and presents a higher degree of difficulty to native speakers of English. Searching the Internet for the translated text returns hits with some degree of overlap, but this is equal to or lower than 50%. 'A' therefore shows a higher likelihood than 'B' to be plagiarism, demonstrating that error checking has a strong discriminatory power to help identify texts that derive from originals in other languages. By this token, a higher index of foreignness is a good indicator that the text derives from literal translation or from wrong language transfer.

[14]This percentage of overlap should be seen as a relative, rather than an absolute value on the grounds that it is not possible at this stage to determine the threshold of overlap at which a text would be considered plagiarism.

PT	A capacidade de se bronzear - (...) - é controlada pela hormona de estimulação dos melanócitos, que se liga a uma proteína que existe no exterior destas células. Esta proteína, que se chama receptor de melanocortina 1, funciona mal em muitas pessoas que têm a pele clara e o cabelo ruivo. É por isso que não se conseguem bronzear, e ainda por cima correm maiores riscos de desenvolver cancro de pele.
PT-EN	**The ability to tan** - (...) - **is controlled by hormone** stimulation of melanocytes, **which binds to** a protein that exists outside these cells. This protein, called **the melanocortin 1 receptor**, malfunctions in **many people** who have **fair skin and red hair**. That is why we can not **tan**, and on top of a higher risk of developing **skin cancer**.
EN	**The ability to tan is** largely **controlled by** a **hormone** called melanocyte-stimulating hormone, **which binds to the melanocortin 1 receptor** (MC1R) on the outside of melanocytes. **Many people** with **red hair and fair skin** have a defect in this receptor, meaning they find it almost impossible to **tan** and are prone to **skin cancer**.

Table 5.4: Example of translingual plagiarism.

In cases where such borrowing might be less evident, it is necessary to compare the source and target texts. From a linguistic perspective, especially if it is considered that linguistics should deal only with textual comparison, leaving aside all instances of plagiarism of ideas (including translation), the only way to compare two texts is having them in the same language. Therefore, based on Selinker's concept of surface structure, it seems plausible to machine-translate a suspect text to the language of the original and then compare the two texts. This method offers several advantages, when compared to the sophisticated procedures described in section 5.2. Firstly, contrary to the methods above, it is very simple and uses tools publicly available, with the advantage of using the same tools as the plagiarists. Secondly, it does not require specific computational resources, such as thesauri and dictionaries, that are costly and of limited availability. Thirdly, since this system works based on crowdsourcing techniques, the system remains updated over time, with no additional investment required. Additionally, although a comparison of the plagiarising text against the plagiarised version is ultimately necessary to explain and justify the reuse, a known source is not nec-

PT	Numa segunda série de experiências os cientistas usaram ratinhos susceptíveis ao cancro, expondo-os ao equivalente a uma a duas horas de Sol na altura do meio-dia solar, diariamente, durante 20 semanas.
PT-EN	**In a second** series of **experiment**s the scientists used **mice** susceptible to cancer, exposing them **to the equivalent of one to two hours of** sunshine at the time of solar noon **each day for 20 weeks**.
EN	**In a second experiment**, a particularly cancer-prone strain of **mice**, also bred to lack effective MC1Rs, were exposed **to the equivalent of 1 to 2 hours of** midday Florida sunlight **each day for 20 weeks**.

Table 5.5: Example of translingual plagiarism.

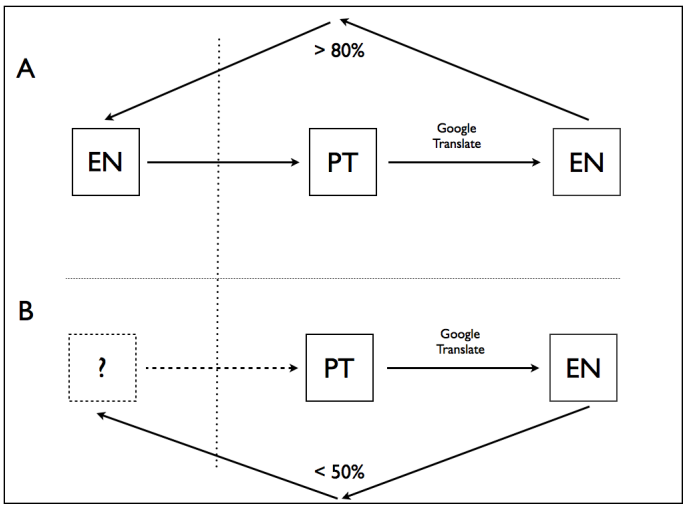

Figure 5.1: Diagram of translingual plagiarism.

essary to initiate the detection procedure, since the machine-translated version of the suspect text is used to search the Internet for similar or identical texts.

The empirical data analysed so far indicates that this method can provide good results, mostly from an investigative, but also evidential perspective. In order to verify the potential of this approach, a comparison of the two texts (the original and the plagiarism) was performed using plagiarism detection software[15]. The analysis showed a word overlap of almost 70%. These results assist the detection procedure, but may be impacted by differences e.g. in lexical density; cases of translated texts that are heavily edited can represent an additional challenge. Based on a small study with collected

[15] *CopyCatch Gold* was used.

texts which showed that handwritten texts tend to be less dense than wordprocessed texts, Grant (2005) strongly argued that, as texts are reworked, they tend to become lexically denser (per thousand words). Supporting his claims with the work of Laviosa (1998), he contends that this density is particularly increased by translation. Laviosa (1998) concluded that, although the percentage of content words in translated narratives is lower than that of grammatical words, the lexical density in translated narratives tends to be higher than in their original versions. Consequently, in these cases, derivative texts tend to distance themselves from the original. However, since: (a) Laviosa's findings resulted from the analysis of professionally, carefully, commercially translated narratives, rather than amateur or machine translation; (b) the likelihood that a text is heavily edited lexically is very low, especially in a case of student plagiarism where the best results are sought with the minimum effort; and (c) although the amount of lexical overlap can decrease as the lexical density increases, that overlap can still provide clues to the lifting; then the performance of the method should not be impacted significantly. This approach therefore supports previous claims that the tools that help plagiarise also help detect plagiarism (Coulthard and Johnson, 2007).

Even if MT is used to plagiarise from a text in a foreign language, the task of lifting while disguising the original authorship involves a considerable effort on the part of the plagiarist to pass the derivative, plagiarising text as their own. Therefore, considering, as argued in chapter 4, that the stronger the effort to obfuscate the authorship, the clearer the plagiarist's intention, cases where a minor, even if insufficient reference is provided can raise the doubt, and may indicate negligence or unintentional plagiarism. Conversely, using translation as a plagiarism strategy, with no reference at all to the original, reveals one of the most intentional attempts to plagiarise.

5.6 Chapter Summary

This chapter discussed the concept of translingual plagiarism, where translation is used to reuse texts from other languages without acknowledgment. It demonstrated that plagiarism of this type can be detected by methods of forensic linguistic analysis. A detection approach was proposed.

The concept of originality and translation over time was discussed, with reference to diverse approaches to translation, from traditional perspectives, where translation implies a transfer of meaning from one language to another, to more recent ones, where translation is considered meaning negotiation between different languages and

cultures, and post-structuralist ones, which argue for a combination of multiple approaches. It was argued that any text, and not just linguistic text, is subject to the principles of originality, and hence can be plagiarised. Very limited prior research had been conducted on plagiarism and translation, and the most relevant was dedicated to the comparison of translations into the same language.

This discussion showed that translingual plagiarism is very difficult to detect, for various reasons. Firstly, not all translated texts are intuitively identified as such. Secondly, it is not possible to make a linguistic comparison of two texts in two different languages. Thirdly, the translation can be diametrically opposed to the original, on the basis that the translator is free to make his/her own choices, both syntactically and lexically. It was therefore argued that it is necessary to rely on 'sentence-based grammars', as well as on the 'discourse grammar', as these impact the informational packaging of the text, as well as the use of cohesion and coherence. We concluded that several linguistic elements (such as morphosyntax, lexis, semantics, pragmatics, informational packaging, discourse grammar and conceptual representations) are part of a speaker's or writer's interlanguage, and get transferred from the source to the target text.

Finally, a method was proposed to detect plagiarism across translated texts. Machine translation allows a comparison between the derivative and the original texts, manually or by using any plagiarism detection software. Additionally, considering that the more the effort to disguise a surreptitious theft, the higher the degree of intention, it is argued that forensic linguistic analysis can provide hints to the plagiarist's intention.

The chapter concluded by showing that linguistic analysis is crucial, in addition to the automatic detection procedure, to investigate, analyse, describe, explain and demonstrate textual reuse.

Chapter 6

Computational Linguistic Approaches to Plagiarism Detection

6.1 Introduction

This chapter discusses computational approaches to plagiarism detection. Firstly, it reviews and examines authorship attribution and the type of evidence that is found to be relevant in attribution cases: external, intrinsic, and stylometric. A tentative explanation is offered of the theoretical assumptions underlying plagiarism detection, with reference to the types of evidence used by attributionists. The core similarities and dissimilarities between authorship attribution and plagiarism detection are identified, in relation to their relevance to computational approaches to both areas. The section concludes with a discussion of stylometric analysis, and the identification of some of the main problems and limitations of machine processing in stylistically analysing textual data.

The chapter then establishes a comparison between human- and machine-attribution methods to make a distinction between markers of authorship traditionally used by attribution studies that can be used effectively by computational approaches, and those markers that are too complex and sophisticated to be handled effectively by machine processing. It reviews research on some of the more relevant computational methods of authorship attribution, classification and identification, to identify the areas where those systems perform well and the areas where they perform worse.

The chapter continues with a discussion of computational responses to plagiarism detection, and discusses the two main types of plagiarism detection, as investigated by computational linguistics: *intrinsic* and *external*. It examines different perspectives,

from more traditional, document comparison methods, to more recent approaches that aim to detect plagiarism based on paraphrasing, or even in translation (*translingual plagiarism*), to argue for a more accurate approach to describe and account for instances of plagiarism.

It then goes on to describe and discuss some common plagiarism detection software tools currently available, with reference to their capacity to handle certain linguistic aspects. Illustrative examples are used to demonstrate the effectiveness of these tools and their limitations for detecting plagiarism strategies, and to offer hypotheses and predictions to overcome these problems.

The chapter concludes by discussing the concepts of *automatic* and *machine assisted* plagiarism detection over time, as well as their relevance in comparison to human detection, and whether the computational approach can be run dissociated from a forensic linguistics approach. The chapter proposes the hypothesis that human interpretation is crucial to confirm or otherwise discard the suspicion that a text has been plagiarised, and that, based on this, computational detection should be used primarily as an investigative tool, whereas a sound forensic linguistics analysis is necessary to provide evidence.

6.2 Authorship Attribution Theory and Computational Practice

Authorship attribution has traditionally used different types of evidence, either to ascertain, based on a comparison of the text of unknown authorship against texts written previously by the candidate authors, whether 'one text, commonly referred to as the query text, is assigned to, or excluded from a group of texts' (Grant, 2007), or as a profiling tool to describe some possible sociolinguistic features of the suspect author, by comparing the text of unknown authorship to texts belonging to different populations. Love (2002), in particular, argued that authorship attribution can be based upon three different types of evidence: *external*, *internal*, and *stylistic* evidence.

External evidence is 'that from the social world within which the work is created, promulgated and read' (Love, 2002: 51), including documents that are not directly related to the actual writing of the document(s) whose authorship is questioned. This is the type of evidence resulting from the extra-textual world, rather than from the text; it represents personal, official and commercial records that can be used independently

or in combination with the aim of 'investigating the circumstances of original individual texts' (Love, 2002: 66), and can consist of documents containing information about the circumstances involved in the writing, biographical data, and historical briefs on previous attributions of the author's work, together with the circumstances of the corresponding attributions. Love argues that contextual information, in particular, is often useful to locate a certain work in a certain historical moment, either to attribute it to a particular class of authors or group of people (in approximately the same way as profiling is used by forensic linguistics), or to confirm or reject the assumptions of someone being the author of a certain document. External evidence of this type has the potential to collect information that cannot be obtained from textual data, which allows narrowing down contextualisation, but does not permit a positive identification (Love, 2002). This identification is more successfully done using linguistic analysis.

Unlike external evidence, *internal evidence* results from the analysis of the work itself (Love, 2002), and includes stylistic evidence, author's 'self-reference and self-presentation within the work' being analysed, and information obtained 'from the themes, ideas, beliefs and conceptions of genre' (Love, 2002: 51) in that work. Evidence of this type requires a close examination of the textual data, and has traditionally been subdivided into *style* and *ideas*. Attribution based on ideas and their potential to represent the beliefs and intellectual stance of the author could only be considered if they were rooted and stable, so that the hypothesis that they could be an influence of another author can be rejected. Similarly, it has to be sufficiently complex not to be shared across many authors, so that the co-presence of two or more unrelated ideas in two different documents gains strength as an indicator of common authorship. This is grounded on the reasoning that ideas, which can be original or derivative, are represented as a part of a coherent world-view in the human mind. Although comparing the ideas of an anonymous work against those in the work of a known author is a difficult task, one can perform it successfully by accumulating individual elements that, albeit not discriminative individually, can work as a whole.

The relevance of the combination of markers for authorship identification was also later demonstrated in forensic contexts by Grant and Baker (2001), who applied Principal Component Analysis (PCA) to select those markers 'which collectively account for the most variance in the texts' (Grant and Baker, 2001: 76). Additionally, it is important to differentiate those elements that remain stable from those that evolve, the former being useful for authorship attribution, and the latter being valuable for dating a

certain work (Love, 2002). Love does not deny the potential of ideas for positive identification of an author, if used with caution. On the contrary, he admits its relevance, but claims that, as a method, it is mainly valuable to exclude candidates for authorship. Given these constraints, internal evidence, especially when built on ideas, was considered weak, and attributing authorship based on internal evidence alone only gained strength with the recent developments of computer-assisted methods and stylometry (Love, 2002).

Stylometric analyses became an accepted method of internal evidence on the grounds that decisions could be made on the basis of probabilities (Love, 2002). The concept of style was earlier defined by Hänlein (1998: 26) with reference to Buffon's premise that style is man himself[1], and consequently it has the potential to reveal the ideas of an author ('an author's way of thinking'). Style is theoretically grounded on the principle that each speaker or writer of a language has his/her own 'idiolect', a unique, distinctive way of expression (Coulthard, 2004), and the analysis of those characteristics, *stylistics*, has been used to determine, confirm, or refute the authorship of queried texts with significant success (Grant, 2007; Love, 2002; McMenamin, 1993, 2002). A stylistic profile can be formed by grouping different style markers together, including elements of lexis, syntax and semantics, but also 'subject matter', 'choice of 'we-groups", 'pragmatic use of inverted commas and quotations', writing tentatively or authoritatively, or 'writing in a detached or an involved manner' (Hänlein, 1998: 29). While different approaches may use diverse analytical methods, the 'systematic examination of the characteristic physical features of a message' (Love, 2002: 99) at the core of stylistics is vital to test the reader's or the analyst's intuition about the authorship of a certain text.

Stylometric analyses do not assess the quality of a certain author's writing. Rather, they aim to determine the consistent and stable use of certain patterns that make a writer's style particularly distinctive. In some cases where the author demonstrates strong individual patterns, based on personal choice, the analytical task is significantly easier, as compared to cases where the writing only subtly varies from a normative, rule-based use, where a closer examination is necessary. Additionally, as Love (2002) argues, it is necessary to consider individual features of style in combination before reaching a conclusion on the authorship of a questioned text, and the quantification of these features. When quantified, stylistic features can be used as powerful discrimi-

[1] In Buffon's words, '*Le style est l'homme même*' (Buffon, 1753).

nants of authorship to support the reader's intuition, by measuring, to some extent, the author's preferences for particular patterns.

This is the approach that lies at the basis of stylometric authorship attribution. Some of these discriminants are relatively simple to compute and quantify. In spite of the increasing use of spell-checkers by word processors, different authors show personal choices at the level of spelling, punctuation, hyphenation, and spacing; text metrics like average sentence, word, paragraph and text lengths, and type/token ratio, can be easily and quickly quantified and compared; and even lexical richness, which might at first seem too complex for a quantitative analysis, can be calculated in a few seconds. Other discriminants can be more challenging to compute. That is the case with the preference for one doublet of a synonym over the other, the use of rare and unusual vocabulary, or even the use of syntactic and grammatical patterns that need a careful intra-writer examination before an inter-writer comparison is made. One can determine the discriminant power of vocabulary by establishing '*how* rare and how unusual' it is (Love, 2002: 108). Word frequencies, keywords in context (KWIC), and word combinations can be provided easily by syntactic data. This output, however, is not independently capable of providing explanations, showing how authors write or how they differ. Other possible discriminants, such as metaphors and prosody, as Love (2002) recalls, are very difficult to quantify, and consequently even more complex to process computationally.

Innumerable approaches to stylometric analysis have, however, been used as the basis of most authorship classification tasks[2], with more or less complex statistical tests and algorithms. But although, as Hänlein (1998: 23) argues, 'there is no general theory or concept of style that would make this field of study a unified area of interest', computational approaches have, in a number of cases, tended to process style, often oversimplistically, as a unified concept. On the contrary, as Grant and Baker (2001) later demonstrated, methodologically analysts have to cope with the problem of determining the validity and reliability of a marker of authorship, since the fact that it works or not as a discriminant in one particular case does not mean that it it works or not in all cases.

Many of these features, including syntactic information (Uzuner *et al.*, 2005), *n*-grams (Hirst and Feiguina, 2007; Kestemont *et al.*, 2011), discourse markers (Rao

[2]The term *classification* is preferred to the terms *attribution*, *identification* or *analysis* when used in the context of computational linguistics because the task in this case consists of matching the questioned text to one of several classes, and hence is considered more accurate.

et al., 2011), and writeprints Abbasi and Chen (2008), have been used by computational linguistic approaches to authorship attribution (Argamon and Juola, 2011; Ayala *et al.*, 2011; Kern *et al.*, 2011; Luyckx, 2011; Mikros and Perifanos, 2011; Recasens and Vila, 2010; Solorio *et al.*, 2011) or profiling (Argamon *et al.*, 2009), as well as plagiarism detection (Barrón-Cedeño and Rosso, 2009; Chester, 2001; Iyer and Singh, 2005; Meyer Zu Eissen and Stein, 2006; Stamatatos, 2009a). However, plagiarism detection and authorship analysis are not identical tasks. Although they share some features of uniqueness of utterance, and as Turell succinctly described (personal communication), they are the two sides of the same coin, they differ in several respects.

Although computational approaches to stylometric analysis in general have shown limitations with respect to the analysis of 'macro' features (e.g. topic choice, 'tentative' vs. 'authoritative', 'detached' or 'involved' writing), they have been proven to handle markers of authorship effectively, and additionally to apply quantifiable and statistically rigorous methods used by forensic linguists (which in a sense is unsurprising, considering that these are the markers that can be computed more easily). One good example of this approach is the one of Mosteller and Wallace described by Grant and Baker (2001), who performed a Bayesian statistical analysis over the rates of relative occurrence of functional and other words, demonstrating the success of this method to determine authorship. Claims by Grant and Baker (2001) that Bayesian statistical analyses have since been very popular as an authorship attribution method are confirmed by the diverse computational approaches to authorship analysis, which actually used Bayesian statistical methods to attribute authorship (e.g. Abbasi and Chen (2008); Argamon *et al.* (2009); Kestemont *et al.* (2011); Luyckx (2011); Manning (2001); Stamatatos (2007, 2009b)).

Claims of authorship can be rejected by establishing, via authorship analysis, that a suspect text was not written by a purported author, but plagiarism can hardly be demonstrated, as will be argued below, without finding additional evidence. Plagiarism detection consists of establishing whether the purported authors of a text have originally written it, and this task becomes particularly complex when the object is academic texts. As discussed in section 1.3, texts written by young academic writers are often expected to include phrases, terms and syntax borrowed from other, senior writers as part of the learning process (Howard, 1995). Consequently, from the perspective of authorship analysis, a high degree of inconsistency, both stylistic and linguistic, would be expected. Moreover, owing to the learning process factor, any

comparison between a writer's suspect and previous texts would need to discount the fact that the author's writing might have evolved.

Notwithstanding the usefulness of authorship analysis to initiate an investigation into plagiarism cases, in a way that is similar to the sense of 'déjà-vu' (Coulthard and Johnson, 2007) and intuition of a reader, further analysis is required. Usually, this involves flagging the suspicious text for plagiarism, finding the original, and finally describing and explaining the derivativeness via a comparative analysis (especially in cases involving sophisticated strategies such as paraphrasing or translation). Conversely, if plagiarism detection is conceived in terms of a 'layered process', starting with the suspicion raised by shifts in style, and subsequently involving the comparative analysis of two or more texts, then the task in both instances consists of determining questioned or disputed authorship. Consequently the most striking differences between authorship attribution and plagiarism detection tend to fade. This is the perspective that seems to appeal the most to computational linguists. The following section discusses the use of stylometric analyses by computational approaches, considering the problems and limitations of stylistic analysis imposed on machine processing.

6.3 Computational Responses to Plagiarism

Computational plagiarism detection, which has attracted the attention of computer scientists in particular (as the work reported by Potthast *et al.* (2009b,a, 2011b) suggests), has been arguably demonstrated to be relevant in cases where human manual analysis proves unfruitful, either due to the large amounts of data for comparison or to human fatigue (Woolls, 2010). Although it has the potential to flag a text as being stylistically inconsistent, as a whole, it is of limited assistance to identify the specific passages where shifts in style occur (Barrón-Cedeño *et al.*, 2008; Maurer *et al.*, 2006; Stamatatos, 2009a; Woolls, 2003). The usefulness of detection software is however unquestionable, not the least because they can process mechanical aspects of text analysis faster (Woolls and Coulthard, 1998), or allow the analysts to perform all kinds of searches and comparisons and to obtain data upon which to draw analytic conclusions. This applies even to software that was not specifically developed for a certain plagiarism detection task, but that may be nonetheless helpful for the analysis. A specific example of this is the text comparison feature of *Microsoft Word*, in addition to other examples such as *WordSmith Tools*, *Vocalyse Tookit*, and *Abridge*, which were cited by Woolls and Coulthard (1998), or *Google*, whose usefulness was later

demonstrated by Coulthard and Johnson (2007) in finding identical strings of text and providing evidence of textual uniqueness.

However, partly owing to the technological developments of recent years, the strategies used for instance by students to plagiarise have also evolved. The traditional 'copying and pasting', verbatim (unacknowledged) borrowing came to be replaced with other more sophisticated strategies, such as paraphrasing and translation. This shift in plagiarism strategies constrained the detection task, impeding detection systems and imposing limits on the true capacity of computational approaches. As explained by Maurer *et al.* (2006: 1052), '[p]lagiarism now is not confined to mere cut and paste; synonymising and translation technologies are giving a new dimension to plagiarism', and although, in theory, such technological developments should be able to cope with those strategies, practice has demonstrated otherwise. Ironically, extensive research on paraphrasing has shown, if anything, that computers are able to correctly generate paraphrases, but their performance tends to drop with paraphrase identification or detection (Madnani and Dorr, 2010). The case of translation, on the other hand, is less challenging. Despite the limited advances on the detection of plagiarism based on translation from other languages, such as those reported by Barrón-Cedeño *et al.* (2008), chapter 5 demonstrated how detecting translingual plagiarism is a feasible task, with potential application in forensic contexts.

Diverse methods of computational plagiarism detection have been proposed in recent years. However, most of these focus on one of the three methods described by Maurer *et al.* (2006): document comparison, Internet search, or stylistic analysis.

