Beyond Goal-Rationality: Traditional Action Can Reduce Volatility in Socially Situated Agents

Chloe M. Barnes, Anikó Ekárt and Peter R. Lewis

Aston Lab for Intelligent Collectives Engineering (ALICE)

Aston University, Aston Triangle, Birmingham, B4 7ET, UK

Email: {barnecm1 — a.ekart — p.lewis}@aston.ac.uk

Abstract

Systems that pursue their own goals in shared environments can indirectly affect one another in unanticipated ways, such that the actions of other systems can interfere with goal-achievement. As humans have evolved to achieve goals despite interference from others in society, we thus endow socially situated agents with the capacity for social action as a means of mitigating interference in co-existing systems. We demonstrate that behavioural and evolutionary volatility caused by indirect interactions of goal-rational agents can be reduced by designing agents in a more socially-sensitive manner. We therefore challenge the assumption that designers of intelligent systems typically make, that goal-rationality is sufficient for achieving goals in shared environments.

1 Introduction

Socio-technical systems are comprised of many interacting components, where decisions are being increasingly delegated from humans to so-called 'intelligent' machines. These machines typically act in reference to a given goal, such as an objective function, utility or goal-state – which may often conflict with other entities in the system [49]. Yet, humans have evolved both social intelligence and social self-awareness to express more complex behaviour than purely goal-rational action in order to succeed in highly social environments, driven by factors such as values,

emotions and traditions [1,23,35,91]. Despite the fact that components in a socio-technical system are inherently socially situated, they are typically limited to performing goal-rational actions and thus only very rudimentary social action. Specifically, they do not perceive or reason about the effect other systems can have on their own ability to learn and evolve, or the capabilities of others around them; the actions that they direct towards others are also not driven by any broader social meaning – unlike in humans [7,30]. These systems are also unaware of their *own* impact on the world around them, which can have a catastrophic and unpredictable effect. In 2010 for example, a \$1 trillion stock market crash occurred in just 36 minutes, caused at least in part by the unforeseen interactions of several automated trading agents [85].

Socio-technical systems such as vehicular networks [31], smart energy grids [62], and trading agents [18] are increasingly being designed to operate in dynamic, uncertain and social environments, where interactions are potentially unanticipated or unknown. Explicit and anticipated interactions can be designed for when integrating systems, however, neglecting to consider potentially *unintended* interactions with others that are co-located can lead to worse performance [89]. [37] argue that as the complexity of these inherently social systems of systems increases, the number of *implicit* and *unintended* interactions between the systems will also increase as a consequence; further, it is not only the interactions with other systems that makes the

1

task of runtime integration challenging, but also as [61] points out – both the intended and unintended interactions with humans as well. The field of self-improving system integration aims to design systems that overcome these issues – without full knowledge, control or authority over the other systems in which they coexist and interact with [6]. The actions of one system can have an unintended effect on the others surrounding it [9]; a transition to more socially intelligent systems that are able to learn about others in their environment is therefore necessary to not only enable systems to self-integrate with others around them at runtime, but to maintain learnt knowledge and goal-achieving behaviour despite interference from others.

More broadly, as we delegate more decisions to intelligent machines, we posit that capturing elements of evolved human social behaviour will be increasingly important; here, social action and social self-awareness are two essential aspects that so far remain largely unexplored. Human societies have no global knowledge or central point of control; how then, do humans interact effectively? Organic Computing approaches this by observing and controlling a group of interacting, self-organising entities [59]; we however take a microsociological approach, as a step towards socially intelligent systems capable of self-awareness at the individual level. Humans have evolved the ability to achieve goals in complex social systems by considering others and acting socially; we thus operationalise social action theory [91] in socially situated agents that pursue individual goals in shared environments with less volatility than current systems. We operationalise two types of social action – traditional action is acting similarly to the rest of the population, and goal-rational action is seen in current systems, where the most effective action to achieve a goal is taken. We also introduce random action, which adds randomness into the behaviour of agents by way of Random Immigrants [19]. The need to act in a socially-sensitive way challenges the assumption that goal rationality is sufficient for agents to achieve goals, when they operate in shared environments with less than complete knowledge.

The experiments use the River Crossing Dilemma as a testbed [4], which was designed to explore arbitrarily complex problems in shared environments. Firstly we explore the impact that interference has on agents that are able to achieve individual goals alone, to assess how learnt knowledge is maintained. We then generalise these results over a large number of experiments for agents that begin evolution with no prior knowledge. We finally draw conclusions to compare the volatility of evolution with each type of social action, and whether goal-achieving behaviour can therefore be learnt and maintained despite interference.

2 Background and Related Work

2.1 Evolution of Sociality in Humans and Animals

We as humans have evolved the ability to navigate and utilise our social environments to our advantage. Social learning can be seen in both humans and the animal world alike, enabling us to learn and do more complex things than we would be able to individually [87]; it has also been shown to be favoured over individual learning when environmental change is slow, or when individual learning or non-social and environmental cues are not useful [55, 65]. Our innate capacity for cooperation over competition and thus our ability to learn from others has been attributed to our success as a species, and to what distinguishes the complexity of our cognitive abilities from that of primates [33, 38, 57].

The cultural intelligence and cultural brain hypotheses posit that species that have evolved to favour social learning consequently evolve improved asocial learning and individual problem-solving skills, and flexibility, adaptability and innovation in learning compared to non-social species; this results in a more intelligent population overall [5, 12, 32, 60, 79, 87]. The social brain hypothesis supports this, postulating that social abilities were favoured and acted as a driving force during evolution; this is said to have led to an increase in brain size to support cognitive ability and intelligence, where neocortex

size correlates with group size in humans and primates [1, 14, 28, 42].

[43] provide supporting evidence for this hypothesis by showing a positive correlation between the size of the mushroom bodies of ants (a higher brain centre like the neocortex in mammalian brains) and colony size. Sociality in ants and other eusocial insects however differs to that of primates; the neural mechanisms driving insect sociality remain relatively unknown [43, 51], but recent efforts explore how this can be addressed in the future [77].

The success of social learning and behaviour in humans and animals alike has inspired many Computer Science researchers. These observations and theories from nature have been widely used to create optimisation algorithms and to design systems with emergent collective intelligence. For example, self-organisation and cooperation in ants and bees have inspired the development of many optimisation algorithms [17, 26, 44]. Swarm robotics is inspired by the cooperation in groups of social animals to achieve goals or complete tasks, where robots cooperate to solve more complex tasks collectively [13, 58]; this has also been used for societally important tasks such as search and rescue [3]. Social learning strategies have also been widely explored [41, 65], taking inspiration from processes such as mimicry, imitation and learning from others [63, 69, 73].

2.2 Sociality in Agents

Individuals in collective systems are inherently socially situated - their actions affect others around them either directly, or indirectly through the environment. Pursuing self-interested action in a social setting can lead to collective irrationality [45]; however, sociality through self-organising institutions can enable groups of self-interested individuals to govern themselves, supporting sustainable management of common pool resources [66]. Social dynamics have been widely explored in areas such as game theory, sociology, economics and evolutionary computation [45, 48, 83, 86].

Agents may pursue common or individual goals when sharing an environment. Argumentation, negotiation and persuasion [80], goal-aware team affiliation [29], norms and obligations [24], social plans and joint intentions [40, 68], and mutual influence detection [71] all promote cooperative behaviour in multi-agent systems, attempting to mitigate interference between goal-pursuing agents. The BDI (beliefs-desires-intentions) agent architecture for example enables agents to reason about norms and obligations that they are explicitly aware of to promote cooperative behaviour [24]. Other research explores deliberative normative agents, which require explicit knowledge of others, as well as the norms that exist in order to deliberately adhere to or violate them [16].

Systems that *intentionally* cooperate, coordinate, or act socially, require social awareness [7], and are capable of perceiving and reasoning about others. The evolution of cooperation has been explored extensively, with social dilemmas such as the Prisoner's Dilemma or the Snowdrift Game used to explore social dynamics and strategies [2, 11, 25, 45]. Although this study explores agents that may evolve to cooperate, *promoting* cooperation is not the focus of this paper. Rather, our questions concern the impact of coexistence and interference on the ability to achieve individual goals, when agents are unable to learn of the existence, goals or intentions of others; knowledge of all others in a system would enable perfect coordination, but is infeasible in dynamic environments due to the infinite power needed for processing and reasoning [40]. Not only this, but information about other systems or entities in dynamic or uncertain environments may not be available at design-time – including their goals, capabilities and how to integrate with them. If a system cannot perceive or learn about potentially unknown others, or the effect that their actions can have on itself, interference can affect how goals are achieved in ways that cannot be understood; without a mechanism to overcome interference and unanticipated events beyond their control, systems will be unable to make appropriate decisions at runtime in accordance with their goals. It is therefore necessary to move towards socially situated systems with the capacity for social awareness, by equipping them first with the ability to maintain goal-achieving behaviour despite unanticipated interference from other systems.

3

2.3 Interference in Multi-Agent Systems

The term *interference* has been used to describe the interaction between entities in a shared environment, arising from the competition for shared resources [72]. Interference is an inherent characteristic of a shared world; it can arise from entities directly interacting with each other, or indirectly by interacting with the environment which acts as a passive entity [15, 27]. The pursuit of goals in a shared world is said to be a core component of *sociality* [15]; social dynamics such as cooperation and competition are the result of interference from entities that help (positive interference) or hinder (negative interference) others with respect to their goals. Other research however states that interference is purely the negative effect on an entity's goal-driven behaviour [53]; here, the term *implicit cooperation* is analogous to positive interference. In this work, we adopt the broader definition of interference which encompasses both the positive and negative effects.

[15] describes *dependence* as a special case of interference, where interference is so strong that entities become dependent on the actions of others to achieve their goals, and thus cannot achieve their goals through their own actions alone. Other researchers have explored the concept of interference with different terminology. [40] describes this as an *interdependence* of actions that arises when the actions of entities are related, or have an impact on others. [27] states that co-existing robots will have social interactions that can be indirect, as their actions will *influence* and affect others around them. [71] explore how to detect *mutual influences* that arise in smart camera networks, and the resulting dependencies. [82] explore interference *within a single agent*; internal interference can arise from conflicting actions necessary to fulfil parallel conflicting goals (e.g. an action is a step towards achieving one goal, but a step away from achieving another, conflicting goal).

These methods require agents to be aware of others in the environment, and potentially their goals or intentions. In unpredictable and dynamic environments, this is not always possible; the environment and the agents within it may change over time, so this information may not always be available or predictable at design-time. [84] identify that subsystems can influence one another either directly or indirectly, whether they are intended to interact together or not; consequently, the 'interwoven' nature of these systems where uncertainty, heterogeneity of entities and interference between said entities prevail means that traditional methods of integration become impractical [8]. Thus, agents in dynamic or uncertain environments, where the presence of others is potentially unknown, need the ability to develop and maintain these models on-the-fly. Endowing systems with the ability to learn from their inherently social environment, as human societies have evolved to do, could therefore enable them to learn and evolve potentially more complex individual behaviours, problem-solving competencies, and goal-achieving abilities.

This work extends that of [4], and explores the volatility and maintenance of goal-achieving behaviour in agents that experience interference when they have no models of others in the environment, as a prerequisite to social awareness in systems. Specifically, we conduct further experiments to explore and generalise the effects that both goal-rational and traditional action, as well as random action, have on expected fitness and volatility in learning on a broader scale. With this, we aim to distinguish the differences between goal-rational and traditional action, and additionally investigate whether these approaches are quantitatively different to introducing Random Immigrants into the population [19] – a mechanism widely in dynamic optimisation problems and those with dynamic environments to add diversity to populations in genetic algorithms [34, 47, 75, 93, 97].

An important distinction between this work and that of the work it is extended from is a change in terminology. [4] explored the instability in evolution caused by interference, whereas in this work we investigate the *volatility* in evolution as a result of interference. This change in terminology was influenced by the additional statistical analysis that has been conducted in this work; three newly defined metrics have been proposed and used to capture different elements of

volatility in evolution, which were inspired and adapted from the established historical volatility metric used in financial modelling and volatility forecasting [67]. These metrics capture the dispersion of values over time. We also conduct an analysis of how the expected fitness changes depending on the type of action used, and whether agents are socially situated or not, as well as a deeper analysis into the results in the original work [4]. Finally, we analyse the trade-off that exists between fitness and volatility for the three different types of action operationalised in this work.