Document comparison consists of comparing a suspect document, locally or remotely, usually word by word, against the possible originals. Local comparison requires that the user feeds the suspect and source documents into a system, usually a client software installed and run on the user's computer, for comparison. Remote comparison requires that the user feeds the suspect documents into a remote server database for comparison against databases containing up to billions of documents, and even against documents indexed by search engines[3], that are available on the Internet. Document source comparison of these types are usually based on computational techniques like word stemming and fingerprinting. The latter in particular, which was first introduced by Manber (1994), is used by a number of systems. It consists of indexing documents' strings of a moderate size and then comparing those indexes

[3]Usually, when documents are not indexed by search engines, only the first few hundred words are shown.

against preprocessed indexes of other documents, and subsequently searching for similarities. As explained by Maurer *et al.* (2006), this apparently simple approach led to the development of more sophisticated querying and ranking algorithms, some of which are claimed to be used by commercially available detection systems. The results of the search query are then returned in the form of a report, and need to be interpreted in order to check whether the instances reported are actually plagiarism or, on the contrary, false positives (as is the case of flagged instances that are actually acknowledged).

The second method, *manual search for characteristic phrases*, consists of taking some strings representing core concepts of a suspect document and searching for them on the Internet. This approach is similar to the one proposed by Coulthard and Johnson (2007), but whereas the latter suggested that the analyst's search might be triggered by intuition, Maurer *et al.* (2006) proposed searching for core concepts, e.g. uncommon terms. Their reasoning is that those concepts are the ones that can be more easily found by means of web searches, and hence more discriminant. A similar approach, which considered lexical items rather than uncommon terms, had been previously demonstrated by Johnson (1997) to be robust in identifying such concepts.

The third method is based on the *stylometric analysis* of texts, and is grounded on the principle that different people have differing, individual and unique writing styles. The comparison of the writing styles can be performed within the same document, or against other documents written previously by the same author (Maurer *et al.*, 2006). The former, which Meyer Zu Eissen and Stein (2006) called 'intrinsic plagiarism detection', does not require external references for document comparison. It identifies inconsistencies within a document, based on the statistical analysis of linguistic features, such as the inconsistent use of singular/plural forms of pronouns (*we* / *I*) or articles (*the*), shifts in vocabulary, common spelling mistakes, use of punctuation and structural features, such as word counts, sentence lengths, and other stylometric measures. Maurer *et al.* (2006) argue (and Sousa-Silva *et al.* (2010) later reached similar findings) that the mismatch in style between the borrowed text and the overall writing style can be indicative of plagiarism. Importantly, however, Maurer *et al.* (2006) also claim that instances of multiple authorship have to be discounted; as attempts to make the text read as single authorship (Love, 2002) in cases of multiple authorship tend to fail, so do attempts, in cases of plagiarism, to make the overall style of the text consistent (Maurer *et al.*, 2006) – which consequently leads to false positives.

Stylometric analysis can also operate based on the 'cloze procedure' (Taylor, 1953), which Maurer *et al.* (2006) argue is the method used by *Glatt Plagiarism Services*[4]. This procedure works by replacing every fifth word of a suspect document with a blank space, and asking the writer of the suspect document to fill in the blanks. The rationale behind this approach is that the percentage of correct answers has the potential to determine if a document has been written by the supposed author. This approach raises a few problems. Firstly, it is not feasible in academic contexts, where large numbers of documents need to be processed, as admitted by Maurer *et al.* (2006). Secondly, as Coulthard (2004) demonstrated, the same person is very unlikely to produce the same utterance on two different occasions.

This classification framework, which considers document comparison, Internet search and stylistic analysis, has been frequently applied to computational approaches to plagiarism detection, based on the assumption that plagiarism consists of identifying 'highly similar sections in texts' (Potthast *et al.*, 2009b: 2). This is similar to the task of finding *duplicates* or *near-duplicates* in computer sciences. However, since detecting similar or identical strings (duplicates and near-duplicates) is manifestly insufficient to account realistically for plagiarism detection, Potthast *et al.* (2009b) suggest a computational approach to plagiarism detection that is based on two other main categories: *external* and *intrinsic* plagiarism detection. These concepts are discussed in the following section.

6.4 *External* and *Intrinsic* Plagiarism Detection

Potthast *et al.* (2009b) strongly argue that, in order to account for plagiarism detection, computational approaches should be structured into *external plagiarism detection* and *intrinsic plagiarism detection*. Although both categories imply some kind of comparison, intrinsic plagiarism detection handles stylistic analysis, whereas external plagiarism detection includes document comparison and Internet searches. Intrinsic and extrinsic detection could be explained in simple terms as to whether they use a corpus of original texts for comparison (external, with reference) or otherwise (i.e. whether the comparison is internal, without reference).

Based on the work of Stein (2005) on the use of fuzzy fingerprints for information retrieval, Meyer Zu Eissen and Stein (2006) propose a model to describe the over-

[4] *Glatt Plagiarism Services* – `http://www.plagiarism.com/`

all strategies used to plagiarise ('plagiarism delicts') and the corresponding detection methods. The 'plagiarism delicts' are identified in figure 6.1.

<table>
<tr><td rowspan="8">Accurate copy</td><td rowspan="4">Large part of document</td><td rowspan="4"></td><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="4">Small part of document</td><td rowspan="2">With reference corpus</td><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="2">Without reference corpus</td><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="8">Modified copy</td><td rowspan="4">Language translation</td><td rowspan="4"></td><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="4">Transformation</td><td rowspan="2">Large part of document</td><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="2">Small part of document</td><td>With reference corpus</td></tr>
<tr><td>Without reference corpus</td></tr>
</table>

Figure 6.1: 'Plagiarism delicts'.

The detection methods proposed originally by Stein (2005) to handle each 'delict' are shown in figure 6.2.

<table>
<tr><td rowspan="8">Accurate copy
Identity analysis</td><td rowspan="4">Large part of document
Global identity analysis</td><td rowspan="4"></td><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="4">Small part of document
Local identity analysis</td><td rowspan="2">With reference corpus</td><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="2">Without reference corpus</td><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="8">Modified copy
Structure analysis
Similarity analysis</td><td rowspan="4">Language translation
Structure analysis</td><td rowspan="4"></td><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="4">Transformation
Similarity analysis</td><td rowspan="2">Large part of document
Global similarity analysis</td><td></td></tr>
<tr><td></td></tr>
<tr><td rowspan="2">Small part of document
Local similarity analysis</td><td>With reference corpus
Fuzzy fingerprinting</td></tr>
<tr><td>Without reference corpus
Style analysis</td></tr>
</table>

Figure 6.2: 'Plagiarism delicts' – detection methods.

The contribution that Meyer Zu Eissen and Stein (2006) made to Stein's model

consisted of adding a method for detecting plagiarism, in cases of both accurate and modified copy, by performing a stylistic analysis, as a means to handle instances of plagiarism where no reference corpus is available. This marked a shift in the trend of computational linguistic approaches to plagiarism detection, which had traditionally given top priority to document comparison.

6.4.1 External Plagiarism Detection

External plagiarism detection attempts to find similar or identical text matches by comparing a suspect document against an open or closed set of documents from which the suspect document might have borrowed. Closed sets consists of a limited number of documents against which the suspect document(s) will be compared. Conversely, open sets consist of other texts that are publicly available, which cannot be sufficiently limited to a specific group. Methods that use closed sets are thus preferred when the analyst knows or suspects that the questioned text might have borrowed from a specific group of texts, whereas open sets are used when the suspect text is known or suspected of having borrowed from other sources, such as the Internet.

The method proposed by Stein *et al.* (2007), which is based on a generic, three-step retrieval process, illustrates the approach to this task. The first stage, *heuristic retrieval*, consists of extracting samples of the suspicious document (*dq*), and searching these against an index of chunks to find matching samples that may be indicative of instances of plagiarism. The second stage, *Detailed Analysis*, consists of comparing the potentially plagiarised sections of each candidate document, selected from a huge database of documents against which the comparison is impractical, to the matching sections in *dq*. Finally, in the third step, *Knowledge-based Post-processing*, the system performs an analysis of the identical passages to investigate whether exclusion criteria apply, e.g. whether there has been a proper citation. The sections remaining after the third stage is completed are reviewed by a human analyst to check for plagiarism.

6.4.2 Intrinsic Plagiarism Detection

Contrary to external plagiarism detection, *intrinsic plagiarism detection* operates by 'reading' the document, as happens often intuitively with humans. Approaches of this type aim to measure similarity and dissimilarity between documents, in much the same way as author classification. Computationally, this means creating algorithms to compare suspicious documents against a reference set of potentially original documents

to identify passages that might have been copied or modified, to a greater or lesser extent – as happens with the human analyst (Meyer Zu Eissen and Stein, 2006). In this case, no reference corpus is provided for comparison; instead, the challenge consists of finding stylistic clues, irregularities or inconsistencies within the document that account for variations in the writing style, that may serve as evidence of different instances of authorship.

In their seminal computational work, Meyer Zu Eissen and Stein (2006) described three possible types of plagiarism: plain, 'one-to-one' copy; borrowing of passages from another original document, with some extent of modification; translation and reuse of documents. They reported that previous methods (Brin *et al.*, 1995; Hoad and Zobel, 2003) handled the first type well, but not the other two (which worked well only when comparing exact matches). Additionally, Meyer Zu Eissen and Stein (2006) argue that these approaches have one constraint in common. Despite the system improvements introduced by Stein and Eissen (2006), which permitted generating fingerprints that are robust against modifications of some type, they require a collection of potential original documents as reference for comparison.

The method that they propose to overcome this limitation is identifying style shifts to detect passages that are suspiciously indicative of plagiarism, in much the same way that human readers do intuitively. Arguing that the quality of the quantification of the linguistic features determines the capacity of the detection approach, Meyer Zu Eissen and Stein (2006) proposed a set of features aimed at capturing stylistic information by measuring the 'customariness of word usage', which they called 'quantification of the writing style' (Meyer Zu Eissen and Stein, 2006). This approach works by dividing a document into parts (sentences, paragraphs, sections), each part containing between 20 and 200 words, in order to analyse variance in style. It measures five categories of stylometric features: *text statistics*, calculated at character level; *syntactic features*, operating at sentence level; *part-of-speech (POS)* features, measuring the classes of words; *closed-class* word sets that quantify the use of 'special words'; and *structural features*, operating at the level of text organisation. To these categories, which are based on previous research on authorship attribution (e.g. Koppel and Schler (2004)), they added another feature, the *averaged word frequency class*. Since this class works independently of text length, it is more robust (Meyer zu Eissen and Stein, 2004), and consequently a more powerful discriminant marker.

Word frequency class, which is related to Zipf's law, is computed by calculating the

frequency and the rank of a word in a list; the more frequent the word, the closer to rank '0', and conversely the less frequent, the further it is from '0'. Their experiment, which was run over a corpus of artificially-created instances of plagiarism, included average sentence length, 18 part-of-speech features, average stopword number and the averaged frequency class. Meyer Zu Eissen and Stein (2006) reported that both a 'classical discriminant analysis' and an SVM (support vector machine) classification returned similar results, with average word frequency class, average preposition number, and average sentence length being the three best discriminants. This method has since been improved by other approaches. A good example of this method is the work of Stamatatos (2009a), which was based on the calculation of character n-gram profiles and a style change function to quantify the style variation within a document. Although the author reported that this method was able to detect roughly half of the plagiarised text sections, he also admitted that it needed improvement, especially in terms of the precision rates (i.e. the fraction of instances that are relevant, out of those retrieved by the system).

6.4.3 *External* and *Intrinsic* Plagiarism Detection and Authorship Attribution

Stylistic analyses like the one underlying 'intrinsic plagiarism detection' had been previously used by authorship attribution studies (Love, 2002). But the part where computational linguistic approaches have made significant improvements over the last decades is *stylometry*, i.e. 'the exact quantitative measurement, tabulation and interpretation of designated aspects of verbal performance' (Love, 2002: 132). This is due mostly to technological developments such as the amount of texts currently available in machine-readable format and the processing capabilities of computers. The availability of growing corpora of data and 'data sets', on the one hand, and the increasing machine processing power, on the other, allow the application of statistical measures in forensic contexts for the comparison of 'questioned and authentic samples' (McMenamin, 1993: 46). Stylometry, nevertheless, is only capable of 'discriminating between forms of language behaviour' (Love, 2002: 160), and this, by itself, is invaluable for research in forensic linguistics, but cannot account for the findings. One can therefore conclude that, although quantitative methods are vital to identify patterns, especially where large amounts of data need to be analysed, by themselves they are unable to explain those patterns and account for them. As Love clearly put it, 'numbers and

ratios can never be fully persuasive when we have no understanding of what elements in the language of the text have given rise to them' (Love, 2002: 101).

Statistical measures, especially those related to markers of authorship, have been applied to textual analyses by computational linguistic approaches with significant success, yet without much consideration for theoretical grounding on authorship attribution. The type of stylometric analysis proposed by Maurer *et al.* (2006), as well as the intrinsic analysis conducted by Potthast *et al.* (2009b), share with the stylometric analysis of forensic linguists and authorship attributionists the fact that they both compare suspect texts to confirm or refute the authorship of questioned texts. But whereas the aim of linguists and attributionists is to identify and explain the findings related to a certain authorship problem, computer scientists are mainly concerned with obtaining the highest possible classification result, and evaluating the performance of that method. Therefore, by disregarding further theoretical considerations, computational approaches are able, at most, to conduct stylometric, rather than stylistic analyses. The problem with computational plagiarism detection is identical. On the one hand, approaches to the detection procedure as a classification problem shift the focus from accurate detection of the plagiarised instances to the calculation of precision and recall rates, while disregarding theoretical considerations. On the other hand, as mentioned above, detecting plagiarism by identifying shifts in style misses the point of multiple voices that can be expected, for example, from academic writing. Besides, it fails to identify instances where the plagiarist has borrowed from one single source, thereby retaining the same style across the whole text. Thirdly, the conceptual nature of the detection procedure needs to be challenged.

Meyer Zu Eissen and Stein (2006) followed strategies like these, more or less closely, to devise computational detection procedures, which they called *automatic plagiarism detection*. Computer scientists tend to investigate fully- or quasi fully-automated plagiarism detection systems that are capable of detecting, with the least human intervention, the instances of plagiarism in a text or set of texts. Good results are reported when they exceed the baseline. However, the concept of *automatic detection* can be elusive, especially because what would be expected from an automatic detection system is its capacity to analyse one or several documents, independently and without human assistance, and determine whether plagiarism has occurred. Although developments in research suggest that this can eventually be the case, reality proves otherwise, as the description of the software in section 6.4.4 below demon-

strates. Moreover, if the linguistic analysis is valued to the detriment of *Precision* and *Recall* curves, statistical analysis of how much is guessed and how much is missed is irrelevant, if the specific instances of borrowing are not identified. On the contrary, from a linguistic perspective, it is ultimately the interpretation of those instances that is most important.

Nonetheless, even if automatic detection cannot be expected to *describe* all instances of plagiarism, and justify *why* they are plagiarism, the potential contribution of computational approaches to plagiarism detection is indisputable, especially if plagiarism detection systems are taken for what they are: tools that can assist human analysis. Woolls and Coulthard (1998) called them 'tools for the trade'. Therefore, it would perhaps be more precise to term this 'machine-assisted plagiarism detection', rather than 'automatic plagiarism detection'. On the one hand, using this term is more precise in describing what it actually does – help human investigators to find clues to – and detect – instances of plagiarism, that will be further submitted to human scrutiny before passing any type of judgement. This avoids failure to meet overexpectations from detection systems, while admitting that, as Maurer *et al.* (2006) argue, a final report cannot be used without human interpretation.

Apart from these theoretical considerations, the challenges described by Maurer *et al.* (2006) remain valid for the most part. One is that electronic data for comparison are not always available, although the future is promising as more and more data become available in electronic form. Another problem is where paraphrasing is used extensively to (intentionally) mislead detection systems, in which case systems tend to fail. In this respect, although systems perform quite well when checking the similarity of a small set of documents using 'deep techniques [to] determine conceptual equivalence even when heavy paraphrasing is used' (Maurer *et al.*, 2006: 1081), those approaches perform poorly when dealing with large volumes of data. Even more challenging is plagiarism detection across different languages, a challenge which they foresaw would remain unresolved for several years.

As Woolls (2003) concluded, and Woolls and Coulthard (1998) had earlier demonstrated, software detection tools work mechanically and computationally. On the one hand, they simulate human behaviour; on the other hand, their development is limited by human expertise. However, those tools are helpful in providing information to experts faster and more efficiently, as long as they are used 'with care' (Woolls, 2003: 112). Part of the solution is, therefore, acknowledging the limitations of these sys-

tems and rejecting their ability to perform automatically. Some of these limitations are discussed in the section that follows.

6.4.4 Software Tools

Diverse plagiarism detection tools are available, including open source, proprietary or commercial solutions, those that operate locally or based on a server, and those that are free, for purchase or offered by subscription. The range of tools is too extensive for a comprehensive description in this thesis. The list included in this section has been selected based on the system's availability, popularity or functionalities, that can contribute to the research.

Article Checker

Similarly to other free online plagiarism detection systems, *Article Checker* is a duplicate text search tool. It allows the user to copy and paste the suspicious text into the search box, or alternatively input a URL in the web address box to scan a website for duplicate content. The output of the detection process returns the potential original source, as well as the percentage of identical text.

Unlike other similar software tools, *Article Checker* supports only two search engines, *Google* and *Yahoo*, but allows the search to be performed using the two systems, rather than ask the user to choose only one of them.

CiteReader

Citereader is another tool developed by CFL Software which aims to find clues to plagiarism by identifying shifts in style of the author. The software is based on the assumption that shifts in style and stylistic inconsistencies can be indicative of plagiarism instances (Woolls, 2003). This shift can be identified by means of in-text comparison, which consists of comparing suspect strings of text with other instances of the same document, or by between-text comparison, which consists of comparing the writing of a suspect document with texts written previously by the same author.

Woolls (2003) argues that currently available computerised plagiarism detection tools need to compare suspect texts against something. In the case of *CopyCatch*, this comparison is made against peer work; in the case of web search tools, this comparison is made against sources available openly on the Internet.

CiteReader functions by reading individual documents and attributing a score to

each sentence. This procedure is said to be independent of sentence length, in that it works at clause level (Woolls, 2003). The two strategies indicated by Woolls (2003) to identify shifts in style are: an analysis of the vocabulary of the discussion section of a student work, which is not usually part of the quoted sections; and passages using language that is distinct from the remaining text. This is based on the principle that the more proficient the writer, the more even the distribution of sentences across the text, regardless of the subject (including the academic style).

CopyCatch

CopyCatch was primarily designed for collusion detection, but can be used to investigate any suspicious instances of non-independent production of documents, including plagiarism (of which the work of Turell (2008) is a good example).

The software has a clear educational motivation, but is able to compare other documents as well. It works as a client software that can be installed in the user's computer. Although it includes a non-comparative element based on stylistic analysis, *CopyCatch* is mainly a document comparison software.

The graphical UI presents marked up versions of the comparison texts, where the user can have access to full text views of sentences or vocabulary. The vocabulary view shows the similar vocabulary used in sentences, whereas the sentence view shows the differences between the source and derivative documents. The information provided by the system includes full cross-referencing of related sentences by paragraph (P) and sentence (S), so that the user knows which sentence of which paragraph in the source text compares to which sentence of which paragraph in the derivative text. The software is also able to provide information on the content words shared by sentences in the files used for comparison, and show the files containing sentences with these words in common. This allows the user to search the Internet for those (suspect) texts. Additionally, the comparison report can be saved in .rtf or .html format, while preserving the colour mark-ups.

As a tool for stylistic analysis, *CopyCatch* scans single files for stylistic shifts that may indicate non-independent production of a text. The system marks up in red the sentences, clauses and sets of words that do not match the remaining text stylistically. The rationale behind this procedure is that the user can copy and paste the suspect text to a search engine and look for possible sources.

Docol@c

Docol@c is a paid software package used widely in Germany, Austria and Switzerland, and by some universities worldwide. It allows the user to search for matching text fragments across various documents and to detect instances of plagiarism and copyright infringement, as well as quotations or other sources available online. The type of plagiarism that it is able to detect is verbatim, or near-verbatim borrowing.

The software works in three steps, namely: logging in using the user name and password provided; uploading document(s) for analysis (or providing the URL, if the text is available online); and generating a report, which can be stored in the user account or alternatively emailed to the user's address.

The report generated contains interactive functions, such as links to the sources found, as well as statistics of the number of text segments analysed and the percentage of matching. This percentage is calculated as a rate of matching fragments compared to the total number of fragments checked. Although the system stores the submitted documents and the corresponding analysis reports anonymously until they are deleted by the user from their account, it does not check documents against other documents stored in the user account.

Finally, although the software is to a great extent automated, it also includes search preferences parameters that can be set up by the user, for instance to specify the sample length.

Dupli Checker

Dupli Checker is another free online plagiarism detection tool that allows users to copy and paste a limited volume of suspicious text into the search box, or alternatively upload a .txt document, and scan it for plagiarism against web pages indexed by one of three search engines: *Google*, *Yahoo* or *MSN*. The user is asked to search for the text 'With quotes' or 'Without quotes' before hitting the search button. The software then breaks up each sentence automatically, subsequently returning the results of the analysis. The output of this analysis presents the suspicious text, with links to the possible original.

Ephorus

Ephorus is a suite of plagiarism detection tools that can be integrated with e-learning platforms like *BlackBoard*. It allows users to check papers for several instances of plagiarism by comparing them against Internet sources and bibliography databases, such as research papers and journals. It also enables the detection of

cases of collusion, where a group of students show a higher than expected degree of overlap. The suite includes a 'database pool', which is a collection of *Ephorus* databases from different educational institutions, and that institutions submitting instances for analysis can join.

However, as Scaife (2007) argues, the system searches only for matches at the level of the source. Then, a report is generated in summary or detailed form. The latter shows the suspect document and the potential sources side-by-side. This report produces an overview of the word count, the number of matching words and the longest reported matching string. Different, potential instances of plagiarism are also shown in different colours, similarly to other software packages. However, unlike these, *Ephorus* claims to be able to identify quotes, which the user can accept manually, hence checking whether the suspect material has been properly acknowledged; cases where the matching material has been properly acknowledged will be removed from the calculated volume of plagiarised material. The system allows the user (usually the lecturer/tutor) to set a threshold, against which the software will 'flag' each case.

Glatt Plagiarism Services

The *Glatt Plagiarism Screening Program (GPSP)* is a plagiarism detection service, rather than a tool. It is described as a 'highly sophisticated Screening Program to detect plagiarism' used traditionally by academic institutions and 'the legal profession'[5] to identify cases of copyright infringement. The system works locally, on the user's computer instead of online, by analysing patterns and writing styles of suspicious texts. However, the method of stylistic analysis used differs from those of other services. GPSP is based on the Cloze procedure (Taylor, 1953), which consists of removing every fifth word of a suspicious document and replacing it with a blank space. The purported writer of the text is then invited to fill in the blank spaces. The plagiarism probability is subsequently calculated considering the number of correct answers and the answering time (Maurer *et al.*, 2006).

Although this system proposes an interesting approach, not the least from a psychological perspective, its power to discriminate cases of plagiarism can be challenged, not the least because it can be misled by cases where the text is memorised. In this case, it is irrelevant to verify the authorship of a text, and even less so to judge instances of plagiarism.

[5] *Glatt Plagiarism Services* Website – `http://www.plagiarism.com/INDEX.HTM`

Google Translate

Machine translation engines are a good example of a tool that was not developed specifically to address plagiarism detection, although they can provide good results in detecting cases of plagiarism involving translation from other languages, as has been demonstrated in 5. *Google Translate*, as any other MT engine, can be used to detect translingual plagiarism. MT engines work by translating a suspect text into another language and scanning the translation for plagiarism in one of the freely available search engines. This latter step of the procedure is similar to the one used to scan any ordinary text for plagiarism online.

Google Translate allows the user to copy and paste the suspicious text into a search box, input a URL into that same box to translate a whole website, or alternatively upload a whole document to translate it into one of the 64 languages currently supported. The system then returns the text translated into the desired language, allowing the user to pick and choose alternatives, when available, and even suggest corrections to the translation offered. This procedure contributes to improving the system collaboratively. Therefore, although this translation engine is based mainly on statistical and crowdsourcing translation, the fact that some of the alternatives offered to the user are common phrases, collocations and/or multi-word units suggests that the collaborative contribution of different people is used to improve the system via a rule-based approach.

The translated text can be copied and pasted into a common search engine to find matching text. Since the system's ability to provide a faultless translation is very limited, a general search is preferred to an 'exact search'. Importantly, the search engine ignores letter case, as well as punctuation, which raises advantages, as well as disadvantages. On the one hand, the detection of instances where minor changes are introduced can benefit with this procedure; conversely, this type of changes imposes limitations on finding exact matches. Additionally, *Google* search queries are limited to 32 words, so lengthy strings of text cannot be easily searched. Another restriction imposed on the users by the system is that translating massive volumes of text on a regular basis is not permitted. This can impact the performance of commercial plagiarism detection systems, that may be unable to benefit from the assistance of *Google Translate*. The violation of this requirement, as described in the terms of use, can result in the system locking access to the MT engine.

Plagiarism Checker

Plagiarism Checker is a free, online plagiarism detection system that allows users to check a document or a webpage for plagiarism. In cases of document search, the software scans a document, or parts of a document, against the Internet to find similar or identical strings. This search is based on a simple, three-step procedure. The first step consists of typing or copying and pasting one or more phrases of the suspicious text into the search box provided. The user should then click the 'Search' button, being subsequently presented with the search engine hits; the user is offered one of two search engines – *Google* or *Yahoo* – to scan the text for (Internet) plagiarism. Thirdly, the user is asked to use the browser's 'Back' button to return to the search window and check other phrases for plagiarism.

Plagiarism Checker bases its search on a minimum threshold of 6 words, because of claims that shorter strings of text are more likely to produce false positives. On the other hand, owing to the fact that search engines like *Google* and *Yahoo* impose a limitation of 32 and 50 words, longer strings of text have to be edited, so that the system can handle them.

The system, however, has some drawbacks. The first is related to the service terms and conditions: the users are warned that searches performed using this website's functionalities by bypassing the system will be blocked to avoid potential abuse. On the other hand, the results are not always the most successful. For example, the string

`the users are asked if they wish to proceed ‘‘with the search’’`

is a paraphrase of the information provided on *Plagiarism Checker*'s website ('asking you whether you wish to continue with the search'), with only the last three words being identical, yet properly quoted. A search performed on Plagiarism Checker informs the user that no results were found, and the search results matching the phrase, without quotes, are presented. This means that the detection procedure is faced with at least three problems. Firstly, the system is unable to search for paraphrased material, as the inability to identify the string 'the users are asked if they wish to proceed' demonstrates. Secondly, the fact that phrases shorter than three words are eliminated means that the phrase 'with the search' would pass undetected if unacknowledged. Thirdly, the phrases introduced by the users in the search box are apparently searched against exact matches. The fact that *Google*'s message

`No results found for ‘users are asked if they wish to proceed with the search’`

is returned, and the user is presented with '[r]esults for users are asked if they wish to proceed with the search (without quotes)' is indicative of the system's operational constraints imposed on the detection of phrases that are reused with minor changes.

Plagiarisma

Plagiarisma is a free, online plagiarism detector that targets mainly University students and researchers, who might like to check their work for plagiarism prior to submission. The software is compatible with over 190 languages, and offers an online version and a free *Windows* client. The online version allows the user to search for plagiarism on the Internet, using *Google*, *Bing* or *Yahoo* APIs; the use of *Google* is restricted to registered users only. Users are offered a search window, where they can input or copy and paste suspicious text, and select between 'exact match' or 'TOR' (terms of reference) before hitting the 'Check Duplicate Content' button. Alternatively, users can upload documents in several file formats, or check websites for duplicate content by providing a URL. In both cases, the search is limited to 50,000 words. Subsequently, the system returns a table indicating whether the string of text is unique or borrowed, the queried text string and, if the string is not unique, the domains where the text can be found. This information is supplemented by a report containing statistics, namely the number of characters, words, unique sentences and percentage originality, as well as the option of saving this report as a PDF file.

Alternatively, *Plagiarisma* allows users to restrict the body of original texts against which to compare the suspicious documents, by limiting the search to *Google Scholar* or *Google Books*[6]. One enables the user to compare the suspicious documents against journal and newspaper articles, patents and editorials available on *Google Scholar*, whereas the other limits the search to resources available on *Google Books*. In this case, the user is not allowed to use the 'Exact match' or 'TOR' options, or to search against other search engines.

Since the software is mainly aimed at University students and researchers, it offers two additional tools: a 'Synonymizer' and a 'Spell Checker'. The Spell Checker, which underlines the misspelled words like a common spell checker, is a simple spelling and grammar check tool that supports 11 languages (Danish, Dutch, English, Finnish, French, German, Italian, Polish, Portuguese, Spanish and Swedish). The 'synonymizer' allows a user to copy and paste text into the text field, or to upload a document, select the desired percentage of rewriting, and the method of rewriting (i.e. 'Manual' or

[6] *Plagiarisma* Website – `http://plagiarisma.net/`

'Auto'). The system then returns an alternative text, which allows the user to select multiple variations to words, phrases and even sentences from a drop down option, and 'pick and choose' the words for the alternative version of the text. The 'synonimizer' operates like an 'article spinning' tool that helps an author rewrite the same article and pass it off as a new one. This raises ethical issues, especially in relation to encouraging what is usually called self-plagiarism.

Plagium

Plagium is a free online plagiarism detection service that searches for documents publicly available on the Internet or within private repositories. The service supports French, English, German, Italian, Portuguese and Spanish, and works by copying and pasting suspicious text into the search box. Although there is no limitation to the length of the text inputted in the search box, the user is warned that the longer strings of text take more processing time.

The software scans the web for matching text after hitting the 'Quick Search' or the 'Deep Search' button. The 'Quick Search' function is free, and allows searches of up to 25,000 characters. Conversely, the 'Deep Search' function allows the user to do more searches, over texts containing several paragraphs, which is not possible using a search engine directly. The system works by applying lexical and grammar rules to divide the paragraphs of each document into chunks of text before searching the Internet for text matches. The search results are then displayed paragraph by paragraph, with the original text on the left hand-side of the screen, and the matching texts retrieved from the Internet on the right hand-side, with links to identical or near-identical matches that allow the original document previewing. Links that are irrelevant can be deleted from the list, before producing a printable report.

Plagium also returns a graphical representation of the text reuse over time: the *Plagium Timeline*. Each textual match identified is represented by a 'bubble' along the timeline, offering the user access to a visual representation of text reuse across different web publications over time. However, this graphical representation is based on the date and time of publication, and not of creation. The acknowledgement of this limitation is important, especially to avoid using the timeline as an indication of the directionality of the borrowing.

The service also has other drawbacks. One is that the 'Deep Search' function is subject to payment. For instance, a payment of 10.00 US Dollars allows up to 220 searches, at an average of 5,000 characters per search, and up to a total of 1,100,000

characters. The other is that it is admittedly indicative only. The search returns chunks of text matching the text input in the search box, and the user is then asked to judge whether the matching text is properly attributed to a legitimate source or not.

Safe Assign

SafeAssign is a web-based service developed by *Mydropbox*, and currently owned by BlackBoard, that integrates with current VLEs. The system, which is supplied at no extra cost with enterprise versions of *BlackBoard*, allows instructors to set up *SafeAssignments* on the *BlackBoard* modules, so that students can submit their own assignments, or upload the assignments directly. As described by Maurer *et al.* (2006), fingerprints of the submitted papers are kept in individual databases of the account holder institution, thus avoiding legal or copyright issues. *SafeAssign* then checks the assignments for matches using these fingerprints, against the following databases:

- The Internet, by comparing it to a comprehensive set of indexed documents that are publicly available.
- The *ProQuest ABI/Inform* database, which has over 1,100 published titles and exclusive access to about 2.6 million articles from the 1990s to the present, and weekly updates.
- Institutional document archives, which contain all assignments that were submitted to *SafeAssign* by users in their institutions.
- A comprehensive index of documents available for public access on the Internet.

After conducting the comparison, the system produces an originality report showing the results of the text matching process. This report is divided into three panes. The first pane, *Paper Information*, includes details of the paper (e.g. submission date), as well as other information, such as the *Matching Index*, which shows the percentage of the paper matching other sources. It also allows the report to be saved as a print-friendly document, or made available to others via a URL. The second pane, *Suspected Sources*, shows a list of URLs containing textual material matching the one of the assignment. A tick box beneath each link allows the source to be removed from the comparison, so that the user can reprocess the assignment and recalculate the matching values. The third pane, *Paper Text*, shows each string of matching text, each source being indicated by a different colour.

The system calculates two different scores: a sentence matching score and an overall score. The sentence matching score indicates the probability of two sentences having 'the same meaning', and the corresponding probability that this matching is due to chance. The overall score, on the other hand, indicates the percentage matching of the suspect document with the sources identified.

- **Below 15 %:** Very low matching with other sources; assignment can include quotes and small amounts of text matching other documents, but does not indicate an instance of plagiarism.
- **15 – 40 %:** Extensive use of quoted or paraphrased text; hence may include instances of plagiarism.
- **Over 40 %:** High probability that the assignment has borrowed from the other sources; may be indicative of plagiarism.

These three scores and their corresponding descriptions suggest that the software is a good performer. However, a deeper consideration of these descriptors suggest that this performance can be misleading. Firstly, the definition of plagiarism used by the system is unclear. For instance, it is indicated that instances can be assessed as plagiarism based on 'quotes' and 'common phrases' matching other documents, but the description does not clarify the procedure applying to properly acknowledged vs. unacknowledged instances, or how it interprets the common phrases. Secondly, this description suggests that the detection is based purely on quantitative aspects, so that papers including 'quoted or paraphrased text' that is considered to be 'in excess' are prime suspects of plagiarism, and require a closer inspection.

One of the main criticisms to this approach is that it does not take into account the fact that these instances might actually be properly attributed, and disregards the fact that, if anything, excessive use of quoted or paraphrased material is more a case of academic quality than an instance of plagiarism, as discussed in chapter 1. Thirdly, it is suggested that the system is able to detect paraphrased text/text with the 'same meaning'. This ability is based on the probability, calculated by the system, that two sentences have the same meaning, and whether this matching is due to chance or to the suspect text containing content from the known source. However, as the detection process is based on document fingerprinting, the system can at most detect similar and/or identical strings of text. Therefore, this probability calculation is based on identifying the changes that have been introduced to the text, whether these changes were

minor or major, and not on determining whether the sentences have the same meaning, as is the case of paraphrasing. Another drawback of *SafeAssign* is that it tends not to manage page transitions properly. As a consequence, it breaks a sentence at the end of the page, even if that sentence continues on to the following page.

Turnitin

Turnitin is a web based service and proprietary software package developed by *iParadigms*, and one of the most popular plagiarism detection software packages worldwide. It is a paid service, which requires one of several account types, and that allows a comparison of documents against three different sources: indexed archive of billions of Internet pages, journals and books in the *ProQuest*(TM) database, and millions of documents already submitted to the *Turnitin* database (Maurer *et al.*, 2006: 1061). The system can be integrated into most common VLE (virtual learning environment) platforms, and works remotely by uploading the documents to the system database to process and compare the documents, and subsequently detect instances of plagiarism. It creates a fingerprint, i.e. a string of a moderate size (Maurer *et al.*, 2006), of those suspect documents, prior to storing this information. The system then produces an originality report, which includes text matches and links to the (supposedly original) sources marked up with different colour codes, indicating the intensity of plagiarism.

Recent developments of the software included a functionality to detect plagiarism resulting from translated text. This functionality works by translating the documents from several languages into English and then comparing them against the *Turnitin* database. Translation into other languages is advertised as a future plan. Since *Turnitin* is a commercial software, no explanation is provided of how this operates.

One of the system's drawbacks is that the report requires a careful interpretation to avoid being misleading. Since highlighted matches are often false positives, the description of the intensity of plagiarism does not mean that the highlighted instances are actually plagiarism. As Maurer *et al.* (2006) argue, some instances flagged as plagiarism are not necessarily so. This is the case of figure 6.3, for example, where quoted sentences are missed as acknowledged instances, and consequently described as plagiarism.

A possible justification for this procedure is the fact that fingerprinting techniques often require that items like punctuation are removed, so as to allow for a more direct comparison in cases where changes have been introduced to the strings. By doing

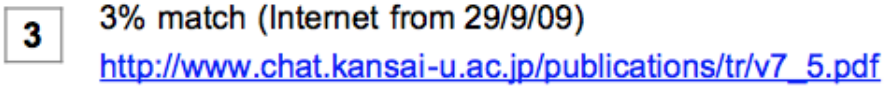

Figure 6.3: Instance of false positive detected by *Turnitin*.

so, the software also removes the quote marks, hence preventing the system from identifying possible cases of attributed authorship.

Urkund

Urkund is used widely in Scandinavia. The software is compatible with VLEs, and has a preventive and a controlling function. On the one hand, it aims to assist lecturers, tutors and examiners, by helping them check the contents of academic papers for plagiarism. On the other hand, it allows students to submit assignments to the system in order to pedagogically support them throughout the writing process, while preventively ensuring that their work will not be plagiarised.

The system works in three steps. Firstly, the student emails the document(s), as an attachment, to the instructor's dedicated e-mail. Subsequently, the documents are submitted to *Urkund* for content check. All text formatting, including bolds and italics, as well as non-textual elements like tables, are removed, and the queried documents are verified against three different types of sources indexed by the system; the Internet, published material and student material. The Internet includes open and password protected webpages, as well as 'Paper Mills' and other sources. Published material includes books (in electronic and conventional format), journals and newspapers. Student material includes assignments submitted to *Urkund* by participating organisations, and which are part of the '*Urkund Archive*'. Upon submission, the user can choose to exclude the document as a future source from the system database, in which case it remains restricted to the user's institution.

The system then screens the suspect documents for plagiarism and generates a report where possible instances of plagiarism are flagged. The report shows similarities between the checked document and the sources indexed by the system. Like other systems (e.g. *Turnitin*), it is the reviewer's role to check those similarities and determine whether they represent instances of plagiarism.

The report generated by the system presents the text submitted for analysis with the percentage of overlap, if any. Unlike other software tools (e.g. *Turnitin*), the comparison between the submitted documents and the sources is made based on blocks

of text, rather than at phrase, sentence or paragraph level. The potential severity level of the matching is indicated by the colours and as a percentage score. The colours range from green to black, where green indicates 'no matches' and black indicates that the whole text is a match. The percentage score, which is calculated based on the ratio between the matching text and the whole document, includes reformulations, changes in word order, replacement with synonyms, among other linguistic modifications that are not identified in the software manual. The report also gives the reviewer access to the main and alternative sources within the same block of text. Only those segments of the text that the system found to be similar to the queried text submitted are shown in the source column. Sensibly, although the system does not offer the option to exclude quotes from the system check, it allows the user to toggle quote highlighting on or off. Like other software tools, the generated report therefore needs to be analysed by the reviewer, in order to interpret the results obtained and verify that the similarities found represent instances of plagiarism.

The software description provided on the software website indicates that the software performs well in identifying borrowed text whose order is shifted, like the following example:

Suspect text: *diving has become very popular, both abroad and at home*
Source: *diving has become extremely popular, both at home and abroad*

The approach described raises some issues, however. Firstly, it does not explain the method used to calculate the degree of similarity. Secondly, the example provided suggests that each block of similar text may include several different matches. Interestingly also, the system keeps a record of the email addresses from which the assignments are sent, in order to avoid, on the one hand, that previous work submitted by the same student is used as a source, and, on other hand, unnecessary matches with earlier versions of the same document. Although this might contribute to improving the system's robustness, especially by avoiding false positives, it also raises the problem of missing the detection of reuse of work previously submitted by the same student, from the same email address.

Overall, the system works more as a machine-assisted, rather than automatic, plagiarism detection tool, which contradicts the claims made when advertising the system. Firstly, a contribution of the reviewer is necessary to confirm or otherwise discard the instances of plagiarism. Secondly, it offers a manual search function that allows the user to copy and paste or type a minimum of 400 characters free text into

the text box provided for comparison. These, together with the fact that the system allows the reviewer to provide Internet sources that will be excluded from the analysis, via source filters, indicates that the system has some degree of customisation, rather than being 'completely automated'.

Viper

Viper is a free *Windows* software client that is available for download from the developer's website, upon registering for a free account. It provides a side-by-side comparison of documents of unlimited length against over 10 billion resources published on the Internet, as well as documents stored locally in the user's computer. The software offers an easy-to-use interface, and returns a comparison highlighting areas of plagiarism. If a match is found when a user submits an essay, other people can see only a percentage match, but not the matching text itself.

The service works by uploading the files to be scanned for plagiarism from the user's computer. The user can choose to compare the essays against documents on the PC, or Internet sources online. Once the system scans the files, it will report matching text in the local files or in the online sources, and indicate exactly where the matching is. The essay can be compared side by side against the source where the information was found, using a dedicated button. When the comparison is finished, it produces a report that can be saved for reference.

This software poses some problems. Firstly, the report can be misleading if misinterpreted, since the system can show false positives; *Viper* identifies a match where there are strings of 5 or more overlapping words. Additionally, the system is not necessarily able to identify material that has been properly attributed to the sources, although it is able to identify material included in quotation marks. On the other hand, the system offers a threshold upon which instances of plagiarism are rated, which once again can be misleading:

- **Overall plagiarism rating 6% or less:** The document is unlikely to have been plagiarised, although careful check is advised for lengthy documents.
- **Overall plagiarism rating 6 - 12%:** The risk of containing plagiarised material is low, but checking the document is advised to ensure that every section has been properly referenced.
- **Overall plagiarism rating 13 - 20%:** The risk of included plagiarised material

is medium, and this may be due to sections that have not been properly referenced, or mistakes with quotation marks (e.g., the system is unable to handle apostrophes rather than inverted commas for quotes, or to identify cases where the inverted commas are open, but not closed).

- **Overall plagiarism rating equal to or higher than 21%+:** The risk of containing plagiarised material is high, indicating a significant number of matching fragments. As in other cases, this might be due to problems with inverted commas or failure to reference or rephrase properly.

Finally, the procedure used by the software can raise ethical issues, as the essays submitted for comparison are added to the developer's database, so that they are available for future comparison against other documents submitted. The developer guarantees the database's confidentiality, and that nobody has access to it. However, by uploading a document for comparison, the user implicitly gives the developer permission to add it to the software database, nine months after submission. This buffer period aims to ensure that the work is only available for comparison long after submission and marking, and hence does not represent any risk that the student user is accused of plagiarism. Additionally, by giving the developer permission to use the essay after this period, the student allows his/her work to be used as an example of good practice, offered via a student essay resource called *EssayCoursework*. The developer also claims that the documents are not indexed by search engines, so only the first 300 words of the essay are shown in Internet searches.

WCopyfind

WCopyfind is a free, open source software, compatible with *Windows* platforms. The software needs to be installed in the user's computer, and works by comparing documents (Scaife, 2007) to find similarities between them, at word and phrase level. It is available in two versions, which differ only in terms of graphical user interface (UI); one, *Copyfind* is provided as a non-graphical, command-line UI; the other, *WCopyfind*, is offered as a graphical user interface for *Windows*. The system starts by loading, hashing and comparing documents, then providing reports on the overlapping text. The first step, *loading*, consists of reading each document provided by the user, one word at a time.

Secondly, each of the words read is converted into a 32-bit hash code, i.e. a simple, numerical representation of the data, that allows the software to store and retrieve the

data faster. At this stage, characteristics of the text like case, numbers and punctuation can be changed or removed, depending on the user's settings. This is helpful when comparing text above the sentence level, where punctuation can interfere with the detection process, but has the disadvantage of returning misleading results in other cases, by producing more false positives. Hashing consists of creating a list of 32-bit hash codes, and giving each word one hash code. The codes are then sorted in the order that they appear in the document. Subsequently, it creates a list of hash codes sorted in numerical order, which corresponds to having an alphabetical list of the words in the document. Then it compares the documents, by scanning them for 'matching phrases'. The numerically sorted hash-code lists allow the software, in this phase, to find the matching pairs faster. When matching hash codes are found in the numerically sorted list, the software looks for matching phrases around those matching hash codes.

The system is compatible with a range of file formats, and allows the user to customise several personal settings. The first one is the opportunity to select, in the 'Most Imperfections to Allow' field, the maximum number of mismatching words in the sequence. Secondly, the user can indicate the 'Shortest Phrase to Match', which allows the system to check if the longest matching phrases found are at least as long as the value indicated in this field, thus operating like a threshold value. By default, words with three or less characters are not used as the starting point for the matching process, since, it is argued, typical documents commonly include words of this length. Finally, marked up versions of the documents with matching phrases are provided in a preliminary report, and an overall comparison report is generated if the software finds a long enough match. The system is advertised as being robust enough to process thousands of documents, at a rate of 30 documents per second.

Some of these features impose limitations on the system. Firstly, it admittedly has some problems in dealing with 'imperfections' that are longer than one word between the two documents, which may impact the detection. Secondly, since the system is unable to count non-matching words in those phrases, it can be misled by minor lexical changes. Thirdly, since the analysis report is based on shared text of considerable length, there is a risk that shorter strings be missed. This risk cannot, however, be ascertained, since that length is not specified. Additionally, since the system works locally, all documents for comparison need to be provided by the user, which prevents the comparison against unknown documents; although Internet searches are not per-

mitted, a comparison can be made against websites, by providing the hyperlinks to those webpages. The advantage of local over server-based systems is that the integrity of the files is not compromised; the disadvantage is that a comparison can only be established among a set of documents provided by the user.

WriteCheck

WriteCheck is a plagiarism detection and grammar checking software tool that was developed by the same team as *Turnitin*. It supports most of the standard file formats, and is available as a pay-as-you-go option. Contrary to *Turnitin*, which identifies the suspect text and indicates matches with the potential sources, *WriteCheck* only identifies the potentially unoriginal text for reviewing, as it is targeted at student users. Unlike *Turnitin*, where papers are added to the programme database, *WriteCheck* does not add the assignments to the database. *WriteCheck* works by uploading a paper to the system and submitting it for comparison against the comprehensive *Turnitin* database. The paper is subsequently screened by the grammar checker for spelling and grammar, style, mechanics and usage. The grammar check uses the *e-rater* grammar analysis tool developed by the Education Testing Service (*ETS*), and provides links to grammar handbooks. The style checker verifies instances of passive voice usage and long sentences. The mechanics checks verifies fused words and compound words, and the usage check verifies the document for faulty comparisons and confused words.

After the analysis, the system produces a report highlighting text matching the database sources and any errors. The user is then able to review the matching instances and the errors reported, and is given the option to hide individual style marks. The software also advertises the use of bibliography and quotation filters, which should allow in-text quotations and bibliography to be excluded from the plagiarism statistics. This feature was unavailable in previous versions of *Turnitin*, as the case in figure 6.3 above demonstrates.

6.4.5 System Comparison

The description of the systems above shows that plagiarism detection methods currently rely, at least in part, on computational assistance, while differing mostly on the level of human intervention required. For example, the *Glatt Plagiarism Detection* could work independently of heavy computational assistance, mostly with human in-

tervention; in fact, although using a computer obviously provides productivity gains, removing every fifth word of a document and replacing that word with a blank space can be done manually. Other systems, such as *CopyCatch*, albeit strongly machine-assisted and involving sophisticated programming algorithms, require a high level of human intervention, starting with selecting and uploading the documents for comparison, and including setting the similarity threshold, and, more importantly, analysing the results. Other systems, like *Viper*, allow the user to select the length of the phrases to be checked, but the remaining matching process is automated. The remaining systems are even more automated, and hence require less human intervention.