3 Inspiration from the Theory of Social Action

Interference means that actions can have unintended or unanticipated consequences on both the actor and others in shared, complex and dynamic environments [54], making shared and individual goals more difficult to achieve. Humans overcome these issues by acting *socially* and not purely *individualistically* in social environments. Many computer science researchers have therefore been inspired by theories of psychology, sociology and cognitive science, such as in organic computing [59], self-awareness [46,50], and social dilemmas, social learning, altruism and agent societies [39, 41, 64, 90]. It thus seems logical to draw parallels between the exploration of human social phenomena in sociology and socially situated agents.

3.1 Outlining Terminology

To explore social concepts in agents, one must first adopt and define the relevant terminology. [68] define 'situated agents' as those that are resource-bound, operate within and continuously interact with dynamic environments, and balance reactivity with deliberation. [52] use the term 'social situatedness' to include both the social and cultural aspects of being situated in an environment, which affect the behavioural and cognitive processes of the agents that are situated within. This is influenced by the theories of [88], such that the social aspects of environments that agents are situated in affect the in-

dividual intelligence of the agents. [21] describes socially situated agents to be those that not only gather information from their physical environment, but from the social component of it as well. We therefore term the agents in this work as *socially situated* due to their operation in a dynamic and shared environment; their behaviour can interfere [15] with the actions and goals of others, and what they learn is influenced by their physical and social environment – whether they are aware of it or not.

As previously stated, socially situated agents experience interference and can consequently be capable of *social action*. [91] define social actions as those being oriented towards and considering the behaviour of others; they hold meaning, which refers to the motivation from the actor's perspective. Actions with inanimate objects do not involve other actors and are thus not social. Additionally, actions that have no meaning, i.e. those without motivation or deliberation, are not actions, but merely behaviours. To further distinguish between these concepts, [81] proposes a hierarchy of social action; we explore social action as defined by these terms. A subset of this hierarchy is outlined below:

- *Behaviour* is automatic, reactive and reflexive.

- *Action* is intentional and purposive, with meaning to the actor.

- *Social Behaviour* holds no meaning to the actor. It is reactive, therefore no deliberation occurs. A behaviour becomes *social* when it is directed or oriented towards another.

- *Social Action* holds meaning to the actor, and is intentional. A rational decision is made to act in a certain way, taking into account different factors such as the actor's emotional state and the current situation. An action becomes *social* when it is directed or oriented towards another.

- *Social Interaction* requires a response to a social action from another actor.

3.2 Ideal Types of Social Action

Social action theory, as outlined by [91], defines four *'idealtypus'* of social action; these ideal types describe the motivations behind social actions in a simplified model, to aid analysis of complex human actions.

Instrumental-Rational Actions, often termed *goal-rational* actions, are those chosen for their effectiveness in achieving a *goal*, and are justifiable from the perspective of the actor. Other goals, the range of possible actions, and the consequences of performing the action are considered to decide the most appropriate action. An example is planning the actions necessary in order to reach long-term goals. Most agents, especially in machine learning, are instrumental-rational by this definition.

Value-Rational Actions are determined by the values or *beliefs* held by an actor, such that performing the action itself carries meaning rather than the result of it. Actions are rationalised in terms of *ethical* or *religious* beliefs, or to any cause valued by the actor. Rationality is a justifiable, conscious decision of how to act, and is understandable when considering the motivations of the actor; actions may seem irrational to outside observers if the motivation is not immediately clear. Pure value-rational action disregards any consequence, such as a soldier sacrificing himself for another; here, the value held outweighs the consequence of the action.

Affective Actions are reactive and impulsive actions in response to an *emotional* state or exceptional stimulus. Affective action appears inherently irrational, as the consequences of the action may not be considered and thus may be difficult to justify. An example is striking someone out of rage.

Traditional Actions are *habitual*, or those in reference to a cultural *custom*; it can therefore be seen as acting in a way that others act. These actions can be described as mindless, automatic, or ritualistic; there is no obligation to act in this way, rather the rationalisation for per-

forming the action is that 'it has always been done this way'. An example is using specific eating utensils; deliberation reduces over time as the action becomes second nature. Traditional actions therefore might not necessarily be optimal or most appropriate in terms of achieving a goal, but can be adequate means to satisfy the end.

3.3 Social Action in Computational Systems

Computational systems with actions determined by error-function-based learning or objective-function-based search are goal-rational; they are engineered to maximise their ability to achieve a particular goal. However, this paper shows that there are unintended consequences associated with the actions of goal-rational agents that coexist in a shared world. This can manifest as volatility in evolution and a loss of ability to achieve one's own goals. Human evolution has favoured social behaviour to deal with issues arising from living in the presence of others [76]; without this, humans struggle to adapt or survive [20]. We explore how computational systems may also begin to overcome these issues by operationalising other ideal types of social action, combined with the goal-rationality already present in current systems.

Social action can be operationalised in different ways. Traditional action could simply be copying what most others are doing, forming traditions over time. Value-rational action would become especially critical when systems make decisions on behalf of humans; this poses the question of how one can trust that the decisions made align with human values. Affective action could be taken when one does not know how to proceed in an unencountered situation; as such, all actions may be irrational unless abstracting previous knowledge.

This work explores the notion of traditional social action and its effect on goal-achievement in socially situated agents. It is operationalised as an action derived from a representative state of the population at specific points in time. As such, to align with the definition of 'social action' stated previously, a social action here is oriented

towards the others in the population in terms of the evolutionary process, and not the other agents that may exist within the environment. We thus operationalise two types of social action as inspired by [91] (goal-rational and traditional), and introduce 'random social action', which adds Random Immigrants to the population [19]. To this end, we can observe whether social action within the evolutionary process is a useful tool to mitigate the effect of interference without an explicit awareness of what is causing it (i.e. other agents in the environment), as a step towards socially self-aware agents – and whether doing so is quantitatively different to adding Random Immigrants.

4 The River Crossing Dilemma: A Shared-Environment Testbed for Social Agents

The experiments in this work use the River Crossing Dilemma as a testbed [4], extended from the River Crossing Task [70]; this is an extensible, gamified grid-world environment, designed to explore how interference affects learning and the ability to achieve goals in socially situated agents.

4.1 The River Crossing Task

The River Crossing Task (RCT) is a 2D grid-world problem designed to explore how agents learn to solve complex tasks with no *a priori* knowledge of the task or environment. It is therefore impossible for agents to plan a sequence of actions – they must react to dynamic elements and previously unseen environments on the fly [70]. An RCT instance comprises an $n \times n$ grid, with a river of *Water*. An agent's goal is to retrieve the *Resource* from the opposite side of the river, which provides a large positive fitness. Conversely, falling in the river provides a large negative fitness. Agents must therefore learn to achieve sub-goals such as building a bridge with *Stones* to cross the river safely to achieve their goal.

4.2 The River Crossing Dilemma

The River Crossing Dilemma (RCD) extends the original RCT to facilitate the exploration of social situations with multiple agents [4]. These may be tractable social dilemmas, such as in this paper, but in general are not constrained in their complexity, since RCD instances may be designed to be arbitrarily complex. The RCD is a 19×19 grid with a two cell deep river of *Water* (Figure 1); a bridge is successfully built if two *Stones* are placed in the same space in the river. An agent's individual goal is to collect both of its allocated *Resources* – one from either side of the river. The RCD reduces the number of objects in the environment compared to the original RCT [70]; extra objects such as *Traps* increase the complexity of the task at hand, without affecting interference. Agents therefore only encounter the objects that are sufficient to achieve their goals.
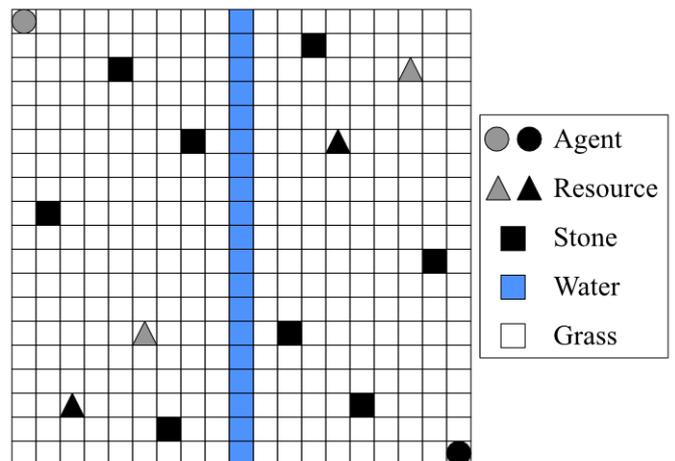

Figure 1: The River Crossing Dilemma: A Gamified Shared Environment for Studying Social Agents [4]

The RCD is gamified through introducing a cost for placing *Stones* in the river, inspired by the Snowdrift Game [56] (also known as the Chicken Game [45, 48, 86] or the Hawk-Dove Game [74]); this is a two-person social dilemma with a cost for cooperation. Gamification adds a subtle complexity to the task as agents must learn to endure a small cost for a large gain. An agent's fitness, or payoff, is calculated with Equation 1.

$$p_i = \frac{r_i}{N} - \left[\frac{C \times s_i}{2}\left(1 + s_i\right) \right] - f_i \quad (1)$$

To calculate the payoff p for agent i with Equation 1, r_i is how many *Resources* that agent i has collected, N is the total number of *Resources* to collect per agent to achieve their individual goal, C is the cost of placing a *Stone* in the river and s_i is the number of *Stones* that agent i has placed in the river. C and N are constants, with $C = 0.1$ and $N = 2$. $f_i = 1$ if the agent falls in the river, and is otherwise 0. The cost of placing *Stones* increases as more are placed, encouraging agents to exert the least effort to achieve their goal. This equation evaluates each agent's payoff individually – independent of others; this also allows agents to learn alone in the environment.

Table 1 shows a simplified payoff matrix using Equation 1. The highest payoff when alone is $p_i = 0.7$ as two *Stones* must be placed. In shared environments, agents can receive a payoff of $p_i = 1.0$ by exploiting the other, who receives $p_i = 0.7$. The overall optimal payoff is $p_i = 0.9$ when agents cooperate by sharing the cost of building a bridge.

Agent 1 \ Agent 2	$S_2 = 0$	$S_2 = 1$	$S_2 = 2$
$S_1 = 0$	0.0 / 0.0	-0.1 / 0.0	0.7 / 1.0
$S_1 = 1$	0.0 / -0.1	0.9 / 0.9	0.7 / 0.9
$S_1 = 2$	1.0 / 0.7	0.9 / 0.7	0.7 / 0.7

Table 1: Payoff Matrix, where S_i is the number of *Stones* placed by agent i

Cooperation in social dilemmas is influenced by knowing of its existence [22], and can be negatively influenced if the dilemma's characteristics are unknown or dynamic [86]; this is an inherent feature of a shared environment. Therefore, cooperation is challenging and cannot be intended when agents are unaware of the existence of others or the consequent interference. It must be noted that the focus of this paper is not the evolution of cooperation, but rather how agents can mitigate the effect of interference when socially situated; cooperation or competition can emerge, however is not intended nor understood by the agents.

4.3 Agent Design

Existing approaches to agent design for the original RCT use a two-layered neural network architecture and neuroevolution as a learning algorithm. These agents are capable of reacting to dynamic environments (such as a change in environment size or configuration) without needing *a priori* knowledge [10, 70, 78]. Our approach for RCD agent design uses a two-layered neural network architecture [4], inspired by [70].

The deliberative layer is a fully-connected neural network that generates high-level sub-goals; the input layer has six neurons, the hidden layer has four neurons, and the output layer has three neurons which represent sub-goals (Figure 2). Inputs correlate to the agent's current state and the state of its environment: *Grass*, *Resource*, *Water* or *Stone*, if it is *Carrying* a *Stone* and if a partial *Bridge* exists in the environment (i.e. one *Stone* in the river). The deliberative network outputs values that correspond to sub-goals, i.e. what to do next. These values are 1 for attraction, 0 for neutrality or -1 for avoidance, and are generated for *Resources*, *Stones* and *Water*.
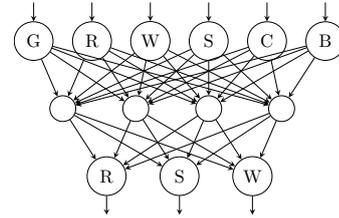

Figure 2: *Deliberative Neural Network for Generating Sub-Goals:* Inputs: Grass, Resource, Water, Stone, Carrying Status, Partial Bridge Built; Outputs: Resource, Stone, Water

The reactive layer generates a dynamic activity landscape through a topologically-organised $n \times n$ lattice of neurons, corresponding to the RCD dimensions; this is used to hill-climb towards current sub-goals. The method of activation propagation is inspired by the shunting model proposed by [94, 95], characterised by a biologically-inspired equation [36]. This is used to calculate the activity of each neuron based on its own activation and the activity surrounding it (Equation 2), for each time-step. A represents the passive decay rate, x_i is the current neuron, I is the Iota value of the neuron (corresponding to the sub-goals from the deliberative layer: for a

value of 1, $I = 15$, for a value of -1, $I = -15$ and $I = 0$ otherwise), w_{ij} is the weight between neurons x_i and x_j where x_j is one of the surrounding cells in x_i's Moore neighbourhood (indicated by $k = 8$), and $[x_j]^+$ is calculated by $max(0, x_j)$ – meaning that negative activity cannot propagate through the network. The Iota value is a large number to create hills and valleys in the activity landscape, as inspired by the original RCT testbed designed by [70]. An example of an activation landscape that is generated by the reactive layer, using Equation 2, is presented in Figure 3.