Most of these software tools involve searching for textual matches online, but others also search against proprietary or institutional databases, or even against digital publications (e.g. *Turnitin* and *SafeAssign*). Additionally, although they tend to vary in some minor details, such as the search engine used in the online plagiarism detection process, these tools restrict human intervention to copying and pasting the text for checking, and/or checking the results for obvious false positives.

The systems analysed also vary in the quality of the results that they return, and which are based mostly on two aspects: the computational approach to the detection procedure (e.g. removing and/or ignoring punctuation may not be irrelevant); and, especially, the linguistic procedure adopted. The description of the systems above indicates that most detection software packages can easily identify instances of *verbatim* plagiarism, even if their approaches vary slightly, e.g. in terms of search length. Their relative success in achieving this analysis is due to the nature of the task, which consists of finding exact matches, or 'duplicates', as they are called in computational linguistics.

Where these systems differ the most is in the level of robustness in finding near-duplicates, i.e. instances of borrowing where minor or major changes are introduced to the original. Some tools are able to find near-duplicates with minor changes, such as adding, editing or removing a few words of the original, while keeping the rest of the text. The greater or lesser success owes, in this case, to the ability that systems have of computing distance similarities. However, the degree of difficulty in finding near-duplicates increases when more sophisticated linguistic strategies are used to plagiarise, such as *paraphrasing* or *translation*. The latter can be resolved via the combination of simple computer tools currently available and manual techniques, such as machine-translating (part of) a suspicious text and then searching for the translated

text on the Internet. The solution to the former, however, is far more complex. As previous research demonstrated, although paraphrase generation can be performed with some degree of success, depending on the system used, paraphrase identification so far has mostly provided disappointing results (Madnani and Dorr, 2010). Therefore, even when systems indicate that they are able to handle paraphrasing, at most they are able to perform a general search of synonymised texts and highlight matching words. This assumption is misleading, since it does not account for more sophisticated types of plagiarism. Paraphrasing is more complex than simply introducing minor changes to exact matches. It should be taken to include also more sophisticated strategies of text rewording and reorganisation.

Table 6.1 below shows a comparison of the systems described above in terms of their linguistic approach. Types of plagiarism included are verbatim plagiarism (word-for-word plagiarism), paraphrasing, translation (translingual plagiarism) and lexical (detection of lexical overlap).

Until now, the most successful plagiarism detection system in finding near-duplicates is perhaps *CopyCatch*. By establishing a lexical comparison between texts, it allows the system to compare texts and identify overlapping lexical items, regardless of their immediate co-occurrence and/or distance in the text. Additionally, the method of lexical analysis underlying this software, not only demonstrates excellent results in finding instances of monolingual plagiarism, but also translingual plagiarism, as discussed in chapter 5. The outputs of the analysis, in particular the statistics and the report where the matching text is marked in red, can then be analysed by the user to assess a textual instance as plagiarism or otherwise.

6.5 Chapter Summary

This chapter discussed different computational approaches to plagiarism detection. I strongly argued that more traditional computational methods tend to encourage an oversimplification of the plagiarism detection, by reducing it to finding *duplicates* or *near-duplicates*. Moreover, I argued that this misrepresents the problem of plagiarism detection, by leaving aside other plagiarism strategies. In addition, most of these software tools have similar capabilities and tend to perform reasonably well in detecting verbatim plagiarism, but their performance tends to drop when screening for other types of plagiarism. Conversely, few systems approach plagiarism detection from a computational linguistics approach that values the contribution of linguistics. The find-

Software tool	Verbatim	Paraphrase	Translation	Lexical
Article Checker	Y	N	N	N
CiteReader	Y	N	N	N
CopyCatch	Y	N	N	Y
Docol@c	Y	N	N	N
Dupli Checker	Y	N	N	N
Ephorus	Y	N	N	N
Glatt	Y	N	N	N
Google Translate	Y	N	Y	Y
Plagiarism Checker	Y	N	N	N
Plagiarisma	Y	N	N	N
Plagium	Y	N	N	N
SafeAssign	Y	N	N	N
Turnitin	Y	N	Y	N
Urkund	Y	N	N	N
Viper	Y	N	N	N
WCopyFind	Y	N	N	N
WriteCheck	Y	N	N	N

Table 6.1: Comparison of linguistic features addressed by plagiarism detection systems.

ings indicated that these are the systems that perform the best.

The findings also indicate that more automatic plagiarism detection systems return more limited results, whereas systems involving more human intervention provide more usable and acceptable results. An alternative term was offered to account more accurately for this type of plagiarism detection: *machine assisted plagiarism detection*. This approach allows the detection procedure to include a combination of selected methods of authorship analysis to the detriment of others for computational plagiarism detection. I strongly argued that more accurate approaches to plagiarism detection need to be linguistically invested, for application in investigative and evidential contexts.

This chapter also demonstrated that extra measures are required to ensure a detection system that is able to identify more sophisticated cases of borrowing, such as paraphrasing and translingual plagiarism. The next chapter discusses some of the

core problems identified in this chapter.

Chapter 7

Improving Plagiarism Detection Methods: A Forensic Linguistics-based Contribution

7.1 Introduction

This chapter discusses the usefulness of current computational linguistics for plagiarism detection, and proposes methods to improve the performance of detection systems. Based on the features, limitations and constraints of existing approaches to plagiarism detection, this chapter claims that detection systems can be improved without developing new computational tools. However, it strongly argues that improvements can be achieved using existing software tools.

Secondly, this chapter claims that one of the problems underlying the poor performance of current systems is the methodology adopted, which consists of testing the software using corpora of artificially-created, and oversimplified plagiarising text. It is hypothesised that using real corpora, ideally of a forensic nature, rather than 'invented' examples provides a more realistic approach to the detection task. A change in the research paradigm is therefore proposed, in order to allow for systems to be tested over natural language texts containing instances, or at least naturally-occurring features of plagiarism. It is suggested that this approach allows the systems to 'learn' from actual examples, consequently bringing additional benefits to detection systems.

The chapter subsequently discusses the main strategies used for plagiarising, in particular verbatim, word-for-word plagiarism, as well as lexico-grammatical analysis and paraphrasing. Linguistic aspects like lexical items, morphology, syntax, semantics and discourse, as well as changes in spelling, all of which are relevant to natural language processing of different language variants, are also discussed in relation to

these strategies. The principle is challenged that overcoming all the limitations identified is realistic, considering the computational requirements involved. By focusing on the identification of obfuscation techniques used by plagiarists, these methods can specifically assist the detection of instances where those strategies are used intentionally.

The chapter concludes by discussing paraphrase-based plagiarism, which is one of the main problems that Maurer *et al.* (2006) identified as being perhaps the most demanding challenge for detecting plagiarism in the coming years. It concludes that paraphrasing remains a demanding challenge, even by adopting a multiphase approach.

7.2 Systems of Detecting Linguistic Strategies

In recent years, diverse tools have been developed and used individually or collectively to assist plagiarism detection, including morphological analysers, parsers, transliterators and part-of-speech (POS) taggers. Each of these tools performs a different task, with different results. However, the development and use of these tools are language-dependent, rather than universal. Consequently, the linguistic and computational effort required for different languages varies, depending on the breadth of research conducted, and possibly on the cultural importance attached to plagiarism. This does not necessarily mean that research on less widely spoken languages is non-existent[1]. But it certainly is comparatively scarcer than, for instance, in English. The application of these tools to different languages means, in some cases, transferring existing tools to new languages; in other cases, the linguistic information underlying the system and some specific algorithms related to processing that information being language-dependent, particular development efforts are required. This research describes how these tools can be developed to improve the detection of each type of plagiarism in Portuguese. As discussed in chapter 1, building a fully operational detection system is far too big an endeavour to be developed in the scope of this project. However, identifying and discussing the main challenges facing the detection procedure can contribute to a more focussed programming of future detection software.

Mozgovoy (2008), and Maurer *et al.* (2006) before him, argued that many instances of plagiarism are only possible to detect by means of human investigation. They there-

[1] For example, on a post to the Corpora List on 23/05/2012, Fatima Zuhra claimed to have developed such tools for Pashto.

fore recognised the limitations currently imposed on plagiarism detection systems, admitting that computer-assisted plagiarism detection can only provide some scores indicating the similarity between documents, but not prove the existence of plagiarism. Analyses of the potential instances of plagiarism beyond the similarity scores require additional manual efforts. However, they also claimed that some other types can be detected using computer tools, either open or hermetic. Hermetic systems are usually based on simple file-file comparison methods. Mozgovoy (2008) gives as an example the UNIX *sdiff* command, which in simple terms consists of comparing two files to find the differences between them, merging the results and outputting them to a new file. These systems are usually not demanding in terms of speed and space requirements, and can be easily implemented in any computer running the *Mac OS*. Conversely, open systems are those more commonly used, especially by commercial plagiarism detection systems. They are a variation of Internet search engines, and require significant computational resources (Mozgovoy, 2008). Capable systems need to be able to handle several challenges at the pre-processing, processing and post-processing phases, from issues of file formatting to output of the results in a form that is clearly understandable.

Firstly, these systems need to be capable of handling data storage properly. Besides legal and ethical issues arising from the retention of information in proprietary databases, there are technical issues resulting from the processing power of computers to handle high volumes of information properly. Secondly, indexing and performance issues are raised, both of which are crucial. It is commonly agreed (Mozgovoy, 2008; Woolls, 2010) that, for the sake of practicality, the comparison algorithms need to operate fast. However, the higher the number of webpages available, the greater the capacity required from processors, not only to conduct the search queries, but also to pre-process (*index*) the files, so that the data can be collected, parsed and stored for information retrieval. Additionally, decisions on whether the computer processors within the same network share the same memory (*parallel computing*), or whether each computer uses its own private memory (*distributed computing*) can impact the performance (Mozgovoy, 2008). In parallel with this, the quality of the results obtained largely depends upon external efforts. A specific example is that of indexing. Despite recent efforts to index the Internet, the exact number of indexed pages is unknown. Estimates such as those provided by the *World Wide Web Size* website[2] indicate that

[2] `http://www.worldwidewebsize.com/`

in April 2012 *Google* had almost 50 billion indexed webpages, compared to 16.2 billion webpages indexed by *Bing* and about 3.2 billion webpages indexed by *Yahoo! Search*. Although this seems a considerable number of webpages, it is far from covering all webpages available. And the fact that *Google*, for example, stopped providing this information several years ago suggests that indexing is a very challenging task, and that indexing all new webpages being created every day is a difficult task.

Additional limitations result from the length of the search strings, as was discussed in chapter 6. *Google* searches, for example, are limited to 32 words, so strings of text longer than this cannot be easily searched. Moreover, systematic queries to search engines, even when using methods like *Google Apps*, are now restricted to a certain limit, which consequently conditions the search for plagiarism.

Moreover, as Mozgovoy (2008) indicated, the results of the NLP task depend significantly on the quality of the linguistic data used to train the system, even when searching for identical text. That subsequently determines the quality of the detection results. Ideally, the data used to train and test the systems should be real forensic data. However, owing to demanding access restrictions imposed on this type of data, on the one hand, and to the volume of data required for analysis, on the other, computational methods have traditionally been tested using artificially-created data. The use of this type of data has the advantage of allowing creative strategies, which may not always be present in real cases, to be tested. As discussed in section 1.3.4, an analysis of instances of plagiarism that is based on descriptive linguistics accounts for the strategies used, and has the potential to explain why a certain instance represents plagiarism. This does not, however, account for other possibilities. The downside to this type of data is that they are usually produced in a controlled environment, using specific instructions of what is to be done to the text. This inevitably results in unnatural representations of plagiarism, and imposes constraints on the performance of the detection systems.

The following sections identify and discuss some linguistic issues that remain a challenge to plagiarism detection systems, and propose methods and common tools to address them. Although ideally these should be tested in practice, using corpora of instances of plagiarism, the nature of the additional computational work required is too vast to receive more than mere mention in this work. Besides, work of this type can be implemented more rapidly and efficiently by a computer scientist.

7.2.1 Verbatim Plagiarism

Verbatim borrowing is the type of plagiarism that most detection software can handle more or less competently, as discussed and illustrated in chapter 6. From a computational perspective, the detection task is quite simple: given a suspect text, the aim of the software is to compare a suspect document either (a) against a document or set of documents that are known to be original, or (b) against an open document source, such as the Internet. The two detection procedures vary also in the amount of preprocessing work involved. Whereas (a) requires that all documents are indexed for comparison, (b) requires only indexing the suspect document for comparison against the already indexed documents available online. The indexing procedure, however, varies in the level of detail and algorithm complexity.

In his description of hermetic plagiarism detection systems, Mozgovoy (2008) explains that the classification task has been approached considering two main classes: *fingerprinting* and *content comparison*. *Fingerprint*-based systems like the ones used by Bernstein and Zobel (2004), Brin *et al.* (1995) and Stein and Eissen (2006) consist of creating a fingerprint. A shortened, machine-readable version of a text, represented as a sequence of bytes, is created and this operates as a descriptor of each document in a collection. This descriptor is based on a set of numerical attributes, including text statistics such as the number of words (average per line and unique words) and the number of sentences, among others. If two fingerprints are close to each other, given a certain mathematical *distance function*, two documents are treated as being similar. The distance function is a mathematical measure that calculates the similarity between two elements (Cha, 2007).

Although Mozgovoy (2008) reports that more recent systems do not follow the fingerprinting technique due to the fact that fingerprints are significantly impacted by even small changes introduced in the document, Stein *et al.* (2007) have shown, using the *Wikipedia* as corpus, that good results can be obtained for similarities above 0.8, especially in terms of recall (i.e. the fraction of relevant plagiarism instances that the system successfully retrieves). This includes *n*-grams. Maia *et al.* (2008) had already discussed the power of *n*-grams as discriminant linguistic elements, but more recently, fingerprinting character *n*-grams has been demonstrated to provide good results in intrinsic plagiarism detection tasks (Kestemont *et al.*, 2011). Specifically, Kestemont *et al.* (2011) divided each text into groups of trigrams and compared each group against each other group for stylistic identity.

The apparently good results of this approach need, however, to be interpreted critically. Firstly, the fact that the system was evaluated as an intrinsic plagiarism detection system means that its performance capacity over the comparison of suspect and potentially original documents has not been tested. Secondly, the system failed to detect plagiarism in most of the instances where manual obfuscation of the text had taken place. Thirdly, admittedly plagiarising sections of short and medium length remain a challenge to the system. Finally, as Kestemont *et al.* (2011) claim, character n-grams have been used in stylistic approaches to authorship attribution, and can also be applied to intrinsic plagiarism detection. However, as discussed in the previous chapter, stylistic analysis for plagiarism detection imposes limitations, both where a document borrows from a single source, and in cases of academic plagiarism where the student fails to competently integrate multiple authorial voices. At most, systems of this type are able to identify different writing styles, but not necessarily plagiarism.

Content comparison techniques are the ones used by most plagiarism detection systems currently available. It consists of comparing the contents of two files, on the grounds that these contents are the ones that best describe the contents of the text structure (Mozgovoy, 2008). Content comparison can be based on *string matching* or on *parse trees*. The latter uses the structural information of a document to build a tree representing that document structure and compare it against the structure of other documents. When parse trees are used, a document is first tokenised, i.e. divided into 'minimal units of meaning' (Mozgovoy, 2008: 28), such as words. Then it is analysed in order to determine the document structure. A graphical comparison of the syntactic organisation of the document is then provided that can be used to compare the structure of two documents, e.g. by analysing the order of the sections of an article, where the document hierarchical structure plays an important role. The downside to this analysis, however, is that identical text genres, such as academic writing, presuppose using identical structures. At this point, the analysis ceases to work as a discriminant element, and comparing two texts in terms of their structure (chapters, sections, subsections, paragraphs and sentences) becomes irrelevant to determine whether two texts have been produced independently. On the other hand, this approach is more complex and slower than string matching (Mozgovoy, 2008).

Conversely, string matching consists of dividing a text into strings for comparison. Different algorithms have been designed to calculate the similarity between documents (Mozgovoy, 2008), using distinct formulae, with distinct results. Some of these algo-

rithms compare tokens shared between two files, or compare two files line-by-line. Other more sophisticated algorithms are based on detecting the strings of text that are used by both texts, but that do not overlap. To this end, the system needs to use as many tokens from both files as possible (Mozgovoy, 2008).

One of the most difficult decisions, in this case, is deciding on the minimum length beyond which a certain similar string is flagged for plagiarism, the plagiarism threshold. Coulthard and Johnson (2007) demonstrated, using a real example which consisted of searching a string on a search engine like *Google*, that any string equal to or longer than 9 words is, in principle, unique. Computationally, experimental results from research conducted by Barrón-Cedeño and Rosso (2009) identified bigrams and trigrams as the best comparison units for plagiarism detection, the former for its Recall potential, and the latter for its Precision capability. Barrón-Cedeño[3] reported an experiment whereby he split a set of phrases into *n*-grams of increasing order, from 0- to 14-grams. He reported that it is unlikely that two sequences of text of 6-grams or higher occur in two independently produced documents.

Empirical data shows, nevertheless, that strings shorter than this can be unique, and the opposite is also true. An analysis of the detection systems available reveals that, especially the commercial ones (e.g. *Turnitin*), do not indicate the search distance used. But other systems allow the users to decide on the search length. For instance, *WCopyfind* searches over a minimum length that is specified by the user, whereas *Plagiarism Checker* operates over a minimum length of 6 words. *Plagium* does not set any limitation in the search box, but the user is warned that, the longer the text, the longer the search will take, and the higher the risk that true positives are missed. Logically, shorter search queries tend to show a higher precision (i.e. when the system identifies a string as being plagiarism, this identification tends to be correct), but their recall tends to drop (i.e. it is highly probable that the system misses many cases). On the contrary, longer search distances allow for the system to search for identical textual elements across longer strings of text, which is particularly useful to identify instances of plagiarism where the plagiarist has inserted textual elements in the borrowed textual string, or possibly where the word order has been changed.

Instances of plagiarism where alterations have been introduced to the text are more difficult to handle.

[3]E-mail sent to the Corpora List on 18/04/2012.

7.2.2 Quasi-verbatim Plagiarism and Lexico-grammatical Changes

A common problem underlying plagiarism detection is where the plagiarist introduces changes to the original, by rewriting the text, replacing some of the original words with synonyms, deleting from or adding to the original, or even reordering the words in a sentence. Usually, when not acknowledged, this results from the plagiarist's attempt to disguise the borrowing.

These alterations, which even if minor also raise problems to the machine-detection procedure (Mozgovoy, 2008; Barrón-Cedeño and Rosso, 2009), cannot strictly speaking be said to represent a case of verbatim, word-for-word plagiarism. But it cannot be termed paraphrasing, either; as explained in section 7.2.3, paraphrasing involves more complex operations, resulting in the expression of the same meaning, but in other words. It is therefore perhaps more accurate to term this type of borrowing quasi-verbatim plagiarism.

Most of the strategies involved in quasi-verbatim plagiarism can be addressed using tools currently available. Deletions and additions can be handled, for example, by making adjustments to the similarity function, as described in the previous section. An issue previously reported when handling larger chunks of text was that systems broke down when matching passages of this type. In this case, matching passages can only be identified by selecting chunks of a smaller size, which represents an increasing effort in terms of fingerprint computation, comparison and storage. To overcome this issue, Stein and Eissen (2006), for example, proposed a method based on 'fuzzy-fingerprinting' to allow detection systems to identify candidates, even when handling larger chunks. This system matches similar strings, even when the exact chain is interrupted by words added towards the middle of the sentence, by increasing the search space while skipping words in the original that are not used in the derivative version. Conversely, borrowed words are also identified.

Similarly, cases of plagiarism where one or a few words are replaced with synonyms can also be addressed realistically, by using thesauri and word sense disambiguation modules. To resolve issues with cases where the meaning of words changes depending on the context, Mozgovoy (2008) suggests applying a word sense disambiguation (WSD) method, prior to running the thesauri. One of the methods that can be used in this case is a semantic analyser. The use of a semantic analyser allows words to be replaced with word classes, rather than the original words, for example.

In the following real example of academic plagiarism[4]

Original	Plagiarism
Pues bien, hace años, debajo de mi casa, había en una pared una frase en spray que decía “Puta Burgos” y eso no era un graffiti, era una pintada.	Pois bem, se numa parede há uma frase em spray que diz “Puta Espanha”, isso não é um graffiti, é uma pintada.

Table 7.1: Real case of plagiarism where one word is replaced with another from another of the same class.

the name of the Spanish city of ‘Burgos’ in the original is replaced with the named entity ‘Espanha’ (‘Spain’), possibly to make the derivative text more coherent. Although the two words do not belong exactly to the same class – the former belongs to the class <CITY>, whereas the latter belongs to the class <COUNTRY> – both classes are part of the same general class, say <PLACE>. The suspect and original documents can then be compared by the system, which looks for matches at the level of class names, instead of searching for the exact words.

Despite the good results of this method reported by Mozgovoy (2008), it also has some disadvantages. Firstly, the performance results depend to a large extent on the processing power of the computer. To enable the comparison of classes between the original and the suspect text, the system needs to replace the original words both in the suspect and in the original texts, which might impact the overall system performance. A second drawback admitted by Mozgovoy (2008) is that in some cases this procedure might produce false positives. He cites the example of the phrases *Vladimir Putin claimed* and *George Bush claimed*, which would both be ‘translated’ to <NAME><NAME><SPEAK>. In a plagiarism case, this could indicate a false positive. Thirdly, this method only works when a set of potential source texts is available. Therefore, it can only be applied in cases of ‘hermetic plagiarism detection’ (Mozgovoy, 2008) or plagiarism detection where reference texts (Stamatatos, 2009a; Barrón-Cedeño and Rosso, 2009) are available. This leaves aside text comparisons against open systems, such as the Internet.

Another strategy used by plagiarists is word and phrase reordering, while keeping

[4] From the CorRUPT corpus.

the same meaning. An example of this is the different use of pronouns in language variant pairs like European Portuguese and Brazilian Portuguese, which hold different places in the sentence in the two variants:

- European Portuguese: verb $>$ pronoun;
- Brazilian Portuguese: pronoun $>$ verb.

At its simplest, word reordering is easily detectable by current systems that are capable of detecting the different words independently. In practice, however, as argued by Mozgovoy (2008) and seconded by the description of the software packages in chapter 6, in order to avoid excessive false positives, most software packages base their operation on a 'shortest length string to match'. In this case, the systems are only able to identify as plagiarism strings that are above a certain threshold. Approaches of this type miss matching words independently, and therefore the ability to detect word reordering.

To resolve this problem, Mozgovoy (2008) proposes using a parser (i.e. a syntactic analyser) to convert the sentences into sequences that reflect the syntactic structure of the strings (sentences, phrases, etc.) without being affected by the word order. In his research, he used a parser that produces an output in the form of a dependency tree. Subsequently, a post-processing tool is used to convert the text into a format in which the words in the sentence are sorted according to the grammatical relation that best describes the dependency between them. The words included in each grammatical relation group are then sorted alphabetically, for ease of comparison.

One important drawback is that the preprocessing phase does not preserve the original structure of the sentences, which prevents the detection system from finding matches in the original documents. In order to overcome this drawback, he suggests either (a) preparing the parser to generate metadata that permits restoring the links between words in the original and parsed files, or (b) programming the system to highlight the whole plagiarising sentences, rather than just word chains (Mozgovoy, 2008). He discards the first alternative, despite admitting that it grants a higher degree of flexibility than the latter, because it would require introducing changes to the parser, which is not realistic. Testing these over corpora of artificially-created plagiarising news reels and essays, he reached several conclusions. Firstly, he reported a considerable increase in the similarity ratios after the parser was run. Secondly, a small value of the 'shortest string length to match' constant reduced the effectiveness of

the parser, whereas larger values increase the probability that these instances of plagiarism, detectable after implementation of the parser, are not identified (Mozgovoy, 2008). Thirdly, he found that using the parser allows a larger constant value to be used, with gains in processing speed. Additionally, it has the potential to provide an indication of the changes introduced, based on the calculation of the difference between similarity scores resulting from the original and the preprocessed collection of documents.

Some of the changes discussed above often result from the manipulation of language variants, which is an additional plagiarism strategy used by plagiarists to adapt the text borrowed from other variants of a language to their own variant, and thereby circumvent detection systems. A typical case, brought to the attention of forensic linguists by Coulthard and Johnson (2007), is the so-called 'Iraq dossier'. The case dates back to 2003, when the UK government released the dossier 'Iraq: Its Infrastructure of Concealment, Deception and Intimidation' to the media. This report on Iraq and its weapons of mass destruction was used by the government to justify the country's involvement in the 2003 invasion of Iraq. An analysis of the document conducted by Glen Rangwala[5] and discussed in the media revealed that the document plagiarised different sources, for the most part an article written by Ibrahim al-Marashi, entitled 'Iraq's Security & Intelligence Network: A Guide & Analysis', which had been previously published in the *Middle East Review of International Affairs*[6]. The document copied whole sections of the original article verbatim, including typographical errors, although some discursive changes were made to strengthen the arguments of the report. There were also some changes in spelling, such as the American spelling 'z' to the English spelling 's' in verbs and nominalisations.

Changes in spelling, used to reflect the standards of a language variant, are also common in instances of plagiarism in academic writing, as the analysis of the data in the CorRUPT corpus demonstrates. In particular, the data include some examples of texts written in European Portuguese that had originally been written in Brazilian Portuguese. In this case, the disguising therefore implied several changes, such as converting

$$cç > ç$$

$$ct > t$$

[5]Available at `http://www.casi.org.uk/discuss/2003/msg00457.html`

[6]The original article was published in Vol. 6, No. 3 (September, 2002).

$$pt > t$$

In this case, a simple preprocessing of the suspect and potential original using free normalisation software tools currently available that function as a dictionary of language variants[7] allows the text to be converted to an identical spelling, thereby making the comparison easier.

The detection of linguistic changes other than spelling is more complex. As the analysis of the instances of plagiarism in the CorRUPT corpus reveals, changes introduced to inflectional and derivational forms of the words, even when minor, result in morphological changes that impact the simplest verbatim plagiarism detection algorithm. A similar situation takes place with changes at the level of verb tenses. The following example from the CorRUPT corpus illustrates this point.

Original	Plagiarism
opção que condicionará a forma e o conteúdo da publicação.	Esta opção irá condicionar a forma e o conteúdo da publicação.

Table 7.2: Example of change of verb tense in real plagiarism cases.

In the example presented in section 7.2, the future tense of the verb *condicionar* – *condicionará* – is shifted to a future tense of the auxiliary verb followed by the infinitive form of the main verb (irá + condicionar). Although this is perhaps a very simplistic example, this plagiarism strategy can potentially be replicated several times, with several phrases, in the same sentence or text. Even if more complex cases may require more sophisticated approaches, morphological analysers like stemmers and lemmatisers can be used to revert lexical items to their root, dictionary form (their lemmas) for comparison and consequently assist the detection procedure. However, since stemming consists of removing the word endings, including derivational affixes, lemmatisation is preferred for its potential to return more reliable results. Table 7.3 gives an example of this procedure, using the word *boss*.

This happens because the rules used in the first phase of *Porter's algorithm* (Porter, 1980)[8] are the following:

[7]The *Lince* tool is available for Portuguese – `http://www.portaldalinguaportuguesa.org/lince.php`

[8]There are other algorithms applied to lemmatisation in English, but Porter's is reported to

Stemming	Lemmatisation
bo	*boss*

Table 7.3: Stemming and lemmatisation of the word *boss*.

$$\text{SSES} > \text{SS}$$
$$\text{IES} > \text{I}$$
$$\text{SS} > \text{SS}$$
$$\text{S} >$$

Lastly, as mentioned in chapter 6, most plagiarism detection systems ignore punctuation, except maybe for strong punctuation marks (such as full stops, question marks and exclamation marks). These are traditionally used to divide the text into sentences for analysis, but are missed by indexing procedures such as the one used by *Google*.

Ignoring the punctuation can bring both advantages and disadvantages to plagiarism detection. The first advantage is that not counting these tokens towards the length of the search string permits running search queries over longer strings of text. Secondly, when querying search engines, the search results are not limited by the boundaries set by the punctuation marks, therefore returning results even if the search terms are in different sentences. Thirdly, ignoring the boundaries of strong punctuation can be beneficial to detecting paraphrasing and summarisation, especially where the rephrasing takes place at a supra-sentential level.

The main drawback of ignoring punctuation is that in writing, spaces and punctuation are used to separate a word from other words that co-occur with it, and this is not irrelevant. For example, a distinction needs to be made between a collocation of two words (*xxx yyy*) and a collocation of the same words separated by a comma (*xxx, yyy*) or a full stop (*xxx. yyy*). Surprisingly, few detection systems have used punctuation marks to reduce the rate of false positives returned. The example presented in figure 7.1, from *Turnitin*, illustrates how using punctuation can contribute the volume of false positives.

A closer analysis of the highlighted passages confirmed that the reused text actually referred to a quotation, which the student cited (two words in the quotation, 'recognizable' and 'characterized', are not identified as plagiarism because the student

be one of the most commonly used – see `http://nlp.stanford.edu/IR-book/html/htmledition/stemming-and-lemmatization-1.html`

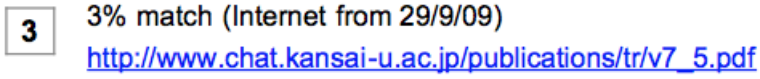

Figure 7.1: Percentage overlap identified by *Turnitin*.

changed the American English spelling to British English):

> 3**definition of genre** is that **"genre is a** recognisable **communicative event** characterised **by a set of communicative purposes identified and mutually understood by the members of the professional or academic community in which it regularly occurs."**

Figure 7.2: Instance of plagiarism detected by *Turnitin*.

This identification of quotations and citations by the system could be resolved easily by means of a simple algorithm that is able to identify text included in quote marks, and discard it from plagiarism statistics. Although the problem remains when searching the Internet for the suspect text using a search engine, owing to the fact that search engines in general are not sensitive to punctuation, it certainly contributes to improving the local comparison of files. A more complex problem is that of detecting paraphrasing.

7.2.3 Approaching Paraphrasing

Howard (1995) convincingly argued and Pecorari (2008) later demonstrated that many instances of academic plagiarism result from improper paraphrasing, either because the students (a) fail to correctly paraphrase original works, or (b) they use paraphrasing as a means to disguise the original authorship. In Howard's terms, the latter, being intentional, is plagiarism. The former, however, results from the students' lack of competence in academic writing, and should not be considered plagiarism, but 'patchwriting' instead. The improper borrowing results, in this case, from the attempt at academic writing resulting from a student's unsuccessful effort to rewrite or explain an author's works or ideas in other words. In her analysis of academic assignments and in the subsequent interviews that she conducted with students, Pecorari (2008) found that the students' inability to properly paraphrase in many cases reflected their own perception of paraphrasing, which illustrates Howard's point that proper academic writing is often a competence issue that not all students are capable of handling properly.

The plethora of existing definitions of paraphrasing reflects a lack of precision at the level of human understanding, which also expands to computational processing. From a computational perspective in particular, the existence of a precise definition is crucial to determine the detection procedure. As Rus *et al.* (2011: 293) argue, a 'more precise operational definition of a paraphrase' is required to competently assess the 'semantic similarity of two texts'.

Paraphrasing has often been considered oversimplistically as the substitution of words with semantic equivalents, to an extent so similar to synonymy that the two are frequently confused. This confusion may be partly due to the fact that traditional approaches to semantics have focused on the meaning of words or morphemes (Oliveira, 2005: 334), to the detriment of other meaning relations, such as at the level of clauses, sentences or even paragraphs. As a consequence, the semantic perspective of paraphrasing has been based on two commonly held assumptions. The first is that, although both paraphrasing and synonymy are part of the object of study of semantics, synonymy operates only at the word level, whereas paraphrasing can also involve linguistic elements such as sentences. The second is that a paraphrase in the strict sense of two propositions (clauses, sentences, etc.) having the same meaning is only successfully achieved if those two propositions represent a case of strict mutual implication; therefore, if the argument of mutual implication does not work, we cannot talk about two propositions having *exactly* the same meaning (Oliveira, 2005).

Although the first assumption is convincing, on the contrary the extent to which two propositions must have exactly the same meaning to be considered a paraphrase can be challenged. For example, in a passive/active scenario, the passive text may not have *exactly* the same mutual implications, strictly speaking, as the active, as is demonstrated e.g. by CDA (Fairclough, 1995). Yet, passivisation is one of several types of lexical and syntactic reformulations that is frequently used legitimately in academic writing to paraphrase someone else's work, or illegitimately to obfuscate the original. In both cases, the outcome does not respect the meaning *exactness* of the two propositions. When operated legitimately, it reflects the meaning of another string of text, even by approximation, and at most it can be considered a 'good' or a 'not-so-good' paraphrase. When used illegitimately, however, the lack of exactness is not sufficient to disregard accusations of plagiarism.

More accurately, therefore, paraphrasing consists of a semantic relation between or among strings of text that, albeit different in form, have similar or even identical

meaning(s). When applied to academic writing, paraphrasing involves restating the thoughts or ideas of another author in one's own words, while using words or a sentence structure that are different from the original and acknowledging the original authorship by means of in-text citation (Pears and Shields, 2008). In other words, it consists of restating in one's 'own words a specific point, finding or argument an author has made' (Williams and Carroll, 2009: 27), 'usually to achieve greater clarity' (Pears and Shields, 2008: 7).

An investigation into paraphrasing therefore first needs to account for how lexical meanings can be combined into an infinite number of sentential or phrasal meanings. But since meaning is context-dependent, lexical meaning cannot, in most cases, be studied in isolation, and reference to the discourse is necessary. The example in table 7.4, which is not a case of paraphrasing in a strict sense, is a good example of referent substitution.

Original	Plagiarism
Pues bien, hace años, debajo de mi casa, había en una pared una frase en spray que decía "Puta Burgos" y eso no era un graffiti, era una pintada.	Pois bem, se numa parede há uma frase em spray que diz "Puta Espanha", isso não é um graffiti, é uma pintada.

Table 7.4: Example of change of extralinguistic referent in real plagiarism cases.

The Spanish city of 'Burgos' in the original is replaced with the country, 'Spain', to adapt the derivative text to the context of a group of Spanish students that are not from the city of Burgos, and thereby avoid a mismatch between the text and the extratextual reality that violates the coherence relation. It is relevant to consider whether alterations of this type, where a city is replaced with a country, retain the exactness of the original proposition, and are, in this sense, a paraphrase; or whether those changes alter the meaning, and consequently the proposition is no longer a paraphrase at all. This example makes a good case for the latter, but computational approaches that can detect these instances would aid in detecting paraphrasing.

Paraphrasing is therefore perhaps more accurately described as the the semantic relationship that operates between propositions, rather than a relationship between

sentences, clauses, etc. Flowerdew (1992) proposed a simple taxonomy, claiming that paraphrases could be sentential (taking the format of a sentence), subordinated (in the form of a subordinate clause or clauses) or supra-sentencial (above the sentence level). However, consensus on a taxonomy of plagiarism has not yet been reached. Approaching plagiarism from a computational linguistics perspective, Rus *et al.* (2011) included in their study six types of paraphrase: *free paraphrase*, *changed words*, *changed structure*, *changed words and structure*, *condensed*, and *improved*.

From an academic writing perspective, Williams and Carroll (2009: 35) suggested that paraphrases can be generated following three basic steps: (1) identifying and grasping the key point of the original; (2) restating the idea(s) of the original in one's own words; and (3) producing a version that is shorter than the original. In theory, a procedure that is logical enough to teach humans how to paraphrase can also be potentially useful to support machine learning. This raises an additional complication: the distinction between *paraphrasing* and *summarisation*.

The Online Writing Lab at Purdue University admits that the distinction between paraphrasing and summarisation is one of detail; '[u]nlike a brief summary, a paraphrase contains more detail'[9]. Pears and Shields (2008) concur with this distinction, by explaining that the two differ mostly on the fact that, whereas paraphrasing consists of retaining the same meaning as the original, summarisation consists of listing the main points, topics or headings, while leaving out any details. In academic writing, this implies that the two are used on a different scale. As Roig (2001) asserts, summarisation is more widely used than paraphrasing, especially when reporting someone else's work, because it allows a writer to briefly explain the main arguments of another author.

On the computational side, approaches to paraphrasing have until recently focused mostly on paraphrase generation, despite the relevance of paraphrase identification for applications like plagiarism detection. Research into computational processing of paraphrasing (Madnani and Dorr, 2010; Rus *et al.*, 2011; Vila *et al.*, 2011) has demonstrated, furthermore, that methods of paraphrase generation and identification investigated so far do not account for its breadth and depth as a complex linguistic phenomenon. Madnani and Dorr (2010) concluded that, although the generation and extraction of semantic equivalents are important tasks, the results are still unsatisfactory. Earlier, Callison-Burch and Osborne (2006) had demonstrated that the use of

[9]Types of Paraphrases – eHow.com – `http://www.ehow.com/`.

paraphrasing in statistical machine translation (SMT) can contribute to improving the coverage and the quality of machine translation systems.

Wubben *et al.* (2010), for example, approached paraphrasing as a process that involved inserting, deleting and substituting words. This is a process that some plagiarism detection software[10], perhaps more accurately, describe as 'linguistic alterations'. Substituting or reordering vocabulary or parts-of-speech is required, not the least to adapt the 'neighbouring' text accordingly, in particular to respond to issues of collocation and 'lexical priming' (Hoey, 2005). It therefore involves changes in vocabulary more prominently, but also reflects alterations in (text) length, parts-of-speech and sentence structure.

Paraphrase detection admittedly has some limitations. Rus *et al.* (2011), despite reporting good results, admitted that their method needed improvements to enhance its performance. In particular, they found that the system should allow each paraphrase to be classified in more than one class, in order to accommodate cases where the paraphrase can fit into more than one class at the same time. But other reasons can account for why paraphrasing seems too complex for computational systems. Firstly, those semantic relations can be of different linguistic nature, including phrases, clauses, sentences or groups of sentences, and may vary in length. Secondly, paraphrasing consists of establishing a relationship between *meanings* of words, and not between the words themselves. Thirdly, the 'mathematical' operation of replacing words in the paradigmatic axis with semantic equivalents in the syntagmatic axis is not always realistic.

Therefore, it is argued that paraphrasing can only be realised using strings of text with identical informational contents, considering that many word pairs share some, but not all meanings. Conversely, instances where only a word or a few words are replaced with semantically-related words are labelled alterations.

7.3 Improving the Detection Procedure

The previous sections have identified the most common strategies used by plagiarists to borrow inappropriately from other sources, and discussed some of the solutions proposed by the literature. This section discusses some of the approaches that can be adopted, using computational linguistics tools, to improve the detection of those

[10] For example, Urkund

instances.

Verbatim plagiarism can be detected quite well using fingerprinting techniques. These techniques allow the comparison of texts, both the source and the suspect, to be indexed for comparison. The strings of identical text are thus identified. Although this procedure works relatively well when comparing texts, the method is not effective to prove instances of plagiarism where no texts for comparison are available. Stylistic analyses have been proposed to address instances of this type. Yet, as was argued, it is very likely that the style is consistent throughout in instances where complete texts are borrowed verbatim from another source. If the hypothesis that students plagiarise mostly from the Internet is taken to be true, the method that is more appropriate to detect instances of this type is the 'open system' comparison. Systems that are able to demonstrate the borrowing in these instances are crucial to investigate also the plagiarist's intention. Specifically, if it is considered that the severity of intention of the plagiarist increases proportionally to the effort made to obfuscate the authorship of the text, then cases of verbatim, unacknowledged text reuse – which are easy to detect – are less severe than cases of obfuscated plagiarism.

On the other hand, even if stylistic inconsistencies are found, especially in academic contexts, this can be due to the inability to competently write in academic contexts more than to plagiarism. A student's failed attempt to integrate the voice of others in their own writing may result in stylistic inconsistencies that do not necessarily account for plagiarism. In this case, a comparison against source texts is necessary to assess the borrowing. If that comparison is not possible, an assessment of the instance of plagiarism that is based on stylistic inconsistencies should be accompanied by an analysis of the plagiarists' intention. Consequently, if the plagiarist omits references to the sources altogether in the derivative texts, then the borrowing is obviously intentional. On the contrary, an improper reuse of the source text that is accompanied by a reference to the sources, indicates a case of unintentional, or at most inadvertent plagiarism. The perceptions of the survey participants in chapter 3 confirm this assumption.

Quasi-verbatim instances of plagiarism impose different challenges on the detection procedure. This involves making morphological, lexical, and syntactical alterations to the text, which may reflect at the level of word endings, replacement/addition/deletion from the sources, or alteration of sentence structures. Each of these strategies can involve different detection procedures. Morphological alterations, which usually consist

of changing some elements of the text, while retaining others that are more core to the text, can be detected by using tools such as lemmatisers and stemmers. These tools operate by removing the word endings in the source and suspect text for comparison of the roots. This is especially useful, for example, in cases of nominalisations.

Lexical alterations, on the other hand, can be detected via searches using fuzzy matches and wild cards. In this case, the suspect text can be searched for plagiarism against other sources to identify textual alterations that allow one to draw conclusions on whether a text has been plagiarised. Text edits like adding, deleting or replacing words may indicate a failed attempt to write academically or an intention to deceive. The latter is illegitimate; the former, however, can represent a legitimate attempt at writing. The use of computational tools that help to accurately identify the changes introduced to the text, together with the attempts made by the plagiarist to attribute or otherwise obfuscate the original, can help in making decisions as to whether the plagiarist acted inadvertently, or attempted to deceive his/her readers.

Thirdly, syntactic alterations can be detected via a lexical or a syntactic analysis of the texts. The type of lexical analysis like the one described by Johnson (1997) demonstrated that this method is very robust, even in instances where sentences are completely rewritten, while the lexical items are maintained. Since lexical items are those that can be used more prominently to express unique text and ideas, lexical overlap is perhaps the most appropriate detection method. In some cases, however, the text can be altered in a way that, although the sentence structure is retained, the amount of lexical overlap is not sufficient to determine that a text has been plagiarised. In this case, a syntactic analyser can be used effectively to annotate texts automatically and correctly. Parse trees can then be compared, in order to determine whether two sentences are similar or identical, or not at all related. For example, the syntactic analyser 'PALAVRAS'[11], part of the VISL project, is able to annotate Portuguese sentences very accurately. Most systems also have the advantage that they can analyse large quantities of text in real time.

The use of parse trees for plagiarism detection imposes, however, some limitations. Firstly, the analyser can only return accurate results when the sentences are syntactically 'correct'. If the standard syntax is not observed, then the analysis tends to break down. In academic writing, where sentences tend to follow a standard structure, this does not usually represent a serious problem. The second issue is more

[11] http://beta.visl.sdu.dk/visl/pt/parsing/automatic/

problematic. Since academic writing usually follows standard patterns, it is very likely that parse trees are not sufficiently discriminative.

In severe cases of intention, where the plagiarists chose to omit a reference to the sources, these tools have the potential to identify the source texts and demonstrate the borrowing. It is very likely that, in extreme cases where standard syntax is to be expected, the syntactic analysis may not be sufficient to demonstrate an instance of plagiarism. In these cases, a closer analysis is required to demonstrate or discard the possibility that the plagiarist may have intentionally plagiarised.

Paraphrasing is the strategy that challenges plagiarism detection systems the most. If the semantic relationships holding between propositions in a paraphrase are realised by different syntactic devices, expanding beyond the boundaries of a fixed, predetermined, well-defined syntax of a text that is divided into sentences or paragraphs, then the first challenge that has to be addressed by computational approaches is determining the units for analysis. Therefore, as discussed in section 7.2.2, a system that separates the strings into sentences or paragraphs for comparison and analysis of vocabulary overlap, as well as semantic overlap, can be illusory. Besides, it raises detection issues in instances where the order of the propositions is shifted, and especially when the relationship between propositions is not of the type one-to-one, where one proposition in one original can be semantically equivalent to several propositions in a derivative text, or vice-versa. The first main challenge to the detection procedure is therefore understanding how paraphrasing is operated, and then determining a classification procedure that approaches the mapping of those operations onto computational systems.

Approaches to paraphrasing that treat it as a classification problem have the potential to increase the system performance. But they have the disadvantage of risking missing important linguistic information. Vila *et al.* (2011) propose a typology of paraphrasing that attempts to account for the linguistic nature of the phenomena involved in the rewording, with expectations that methods of annotation based on such typology may assist plagiarism detection. The method consists of grouping the paraphrases into a two-level typology of nine types of paraphrasing, according to five classes. The rationale behind this typology is that morphological, syntactic and discursive changes retain the *same propositional content*, whereas lexical and semantic changes are *equivalent in their propositional content*. The authors report good results. The main advantage of this application is that it represents a considerable advance in the field of

the natural processing of paraphrases, especially by allowing the machine to learn how to classify new examples. But it also brings some disadvantages. Firstly, a typology of this nature is complex, time-consuming and difficult to implement, besides requiring considerable human and technical resources. Secondly, although this typology is comprehensive, the two-level approach does not cover all possibilities. For example, *deletion* is considered only in the group of *lexicon based changes*, which, for instance, does not take into account cases where entire phrases may be deleted from another original.

These issues could be resolved by basing the system of annotation on (at least) three different axes: *type*, *level* and *strategy*. This taxonomy is based on the instances found on the empirical analysis of the CorRUPT corpus texts, as well as examples from the literature (including those suggested by Vila *et al.* (2011)).

The *type of paraphrase* consists of describing the semantic relation(s) established between strings of text with related meanings, e.g. synonyms, antonyms, superordinates, co-hyponyms, hypernyms, instance of, generic, etc. The second axis, *level of paraphrase*, aims to account for the level at which the relationship is established, e.g. word, phrase, clause, sentence, or above the sentence level (supra-sentential). Thirdly, the *strategy* aims to account for the operation performed to create the paraphrase, e.g. substitution, reordering, addition/expansion, deletion/reduction. This system of classification represents an annotation procedure based on a method of triangulation, whose analysis includes aspects that account for the interaction among three different classes, instead of classifying a paraphrase first according to a class, and then as a sub-class of that class.

This procedure is complex, and requires significant resources, especially in the initial annotation phase, but the subsequent gains can be significant. In particular, this classification can start with manual annotation of cases of paraphrasing, which are then used to teach the system how to annotate subsequent cases, using a method of classification such as Support Vector Machines (SVM) (Joachims, 1998). In a second phase, the system is asked to annotate instances of paraphrasing automatically, and this annotation is verified by human annotators. Once the system demonstrates a considerable level of confidence, by annotating correctly the instances provided, it will be used for automatic annotation. By gradually annotating the examples of paraphrasing automatically, the system will eventually be able to provide a correct annotation, without the fatigue inherent to human annotation.

In order to be able to cope with real-life scenarios, the training procedure should ideally be performed using real and, if possible, forensic examples of plagiarised paraphrases. However, at a training stage, this is not absolutely necessary if we consider that, as discussed in chapter 2, real cases demonstrate *what really happens* in instances of plagiarism, but not *what can happen*, i.e. they do not necessarily account for all possibilities of plagiarism. In this sense, other plagiarism corpora, such as a corpus of news feeds, can potentially offer a wider range of possibilities to train the system.

It is therefore hypothesised that a system of classification that considers information from these three classes, rather than mutually-exclusive classes, has the potential to provide better results than one class individually. Additionally, this procedure helps resolve issues arising from the selection of only one class, where several different elements coexist. However, the development and implementation of this work computationally is far too extensive for more than a mere proposal in this study.