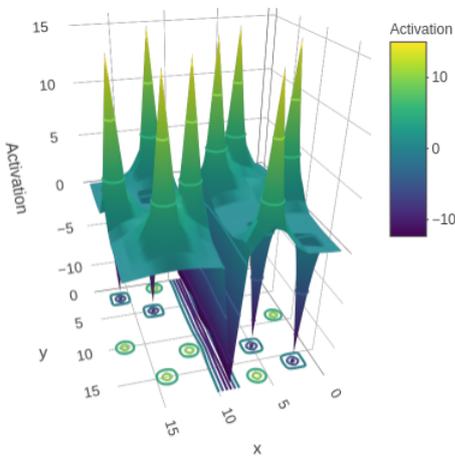$$\frac{dx_i}{dt} = -Ax_i + I_i + \sum_{j=1}^{k} w_{ij}[x_j]^+ \qquad (2)$$

Figure 3: *Activation Landscape Generated by the Reactive Layer:* The reactive layer generates dynamic activity landscapes based on the current sub-goals; here, the sub-goals are $[-1, 1, -1]$, meaning the agent is attracted to *Stones* and avoids *Resources* and *Water*. The activity landscape maps to the physical landscape (Figure 1), and agents are thus able to hill-climb towards their sub-goals whilst avoiding repulsive objects, by traversing the activity landscape and moving to the adjacent cell with the highest value.

4.4 Experimental Setup

Previous work using the RCT testbed has demonstrated that agents are able to solve tasks in *dynamic* configurations of the environment, by expressing both reactive and deliberative behaviours [70]. Whilst the experiments outlined in this paper use a static configuration of the RCD testbed depicted in Figure 1, the agent architecture used is inspired by [70], and thus agents are

expected to be capable of solving dynamic configurations of the RCD environment as well. The experiments in this paper use the following common parameters to operationalise goal-rational action, inspired by previous work [4, 10, 70].

A population of 25 randomly initialised agents was evolved with a Steady State Genetic Algorithm. At each generation, three agents from the population were randomly selected to compete in a tournament, where they were given 500 time-steps to solve the task in the RCD. The worst-performing agent in the tournament was replaced by the offspring of the winners. The offspring that was generated had a probability of $P_{one} = 0.95$ to inherit a whole chromosome (each neuron's connection weights) from a random parent, otherwise single-point crossover was used; this was repeated for each chromosome. Each weight w in the resulting offspring was then mutated by a random value from a Gaussian distribution with mean $\mu = w$ and variance $\sigma = 0.01$. This represents the goal-rationality seen in current systems. Thus, agents do not learn within their lifetime, but on an evolutionary basis.

4.5 Experimental Design

Agents were evolved for either 500,000 generations when alone, or 1,500,000 generations when socially situated. This is to give each agent an adequate amount of time to explore the search space; the effects of interference tend to be seen over a long period of time, thus socially situated agents are evolved for longer than those that evolve alone. Where agents are evolved together, two separate populations are maintained.

The first part of the study explores the effects of both goal-rational and traditional action on evolution in agents that exist either alone or in a shared environment. The first set of experiments in Section 5 explored whether ten randomly-initialised agents were able to achieve individual goals with goal-rational action; these evolved alone for 500,000 generations.

The effect of interference was then explored in a further two sets of experiments, which observed whether agents were able to continue to achieve goals either with or without continued evolution in a shared environment. In Section 6.1, the ten evolved agents from the first set of exper-

Section	Action Type	Learning Type	Agents	Generations	Runs
5	Goal-Rational	Learning Alone	1	500,000	10
6.1	Goal-Rational	No Ongoing Learning	2	500,000	10
6.2	Goal-Rational	Ongoing Learning	2	1,500,000	30
7.1	Goal-Rational & Traditional	Ongoing Learning	2	1,500,000	30

Table 2: Experiment breakdown for the initial part of the study, outlining the section number the experiments are presented in, the type of action used, the type of evolution, the number of agents in the environment, the number of generations, and the number of times the experiment was repeated.

iments were randomly arranged into ten pairs to observe the effects of interference over the course of 500,000 generations; these agents initially evolved alone for 500,000 generations (Section 5), and continue to act in a shared environment for a further 500,000 without further evolution. Section 6.2 then arranged the same ten agents into 30 random pairs; these pairs were evolved together in shared environments for 1,500,000 generations after their initial period of evolving alone, to observe whether goal-achieving behaviour can be maintained despite interference.

Section 7.1 compares goal-rational action (Section 6.2) with the introduction of traditional action in the same 30 pairs of agents. These were evolved for 1,500,000 generations after the initial period of evolving alone. The design for the experiments thus far is summarised in Table 2.

Traditional action is then explored on a broader scale in Section 7.3; 100 goal-rational agents were initialised with random weights, and were evolved with and without traditional action in both individual and social environments. To ensure the results in shared environments were not biased against a particular agent, agents were assigned a randomly-generated partner. The agents were evolved for 500,000 generations if they were alone, and 1,500,000 generations if socially situated.

Section 8 finally ascertains whether traditional action is quantitatively different to the introduction of Random Immigrants [19], as adding randomness is a common way to improve population diversity during evolution. The same 100 agents were evolved with random action in both individual and social environments. Finally, we analyse the trade-off between fitness and volatility observed for the three types of social action.

Sec	Action Type	Agents	Gens	Runs
7.3	Goal-Rational	1 (1G)	500000	100
7.3	Goal-Rational	2 (2G)	1500000	100
7.3	Goal-Rational & Traditional	1 (1GT)	500000	100
7.3	Goal-Rational & Traditional	2 (2GT)	1500000	100
8	Goal-Rational & Random	1 (1GR)	500000	100
8	Goal-Rational & Random	2 (2GR)	1500000	100

Table 3: Experiment breakdown for the final part of the study, outlining the section number the experiments are presented in, the type of action used, the number of agents in the environment, the number of generations, and the number of times the experiment was repeated. Each experiment starts with the same 100 randomly initialised agents.

The design for these experiments is summarised in Table 3.

4.6 Operationalising Traditional and Random Action

The notion of social action has been operationalised in this work within the evolutionary process. Here, a social action is oriented towards others in the population in terms of the evolutionary algorithm, rather than other agents that may exist in the environment. Specifically, the offspring produced at each generation is dependent on the type of action used.

Traditional action is operationalised as follows; at each generation, there is a 90% chance for the worst-performer of each tournament to be replaced by the current goal-rational offspring of the best two parents, and a 10% chance for it to be replaced with an offspring that is a representative state of the population. This is captured with the calculated median of all weights in the deliberative layer across all agents in the population; traditions can potentially be established that fluctuate as the population evolves.

Random action is operationalised similarly to traditional action. At each generation, there

10

is a 10% chance for the worst-performer of the tournament to be replaced with an agent with a randomly-initialised set of weights, and a 90% chance for it to be replaced with the standard goal-rational offspring as defined in Section 4.4.

The replacement probability of 10% in the case of both traditional and random action is in line with the method used by [4]. Preliminary investigation of varying replacement probabilities showed that little effect or benefit was seen with a probability lower than 10%; forming and maintaining traditions is more difficult with lower replacement probabilities, as solutions have little time to influence the population. Further, higher probabilities such as 20% or 30% showed that the population became saturated and diversity was reduced, meaning that it was increasingly difficult for evolution to explore the fitness landscape. Especially in the case of random action, higher probabilities of replacement start to resemble random search.

4.7 Metrics and Statistical Tests

Three metrics have been devised to analyse the evolution, fitness and learning in agents using the three types of action; each metric captures different knowledge, and therefore complements the others.

Historical volatility is a common metric used in financial modelling and volatility forecasting that captures the dispersion of values over time, calculated most commonly by the sample standard deviation over a defined time period [67]. A higher volatility here means that there is a higher variability and dispersion of the values [92]. This is a useful metric to determine the expected volatility in fitness over time across all 100 agents in Sections 7.3 and 8; to differentiate this from the other volatility metrics used, we term this the Standard Deviation over Time (SDoT).

The Cumulative Absolute Change over Time (CACoT) metric captures the magnitude of the changes in fitness across all 100 agents; the count is incremented by the absolute change in fitness between generation g_{i-1} and g_i. A high CACoT therefore indicates that fitness fluctuates by large amounts.

The Count of Change over Time (CCoT) metric captures how often the fitness changes

across all 100 agents; this increments by one only if the fitness in generation g_i is not equal to g_{i-1}. A high CCoT indicates frequent fluctuations in fitness, without capturing the magnitude. Both CACoT and CCoT therefore capture knowledge gain and loss (fitness increase/decrease).

According to [96], the Shapiro-Wilk normality test is powerful for distributions that are suspected to be asymmetric, as well as those that are symmetric with both high and low kurtosis values; for this reason, this test was chosen over other normality and goodness-of-fit tests such as the Cramr-von-Mises and chi-squared tests. As such, a Shapiro-Wilk normality test was conducted for each of the SDoT, CACoT and CCoT distributions for each experiment in Sections 7.3 and 8. Results were significant for each with a p-value below 0.05, indicating that the distributions are non-normal. Wilcoxon Signed Rank tests were therefore used to compare the effect each action type had on evolution. These are non-parametric tests used to compare the medians of two paired distributions; the null hypothesis for two-tailed tests is that the medians are identical, whereas one-tailed tests compare the directional difference in the distribution means. One two-sided ($x \neq y$) and two one-sided tests ($x > y$, and $x < y$) were conducted. The four statistical moments (mean, variance, skewness and kurtosis) and the median were also calculated for each distribution for further comparison.

5 Learning Alone with Goal-Rationality

Agents do not experience interference when they are in an environment alone. As such, they can achieve individual goals with independent asocial learning with goal-rational action. Figure 4 depicts the fitness of a single agent during evolution, showing that the agent can learn to achieve its goal.

5.1 Results

Whilst not impossible to achieve, this task initially appears difficult to solve simply because
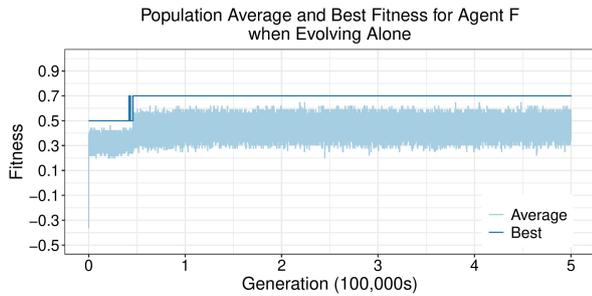
Figure 4: Agent F can achieve its goals when alone in an environment; it initially learns to collect one *Resource* to get a fitness of 0.5, then to build a bridge to achieve its goal around generation 50,000, giving it a fitness of 0.7
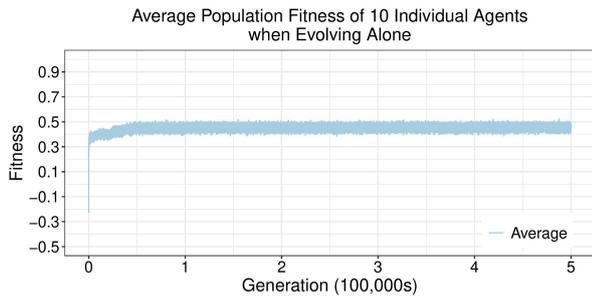

Figure 5: *Average Population Fitness of 10 Agents Evolving Alone:* All ten agents that evolved alone sustained the behaviours necessary to achieve their goal by generation 50,000

the fitness function does not 'lead' agents towards their goals with incremental rewards; agents encounter a very large, neutral network landscape during evolution as a result. In each experiment, goal-achieving behaviours were maintained once learnt. Figure 5 shows the average fitness of ten agent populations during evolution. Random mutations during breeding periodically create agents with lower fitnesses, which may fall in the river for example, thus reducing the fitness average. These ten, individually evolved agents are henceforth labelled Agents A through J.