Paraphrasing remains one of the most challenging tasks of plagiarism detection. This is due, particularly, to the fact that, rather than consisting of alterations at the level of morphology, syntax or vocabulary, it can consist of a combination of several of these. Detection, in these cases, can be especially difficult owing to the complexity of the task and to the resources involved.

It also represents a challenge to detecting intention because, like other alterations, it can result from a failed attempt at paraphrasing someone else's words, works or ideas. Judgements of intention, in this case, can perhaps be more accurately passed by ascertaining whether the plagiarist provided or otherwise omitted the reference to the sources. The former suggests an unintentional case of plagiarism; the latter suggests an intention to deceive.

Another challenge to plagiarism detection, as discussed in chapter 5, is translingual plagiarism. However, it was demonstrated that this type of plagiarism can be detected by using a simple MT approach, followed by a comparison against the Internet. This procedure discards the need for comparison against known sources. As was strongly argued, MT can have some limitations at the syntactic level, but has the capacity to handle lexical items very well. These are the ones that can be more discriminating in detecting translingual plagiarism. Although a comparison cannot be fully established with non-indexed or access-restricted sources, search engines allow searches of the first few words. This permits the borrowing to be investigated further.

Translingual plagiarism is perhaps the one that raises most concerns when determining the authors' intention to deceive. As in other cases, acknowledging the sources is sufficient to determine that the borrowing is, at most, inadvertent. However, omitting the sources, albeit not representing the appropriation of someone else's words, equates with the intentional appropriation of someone else's works or ideas. This therefore represents intentional plagiarism.

The use of computational tools can therefore contribute to assess instances of plagiarism that can represent a violation of the rule, a violation of the principle, or both. Especially in cases involving large volumes of data, these tools can highlight instances of borrowing that could otherwise be missed.

7.4 Chapter Summary

This chapter discussed a set of methods and tools that may contribute to improving existing plagiarism detection methods, as well as the limitations imposed on the different types of plagiarism detection.

The chapter started by identifying some of the main linguistic issues that affect current detection tools. Then it discussed some issues underlying computational processing, including processing power, pre-processing and the comparison procedure.

The findings indicated that fingerprinting techniques have a good performance in detecting verbatim plagiarism, whereas content comparison methods like parse trees perform well when the vocabulary is altered, but the syntax is retained. Alternatively, 'word sense disambiguation' methods can be used to make a comparison between an original and a suspect instance of plagiarism.

It was argued that, on balance of probabilities, 'fuzzy fingerprinting' represents a good trade-off in terms of performance, unlike lexis-based detection systems, which are significantly affected by lexical changes. The use of stemmers and, particularly, lemmatisers was also demonstrated as having the potential to offer good results, especially with languages that are highly inflectional or derivational, or in cases of plagiarism involving different language variants.

This chapter then discussed possible improvements to paraphrase detection systems. A taxonomy was proposed to annotate paraphrases on three different levels simultaneously: *level*, *type* and *strategy*. It was hypothesised that this method should improve the machine learning procedure, and consequently the paraphrase identification task.

It was argued that the software tools can assist the linguist in making more informed decisions on the degree of intention of the plagiarist, by determining whether each instance represents a violation of the rule or a violation of the principle.

The chapter concluded by strongly arguing that currently there is not one single method of machine-assisted plagiarism detection that is able to perform better than all other methods in all possible situations. Ideally, systems should combine all possible methods.

Chapter 8

Conclusions & Implications

8.1 Introduction

This study presents the findings of the research project on plagiarism detection that was conducted over the last four years. This chapter concludes, starting with a brief description of the contribution to the literature, and some reflections on the methodology adopted. A summary of the findings then follows, that describes the results obtained in this research, in relation to the research questions described in chapter 1. This section is followed by an evaluation of the findings, that identifies the main results, as well as the shortcomings of this research. In particular, the main limitations are identified and an explanation is provided for them. A section follows that describes the main implications of this research, and their relevance. The chapter concludes with some final remarks, and an indication of future research work on this area.

8.2 Contributions to the Literature

This study covered an extensive body of research in the field of plagiarism in general, as well as plagiarism detection in particular, across different disciplines, including linguistics, authorship, learning and teaching, law and computational linguistics. However, until now research has tended to focus on one single area, maybe with the exception of forensic linguistics and computational linguistics; the former by applying linguistics research in forensic contexts, and the latter by combining the work of linguistics and computer sciences. Research on computational linguistics has, however, been conducted mostly by computer scientists, based on their knowledge of how language works, with little contribution from linguists. Maybe as a result of this approach,

the results are highly computational, and minimally linguistic, although the joint work of linguists and computer scientists has been shown to provide good results (e.g. the work described by Woolls (2010)). However, in plagiarism detection, for example, studies abound of cases where the detection procedure is based on string matching algorithms that consider no linguistic information at all.

On the other hand, forensic linguistics research into plagiarism detection, albeit taking into account legal aspects, has focused on the linguistic description of the instances of plagiarism, applied to legal contexts. Consequently, considerations of other aspects tend to be left outside the work of the forensic linguist. The evaluation of the plagiarist's competence to write academically is left to lecturers/tutors, assessments of perceptions of plagiarism tend to be dismissed, and labelled as a task of sociologists, judgements of whether the plagiarist has intended to plagiarise, or whether s/he acted inadvertently, tend to be left to psychologists, and computational detection is left in the hands of computer scientists, who eventually overwhelm linguists with their ability to handle large volumes of data.

This thesis has made some original contributions to deepening the research on plagiarism and its detection, by approaching these areas, while retaining a focus on forensic linguistics. Firstly, it discussed and developed the cross-cultural contributions to research into plagiarism. The analysis of the findings of the survey conducted in the UK and in Portugal allowed for a comparison of perceptions of plagiarism in the two countries, and concluded that those perceptions can be strikingly different in some cases, even in two countries that are geographically close. Studies of this type, which contribute to deepening the understanding of plagiarism across different cultures and countries, are scarce. At most, until now studies have tended to investigate academic ethics in general, and not plagiarism in particular; or they have focused more in one country than others. None has however been conducted that: (a) focused specifically on plagiarism; (b) dedicated the same attention to two countries, rather than focusing on one country to the detriment of others; (c) established a direct comparison between two countries; and (d) compared these two countries intrinsically. Additionally, a study of this type involving Portugal is novel.

The findings from this study, together with previous findings from the literature on pedagogical approaches to plagiarism, as well as with the findings from the analysis of institutional approaches to plagiarism, contributed to understanding plagiarism cross-culturally, in particular why plagiarism may represent a problem of academic writing,

or more liberally an academic writing strategy, whether universities tend to consider plagiarism uniformly, or whether these approaches tend to vary depending on the corresponding institution, and why approaching plagiarism detection from a purely punitive perspective might miss the objectives of the detection procedure.

Secondly, it discussed the role of intention in relation to plagiarism, paying particular attention to the legal context. The findings from the literature on law, ethics, morals and philosophy, as well as the findings from the survey, indicated that a distinction needs to be made between judgements of intentional and unintentional instances of plagiarism, that this degree of intentionality needs to be considered in a continuum, rather than dichotomically, to include more degrees than a simple binary classification, and that the penalties should vary correspondingly. I found that, despite the tendency of linguists to avoid interfering with judgements of intention, this is an issue that is not irrelevant in legal contexts. Using illustrative examples from the CorRUPT corpus (my own Corpus of Reused and Plagiarised Texts), I proposed a taxonomy of different degrees of intention, and argued that different degrees of intention can be determined by examining the linguistic strategies used to plagiarise.

Thirdly, I proposed a method to detect instances of plagiarism that consist of translating texts from foreign languages. Plagiarism detection had, until now, concentrated mostly on the linguistic analysis of texts written in the same language. The examples from the CorRUPT corpus demonstrated, however, that using translation to plagiarise (i.e. translingual plagiarism) is a relevant issue. In this thesis, I proposed an empirically-oriented and theoretically-grounded method to detect translingual plagiarism, and demonstrated that this method has the potential to provide effective results across different languages.

Additionally, this thesis made a contribution to the development of plagiarism detection software. The evaluation of detection software has been a common procedure in recent years, but it has tended to focus on the assessment of the performance of the detection systems from a computational perspective; i.e., detection systems have been assessed based on their performance, precision and recall, rather than concentrating on linguistic aspects. I examined current detection software to identify their strengths and weaknesses, using a set of linguistic parameters and identified several points for detection software improvement, specifying for each linguistic strategy the corresponding necessary improvements.

Finally, I approached plagiarism detection in texts written in Portuguese. Plagia-

rism detection in Portuguese has until now attracted little research attention, if any, although Portuguese is a widely-spoken language. This study has focused mostly on plagiarism detection in Portuguese, and examples from the CorRUPT corpus were used throughout the thesis to illustrate my arguments. I contended, with reference to the illustrative examples from the CorRUPT corpus, that detection systems need to consider particular aspects to be able to effectively detect plagiarism in texts written in Portuguese.

Overall, these original contributions demonstrated the relevance of research on plagiarism detection, and moreover that approaches to plagiarism detection need to consider the perspectives of several different disciplines, all of which are terms in the plagiarism equation, and consequently ideally involved in the discussion and application of the detection process. I argued that plagiarism detection, especially when approached from a forensic linguistic perspective, involves a range of internal and external circumstances, that may increase the severity of the plagiarism instance, or on the contrary contribute extenuating circumstances to it, and consequently needs to account for more than a binary, 'yes'–'no' classification.

8.3 Reflections on the Methodology

The research questions and hypothesis studied in this thesis were addressed based on mainly two types of data: (a) the survey data; and (b) the textual data. On the one hand, an online survey was conducted, so as to understand the perceptions that students and lecturers/tutors have of plagiarism, as well as the circumstances involving the elements that enter the plagiarism equation. On the other hand, in order to identify the linguistic strategies used by university students to plagiarise, especially in Portugal, a linguistic analysis was conducted of a small number of real plagiarism cases.

8.3.1 Reflections on the Survey

The aim of the survey was to attempt to understand the perceptions that students and lecturers/tutors, of two different countries – Portugal and the UK – have of plagiarism. In order to approach a topic that is not uncommonly considered to be highly sensitive, interviewing methodologies were initially discarded. Also, in order to encourage the discussion of different issues related to plagiarism, an open question survey was subsequently rejected. It was found that the best approach to this consisted of a survey

based on vignettes. These had the advantage that particular aspects could be introduced naturally, as fictional, but realistic stories, and hence elicit the opinion of the participants.

Although the number of responses to the final survey was higher than anticipated, the pilot test took longer than initially expected, owing to the fact that the number of vignettes was very high, consequently taking 1 to 2 hours to complete. To avoid this problem, and the delays that it caused, the total number of vignettes should have been lower, or alternatively the survey should have been advertised more widely, across other institutions, especially in the UK.

In the final survey a field was included after each scenario so as to allow participants to justify their choice, if they wished. Providing this type of feedback was optional; yet, the volume of feedback provided by the participants was extremely high, and the responses are engaged. This suggests that the aim to have the participants engaged in the survey was also achieved.

A quantitative analysis of the responses was conducted, later followed by a qualitative analysis of the participants' feedback, in order to be able to contextualise those responses.

8.3.2 Reflections on the Linguistic Analysis

A study of plagiarism detection should necessarily include a linguistic analysis of real cases of plagiarism. However, owing to confidentiality and anonymity issues, getting access to data of this nature (usually very restricted), and obtaining the corresponding permission to use them for research purposes, is often a problem.

The volume of data is not, however, the most important element in an analysis of this type, where the quality of the data is more important than the quantity. Therefore, despite the apparently small number of texts containing plagiarism, these include a diverse range of linguistic strategies, whose analysis contributed to demonstrating that linguistic data can be used to investigate the degree of intention behind the instances of plagiarism; that these data allow for the detection of translingual plagiarism; and furthermore that their analysis can contribute to improving the detection systems.

In the future, however, other cases of plagiarism to which access is granted will be added to the CorRUPT corpus.

8.4 Summary of the Findings

At the beginning of this study (section 1.1.2), five research questions were posed. This section presents a summary of the answers revealed by the data obtained during this research project.

8.4.1 Identifying the Strategies Used by University Students to Plagiarise

The first research question was concerned with how university students use information available (especially online) to plagiarise in their academic work. The data collected indicated that one of the most typical forms of plagiarism, which consists of borrowing literally, word-for-word from other texts, by simply copying and pasting the text, is still used by academic students to plagiarise. However, use of this strategy shows a tendency to decrease in parallel with the awareness that, not only does this represent the most obvious form of plagiarism, but also it is a strategy that is now, in an era of mass information, easy to detect using any simple string-matching technique (from simple plagiarism detection tools, to even the document comparison function of *Microsoft Word*).

Conversely, the data demonstrated that other more sophisticated strategies are increasingly used in the academy to plagiarise. One of these strategies consists of reordering the textual elements, one obvious example of which is passivisation. However, text reordering can also involve other linguistic strategies as well, including reordering of clauses, which consequently involves adjustments in grammar that create 'noise' in the text, thus making the detection procedure more difficult. In parallel, word substitution is also operated, which involves both replacing one or two lexical items with synonyms, but also cases where verbs, for instance, can be replaced with prepositions. In both cases, string matching techniques, that consist of finding levels of identity or similarity between two or more texts, tend to break down owing to the fact that the word sequences of the original texts are interrupted by new textual elements. In these cases, an analysis of lexical overlap has also demonstrated good results in identifying and demonstrating the borrowing.

Although paraphrasing is both a concern and a challenge to plagiarism detection, the cases of paraphrasing in this data set are very scarce and, when existing, they tend to be more consistent with the classification of word substitution, than actually para-

phrasing. One possible explanation for this is that true, proper paraphrasing involves a complex and time-consuming process (often more complex than creating an original) of text interpretation and re-writing that are incompatible with the most typical cases of plagiarism, where a text is borrowed as a result of laziness or lack of time. Use of paraphrasing to plagiarise can become a real problem in the future, if paraphrase generation tools come to be sophisticated enough to allow plagiarists to paraphrase automatically. However, even in this case the plagiarism detection systems should develop at least at the same pace.

Surprisingly, the data revealed that translation is increasingly used to plagiarise. Given the recent developments in automatic translation technology, anyone can nowadays translate any text immediately online and for free, and subsequently copy and paste the result of the translation to any assignment without acknowledgment. The analysis of the data revealed that machine translation systems nowadays perform quite well, especially in some language pairs, and in particular those involving English. This is not surprising considering that the most popular machine translation systems use 'crowdsourcing' as a method of machine-learning, and since English is the Internet's 'lingua franca', other languages can vary across different documents, but English tends to remain a constant. The method proposed consisted of using freely available machine-translation tools, as these are the ones that are at the disposal of the common student user to translate the suspect text to the language of the suspect original. Unsurprisingly, machine translation shows a good performance when dealing with lexical vocabulary, but performs less well in processing grammar and syntax. These are therefore the linguistic elements that plagiarists are compelled to correct when they borrow from other languages, not only to make the text readable, but also to attempt to make it read it as a 'natural language', as opposed to machine-generated text. Once more, the analysis of lexical overlap demonstrated excellent results in detecting this type of plagiarism. Considering that these are the textual elements that tend to be machine-translated more correctly across different languages, unlike grammatical and syntactic elements, an analysis of the overlap of those elements helps to detect, as well as prove or disprove the borrowing.

8.4.2 Improving the Computational Detection of the Different Instances of Plagiarism

The second research question was concerned with improving the computational plagiarism detection methods and tools, so as to increase the accuracy and the robustness of (semi)automatic detection of the different instances of plagiarism revealed by the analysis of the data, and consequently improve the detection results, not only in terms of the number of the instances detected, but also – and most importantly – of the quality of the detection procedure.

The review of some of the most popular plagiarism detection software packages, with reference to the textual data analysed, revealed that existing software performs relatively well in detecting some types of plagiarism, in particular textual identity and, with some limitations, textual similarity. Textual identity is the one that computational systems detect more easily, considering that this operates based on the principle of exact string match. On the contrary, textual similarity poses some problems to computational plagiarism detection. The former is a problem of determining the distance comparison. By specifying shorter distances for comparison, the system is able to detect string matching more easily, with the downside that it increases the chance of identifying false positives; i.e. the number of instances identified as plagiarism is higher, but a high percentage of these are false, owing to the fact that many of these words are grammatical and consequently not plagiarising. In the latter situation, selecting longer distances for comparison causes the system to search for identical or highly similar strings that are too long, thus increasing the likelihood that true positives are missed.

Consequently, it was concluded that 'fuzzy fingerprinting' represents a good trade-off in terms of performance, since it is robust in relation to changes that are missed by exact matches, such as small additions, deletions and reordering, and has the advantage, when compared to lexis-based detection systems, that they are not significantly affected by changes at the level of the lexical items. At the same time, systems of this type allow for searches of longer chunks of text.

The data revealed that these strategies are unable to detect most cases of non-literal plagiarism, and that the implementation of other existing computational linguistics tools could help improve the number of 'correct guesses' of automatic detection systems. In particular, using stemmers and lemmatisers to reduce words to their root version, thus resisting changes that are purely inflectional or derivational, can offer

promising results, especially with languages that are highly inflectional or derivational. Additionally, using simple methods like 'dictionaries' to standardise the spelling before establishing the comparison offers the ability to detect plagiarism across language variants.

Although no particular instances of paraphrasing were reported from the analysis of the data, this is a known strategy that can and sometimes, although not often, is used in cases of plagiarism. This study suggested that, in order to improve paraphrase detection systems, a taxonomy is required to annotate paraphrases, and, moreover, that this taxonomy should enable a non-mutually exclusive classification. The solution proposed in this classification to improve the machine-learning procedure, that will eventually permit the automatic paraphrase identification task, is based on an annotation of three different layers simultaneously: level, type and strategy. Conversely, it was concluded that co-reference identification, despite being one aspect of semantics that is clear to the human reader, is difficult to identify by computational systems, and therefore will remain a challenge to computational systems.

8.4.3 Identifying Clues to the Degree of Intentionality behind Instances of Plagiarism

The third research question was concerned with identifying the degrees of intentionality that underlie the instances of plagiarism. Initially, the survey data were analysed to determine whether the participants make any distinction between intentional and unintentional instances of plagiarism. The data revealed that the perception of the participants of whether all cases of textual borrowing should be considered plagiarism, or alternatively whether only those cases where the plagiarist intends to deceive, is not unanimous. The perceptions of the participants clearly reveal, however, common agreement that intentional and unintentional plagiarism involve different degrees of violation of the moral principle of not plagiarising, and should accordingly imply different severity measures.

Based on these perceptions, a taxonomy of degrees of intentionality was proposed that take into account the paradigms of criminal law. According to this taxonomy, the most serious degree of intention consists of intentional plagiarism, which reflects linguistically on the instances where the text is manipulated to disguise the sources and obtain undeserved credit.

On the other end of the scale is unintentional 'plagiarism', which it was argued

should not be considered plagiarism. Examples of unintentional plagiarism include 'patchwriting' and borrowing that results from poor academic writing skills, which are a problem of the students' education and training, rather than an 'intention to deceive'. Linguistically, this is usually reflected in the poor quality of the writing, which is accompanied by a reference to the sources – even in cases where the referencing is not correctly done.

In the middle of the scale, in between the two extremes, is oblique intentional plagiarism, which includes cases of knowingly borrowing without acknowledgement, and hence culpable recklessness. This includes instances where the plagiarists intend the effects of the instance of plagiarism, but only indirectly: they foresee the effects of their actions but do not impede their actions; in other words, despite not agreeing with the means altogether, they desire the outcome. The textual data analysed suggest that most of the instances of plagiarism would fit in this class.

Admittedly, one of the main difficulties consists of determining the plagiarism cases deriving from an attempt to be creative, and where the effort to edit the original can result either from the intention to deceive, by passing the text off as his/her own, or from a legitimate attempt to write academically, hence oscillating between that which is 'highly probable', and the 'virtually certain'. In these cases, determining whether references are provided in any form, or on the contrary omitted altogether can function as a clue to the Forensic Linguist, despite not one of absolute certainty.

It was concluded that a forensic linguistic analysis of improper textual borrowing has the potential to identify and provide clues or, in extreme cases, evidence of intention in plagiarism. These considerations can lead to a fairer judgement of cases of plagiarism and, on the other hand, admit the existence of extenuating circumstances that may exempt a suspect plagiarist from more severe accusations of plagiarism.

8.4.4 Assessing the Investigative and Evidential Value of the Linguistic Analysis in Cases of Plagiarism in Forensic Contexts

The last, interrelated research questions, were concerned, on the one hand, with the contribution of the linguistic analysis as an investigative tool to a more reliable and accurate plagiarism detection procedure, while, on the other hand, considering whether these findings can also be used as evidence of plagiarism.

In forensic contexts, both the investigative and evidential nature of linguistic analyses can make a significant contribution. This is particularly true in cases of plagiarism,

where this contribution, based on the identification of the linguistic strategies used by suspected plagiarists to borrow similar or identical text from external sources, can be used to improve the detection methods and procedures.

Linguistic analyses can, however, go beyond the investigative phase, to provide evidence that is able to prove or disprove the thesis that a certain text has been plagiarised. In this case, more than a simple comparison of the suspect and the original text(s) that presents a description of the linguistic operations performed is required. In addition to the description of the linguistic operations performed, an explanation of those operations that details the steps adopted is necessary. Furthermore, ideally those explanations need to be theoretically grounded, so as to justify why they represent instances of plagiarism.

8.5 Evaluation of the Findings

Invariably, a study of this nature presents both strengths and limitations. This section considers some of these, weighing the strengths against the limitations that are imposed mainly by technological advances.

A multidisciplinary approach such as the one conducted in this study contributes to this approach to plagiarism detection. By considering the complexity involving the issues of authorship, as well as the contribution of pedagogical approaches to learning and teaching in higher education, and setting these in the context in which instances of plagiarism usually occur – i.e. legal or legally-related contexts – rather than confining itself to a linguistic analysis that simply describes *what* happens in the text while dismissing the *how* and the *why*. This multidisciplinary approach represents one small step towards a critical forensic linguistics approach.

At the same time, this small step should contribute, even if modestly, to pointing towards new directions in computational linguistic approaches to plagiarism detection that, being more accurate, can assist the forensic linguist both in the investigative and the evidential process. Although an extraordinary amount of research on computational approaches to plagiarism detection has been conducted in recent years, these have concentrated, with minor differences, on the detection of identical or similar strings, which is only useful, when effective, to detect instances of literal plagiarism. On the contrary, computational plagiarism detection tools that rely on lexical analysis have been shown to provide better, more accurate and reliable results. A third alternative was proposed, which consists of using a combination of lexical methods and

fingerprinting techniques, so as to allow computational tools to, on the one hand, be able to detect instances of literal, word-for-word borrowing, while at the same time avoiding compromising lexically-grounded methods of detection.

A method was proposed also to detect plagiarism resulting from the translation of texts from other languages. To my knowledge, a solution to this type of plagiarism has never been proposed[1], and even views of this strategy as a problem are rare. The analyses of the data that were presented and discussed show promising results.

However, these findings have to be evaluated considering the following limitations. Firstly, although an effective method to approach 'translingual' plagiarism was proposed that demonstrated good results, this method has not yet been implemented computationally. This implementation can, however, be faced with the following problems. On the one hand, the machine translation system used in this study was *Google Translate*, not only because it is one of the most popular translation systems (and the one that students speculatively use when they need to translate a text), but also because it is one of the systems that has evolved the most in recent years. Initially, *Google Translate* allowed the system to be integrated in proprietary applications. Later, however, as a result of translation companies heavily using – and abusing – *Google*'s machine translation system, *Google* suspended the permission to integrate the system in other tools, so that nowadays it is only available via its web interface. Although this does not impact the manual use of the tool, it poses problems to the integration of the system into proprietary solutions. In the best scenario, the workaround to this consists of invoking *Google*'s translation system, retrieving the results and establishing the comparison.