5.2 Implications

Goal-rational action enables agents to evolve and maintain goal-achieving behaviour when they evolve alone; this can therefore be compared to evolving socially situated agents, to observe the effect of interference on learning.

6 Coexistence and Volatility

Interference from the actions of others causes learnt knowledge to become unreliable, resulting in goal-rational agents behaving differently. The ten agents in the previous section are then observed in shared environments both without (Section 6.1) and with (Section 6.2) continued evolution.

6.1 Results: Coexistence Without Ongoing Evolution

Three emergent dynamics are observed in these non-evolutionary experiments: one agent exploits the other for a higher payoff (Figure 6a); both agents co-exist and achieve their goals similarly to when they are alone (Figure 6b); one or both agents cannot achieve their goals (Figure 6c). Interference can change the world view of the agents, and consequently change their expected behaviour and fitness. As the results are dependent on the interactions with a specific partner, graphing an average of all ten experiments would mask the specific interactions. A representative sample of three experiments showing average population fitnesses are therefore shown in Figure 6.

The fitnesses fluctuate based on the actions of both individuals at every generation, and as such appear very volatile. Emergent exploitation has no long-term implication here as agents do not learn; agents would continue to achieve their goals as before if they were again alone in an environment. The critical observation here is that interference causes the agents' world view to change in unanticipated ways, thus affecting their ability to achieve goals.

6.2 Results: Coexistence with Ongoing Evolution

Interference can affect learnt and reliable knowledge, and thus an agent's ability to solve tasks independently – even when interaction is not direct. The most commonly observed emergent dynamic when agents continue to evolve together was exploitation: an agent would receive a higher
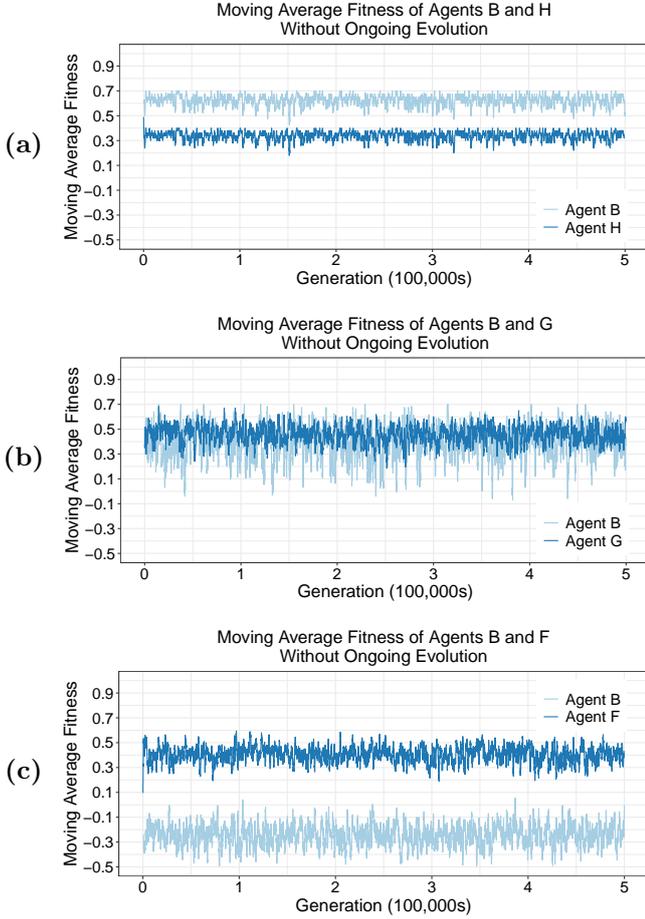
(a) Best in Population Fitness of Agents B and F
with Ongoing Evolution

(b) Best in Population Fitness of Agents B and F
with Ongoing Evolution and Traditional Action

Figure 7: Agents B and F continuing to evolve together (a) without and (b) with traditional action. (a) Without traditional action, Agent B exploits Agent F and cannot achieve its goals alone; evolution is volatile as agents evolve to be codependent. Cooperation gives 0.9 fitness. If fitness < 0.7, agents do not achieve their goals. (b) With traditional action, the agents endure a smaller period of low fitness, and get a better fitness of 0.9 by cooperating.

(a) Moving Average Fitness of Agents B and H
Without Ongoing Evolution

(b) Moving Average Fitness of Agents B and G
Without Ongoing Evolution

(c) Moving Average Fitness of Agents B and F
Without Ongoing Evolution

Figure 6: The moving average fitness of Agent B and Agents (a) H, (b) G and (c) F, without ongoing evolution. When socially situated with different pairs, Agent B (a) exploits Agent H to receive a higher average fitness from not exerting as much effort, (b) is more unpredictable than when alone, and is often unable to achieve its goal, and (c) cannot achieve its goal. In each experiment, Agent B's pair performs similarly to when it is alone.
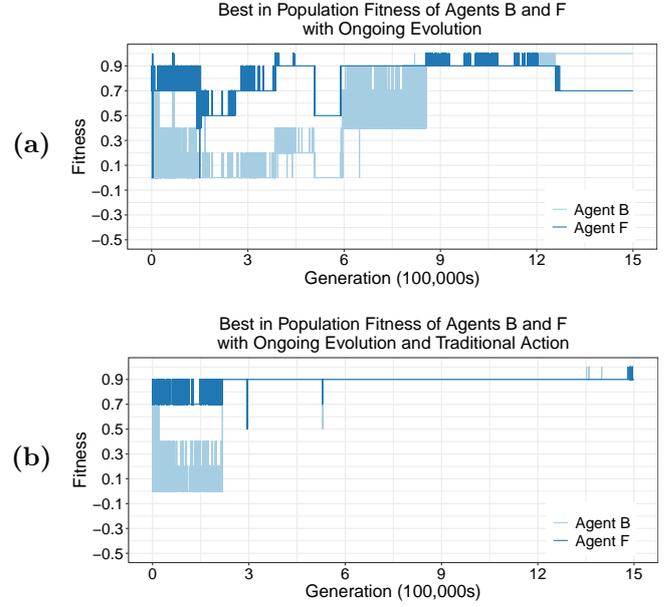
Resources \ Stones	$S = 0$	$S = 1$	$S = 2$
$R = 0$	0.0	-0.1	-0.3
$R = 1$	0.5	0.4	0.2
$R = 2$	1.0	0.9	0.7

Table 4: Commonly observed agent behaviours and the resulting fitnesses

payoff at another's expense. No pair of agents were observed to both maintain their ability to achieve goals alone – at least one agent was observed to change its learnt behaviour. Figures 7a to 11a depict the best fitness in each population for the named agents, to show that learnt knowledge can be lost with coevolution. Table 4 shows common fitnesses and their associated behaviours.

Figure 7a shows the most extreme case of interference observed; agents evolved to be codependent, as what agents learn depends on the actions of the other. Both endure periods of being unable to achieve their goals (a fitness below 0.7), with Agent B being more negatively affected than Agent F. The brief period of cooperation is volatile as Agent F is unreliable; Agent

13

B thus evolves to rely on itself to achieve its goals. The highest fitness for Agent B continuously fluctuates between 0.9 (cooperation, placing one *Stone* and collecting two *Resources*) and 0.4 (placing one *Stone* but only collecting one *Resource*); this volatility leads to exploitation. One can thus postulate that the effect of interference can be great, complex and uncertain as the world changes in unanticipated ways.

Periodic dips in fitness lead to exploitation (Figure 8a). Exploitation in early evolution leads to a mutual loss of fitness and inability to achieve goals, which evolves into an exploitative relationship. The spikes observed in Agent D's fitness indicate that occasionally Agent F will cooperate, but is unreliable and predictably self-interested. As a result, Agent D sustains its independent behaviour, whilst Agent F evolves to capitalise on this.

Figure 9a shows an initial period of cooperation between Agents D and G. Agent G evolves to exploit, and Agent D evolves to be independent as the actions of Agent G are not predictable enough to make cooperation beneficial.

Figures 10a and 11a show exploitative relationships; the remainder of the 30 experiments are similar. As with Figure 8a, peaks in the exploited agent's fitness indicate that the exploitative agent occasionally helps to build a bridge; it however evolves to be independent as the other agent is not predictable enough to rely on. This, again, evidences that interference affects how and what behaviours are evolved and maintained during evolution.

6.3 Implications

Learnt knowledge becomes unreliable when agents pursue individual goals with goal-rational action in shared environments. Interference from the actions of others changes each agent's perception of the world, and often leads to volatile evolution and knowledge loss; if agents cannot perceive the cause of the environmental changes, they will attempt to adapt and often lose learnt knowledge as a result. Exploitative agents for example depend on the other to achieve their goals, and thus 'forget' how to interact with *Stones* over time; this would be detrimental if they became suddenly alone in an environment, as they have
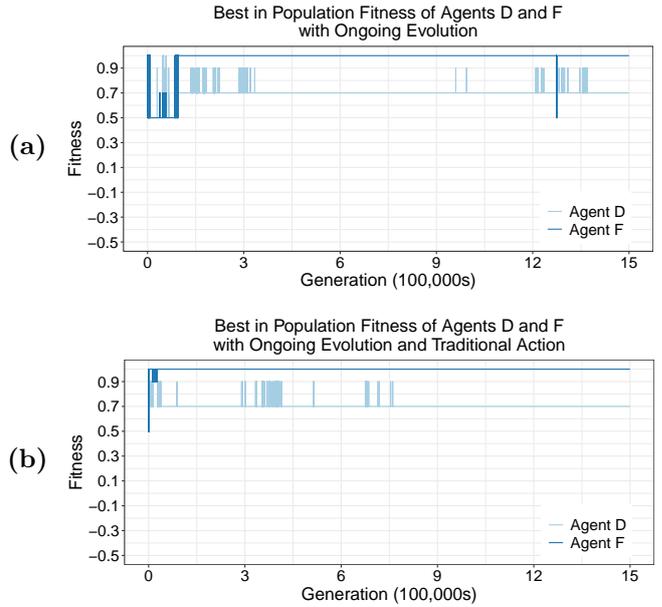

Figure 8: Agents D and F continuing to evolve together (a) without and (b) with traditional action. (a) Without traditional action, both agents initially lose their ability to achieve goals, with a fitness of 0.5; Agent F then evolves to exploit Agent D. (b) With traditional action, the agents endure less knowledge loss with traditional action and maintain their exploitative relationship in less time.

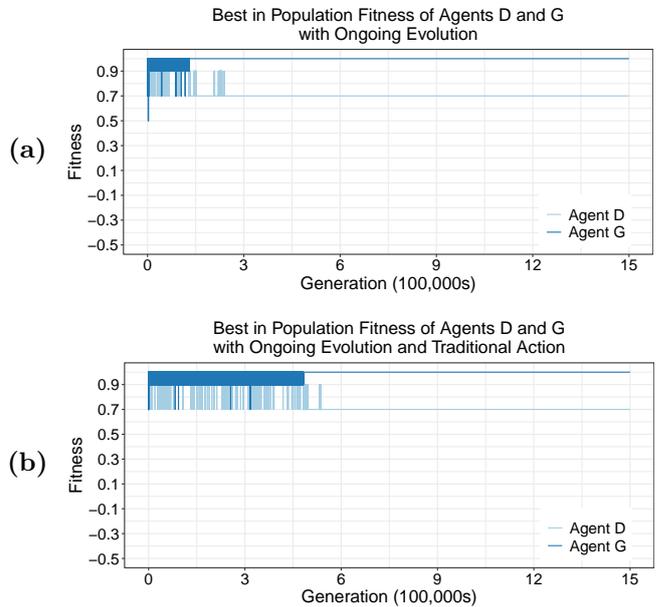**Figure 9:** Agents D and G continuing to evolve together (a) without and (b) with traditional action. (a) Without traditional action, Agent G cannot always rely on Agent D to exploit it, so it sometimes cooperates; Agent D cannot rely on G to cooperate, so evolves to achieve its goals alone. (b) With traditional action, the agents are able to cooperate and receive a higher fitness for longer.
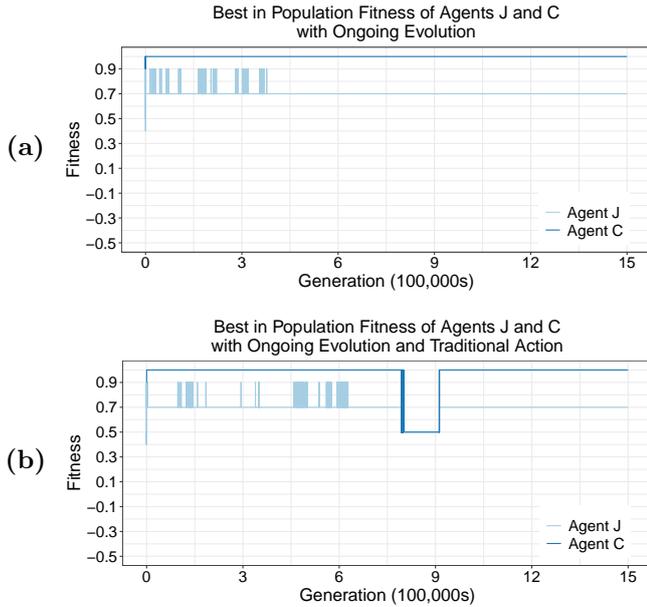
Figure 10: Agents J and C continuing to evolve together (a) without and (b) with traditional action. (a) Without traditional action, Agent C exploits Agent J, which evolves to achieve goals alone. (b) With traditional action, the agents endure low fitness temporarily with traditional action, but overcome this before maintaining an exploitative relationship.
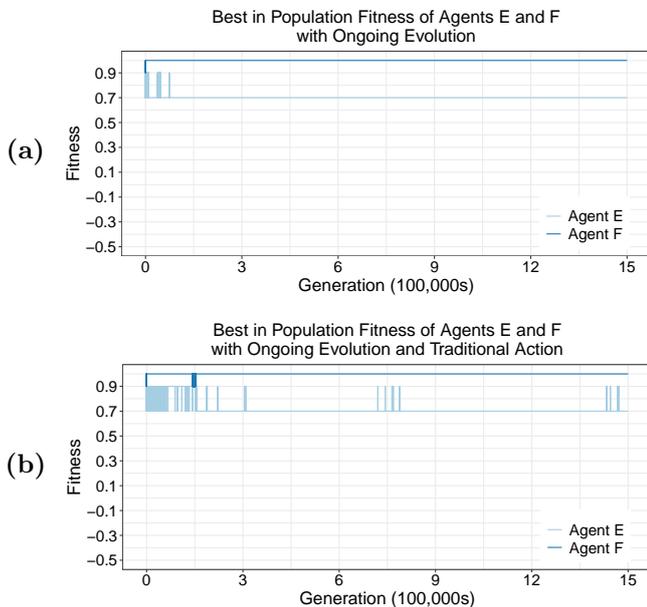

Figure 11: Agents E and F continuing to evolve together (a) without and (b) with traditional action. (a) Without traditional action, Agent F exploits Agent E, which evolves to achieve goals alone. (b) With traditional action, Agent F still evolves to exploit Agent E.

evolved to be codependent rather than independent.

Further, the unanticipated and imperceptible actions of the other agent change the state of the environment, and can cause agents to alter their knowledge and behaviour in an attempt to adapt to the change in circumstances. This can lead to a change in the mapping of inputs to sub-goals in the deliberative layer, and may result in drastic changes of behaviour. Specific examples could be an agent suddenly falling into the river due to an inability to cope with the environmental changes appropriately, or an agent that simply places many *Stones* into the river as its sub-goals have been altered to be repulsed from *Resources*. Interference can therefore impair an agent's ability to achieve its goal, even though it was once capable of being successful. An inability to overcome unexpected situations and unanticipated changes means agents are susceptible to knowledge loss and changes in behaviour. We operationalise traditional action in an attempt to mitigate the effect of interference in goal-rational, socially situated agents.

7 Traditional Action Can Promote Less Volatile Learning

Motivated by the behavioural changes observed in socially situated agents, traditional action was compared to goal-rational action to explore how coevolution and interference were affected. As a result, socially situated agents evolved in a more predictable and less volatile manner.

7.1 Results

Figures 7a and 7b depict the same pair of goal-rational agents with ongoing evolution, without and with traditional action respectively. The volatility seen in Figure 7a is drastically reduced with traditional action (Figure 7b). Further, the period of knowledge loss endured by Agent B as it learns about its new environment is much shorter; additionally, both agents receive a mutual benefit from the presence of the other such that they evolve and maintain emergent cooper-

ative behaviours. This enables both agents to achieve their goals with the best overall payoff.

Figures 8a and 8b also present the same pair of agents. Traditional action significantly reduces the initial period that agents are unable to achieve their goals; instead of a dip in fitness, agents briefly cooperate instead. Agent F then evolves to exploit Agent D, which is observed both with and without traditional action. Traditional action is beneficial overall for these agents.

Figures 9a and 9b demonstrate that cooperation can exist with exploitation. Fluctuations between fitnesses of 0.9 and 1.0 in Agent G indicate that it cannot rely on the other agent, and therefore must still maintain knowledge of how to build a bridge. Agent D is more independent, and fluctuates between a fitness of 0.7 and 0.9; Agent D cannot always rely on Agent G to cooperate, so maintains its ability to achieve its goal alone. Cooperation is maintained for longer with traditional action, so agents are therefore generally much better off.

Figures 10a and 10b show exploitative behaviour; however, employing traditional action causes a period of around 100,000 generations where both agents lose fitness from losing the ability to build bridges, thus only collecting one *Resource* each with a payoff of 0.5. Agent C evolves to exploit Agent J; when its behaviour changes such that it cannot achieve its goal, neither can Agent C. This anomalous evolutionary event causes a temporary dip in fitness that is not observed without traditional action. Furthermore, traditional action enables the agents to overcome this dip in fitness, indicating that agents can potentially learn how to cope with unexpected events.

Figures 11a and 11b show that traditional action can have a minimal effect on evolution. In the remaining experiments, if pure goal-rational agents evolve to be exploitative, traditional agents will as well. Exploitative agents essentially unlearn their previously dependable knowledge; this change in behaviour is simply caused by the interference that arises from coexistence.

7.2 Implications

Traditional action can reduce the effect of interference on agent evolution and goal-achievement. Whilst intentional cooperation requires an explicit awareness of others [15], *unintentional*, emergent cooperation can be observed when agents receive a mutual benefit from the actions of the other; this enables them to pursue their own goals and achieve a better fitness overall. Additionally, coexisting agents are observed to evolve and maintain their expected behaviours faster with traditional action, allowing them to recover from unforeseen events or states.

7.3 Analysing the Effect of Traditional Action on Interference

The effect of interference can be explored on a broader scale. As outlined in Section 4.5, 100 agent populations were initialised with random weights, used to explore how goal-rational agents evolved both with and without traditional action, in individual and social environments.

Over the 100 experiments, agents using traditional action were observed to have the same or a lower median SDoT, CACoT and CCoT than their purely goal-rational counterparts, which can be seen in Figure 12. Agents using traditional action thus typically have a lower volatility in evolution than purely goal-rational agents. The difference in volatility measured by the three metrics is much larger in socially situated agents than agents that are alone in an environment; fitness with or without traditional action is not significantly different when agents do not experience interference.

Agents using traditional action typically have an SDoT with less variability than those that do not (Table 5); this is more evident when agents exist alone, than when they are socially situated. Combined with a lower expected SDoT, and more skew and excess positive kurtosis ($k_{excess} = k - 3$), agents using traditional action are expected to have a lower and more predictable SDoT. Kurtosis is used to analyse the tailedness and probability of outliers.

When alone, agents that use traditional action have a CACoT and CCoT with a higher variability and kurtosis, as well as a lower expected

Moment	Alone			Together		
	G	GT	GR	G	GT	GR
mean	0.01545549	0.01254882	0.04222212	0.05888082	0.04236362	0.03868761
variance	0.001083114	0.0008872677	0.0003943957	0.006211676	0.00474367	0.000329736
skewness	1.835432	2.166065	0.1707893	1.286178	1.752168	0.329795
kurtosis	4.562028	5.946412	2.740262	3.351055	5.035609	3.637307
median	0.0	0.0	0.041909	0.010199	0.003317	0.040580

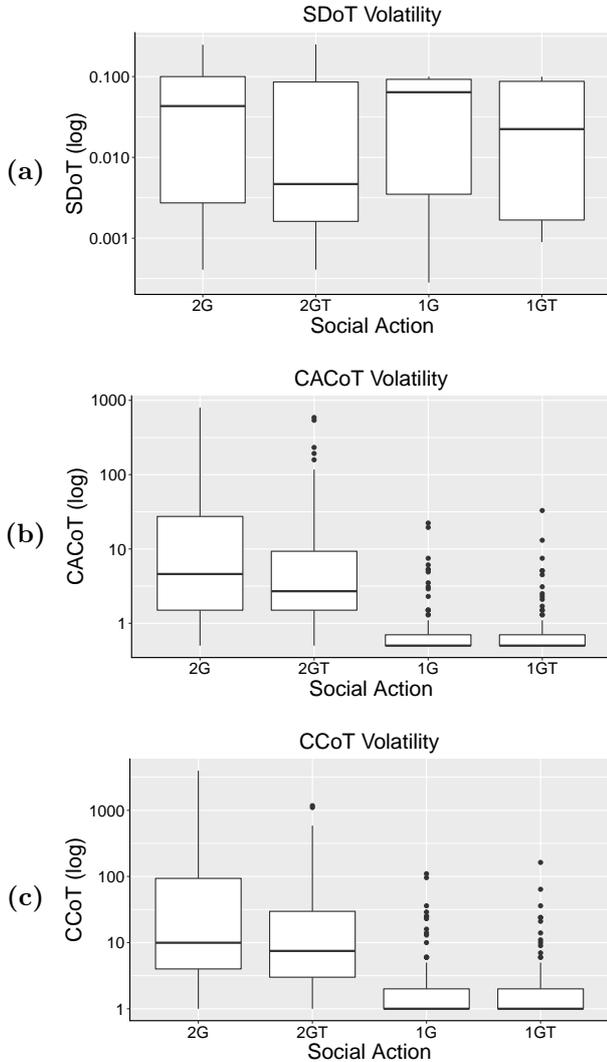**Table 5:** SDoT Statistical Moments and Median

Figure 12: Log-scale box plots of the (a) SDoT, (b) CA-CoT and (c) CCoT, of agents using goal-rational or traditional action, evolved together or alone (2G, 2GT, 1G and 1GT respectively). All points above $(Q_3 + 1.5IQR)$ are classed as mild outliers, whereas those above $(Q_3 + 3IQR)$ are extreme outliers; in all cases, there were no outliers below $(Q_1 - 1.5IQR)$. All data points – including all outliers, indicated with a black circle – are used in the analysis. Note that the boxplots are on a log scale. Traditional action makes evolution less volatile than goal-rational action when agents are together. Volatility is similar when agents are alone, so the type of action used makes little difference.

CACoT and CCoT than those that do not (Tables 6 and 7). This changes when agents are socially situated, with traditional action reducing the variability and kurtosis in CACoT and CCoT. This shows that traditional action can reduce volatility in evolution when agents experience interference. Agents that use traditional action typically have a lower CACoT than those that do not, indicated by a greater skewness; this is the same for the CCoT in agents that use traditional action when alone. The expected CCoT and skewness of the CCoT are lower in socially situated agents that use traditional action than those that do not, meaning fluctuations in fitness are typically reduced.

It is not surprising that the type of action used does not significantly affect the expected SDoT, CACoT and CCoT when agents are alone; they do not experience interference and thus the environment is more predictable (Table 8). However, when agents are socially situated and experience interference, goal-rational agents have a significantly higher SDoT and CCoT; traditional action in socially situated agents therefore reduces volatility in evolution compared to goal-rational action alone.

Table 9 presents further information about the SDoT, CACoT and CCoT of each experiment, including the minimum and maximum value as well as the quartiles. Table 10 then presents an analysis of the outliers in the data, including the median and mean of the outliers. The median and mean of the outliers are generally higher in agents using goal-rational action – both when agents act alone and when together – compared to those that also use traditional action. This is especially evident for the CACoT and CCoT for when agents act together; saying this, the median of the outliers of the CACoT and CCoT is marginally lower when agents use goal-rational action in an environment alone.