This implementation is absolutely necessary if a system is required to systematically screen student assignments – and other text genres – to be screened for plagiarism. Another pragmatic issue involves determining the expected language of the original. It has been demonstrated, with reference to the theory of interlanguage, that trained humans are able to find clues in translated texts that suggest the language of the source language. Computers are not able to understand these clues, or at least not in the same terms. Specific software tools, nonetheless, are nowadays able to make intelligent guesses regarding the language of an original, and these could potentially be used to detect the language of the alleged original texts. However, some

[1] This topic was first approached in a presentation that I delivered, in co-authorship with Tim Grant and Belinda Maia, in 2009, at the IAFL9, the 9th Conference of the International Association of Forensic Linguists, in Amsterdam.

of these deployments and assumptions, albeit apparently simple theoretically, cannot be confirmed before practical implementation is made, and this can take several years to complete.

Secondly, new directions have been discussed to improve computational plagiarism detection systems. In particular, this improvement suggested the implementation of already existing computational linguistics tools to plagiarism detection systems. However, it is a known limitation of computational systems that, despite the constantly increasing processing power and storage capacities of computers, computers still have limitations regarding the implementation of several different modules in a single package. Although humans, especially when trained, can easily and almost instantly determine the type of plagiarism used, the strategies involved, and the best tools to demonstrate it, computers follow a pattern of logics that does not allow skipping unnecessary steps.

8.6 Implications of this Research

The first implication is that technological advances still determine, to a great extent, the research possibilities in this area. Although, in most cases, great improvements could be achieved simply from the implementation of existing natural language processing tools, in practice computers, unlike the human mind, follow a logic of 'one thing at a time', which not only impacts the computational processing capabilities, but also brings implications in terms of the processing time required to detect instances of plagiarism. This supports previous findings (Woolls, 2010) that computer systems can only approximately undertake tasks that humans can recognise and handle in simple terms. Additionally, computer systems are not yet able to detect paraphrasing automatically. Although they have an extraordinary capacity to handle large volumes of data considerably faster than humans, they are very limited in recognising semantic information and making correct guesses when the same meanings are expressed in different forms. More research is required on this area before computers can successfully detect paraphrasing automatically. Therefore, current plagiarism detection is determined by the technology available, without which large volumes of data cannot be processed, and only human intuition can be used. This explains, at least in part, the fact that until now translation as a plagiarism strategy has been hardly studied – the other possible main reason being that, since most research on plagiarism is conducted on English, and since most scientific production worldwide is in English, this

has not so far been considered a relevant area.

The second implication is that conceptions of plagiarism had not consistently considered intention as a component of plagiarism. Even forensic linguistics approaches have concentrated on the description of the linguistic events involved in the texts, leaving aside any concerns regarding the plagiarist's intentions. In the academic context, forensic linguistics has limited itself to identifying the instances of matching or overlapping text, leaving considerations of intentionality to university disciplinary boards. However, as the survey data demonstrated, all the groups of participants in the survey consider that plagiarism should be judged more severely when intentional, but not so strictly when unintentional. Accordingly, the textual data demonstrated that linguistic clues can be found in the text that suggest whether the suspect plagiarist acted intentionally, knowingly or unintentionally.

The third implication is that academic plagiarism cannot be assessed independently of academic honesty policies, a clear definition of plagiarism, cultural context and, more importantly, pedagogical aspects. The idea that definitions of plagiarism were culturally variable, so that Asian university students were in an unfavourable situation when forced to apply Western academic writing conventions, was endorsed by (Scollon, 1994, 1995). However, as the survey data demonstrated, plagiarism is not equally understood by everyone, even in two Western, not very geographically distant countries, which suggests that different perspectives of plagiarism are firstly a problem of pedagogy (Howard, 1995; Robillard, 2008).

The fourth implication is that not only is plagiarism metaphorically and metonymically associated with concepts of crime, it also *is* a legal matter. As is illustrated by the Portuguese copyright law, which like in many countries was influenced by the Berne Convention and subsequent international law (including the European Union Directive on Copyright), plagiarism can involve the violation of both moral and/or financial rights, both involving financial penalties and imprisonment up to three years. This demonstrates that plagiarism, albeit a moral issue, is also a legal problem.

The fifth implication, in connection with this, is that research on plagiarism analysis and detection that considers a multidisciplinary approach contributes to extending the scope of forensic linguistics beyond a descriptive linguistic analysis, to consider the legal, social and educational circumstances involving the textual borrowing. As a consequence, such an approach can not only contribute to plagiarism detection as an investigative, but also as an evidential tool.

Another implication is that a computational forensic linguistics approach to plagiarism detection has the ability to resolve many of the problematic issues that forensic linguists often need to resolve manually. A computational linguistics approach that is informed by the input of forensic linguistics can contribute to improving the detection methods, procedures and tools, by combining the strengths of linguistics, in particular the knowledge of how language works and the ability to explain why a certain instance is or is not unique, when compared to another one, with the strengths of computational linguistics. This supports the claims by Woolls (2010) that, despite their limitations, computers have the ability to handle large volumes of data, quickly, consistently and without the human mental stress.

The Plagiarism Handling Flowchart presented in figure 8.1 provides a succinct description of the workflow that can be used to handle instances of academic plagiarism.

The flowchart starts by making a distinction between plagiarism and other cases of academic dishonesty, which reflects the discussion presented in section . Plagiarism is then divided between *linguistic plagiarism* and *plagiarism of ideas*, each reflecting the respective linguistic strategies used for improper text reuse. A forensic linguistic analysis of the suspect instances of plagiarism is able to identify the existence – or otherwise the absence – of clues in the plagiarising text that may suggest that it was appropriated – and repurposed – from another source.

The existence of these clues suggests that the suspect plagiarist attempted to attribute the text, even if incorrectly, to the original source. The failure to meet the established rules and conventions of academic writing reflects, in this case, a student's failed attempt at writing academically. In other words, this represents a violation of the rules and indicates that the suspect plagiarist did not try to pass off the text as his/her own, and hence is very likely to be unintentional. These problems with academic writing can be overcome by retraining the student on academic writing. Only when the case is not resolved – e.g. the suspect plagiarist repeats the violation of the academic writing rules – should the case be submitted to the scrutiny of a disciplinary board.

Conversely, the absence of clues to the original indicates a possible attempt of the student at deceiving the reader, by making the text pass as his/her own. This, in itself, represents a violation of the principle of ownership of the original text (especially moral ownership), i.e. the suspect plagiarist knowingly and improperly lifted someone else's text and tried to pass it off as his/her own. This behaviour is, therefore, intentional, since the suspect not only knew what s/he was doing, but also consciously obfuscated

the original authorship. In this case, a forensic linguistic analysis can be submitted with the case to the scrutiny of a disciplinary board for the possible application of disciplinary action.

Although this flowchart describes a common procedure that can be adopted to handle instances of academic plagiarism, it is by no means the only possibility available. Rather on the contrary, it should be taken for what it is: a tentative first attempt at describing the plagiarism detection workflow, which can be improved and/or updated if and when necessary. Additionally, this flowchart can also be used with other types of academic dishonesty that are expressed linguistically, such as self-plagiarism or salami slicing. Finally, although the procedure described focused on academic plagiarism, it can also be used with other types of plagiarism and copyright infringement where linguistic analysis is able to establish that an author's moral and/or financial rights were violated.

8.7 Concluding Remarks

Plagiarism detection is arguably one of the most challenging tasks in forensic linguistics. On the one hand, even when observing strict confidentiality and anonymity procedures, access to real academic plagiarism data is often restricted; if access is restricted, permission is not given to use the data, even for research purposes. On the other hand, unlike other research areas, the identification of linguistic strategies for plagiarism detection cannot rely solely on the findings of particular sets of data, considering that there is always the chance that those sets are not comprehensive enough to cover all the possibilities. As Finegan argued very clearly, 'descriptive [linguistic] practices do not fully or adequately serve the entire gamut of linguistic needs, especially in terms of social and cultural demands' (Finegan, 2003: 222). If we consider the legal implications of plagiarism in most countries, it is crucial that it is a role for the forensic linguistics expert to frame a linguistics-based approach that takes into account the legal definition of plagiarism, and the implications of this definition. But more importantly, a critical approach that assesses the clarity of the applicable definition, and that admits the possibility that it might not be the most appropriate is vital to analyse instances of plagiarism, and subsequently apply the findings of that analysis accordingly. Pending, of course, the ethical considerations of the linguistics expert.

Moreover, with the technological advances and the evolution of detection techniques, new strategies tend to be invented, which lead to the need to develop new

detection methods. As a result of this dynamic, a constant update of the linguistic patterns and strategies used across different instances of plagiarism is required. On the computational detection side, there is not one single method of machine-assisted plagiarism detection that is able to perform better than all other methods in all possible situations; each method may perform better than another in detecting one particular strategy – the one which it was developed to address. In practice, however, any plagiarism strategy can be used, independently or in combination with other strategies, to produce a plagiarising version of a document that integrates multiple forms of plagiarism. Ideally, a functional system should combine as many methods as possible to improve the detection procedure. Additionally, computer systems will continue to face several of the current challenges for the coming years. One of them is paraphrase detection; the other is co-reference identification. This is crucial aspect in plagiarism detection, considering that, when misused, it easily allows a human reader to suspect that a text has been plagiarised, but not when properly used. In this case, it makes the detection task more difficult. However, owing to technological limitations, in particular the processing capacity of the systems, it is unlikely that the ideal system will be achieved over the next few years.

This is the first main aim of future work in this area: implement this research, by developing a computational system that is able to detect each plagiarism strategy individually. One of these modules needs to consider and be able to handle specifically paraphrase detection. The different individual modules subsequently will be combined into one same machine detection system.

The second aim is to take the contribution of computational forensic linguistics to plagiarism detection even further. Academic plagiarism detection has traditionally adopted a common procedure of simple identification of shared texts (increasingly using one of the many plagiarism detection software available), followed by a submission of cases of a substantial amount of textual overlap to the scrutiny of a disciplinary board. Forensic linguistics, as this work demonstrated, can make a special contribution to plagiarism detection. It has the potential to go further than the simple identification of the most evident instances of matching text, to investigate also obfuscated textual reuse, and the possible intention of the plagiarist behind this reuse, to establish cases of plagiarism of ideas linguistically, and to demonstrate and explain why a certain instance (even those apparently less evident) can be rather severe. I demonstrated that this task cannot be done independently of an analysis of cultural

assumptions. But, perhaps more importantly, this research indicated a new turn in forensic linguistics research: the turn from the purely legal to the ethical aspects of plagiarism detection. It is, therefore, the ultimate aim of this research to extend this work in other forensic contexts, starting by detecting plagiarism in this new forensic linguistics turn.

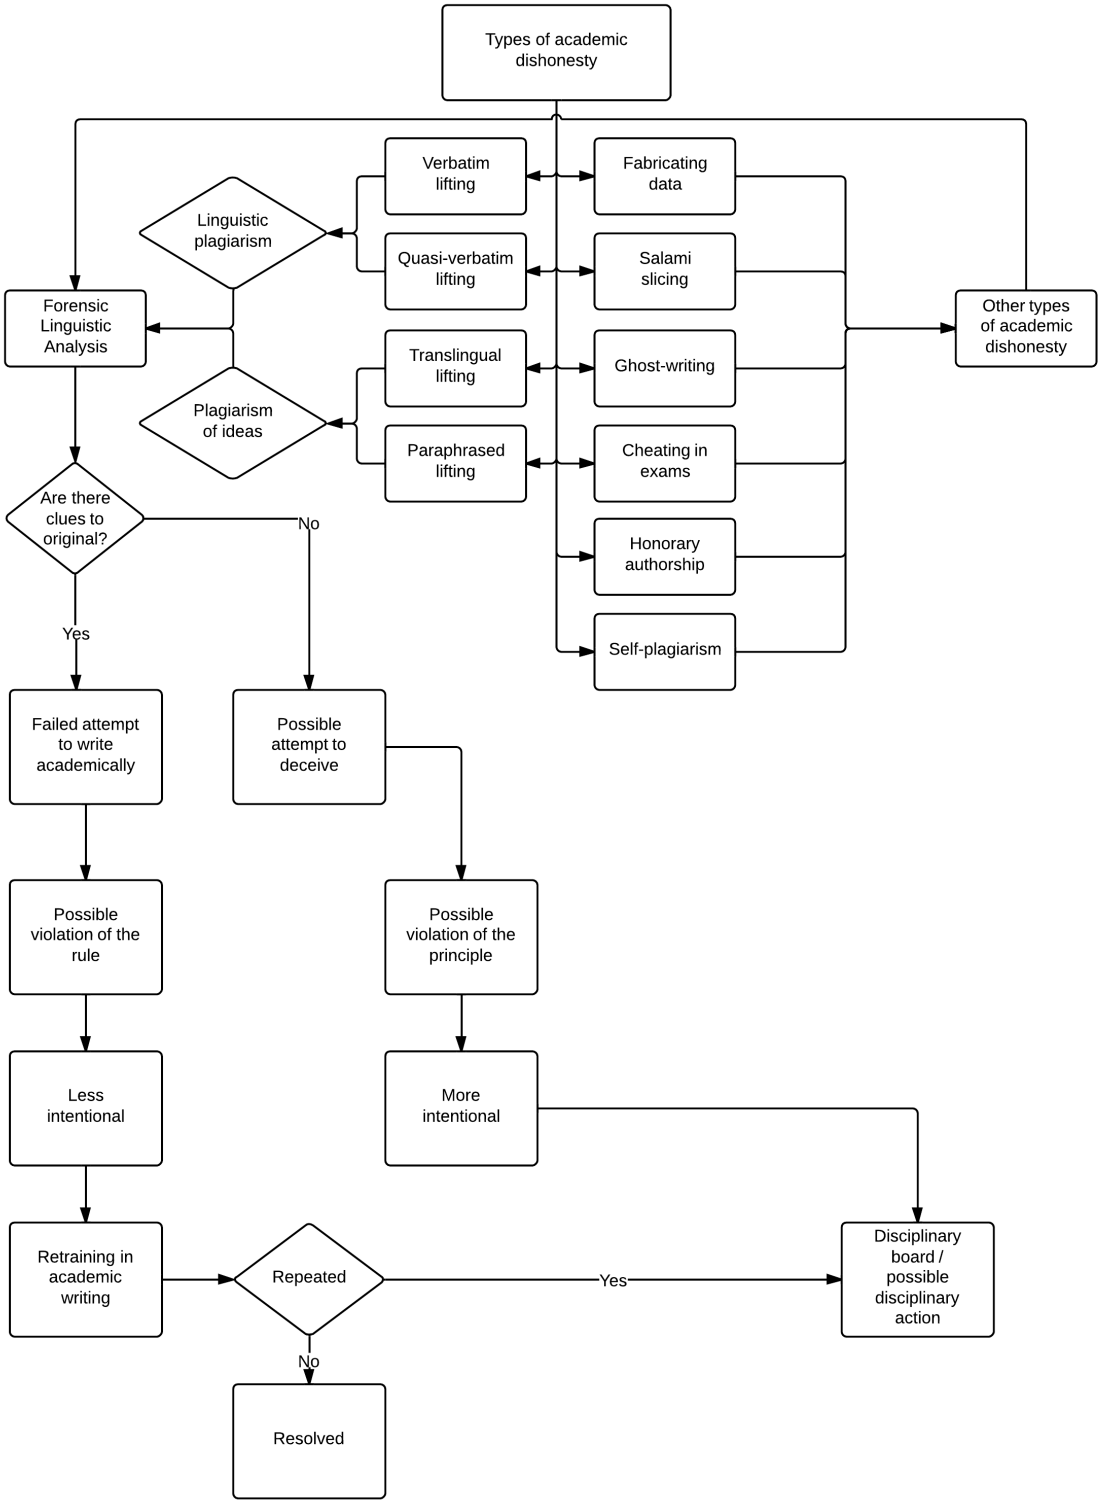

Figure 8.1: Plagiarism Handling Flowchart.

Bibliography

Abbasi, A. and Chen, H. (2008). Writeprints: A stylometric approach to identity-level identification and similarity detection in cyberspace. *ACM Transactions on Information Systems*, 26(2).

Almeida, C. P. (2010). *Código Penal*. Coimbra: Almedina.

Alvarez, M. (2009). Acting intentionally and acting for a reason. *Inquiry*, 52(3), 293–305.

Anderson, J. (1998). *Plagiarism, Copyright Violation and other Thefts of Intellectual Property: An Annotated Bibliography with a Lengthy Introduction*. Jefferson, North Carolina and London: McFarland & Company, Inc.

Angèlil-Carter, S. (2000). *Stolen Language? : Plagiarism in Writing*. Real Language Series. Harlow: Longman.

Anson, C. M. (2008). We never wanted to be cops: Plagiarism, institutional paranoia, and shared responsibility. In R. M. Howard and A. E. Robillard, Eds., *Pluralizing Plagiarism: Identities, Contexts, Pedagogies*, 140–157. Portsmouth, NH: Boynton/Cook.

Argamon, S. and Juola, P. (2011). Overview of the international authorship identification competition at pan-2011. In V. Petras, P. Forner and P. D. Clough, Eds., *CLEF (Notebook Papers/Labs/Workshop)*.

Argamon, S., Koppel, M., Pennebaker, J. W. and Schler, J. (2009). Automatically profiling the author of an anonymous text. *Commun. ACM*, 52(2), 119–123.

Argamon, S., Whitelaw, C., Chase, P., Hota, S. R., Garg, N. and Levitan, S. (2007). Stylistic text classification using functional lexical features: Research articles. *Journal of the American Society for Information Science and Technology*, 58(6), 802–822.

Ascensão, J. O. (1992). *Direitos de Autor e Direitos Conexos*. Coimbra: Coimbra Editora.

Austin, J. (1962). *How to do Things with Words*. Oxford: Clarendon Press.

Ayala, D. V., Castillo, E., Pinto, D., León, S. and Tovar, M. (2011). Baseline approaches for the authorship identification task – notebook for PAN at CLEF 2011. In *CLEF (Notebook Papers/Labs/Workshop)*.

Azuelos-Atias, S. (2007). *A Pragmatic Analysis of Legal Proofs of Criminal Intent*. Amsterdam and Philadelphia: J. Benjamins Pub. Co.

Barrón-Cedeño, A. and Rosso, P. (2009). On automatic plagiarism detection based on n-grams comparison. In *Proceedings of the 31th European Conference on IR Research on Advances in Information Retrieval*, ECIR '09, 696–700, Berlin, Heidelberg: Springer-Verlag.

Barrón-Cedeño, A., Rosso, P., Agirre, E. and Labaka, G. (2010). Plagiarism detection across distant language pairs. In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING '10, 37–45, Stroudsburg, PA, USA: Association for Computational Linguistics.

Barrón-Cedeño, A., Rosso, P., Pinto, D. and Juan, A. (2008). On cross-lingual plagiarism analysis using a statistical model. In *Proceedings of PAN-08*.

Barthes, R. (1977). *Image-Music-Text*. London: Fontana Press.

Bassnett, S. (1998). When is a translation not a translation? In S. Bassnett and A. Lefevere, Eds., *Constructing Cultures: Essays on Literary Translation*, 25–40. Clevedon: Multilingual Matters.

Bassnett, S. (2002). *Translation Studies*. London and New York: Routledge.

Becker, H. S. (1963). *Outsiders: Studies in the Sociology of Deviance*. New York: Free Press.

Bennett, R. (2005). Factors associated with student plagiarism in a post-1992 university. *Assessment & Evaluation in Higher Education*, 30(2), 137–162.

Bently, L. and Sherman, B. (2009). *Intellectual Property Law*. Oxford: Oxford University Press.

Bernstein, Y. and Zobel, J. (2004). A scalable system for identifying co-derivative documents. *String Processing and Information Retrieval*, 55–67.

Biggs, J. (1987). *Student Approaches to Learning and Studying*. Melbourne: Australian Council for Educational Research.

Bloch, J. (2012). *Plagiarism, Intellectual Property and the Teaching of L2 Writing*. Bristol: Multilingual Matters.

Brin, S., Davis, J. and Garcia-Molina, H. (1995). Copy detection mechanisms for digital documents. In *ACM International Conference on Management of Data (SIGMOD 1995)*.

Buffon, M. d. (1753). *Discours sur le Style Prononcé a l'Académie Française*. Paris: Librairie Classique Eugène Belin.

Burns, R. P. and Burns, R. (2009). Business research methods and statistics using spss. Online: http://tinyurl.com/bf93cyx – Date accessed: 30/11/2012.

Callison-Burch, C. and Osborne, M. (2006). Improved statistical machine translation using paraphrases. In *Proceedings of HLT/NAACL-2006*, 17–24.

Carroll, J. and Appleton, J. (2001). Plagiarism: A good practice guide. Online: http://www.altcexchange.edu.au/plagiarism-good-practice-guide – Date accessed: 20/06/2009.

Ceska, Z., Toman, M. and Jezek, K. (2008). Multilingual plagiarism detection. In D. Dochev, M. Pistore and P. Traverso, Eds., *Artificial Intelligence: Methodology, Systems, and Applications*, volume 5253 of *Lecture Notes in Computer Science*, 83–92. Springer Berlin Heidelberg.

Cha, S.-h. (2007). Comprehensive survey on distance / similarity measures between probability density functions. *International Journal of Mathematical Models and Methods in Applied Sciences*, 1(4), 300–307.

Chester, G. (2001). Pilot of free-text electronic plagiarism detection software. Online: http://tinyurl.com/agbezc2 – Date accessed: 20/10/2009.

Cohen, S. (1972). *Folk Devils and Moral Panics*. Oxon and New York: Routledge.

Corezola Pereira, R., Moreira, V. P. and Galante, R. (2010). A new approach for cross-language plagiarism analysis. In M. Agosti, N. Ferro, C. Peters, M. Rijke and A. Smeaton, Eds., *Multilingual and Multimodal Information Access Evaluation*, volume 6360 of *Lecture Notes in Computer Science*, 15–26. Springer Berlin Heidelberg.

Coulthard, M. (2004). Author identification, idiolect and linguistic uniqueness. *Applied Linguistics*, 25(4), 431–447.

Coulthard, M. and Johnson, A. (2007). *An Introduction to Forensic Linguistics: Language in Evidence*. London and New York: Routledge.

Coulthard, M., Johnson, A., Kredens, K. and Woolls, D. (2010). Four forensic linguists' responses to suspected plagiarism. In M. Coulthard and A. Johnson, Eds., *The Routledge Handbook of Forensic Linguistics*, 523–538. Milton Park, Abingdon, Oxon; New York, NY: Routledge.

Cronin, M. (2000). *Across the Lines: Travel Language and Translation*. Cork: Cork University Press.

Dainow, J. (1966). The civil law and the common law: Some points of comparison. *The American Journal of Comparative Law*, 15(3), 419–435.

de Nahlik, C. (2011). Ethical behaviour and academic integrity: contemporary practice. Online: http://tinyurl.com/a46tcqs – Date accessed 01/06/2012.

Duff, R. A. (2007). *Answering for Crime: Responsibility and Liability in the Criminal Law*. Oxford and Portland, Oregon: Hart.

Dõrnyei, Z. (2007). *Research Methods in Applied Linguistics*. Oxford: Oxford University Press.

D'Andrea, V.-M. (1999). Organizing teaching and learning: Outcomes-based planning. In H. Fry, S. Ketteridge and S. Marshall, Eds., *A Handbook for Teaching and Learning in Higher Education: Enchancing Academic Practice*. London: Kogan Page.

Eggington, W. G. (2008). Deception and fraud. In J. Gibbons and M. T. Turell, Eds., *Dimensions of Forensic Linguistics*, 249–264. Amsterdam and Philadelphia: John Benjamins.