Moment	Alone			Together		
	G	GT	GR	G	GT	GR
mean	1.374	1.33	2.276	38.348	27.173	1001.674
variance	9.46962	12.8302	2.365479	10840.82	7419.292	1694746
skewness	5.430996	7.385912	1.433599	5.066438	5.186799	1.073078
kurtosis	34.39726	62.88242	5.654775	32.85393	31.60225	2.824496
median	0.5	0.5	1.9	4.6	2.7	85.1

Table 6: CACoT Statistical Moments and Median

Moment	Alone			Together		
	G	GT	GR	G	GT	GR
mean	5.37	5.15	9.88	150.08	64.44	4556.34
variance	236.7405	320.7551	59.13697	215319.4	33005.91	37515029
skewness	5.430996	7.385912	1.433599	6.216094	4.630009	1.108448
kurtosis	34.39726	62.88242	5.654775	48.31377	26.21253	2.847056
median	1.0	1.0	8.0	10.0	7.5	187.5

Table 7: CCoT Statistical Moments

Test	SDoT		CACoT		CCoT	
	Alone	Together	Alone	Together	Alone	Together
$G \neq GT$	0.4282	**0.04565**	0.7669	0.192	0.7426	0.0978
$G < GT$	0.7885	0.9774	0.62	0.9046	0.6322	0.9515
$G > GT$	0.2141	**0.02283**	0.3835	0.09601	0.3713	**0.0489**

Table 8: Wilcoxon Signed Rank Statistical Tests Comparing Goal-Rational with Goal-Rational and Traditional Action (G-GT). Results are p-values; significant results ($p < .05$) are **bold**

	Metric	Min	Q_1	Median	Mean	Q_3	Max	IQR
SDoT	2G	0	0.001122	0.010199	0.058881	0.088579	0.24805	0.087457
	2GT	0	0.00108	0.003317	0.042364	0.069555	0.24989	0.068475
	1G	0	0	0	0.015455	0.00301	0.099989	0.00301
	1GT	0	0	0	0.012549	0.001523	0.099676	0.001523
CACoT	2G	0.5	1.5	4.6	38.348	27.45	798.9	25.95
	2GT	0.5	1.5	2.7	27.173	9.3	588	7.8
	1G	0.5	0.5	0.5	1.374	0.7	22.3	0.2
	1GT	0.5	0.5	0.5	1.33	0.7	32.9	0.2
CCoT	2G	1	4	10	150.08	94	3967	90
	2GT	1	3	7.5	64.44	29.75	1176	26.75
	1G	1	1	1	5.37	2	110	1
	1GT	1	1	1	5.15	2	163	1

Table 9: The minimum, maximum, first and third quartiles, median, mean and interquartile range (IQR) for the SDoT, CACoT and CCoT of each experiment (rounded to 5SF).

Metric		Mild	Extreme	Total	Median of Outliers	Mean of Outliers
SDoT	2G	8	0	8	0.24229	0.24503
	2GT	10	0	10	0.21766	0.21368
	1G	0	18	18	0.083216	0.091756
	1GT	3	16	19	0.065154	0.080340
CACoT	2G	4	8	12	238.46	171.4
	2GT	2	14	16	149.26	93.85
	1G	5	15	20	4.73	2.6
	1GT	4	13	17	5.1588	2.3
CCoT	2G	3	10	13	936.54	643
	2GT	3	12	15	365.73	236
	1G	3	17	20	22.15	11.5
	1GT	2	15	17	24.294	10

Table 10: An analysis of the outliers for the SDoT, CACoT and CCoT of each experiment (rounded to 5SF). Mild outliers are above $(Q_3 + 1.5IQR)$, whereas those above $(Q_3 + 3IQR)$ are extreme outliers. In all cases, there were no outliers below $(Q_1 - 1.5IQR)$.

8 Analysing the Trade-Off Between Fitness and Volatility

As outlined in Section 4.5, this section ascertains the difference between random and traditional action when combined with goal-rational action, and analyses the trade-off between fitness and volatility.

8.1 Random Action: High Expected Fitness, High Volatility

Goal-rational agents that use random action have a higher expected SDoT, CACoT and CCoT than those that use the other two types of action; the exception to this is that the expected SDoT is lower when agents experience interference with random action, despite the median still being larger than the other types of action (Table 5). Additionally, agents that use random action have an SDoT with lower variability, and those that act alone have a CACoT and CCoT with lower variability than those that use goal-rational or traditional action. Therefore, when agents do not experience interference, using random action means the SDoT, CACoT and CCoT will be predictably high due to low dispersion around the mean. This changes drastically when agents are socially situated agents (Tables 6 and 7), as interference can cause fitness in agents using random action to fluctuate often, indicated by a CACoT

and CCoT with very high variability.

In all experiments, the SDoT, CACoT and CCoT are positively skewed, with random action being the least skewed; this means that the CACoT and CCoT are typically higher than the other types of action, as the values are more symmetrical around the higher means (Tables 6 and 7). It can therefore be said that evolution with random action is more volatile than goal-rational or traditional action.

The SDoT, CACoT and CCoT in both goal-rational and traditional action are highly leptokurtic, meaning there is positive excess kurtosis (Tables 5, 6 and 7). The SDoT, CACoT and CCoT are expected to be less extreme in agents using random action, but are generally higher overall.

Agents that use random action have a significantly higher expected CACoT and CCoT than those that use the other two types of action, as well as a significantly higher expected SDoT when agents are alone ($p < .05$, Tables 11 and 12). This indicates that the evolution of agents using random action is significantly different to the other types of action, and agents are affected significantly more by interference.

To summarise, goal-rational agents that use random action typically have a significantly higher SDoT, CACoT and CCoT during evolution than those that use traditional or pure goal-rational action; this is due to a higher mean and median, combined with less skew and kurtosis. There is not enough evidence to say that the SDoT is significantly greater when agents experience interference with random action, how-

19

Test	SDoT		CACoT		CCoT	
	Alone	Together	Alone	Together	Alone	Together
$G \neq GR$	**1.415E-08**	0.4993	**6.814E-09**	**6.353E-12**	**6.146E-09**	**5.557E-11**
$G < GR$	**7.077E-09**	0.7515	**3.407E-09**	**3.177E-12**	**3.073E-09**	**2.778E-11**
$G > GR$	1.0	0.2496	1.0	1.0	1.0	1.0

Table 11: Wilcoxon Signed Rank Statistical Tests Comparing Goal-Rational with Goal-Rational and Random Action (G-GR). Results are p-values; significant results ($p < .05$) are **bold**

Test	SDoT		CACoT		CCoT	
	Alone	Together	Alone	Together	Alone	Together
$GT \neq GR$	**2.643E-11**	0.1816	**9.769E-11**	**1.797E-15**	**9.594E-11**	**1.473E-14**
$GT < GR$	**1.322E-11**	0.09081	**4.884E-11**	**8.986E-16**	**4.797E-11**	**7.363E-15**
$GT > GR$	1.0	0.9098	1.0	1.0	1.0	1.0

Table 12: Wilcoxon Signed Rank Statistical Tests Comparing Goal-Rational and Traditional with Goal-Rational and Random Action (GT-GR). Results are p-values; significant results ($p < .05$) are **bold**

ever both the CACoT and CCoT are significantly higher. When agents exist alone, the SDoT, CACoT and CCoT have less variability, but in social environments, the variance is extremely large. This indicates that evolving and retaining knowledge is much more volatile when agents experience interference. It can therefore be concluded that evolution with random action is extremely volatile compared to the other types of actions implemented, especially in social environments.

8.2 Balancing the Fitness and Volatility Trade-Off

The expected fitness after evolution between goal-rational agents that use traditional action and those that do not is not significantly different (Tables 13 and 14). Goal-rational agents using random action however evolve to have a significantly higher expected fitness than the other two types of action. All 100 agents achieve a goal-achieving fitness; a median fitness of 1.0 indicates socially situated agents are expected to evolve exploitative behaviour, and thus lose the knowledge and ability to achieve their goals on their own. Although a higher expected fitness is desirable, dependency on others would not be desirable if the existence of others is unpredictable, or the environment is extremely dynamic.

There appears to be a trade-off between expected fitness and the volatility in evolution in socially situated agents. The addition of traditional action significantly reduces the volatil-

ity in evolution in otherwise purely goal-rational agents, without sacrificing fitness. The addition of random action however significantly increases the expected fitness in otherwise purely goal-rational agents, but this is at the cost of a significantly increased volatility in evolution; this can be seen in the kernel density estimations for the fitness received in each experiment, presented in Figure 13, and for the SDoT, CACoT and CCoT of each experiment, presented in Figure 14. It can therefore be said that incorporating traditional action into otherwise goal-rational agents is a logical way to achieve high fitness and low volatility in evolution, in both agents that are socially situated and those that are alone.

9 Conclusions, Discussion and Direction

We have used the River Crossing Dilemma testbed introduced by [4] to analyse the effects that various types of social action have on the way interference is experienced by socially situated agents. Our results demonstrate that endowing goal-rational agents with the capability of traditional action reduces the volatility in the evolutionary process caused by interference. Conversely, when agents evolve alone, no significant difference in fitness has been found between purely goal-rational agents and those that also use traditional action. The main contributions therefore demonstrate that traditional

Moment	Alone			Together		
	G	GT	GR	G	GT	GR
mean	0.546	0.534	0.7	0.676	0.669	0.87
variance	0.007155556	0.00570101	0	0.0416404	0.04397879	0.02151515
skewness	1.283171	1.757035	NaN	0.6358369	0.7252645	-0.2817959
kurtosis	2.646527	4.087172	NaN	1.783025	1.830542	1.113205
median	0.5	0.5	0.7	0.6	0.5	1

Table 13: Fitness Statistical Moments and Median

Test	G-GT		G-GR		GT-GR	
	Alone	Together	Alone	Together	Alone	Together
$x \neq y$	0.3217	0.8379	**2.2e-16**	**2.453e-09**	**2.2e-16**	**4.243e-09**
$x < y$	0.8435	0.5836	**2.2e-16**	**1.226e-09**	**2.2e-16**	**2.122e-09**
$x > y$	0.1608	0.419	1	1	1	1

Table 14: Wilcoxon Signed Rank Statistical Tests Comparing Fitness of Goal-Rational, Goal-Rational and Traditional, and Goal-Rational and Random Action. Results are p-values; significant results ($p < .05$) are **bold**
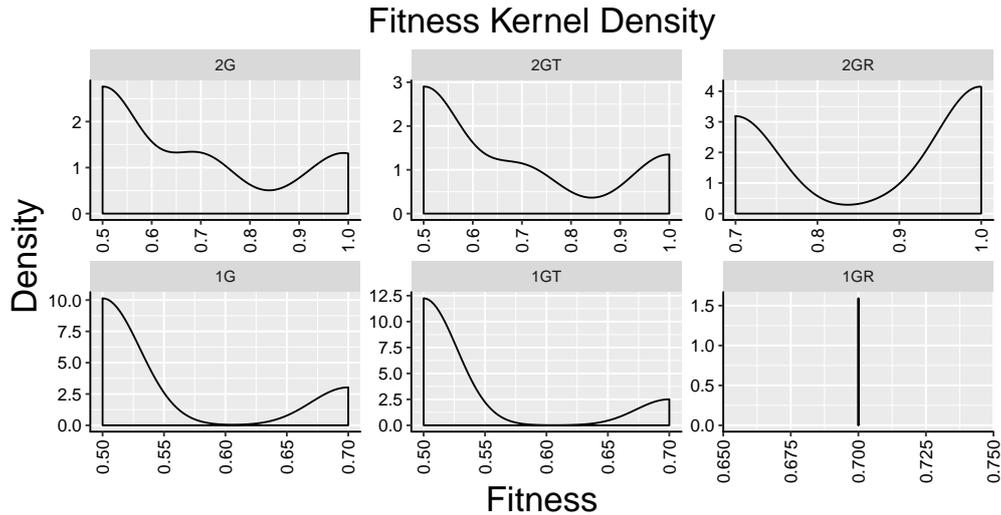

Figure 13: Kernel density estimation of the fitness of agents using goal-rational, traditional and random action, evolved together or alone (*2G*, *2GT*, *2GR*, *1G*, *1GT* and *1GR* respectively). Note that the plot scales are not comparable
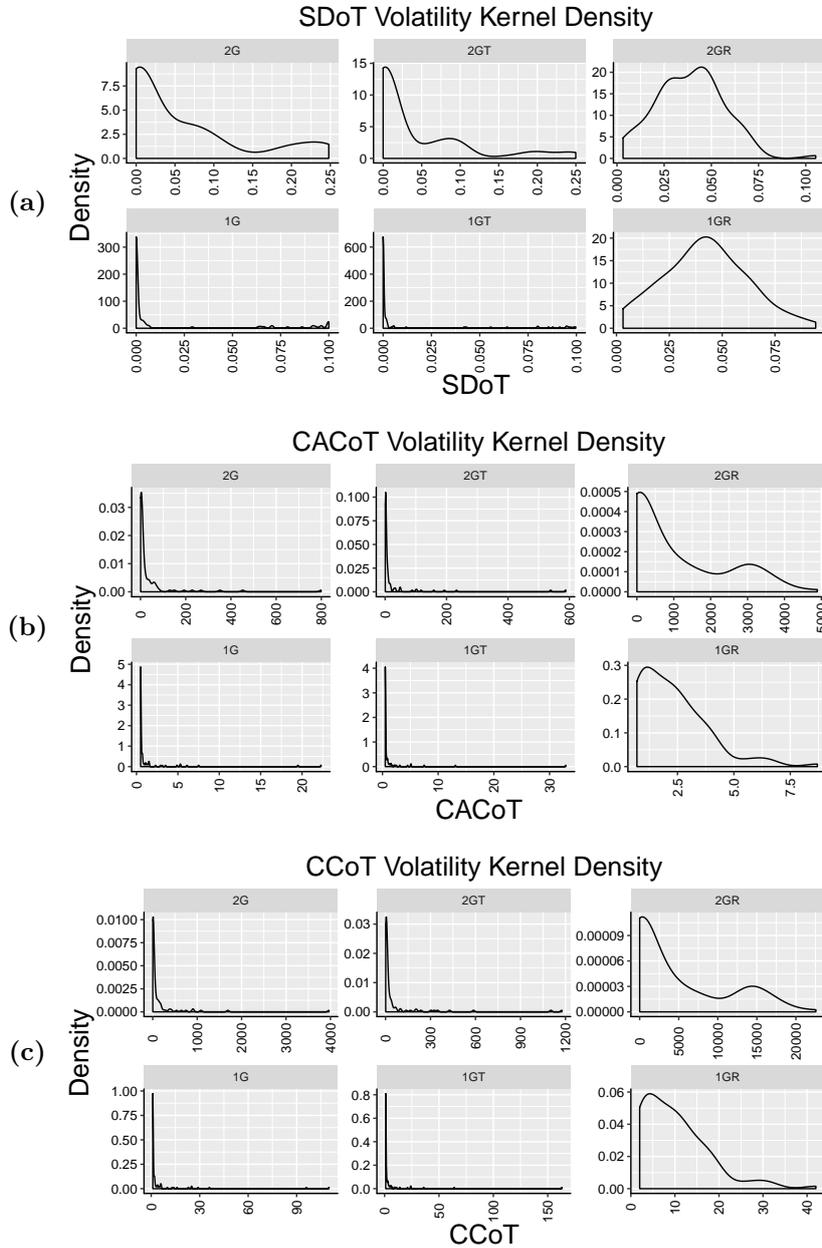
Figure 14: Kernel density estimation of the (a) SDoT, (b) CACoT and (c) CCoT, of agents using goal-rational, traditional or random action, evolved together or alone (2G, 2GT, 2GR, 1G, 1GT and 1GR respectively). Note that the plot scales are not comparable

action increases resilience to interference when agents evolve in dynamic, social environments compared to purely goal-rational agents. Traditional action was found to reduce the amount of volatility in the evolutionary process compared to agents that just use pure goal-rational action – without sacrificing fitness.

The final contribution shows that traditional action is indeed conceptually different to simply adding an element of randomness to the evolution of goal-rational agents. In fact, random action was found to increase the expected volatility measured by each of the three new metrics proposed. Volatility was significantly higher in socially situated agents, indicating that random action is very sensitive to interference and unknown changes in the environment. Whilst the expected fitness of agents that use random action is higher than those that use either goal-rational or traditional action, the expected volatility also drastically increases. As a result, agents that use traditional action are expected to receive a similar fitness to those that just use goal-rational action, as well as a significant decrease in volatility compared to those that use goal-rational or random action.

As socio-technical systems grow in size and complexity, it is inevitable that systems will indirectly interfere with one another by interacting with the environment. Related work on interference utilises knowledge of the existence of interference itself or others, their goals, intentions or abilities; in reality, this knowledge might not always be available. As the field of integration science flourishes, attention is being drawn to the inherently social nature of today's systems, and both the explicit and – more importantly in this work – the *implicit* interactions in which they participate. Consequently, it is becoming increasingly necessary to enable systems to *self*-integrate with others, despite the challenges presented by an incomplete knowledge of others and the environment, uncertainty, and the task of interacting in a multitude of ways with other systems that are heterogeneous in nature. A precursor to learning about others in one's environment and how to interact with them effectively is to learn about unanticipated changes and how to manage them; this work is therefore novel by

exploring how agents can mitigate interference without explicitly knowing about it.

Future work will explore different operationalisations of traditional action, and other combinations of social action in order to observe their effects on learning in socially situated agents. Further analysis using other metrics such as those that capture the change in entropy of fitness values can be utilised to gain more insight into the effect of social action on the evolutionary process.

[37] argue that, as the complexity of a system increases, the number of implicit and unintended interactions will rise as a result; we would therefore expect that it will become more challenging for agents to remain resilient to interference as the number of other entities in which they may potentially interact with increases. The results presented in this work show the impact that social situatedness can have on an agent's ability to achieve goals in the most minimal sense – i.e. there is only one other agent to interact with. However, whilst these results provide an indication of the effect that interference can have on evolution, this may change drastically depending on the environment, task, and the number of agents within the environment; this will be the topic of future investigation.

We hypothesise that endowing future systems with social self-awareness will mean systems with many co-located components are more capable of dealing with unanticipated and unknown interactions, due to their ability to detect and reason about others and the impact their actions can have on themselves. According to [7], systems that showcase social self-awareness are able to *intentionally* coordinate with one another; the agents in this work are therefore not socially self-aware as they have no capacity to perceive one another, but are instead tolerant of the consequences of interference to varying degrees. Further, traditional action has proven to be a beneficial step towards realising socially self-aware systems that can manage the impact of the actions of others on themselves in a socially acceptable way. This work challenges the assumption that goal-rationality is sufficient for achieving goals in social environment, as traditional action significantly reduces volatility in the evolutionary process without sacrificing fitness.

References

[1] Ralph Adolphs. Social cognition and the human brain. *Trends in cognitive sciences*, 3(12):469–479, 12 1999.

[2] R Axelrod and W D Hamilton. The evolution of cooperation. *Science (New York, N.Y.)*, 211(4489):1390–6, 3 1981.

[3] M. Bakhshipour, M. Jabbari Ghadi, and F. Namdari. Swarm robotics search & rescue: A novel artificial intelligence-inspired optimization approach. *Applied Soft Computing*, 57:708–726, 8 2017.

[4] C. M. Barnes, A. Ekárt, and P. R. Lewis. Social action in socially situated agents. In *Proceedings of the IEEE 13th International Conference on Self-Adaptive and Self-Organizing Systems*, pages 97–106, 2019.

[5] Louise Barrett, Peter Henzi, and Drew Rendall. Social brains, simple minds: does social complexity really require cognitive complexity? *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480):561, 4 2007.

[6] Kirstie Bellman, Jean Botev, Ada Diaconescu, Lukas Esterle, Christian Gruhl, Chris Landauer, Peter R Lewis, Anthony Stein, Sven Tomforde, and Rolf P Würtz. Self-improving system integration-status and challenges after five years of sissy. In *2018 IEEE 3rd International Workshops on Foundations and Applications of Self* Systems (FAS* W)*, pages 160–167. IEEE, 2018.

[7] Kirstie Bellman, Jean Botev, Hanno Hildmann, Peter R. Lewis, Stephen Marsh, Jeremy Pitt, Ingo Scholtes, and Sven Tomforde. Socially-Sensitive Systems Design: Exploring Social Potential. *IEEE Technology and Society Magazine*, 36(3):72–80, 2017.

[8] Kirstie Bellman, Sven Tomforde, and Rolf P Würtz. Interwoven systems: Self-improving systems integration. In *2014 IEEE Eighth International Conference on Self-Adaptive and Self-Organizing Systems Workshops*, pages 123–127. IEEE, 2014.

[9] Kirstie L Bellman, Christian Gruhl, Chris Landauer, and Sven Tomforde. Self-improving system integration-on a definition and characteristics of the challenge. In *2019 IEEE 4th International Workshops on Foundations and Applications of Self* Systems (FAS* W)*, pages 1–3. IEEE, 2019.

[10] James M Borg, Alastair Channon, Charles Day, et al. Discovering and maintaining behaviours inaccessible to incremental genetic evolution through transcription errors and cultural transmission. In *Advances in Artificial Life, ECAL 2011: Proceedings of the Eleventh European Conference on the Synthesis and Simulation of Living Systems*, pages 101–108. MIT Press, 2011.

[11] Robert Boyd and Peter J. Richerson. Culture and the evolution of human cooperation, 2009.

[12] Robert Boyd, Peter J Richerson, and Joseph Henrich. The cultural niche: why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences of the United States of America*, 108 Suppl 2(Supplement 2):10918–25, 6 2011.

[13] Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 3 2013.

[14] Richard W. Byrne. Machiavellian intelligence. *Evolutionary Anthropology: Issues, News, and Reviews*, 5(5):172–180, 1 1996.

[15] Cristiano Castelfranchi. Modelling social action for AI agents. *Artificial Intelligence*, 103(1-2):157–182, 8 1998.

[16] Cristiano Castelfranchi, Frank Dignum, Catholijn M. Jonker, and Jan Treur. Deliberative Normative Agents: Principles and Architecture. In N.R. Jennings and Lespérance Y., editors, *Intelligent Agents VI. Agent Theories, Architecture, and Languages*, pages 364–378. Springer, Berlin, Heidelberg, 2000.

[17] Sherry Chalotra, Sumeet Kaur Sehra, and Sukhjit Singh Sehra. A systematic review of applications of Bee Colony Optimization. In *2016 International Conference on Innovation and Challenges in Cyber Security (ICICCS-INBUSH)*, pages 257–260. IEEE, 2 2016.

[18] Dave Cliff and Linda Northrop. The global financial markets: An ultra-large-scale systems perspective. In *Monterey workshop*, pages 29–70. Springer, 2012.

[19] Helen G. Cobb and John J. Grefenstette. Genetic algorithms for tracking changing environments. In *Proceedings of the 5th International Conference on Genetic Algorithms*, pages 523–530, San Francisco, CA, USA, 1993. Morgan Kaufmann Publishers Inc.

[20] Kerstin Dautenhahn. Ants don't have friends– thoughts on socially intelligent agents. *Socially intelligent agents*, pages 22–27, 1997.

[21] Kerstin Dautenhahn, Bernard Ogden, and Tom Quick. From embodied to socially embedded agents – implications for interaction-aware robots. *Cognitive Systems Research*, 3(3):397 – 428, 2002. Situated and Embodied Cognition.

[22] Robyn M. Dawes. Social Dilemmas. *Annual Review of Psychology*, 1980.

[23] Edward Diener and Thomas K. Srull. Self-awareness, psychological perspective, and self-reinforcement in relation to personal and social standards. *Journal of Personality and Social Psychology*, 37(3):413–423, 1979.

[24] F. Dignum, D. Morley, E.A. Sonenberg, and L. Cavedon. Towards socially sophisticated BDI agents. In *Proceedings Fourth International Conference on Multi Agent Systems*, pages 111–118. IEEE Comput. Soc, 2000.

[25] Michael Doebeli and Christoph Hauert. Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game, 2005.

[26] M. Dorigo, V. Maniezzo, and A. Colorni. Ant system: optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, 26(1):29–41, 1996.

[27] Brian Duffy. Robots social embodiment in autonomous mobile robotics. *International Journal of Advanced Robotic Systems*, 1(3):17, 2004.

[28] Robin I. M. Dunbar. The social brain hypothesis. *Evolutionary Anthropology: Issues, News, and Reviews*, 6(5):178–190, 1 1998.

[29] Lukas Esterle. Goal-Aware Team Affiliation in Collectives of Autonomous Robots. In *2018 IEEE 12th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pages 90–99. IEEE, 9 2018.

[30] Lukas Esterle and John NA Brown. Levels of Networked Self-Awareness. In *2018 IEEE 3rd International Workshops on Foundations and Applications of Self* Systems (FAS*W)*, pages 237–238. IEEE, 9 2018.

[31] Xiang Fei, Hani S Mahmassani, and Pamela Murray-Tuite. Vehicular network sensor placement optimization under uncertainty. *Transportation Research Part C*, 29:14–31, 2013.