Eiras, H. and Fortes, G. (2010). *Dicionário de Direito Penal e Processo Penal*. Lisbon: Quid Juris.

Entwistle, N. and Ramsden, P. (1983). Contrasting forms of understanding for degree examination: the student experience and its implications. *Higher Education*, 22, 205–27.

Estrada, F. (2001). Juvenile violence as a social problem: Trends, media attention and societal response. *British Journal of Criminology*, 41, 639–655.

Fairclough, N. (1995). *Critical Discourse Analysis: The Critical Study of Language*. London: Longman.

Fairclough, N. and Wodak, R. (1997). Critical discourse analysis. In T. A. v. Dijk, Ed., *Discourse Studies: A Multidisciplinary Introduction - Discourse as Social Interaction*, volume 2, 258–284. London: SAGE Publications Ltd.

Fillmore, L. (1991). When learning a second language means losing the first. *Early Childhood Research Quarterly*, 6(3), 323–346.

Finegan, E. (1993). Ethical considerations for expert witnesses in forensic linguistics. *Issues in Applied Linguistics*, 4(2), 179–187.

Finegan, E. (2003). Linguistic prescription: Familiar practices and new perspectives. *Annual Review of Applied Linguistics*, 23(1), 213–224.

Finegan, E. (2009). Expert linguists and the whole truth. *International Journal of Speech, Language and the Law*, 16(2), 267–277.

Finegan, E. (2012). *Language: Its Structure and Use*. Australia; United Kingdom: Wadsworth, 6th, international ed. edition.

Finnis, J. (1991). Intention and side-effects. In R. Frey and C. Morris, Eds., *Liability and responsibility: Essays in law and morals*, chapter 2, 32–64. Cambridge: Cambridge University Press.

Flowerdew, J. (1992). Definitions in science lectures. *Applied Linguistics*, 13(2), 202–221.

Foucault, M. (1979). Authorship: What is an author? *Screen*, 20(1), 13–34.

Garland, D. (2001). *The Culture of Control: Crime and Social Order in Contemporary Society*. Oxford: Oxford University Press.

Garner, B. A. (2009). *Black's Law Dictionary*. St. Paul, MN: West.

Gentzler, E. (1998). Foreword. In S. Bassnett and A. Lefevere, Eds., *Constructing Cultures: Essays on Literary Translation*, chapter Foreword, ix–xxii. Clevedon: Multilingual Matters.

Gentzler, E. (2011). *Contemporary Translation Theories*. Clevedon: Multilingual Matters, 2nd edition.

Goffman, E. (1975). *Frame Analysis: An Essay on the Organization of Experience*. Harmondsworth: Penguin.

Goffman, E. (1981). *Forms of Talk*. Philadelphia: University of Pennsylvania Press.

Goldstein, P. (2003). *Copyright's Highway: From Gutenberg to the Celestial Jukebox*. Stanford: Stanford University Press.

Grant, T. (2005). *Authorship attribution in a forensic context*. PhD thesis, University of Birmingham, Birmingham.

Grant, T. (2007). Quantifying evidence in forensic authorship analysis. *The International Journal of Speech, Language and the Law*, 14(1), 1–25.

Grant, T. (2008). Approaching questions in forensic authorship analysis. In J. Gibbons and M. T. Turell, Eds., *Dimensions of Forensic Linguistics*, 215–229. Amsterdam and Philadelphia: John Benjamins.

Grant, T. and Baker, K. (2001). Identifying reliable, valid markers of authorship: a response to Chaski. *Forensic Linguistics*, 8(1), 66–79.

Grice, H. P. (1969). Utterer's meaning and intention. *Philosophical Review*, 78(2), 147–177.

Grosz, B. J. and Sidner, C. L. (1986). Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3), 175–204.

Halliday, M. A. K. (1994). *An Introduction to Functional Grammar*. London: Edward Arnold.

Hänlein, H. (1998). *Studies in Authorship Recognition – A Corpus-based Approach*. Frankfurt: Peter Lang.

Hansen, B. C. and Hansen, K. D. (1995). Academic and scientific misconduct: Issues for nursing educators. *Journal of Professional Nursing*, 11(1), 31–39.

Hart, W. and Albarracín, D. (2010). Learning about what others were doing. *Psychological Science*, 22(2), 261–266.

Hirst, G. and Feiguina, O. (2007). Bigrams of syntactic labels for authorship discrimination of short texts. *Literary and Linguist Computing*, 22(4), 405–417.

Hoad, T. C. and Zobel, J. (2003). Methods for identifying versioned and plagiarized documents. *Journal of the American Society for Information Science and Technology*, 54(3), 203–215.

Hoey, M. (2005). *Lexical Priming: A New Theory of Words and Language*. London and New York: Routledge.

Howard, R. (1995). Plagiarisms, authorships, and the academic death penalty. *College English*, 57(7), 788–806.

Howard, R. and Robillard, A. (2008). Plagiarisms. In R. Howard and A. Robillard, Eds., *Pluralizing Plagiarism: Identities, Contexts, Pedagogies*, chapter Plagiarisms, 1–7. Portsmouth: Boynton/Cook.

Howard, R. M. (1999). *Standing in the Shadow of Giants: Plagiarists, Authors, Collaborators*. Stamford, CT: Ablex Publising Corporation.

Howard, R. M. (2001). Forget about policing plagiarism. Just teach. *The Chronicle of Higher Education*, B24.

Hughes, R. (1998). Considering the vignette technique and its application to a study of drug injecting and HIV risk and safer behaviour. *Sociology of Health and Illness*, 20(3), 381–400.

IEEE, (2006). IEEE publication services and products board operations manual. Online: Online: http://www.ieee.org/web/publications/PSPB/index.html – Date accessed: 15/02/2009.

Iyer, P. and Singh, A. (2005). Document similarity analysis for a plagiarism detection system. In *IICAI'05*, 2534–2544.

Jameson, D. A. (1993). The ethics of plagiarism: How genre affects writers' use of source materials. *Bulletin of the Association for Business Communication*, 56(2), 18.

Joachims, T. (1998). Text categorization with support vector machines: learning with many relevant features. In C. Nédellec and C. Rouveirol, Eds., *ECML-98, 10th European Conference on Machine Learning*, 137–142, Heidelberg: Springer Verlag.

Johnson, A. (1997). Textual kidnapping – a case of plagiarism among three student texts? *The International Journal of Speech, Language and the Law*, 4(2), 210–225.

Johnstone, B. (2008). *Discourse Analysis*. Oxford: Blackwell.

Jones, M. (2009). Back-translation: the latest form of plagiarism. In *The 4th Asia Pacific Conference on Educational Integrity*, 1–7, Wollongong: University of Wollongong.

Kaplan, T. and Torbati, J. (2007). Plagiarism rampant at peking university, Yale professor says. *Yale Daily News*.

Kern, R., Seifert, C., Zechner, M. and Granitzer, M. (2011). Vote/veto meta-classifier for authorship identification – notebook for PAN at CLEF 2011. In *CLEF (Notebook Papers/Labs/Workshop)'11*.

Kestemont, M., Luyckx, K. and Daelemans, W. (2011). Intrinsic plagiarism detection using character trigram distance scores – notebook for PAN at CLEF 2011. In V. Petras, P. Forner and P. D. Clough, Eds., *CLEF (Notebook Papers/Labs/Workshop)*.

Koehn, P. (2010). *Statistical Machine Translation*. Cambridge: Cambridge University Press.

Koppel, M. and Schler, J. (2004). Authorship verification as a one-class classification problem. In *Proceedings of the twenty-first international conference on Machine learning*, ICML '04, New York, NY: ACM.

Kress, G. (2001). Critical sociolinguistics. In R. Mesthrie, Ed., *Concise Encyclopedia of Sociolinguistics*, 542–545. Oxford: Elsevier.

Kugler, I. (2002). *Direct and Oblique Intention in the Criminal Law: An Inquiry into Degrees of Blameworthiness*. Hampshire: Ashgate.

Laviosa, S. (1998,). Core patterns of lexical use in a comparable corpus of English narrative prose. *Meta*, 43(4), 557–570.

Lefevere, A. (1998a). Chinese and Western thinking on translation. In S. Bassnett and A. Lefevere, Eds., *Constructing Cultures: Essays on Literary Translation*, 12–24. Clevedon: Multilingual Matters.

Lefevere, A. (1998b). The gates of analogy: the kalevala in English. In S. Bassnett and A. Lefevere, Eds., *Constructing Cultures: Essays on Literary Translation*, 76–89. Clevedon: Multilingual Matters.

Lefevere, A. and Bassnett, S. (1998). Introduction: Where are we in translation studies? In S. Bassnett and A. Lefevere, Eds., *Constructing Cultures: Essays on Literary Translation*, 1–11. Clevedon: Multilingual Matters.

Leitão, L. M. T. M. (2011). *Direito de Autor*. Coimbra: Almedina.

Lindey, A. (1952). *Plagiarism and Originality*. New York: Harper & Brothers.

Love, H. (2002). *Attributing Authorship: An Introdution*. Cambridge: Cambridge University Press.

Luyckx, K. (2011). Authorship identification of e-mail as a multi-class task – notebook for PAN at CLEF 2011. In *CLEF (Notebook Papers/Labs/Workshop)'11*.

Macherey, K., Och, F. J. and Ney, H. (2001). Natural language understanding using statistical machine translation. In *European Conference on Speech Communication and Technology*, 2205–2208.

Madnani, N. and Dorr, B. J. (2010). Generating phrasal and sentential paraphrases: A survey of data-driven methods. *Computational Linguistics*, 36(3), 341–387.

Maia, B., Sousa-Silva, R., Barreiro, A. and Fróis, C. (2008). N-grams in search of theories. In B. Lewandowska-Tomaszczyk, Ed., *Corpus Linguistics, Computer Tools, and Applications - State of the Art (PALC 2007)*, volume 17. Frankfurt am Main, Berlin, Bern, Bruxelles, New York, Oxford, Wien: Peter Lang.

Major, R. C. (1992). Losing English as a first language. *The Modern Language Journal*, 2, 190–208.

Mallon, T. (2001). *Stolen Words: The Classic Book on Plagiarism*. San Diego: Harvest/HBJ Book.

Manber, U. (1994). Finding similar files in a large file system. In *USENIX Winter'94*, 1–10.

Manning, C. D. (2001). Probabilistic approaches to syntax. Online: http://tinyurl.com/bcz9vzg – Date accessed 04/02/2010.

Maurer, H., Kappe, F. and Zaka, B. (2006). Plagiarism – a survey. *Journal of Universal Computer Science*, 12(8), 1050–1084.

McCawley, J. D. (1999). Participant roles, frames, and speech acts. *Linguistics and Philosophy*, 22, 595–619.

McMenamin, G. R. (1993). *Forensic Stylistics*. Amsterdam: Elsevier.

McMenamin, G. R. (2002). *Forensic Linguistics: Advances in Forensic Stylistics*. Boca Raton and New York: CRC Press.

Merton, R. (1968). Foreword. In E. Garfield, Ed., *Citation Indexing – Its Theory and Application in Science, Technology, and Humanities*. Philadelphia: Isi Press.

Meyer zu Eissen, S. and Stein, B. (2004). Genre classification of web pages: User study and feasibility analysis. In S. Biundo, T. Frühwirth and G. Palm, Eds., *Proceedings of the 27th Annual German Conference on Artificial Intelligence (KI 04)*, 256–269, Ulm, Germany: Springer.

Meyer Zu Eissen, S. and Stein, B. (2006). Intrinsic plagiarism detection. In *Proceedings of the European Conference on Information Retrieval (ECIR-06)*.

Mikros, G. K. and Perifanos, K. (2011). Authorship identification in large email collections: Experiments using features that belong to different linguistic levels – notebook for PAN at CLEF 2011. In V. Petras, P. Forner and P. D. Clough, Eds., *CLEF (Notebook Papers/Labs/Workshop)*.

Mozgovoy, M. (2008). *Enhancing Computer-Aided Plagiarism Detection*. Saarbrücken: VDM Verlag Dr. Müller.

Muncie, J. (2008). The 'punitive turn' in juvenile justice: Cultures of control and rights compliance in Western Europe and the USA. *Youth Justice*, 8(2), 107–121.

Oberreuter, G., L'Huillier, G., Ríos, S. A. and Velásquez, J. D. (2011). Approaches for intrinsic and external plagiarism detection – notebook for PAN at CLEF 2011. In V. Petras, P. Forner and P. D. Clough, Eds., *CLEF (Notebook Papers/Labs/Workshop)*.

Och, F. J., Tillmann, C., Ney, H. and Informatik, L. F. (1999). Improved alignment models for statistical machine translation. In *University of Maryland, College Park, MD*, 20–28.

Oliveira, F. (2005). Semântica. In I. H. Faria, E. R. Pedro, I. Duarte and C. A. M. Gouveia, Eds., *Introdução à Linguística Geral e Portuguesa*, chapter 7, 333–379. Lisbon: Caminho.

Pataki, M. (2012). A new approach for searching translated plagiarism. In *Proceedings of the 5th International Plagiarism Conference*, Newcastle upon Tyne.

Pavlenko, A. (2000). L2 influence on L1 in late bilingualism. *Issues in Applied Linguistics*, 11(2), 175–205.

Pears, R. and Shields, G. (2008). *Cite them Right: The Essential Referencing Guide*. Newcastle upon Tyne: Pear Tree Books.

Pecorari, D. (2008). *Academic Writing and Plagiarism: A Linguistic Analysis*. London: Continuum.

Pereira, A. L. D. (2003). Problemas actuais da gestão do direito de autor: gestão individual e gestão colectiva do direito de autor e dos direitos conexos na sociedade da informação. In *Estudos em Homenagem ao Professor Doutor Jorge Ribeiro de Faria – Faculdade de Direito da Universidade do Porto*, 17–37. Coimbra Editora.

Pereira, V. S. and Lafayette, A. (2008). *Código Penal Anotado e Comentado*. Lisbon: Quid Juris.

Porter, M. F. (1980). An algorithm for suffix stripping. *Program*, 14(3), 130–137.

Potthast, M., Barrón-Cedeño, A., Stein, B. and Rosso, P. (2011a). Cross-language plagiarism detection. *Language Resources and Evaluation*, 45, 45–62.

Potthast, M., Eiselt, A., Barrón-Cedeño, A., Stein, B. and Rosso, P. (2011b). Overview of the 3rd international competition on plagiarism detection. In V. Petras, P. Forner and P. D. Clough, Eds., *CLEF (Notebook Papers/Labs/Workshop)*.

Potthast, M., Eiselt, A., Stein, B., Barrón-Cedeño, A. and Rosso, P. (2009a). Pan plagiarism corpus pan-pc-09. Online: http://www.uni-weimar.de/cms/medien/webis/research/corpora/pan-pc-09.html – Date accessed: 02/03/2010.

Potthast, M., Stein, B., Eiselt, A., Barrón-Cedeño, A. and Rosso, P. (2009b). Overview of the 1st international competition on plagiarism detection. In B. Stein, P. Rosso, E. Stamatatos, M. Koppel and E. Agirre, Eds., *SEPLN 2009 Workshop on Uncovering Plagiarism, Authorship, and Social Software Misuse (PAN 09)*, 1–9: CEUR-WS.org.

Rao, S., Gupta, P., Singhal, K. and Majumder, P. (2011). External & intrinsic plagiarism detection: VSM & discourse markers based approach – notebook for PAN at CLEF 2011. In V. Petras, P. Forner and P. D. Clough, Eds., *CLEF (Notebook Papers/Labs/Workshop)*.

Recasens, M. and Vila, M. (2010). On paraphrase and coreference. *Computational Linguistics*, 36(4), 639–647.

Robillard, A. E. (2008). Situating plagiarism as a form of authorship: The politics of writing in a first-year writing course. In R. Howard and A. E. Robillard, Eds., *Pluralizing Plagiarism: Identities, Contexts, Pedagogies*, 27–42. Portsmouth, NH: Boynton/Cook.

Robillard, A. E. and Howard, R. M. (2008). Plagiarisms. In R. M. Howard and A. E. Robillard, Eds., *Pluralizing Plagiarism: Identities, Contexts, Pedagogies*, 1–7. Portsmouth, NH: Boynton/Cook.

Robson, C. (2002). *Real World Research*. Oxford: Blackwell Publishing.

Roig, M. (2001). Plagiarism and paraphrasing criteria of college and university professors. *Ethics and Behavior*, 11(3), 307–323.

Rus, V., Feng, S., Brandon, R., Crossely, S. and McNamara, D. S. (2011). A linguistic analysis of student-generated paraphrases. In *Proceedings of the Twenty-Fourth International Florida Artificial Intelligence Research Society Conference*, 293–298.

Saville-Troike, M. (1982). *The Ethnography of Communication*. Oxford: Basil Blackwell.

Scaife, B. (2007). Evaluation of plagiarism detection software. Online: http://www.plagiarismadvice.org/nccgroupreport.php – Date accessed: 30/11/2010.

Scanlan, C. (2004). The poynter online published. Online: http://www.poynter.org/column.asp?id=52&aid=70742&view=print – Date accessed: 16/02/2009.

Scollon, R. (1994). As a matter of fact: The changing ideology of authorship and responsibility in discourse. *World Englishes*, 13(1), 33–46.

Scollon, R. (1995). Plagiarism and ideology: Identity in intercultural discourse. *Language in Society*, 24, 1–28.

Seliger, H. and Vago, R. (1991). The study of first language attrition. In H. Seliger and R. Vago, Eds., *The Study of First Language Attrition: An overview*, 3–15. Cambridge University Press.

Selinker, L. (1972). Interlanguage. *International Review of Applied Linguistics*, 10, 209–241.

Selinker, L. and Lakshmanan, U. (1992). Language transfer and fossilization: The "multiple effects principle". In S. Gass and L. Selinker, Eds., *Language Transfer in Language Learning*. Amsterdam and Philadelphia: John Benjamins.

Seneff, S. (1992). Tina: a natural language system for spoken language applications. *Computational Linguistics*, 18(1), 61–86.

Shuy, R. W. (2001). Discourse analysis in the legal context. In D. Schiffrin, D. Tannen and H. E. Hamilton, Eds., *The Handbook of Discourse Analysis*, chapter 22, 437–452. Malden: Blackwell.

Slocum, J. (1984). Machine translation: its history, current status, and future prospects. In *Proceedings of the 10th international conference on Computational linguistics*, COLING '84, 546–561, Stroudsburg, PA, USA: Association for Computational Linguistics.

Smith, W. (1994). Computers, statistics and disputed authorship. In J. Gibbons, Ed., *Language and the Law*, chapter 16, 374–413. Harlow: Longman.

Solorio, T., Pillay, S. and y Gómez, M. M. (2011). Authorship identification with modality specific meta features – notebook for PAN at CLEF 2011. In *CLEF (Notebook Papers/Labs/Workshop)*.

Sousa-Silva, R., Grant, T. and Maia, B. (2010). 'I didn't mean to steal someone else's words!': A forensic linguistic approach to detecting intentional plagiarism. In *4th International Plagiarism Conference: "Towards an Authentic Future" - 2010 Conference Proceedings & Abstracts*: PlagiarismAdvice.org.

Stamatatos, E. (2006). Ensemble-based author identification using character n-grams. *ReCALL*, 41–46.

Stamatatos, E. (2007). Author identification using imbalanced and limited training texts. In *Proceedings of the 18th International Conference on Database and Expert Systems Applications*, 237–241, Washington, DC, USA: IEEE Computer Society.

Stamatatos, E. (2009a). Intrinsic plagiarism detection using character n-gram profiles. *Knowledge Creation Diffusion Utilization*, 2, 38–46.

Stamatatos, E. (2009b). A survey of modern authorship attribution methods. *Journal of the American Society for Information Science and Technology*, 60(3), 538–556.

Stein, B. (2005). Fuzzy-fingerprints for text-based information retrieval. *Knowledge Management*, Special Issue, 572–579.

Stein, B. and Eissen, S. (2006). Near similarity search and plagiarism analysis. *From Data and Information Analysis to Knowledge Engineering*, 430–437.

Stein, B., zu Eissen, S. M. and Potthast, M. (2007). Strategies for retrieving plagiarized documents. In *SIGIR'07*, 825–826.

Sunderland, J. (2004). *Gendered Discourses*. Basingstoke: Palgrave Macmillan.

Swales, J. M. (1990). *Genre Analysis: English in Academic and Research Settings*. Cambridge: Cambridge University Press.

Taylor, W. (1953). Cloze procedure: A new tool for measuring readability. *Journalism Quarterly*, 30, 415–433.

Tonry, M. (2001). Symbol, substance, and severity in Western penal policies. *Punishment & Society*, 3, 517–536.

Toury, G. (1995). *Descriptive Translation Studies and Beyond*. Amsterdam: John Benjamins Publishing Company.

Turell, M. T. (2004). Textual kidnapping revisited: the case of plagiarism in literary translation. *The International Journal of Speech, Language and the Law*, 11(1), 1–26.

Turell, M. T. (2007). Plagio y traducción literaria. *Vasos Comunicantes*, 37(1), 43–54.

Turell, M. T. (2008). Plagiarism. In J. Gibbons and M. T. Turell, Eds., *Dimensions of Forensic Linguistics*, 265–299. Amsterdam and Philadelphia: John Benjamins.

Uzuner, O., Katz, B. and Nahnsen, T. (2005). Using syntactic information to identify plagiarism. In *Proc. 2nd Workshop on Building Educational Applications using NLP. Ann Arbor*.

Vila, M., Martí, M. A. and Rodríguez, H. (2011). Paraphrase concept and typology. A linguistically based and computationally oriented approach. *Procesamiento del Lenguaje Natural*, 46, 83–90.

Vogel, S., Ney, H. and Tillmann, C. (1996). Hmm-based word alignment in statistical translation. In *Proceedings of the 16th conference on Computational linguistics*, volume 2 of *COLING '96*, 836–841, Stroudsburg, PA, USA: Association for Computational Linguistics.

Wacquant, L. (1999). How penal common sense comes to Europeans: Notes on the transatlantic diffusion of the neo-liberal doxa. *European Societies*, 1(3), 319–352.

Wakeford, R. (1999). Principles of assessment. In N. H. Fry, S. Ketteridge and S. Marshall, Eds., *A Handbook for Teaching and Learning in Higher Education: Enchancing Academic Practice*. London: Kogan Page.

Weinreich, U. (1953). *Languages in Contact*. London/The Hague/Paris: Mouton.

White, C. (2004). The cope report 2003: Annual report of the committee on publication ethics. Online: http://publicationethics.org/annualreport/2003 – Date accessed: 15/10/2010.

Williams, K. and Carroll, J. (2009). *Referencing and Understanding Plagiarism*. Basingstoke: Palgrave Macmillan.

Woolls, D. (2003). Better tools for the trade and how to use them. *International Journal of Speech, Language and the Law*, 10(1), 102–112.

Woolls, D. (2006). Plagiarism. In K. Brown, Ed., *The Encyclopedia of Language and Linguistics*, volume 9. Oxford: Elsevier.

Woolls, D. (2010). Computational forensic linguistics: Searching for similarity in large specialised corpora. In M. Coulthard and A. Johnson, Eds., *The Routledge Handbook of Forensic Linguistics*, 576–590. Milton Park, Abingdon, Oxon; New York, NY: Routledge.

Woolls, D. (2012). Detecting plagiarism. In L. Solan and P. Tiersma, Eds., *The Oxford Handbook of Language and Law*. Oxford: Oxford University Press.

Woolls, D. and Coulthard, M. (1998). Tools for the trade. *International Journal of Speech, Language and the Law*, 5(1), 33–57.

Wubben, S., van den Bosch, A. and Krahmer, E. (2010). Paraphrasing headlines by machine translation: Sentential paraphrase acquisition and generation using Google News. In E. Westerhout, T. Markus and P. Monachesi, Eds., *Proceedings of the 20th Meeting of Computational Linguistics in the Netherlands*.