[32] Sofia I. F. Forss, Erik Willems, Josep Call, and Carel P. van Schaik. Cognitive differences between orang-utan species: a test of the cultural intelligence hypothesis. *Scientific Reports*, 6(1):30516, 9 2016.

[33] Uta Frith and Chris Frith. The social brain: allowing humans to boldly go where no other species has been. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1537):165–176, 1 2010.

[34] Xiaowei Fu and Xiaoguang Gao. Genetic algorithm with adaptive immigrants for dynamic flight path planning. In *2010 IEEE International Conference on Intelligent Computing and Intelligent Systems*, volume 1, pages 630–634. IEEE, 2010.

[35] Gordon G. Gallup. Self-awareness and the evolution of social intelligence. *Behavioural Processes*, 1998.

[36] Stephen Grossberg. Nonlinear neural networks: Principles, mechanisms, and architectures, 1988.

[37] Joerg Haehner, Uwe Brinkschulte, Paul Lukowicz, Sanaz Mostaghim, Bernhard Sick, and Sven Tomforde. Runtime self-integration as key challenge for mastering interwoven systems. In *ARCS 2015-The 28th International Conference on Architecture of Computing Systems. Proceedings*, pages 1–8. VDE, 2015.

[38] Esther Herrmann, Josep Call, Maráa Victoria Hernàndez-Lloreda, Brian Hare, and Michael Tomasello. Humans have evolved specialized skills of social cognition: the cultural intelligence hypothesis. *Science (New York, N.Y.)*, 317(5843):1360–6, 9 2007.

[39] Chung Yuan Huang, Sheng Wen Wang, and Chuen Tsai Sun. Self-aware intelligent agents in the prisoner's dilemma. In *Proceedings - 2011 International Conference on Future Computer Sciences and Application, ICFCSA 2011*, pages 127–131, 2011.

[40] N.R. Jennings. Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75(2):195–240, 6 1995.

[41] Ben P. Jolley, James M. Borg, and Alastair Channon. Analysis of social learning strategies when discovering and maintaining behaviours inaccessible to incremental genetic evolution. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 9825 LNCS, pages 293–304, 2016.

[42] Alison Jolly. Lemur Social Behavior and Primate Intelligence. *Science*, 153(3735):501–506, 7 1966.

[43] J. Frances Kamhi, Wulfila Gronenberg, Simon K. A. Robson, and James F. A. Traniello. Social complexity influences brain investment and neural operation costs in ants. *Proceedings of the Royal Society B: Biological Sciences*, 283(1841):20161949, 10 2016.

[44] Dervis Karaboga and Bahriye Basturk. A powerful and efficient algorithm for numerical function optimization: artificial bee colony (abc) algorithm. *Journal of global optimization*, 39(3):459–471, 2007.

[45] Peter Kollock. Social Dilemmas: The Anatomy of Cooperation. *Annual Review of Sociology*, 1998.

[46] Samuel Kounev, Peter Lewis, Kirstie L. Bellman, Nelly Bencomo, Javier Camara, Ada Diaconescu, Lukas Esterle, Kurt Geihs, Holger Giese, Sebastian Götz, Paola Inverardi, Jeffrey O. Kephart, and Andrea Zisman. The Notion of Self-aware Computing. In *Self-Aware Computing Systems*, pages 3–16. 2017.

[47] Seung-Kyu Lee and Byung-Ro Moon. Genetic algorithm with adaptive elitism-based immigrants for dynamic optimization problems. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 1865–1866, 2009.

[48] Joel Z. Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent Reinforcement Learning in Sequential Social Dilemmas. *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017.

[49] Peter R. Lewis. Self-aware computing systems: From psychology to engineering. In *Proceedings of the 2017 Design, Automation and Test in Europe, DATE 2017*, pages 1044–1049, 2017.

[50] Peter R. Lewis, Arjun Chandra, and Kyrre Glette. Self-awareness and Self-expression: Inspiration from Psychology. In Peter R. Lewis, Marco Platzner, Bernhard Rinner, Jim Torresen, and Xin Yao, editors, *Self-aware Computing Systems*, chapter 2, pages 9–21. Springer International Publishing, 1 edition, 2016.

[51] Mathieu Lihoreau, Tanya Latty, and Lars Chittka. An Exploration of the Social Brain Hypothesis in Insects. *Frontiers in Physiology*, 3:442, 11 2012.

[52] Jessica Lindblom and Tom Ziemke. Social situatedness: Vygotsky and beyond. *Proceedings of the Second International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems*, pages 71–78, 2002.

[53] Maja J Matarić. *Interaction and Intelligent Behavior*. PhD thesis, 1994.

[54] Robert K. Merton. The Unanticipated Consequences of Purposive Social Action. *American Sociological Review*, 1936.

[55] Alex Mesoudi, Lei Chang, Sasha R.X. Dall, and Alex Thornton. The Evolution of Individual and Cultural Variation in Social Learning. *Trends in Ecology & Evolution*, 31(3):215–225, 3 2016.

[56] Krzysztof Mogielski and Tadeusz Płatkowski. A mechanism of dynamical interactions for two-person social dilemmas. *Journal of Theoretical Biology*, 2009.

[57] Henrike Moll and Michael Tomasello. Cooperation and human cognition: the Vygotskian intelligence hypothesis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1480):639–648, 4 2007.

[58] Francesco Mondada, Luca Maria Gambardella, Dario Floreano, Stefano Nolfi, Jean-Louis Deneubourg, and Marco Dorigo. The cooperation of swarm-bots: physical interactions in collective robotics. *IEEE Robotics Automation Magazine*, 12(2):21–28, June 2005.

[59] C Müller-Schloer. Organic Computing – On the Feasibility of Controlled Emergence. *Proceedings of the 2nd IEEE/ACM/IFIP international conference on Hardware/software codesign and system synthesis*, 2004.

[60] Michael Muthukrishna, Michael Doebeli, Maciej Chudek, and Joseph Henrich. The Cultural Brain Hypothesis: How culture drives brain expansion, sociality, and life history. *PLoS computational biology*, 14(11):e1006504, 2018.

[61] Phyllis R Nelson. Ensuring appropriate use of computational resources in interwoven systems. In *2017 IEEE 2nd International Workshops on Foundations and Applications of Self* Systems (FAS* W)*, pages 139–143. IEEE, 2017.

[62] Daphne Ngar-yin Mah, Johannes Marinus van der Vleuten, Jasper Chi-man Ip, and Peter Ronald Hills. Governing the transition of socio-technical systems: A case study of the development of smart grids in Korea. *Energy Policy*, 45:133–141, 6 2012.

[63] Mark Nielsen. Copying actions and copying outcomes: Social learning through the second year. *Developmental Psychology*, 2006.

[64] Sascha Ossowski and Ana García-Serrano. Social Structure in Artificial Agent Societies: Implications for Autonomous Problem-Solving Agents. In *Intelligent Agents V. ATAL'98: Proc. 5th Int'l Workshop on Agent Theories, Architectures and Languages*, pages 133–148, 1998.

[65] Charles Perreault, Cristina Moya, and Robert Boyd. A Bayesian approach to the evolution of social learning. *Evolution and Human Behavior*, 33(5):449–459, 9 2012.

[66] Jeremy Pitt, Julia Schaumeier, and Alexander Artikis. The Axiomatisation of Socio-Economic Principles for Self-Organising Systems. In *2011 IEEE Fifth International Conference on Self-Adaptive and Self-Organizing Systems*, pages 138–147. IEEE, 10 2011.

[67] Ser Huang Poon and Clive W.J. Granger. Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature*, 41(2):478–539, 2003.

[68] Anand S Rao, Michael P Georgeff, and Elizabeth A Sonenberg. Social Plans: A Preliminary Report. Technical report, 1992.

[69] Simon M. Reader. Innovation and social learning: Individual variation and brain evolution. *Animal Biology*, 2003.

[70] Edward Robinson, Timothy Ellis, and Alastair Channon. Neuroevolution of agents capable of reactive and deliberative behaviours in novel and dynamic environments. In *Advances in Artificial Life*, pages 1–10. Springer, 2007.

[71] Stefan Rudolph, Sven Tomforde, Bernhard Sick, and Jorg Hahner. A Mutual Influence Detection Algorithm for Systems with Local Performance Measurement. In *2015 IEEE 9th International Conference on Self-Adaptive and Self-Organizing Systems*, pages 144–149. IEEE, 9 2015.

[72] M. Schneider-Fontán and M.J. Mataric. Territorial multi-robot task division. *IEEE Transactions on Robotics and Automation*, 14(5):815–822, 1998.

[73] Daniel Joel Shaw, Kristína Czekóová, Jakub Chromec, Radek Mareček, and Milan Brázdil. Copying you copying me: Interpersonal motor coordination influences automatic imitation. *PLoS ONE*, 2013.

[74] Karl Sigmund and Martin A Nowak. Evolutionary game theory. *Current Biology*, 1999.

[75] Anabela Simões and Ernesto Costa. An immune system-based genetic algorithm to deal with dynamic environments: diversity and memory. In *Artificial Neural Nets and Genetic Algorithms*, pages 168–174. Springer, 2003.

[76] Jeffry A. Simpson and Douglas T. Kenrick. *Evolutionary social psychology*. Lawrence Erlbaum Associates, 1997.

[77] Frank M. J. Sommerlandt, Axel Brockmann, Wolfgang Rössler, and Johannes Spaethe. Immediate early genes in social insects: a tool to identify brain regions involved in complex behaviors and molecular processes underlying neuroplasticity. *Cellular and Molecular Life Sciences*, 76(4):637, 2019.

[78] Adam Stanton and Alastair Channon. Incremental Neuroevolution of Reactive and Deliberative 3D Agents. *European Conference on Artificial Life (ECAL)*, pages 341–348, 2015.

[79] Sally E Street, Ana F Navarrete, Simon M Reader, and Kevin N Laland. Coevolution of cultural intelligence, extended life history, sociality, and brain size in primates. *Proceedings of the National Academy of Sciences of the United States of America*, 114(30):7908–7914, 7 2017.

[80] Katia P. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence - Volume 1*, IJCAI'89, pages 517–523, San Francisco, CA, USA, 1989. Morgan Kaufmann Publishers Inc.

[81] Piotr Sztompka. *Socjologia*. Znak, Kraków, 2002.

[82] John Thangarajah, Lin Padgham, and Michael Winikoff. Detecting & Avoiding Interference Between Goals in Intelligent Agents. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, 2003.

[83] John Thøgersen. Social norms and cooperation in real-life social dilemmas. *Journal of Economic Psychology*, 2008.

[84] Sven Tomforde, Stefan Rudolph, Kirstie Bellman, and Rolf Würtz. An organic computing perspective on self-improving system interweaving at runtime. In *2016 IEEE International Conference on Autonomic Computing (ICAC)*, pages 276–284. IEEE, 2016.

[85] U.S. CFTC and U.S. SEC. Findings regarding the market events of May 6, 2010. Technical report, 2010.

[86] Paul A.M. van Lange, Jeff Joireman, Craig D. Parks, and Eric Van Dijk. The psychology of social dilemmas: A review. *Organizational Behavior and Human Decision Processes*, 2013.

[87] Carel P van Schaik and Judith M Burkart. Social learning and evolution: the cultural intelligence hypothesis. *Philosophical transactions of the Royal Society of London. Series B, Biological Sciences*, 366(1567):1008–16, 4 2011.

[88] Lev Semenovich Vygotsky. *Mind in society: The development of higher psychological processes*. Harvard university press, 1978.

[89] Guy H. Walker, Neville A. Stanton, Dan Jenkins, Paul Salmon, Mark Young, and Amerdeep Aujla. Sociotechnical Theory and NEC System Design. In *Engineering Psychology and Cognitive Ergonomics*, pages 619–628. Springer, 2007.

[90] Jane X. Wang, Edward Hughes, Chrisantha Fernando, Wojciech M. Czarnecki, Edgar A. Duéñez Guzmán, and Joel Z. Leibo. Evolving intrinsic motivations for altruistic behavior. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '19, pages 683–692, Richland, SC, 2019. International Foundation for Autonomous Agents and Multiagent Systems.

[91] Max Weber, Guenther. Roth, and Claus. Wittich. *Economy and society: an outline of interpretive sociology*. University of California Press, 1978.

[92] Holger Wolf. Volatility: Definitions and Consequences. In Joshua Aizenman and Brian Pinto, editors, *Managing Economic Volatility and Crises: A Practitioner's Guide*, pages 45–64. Cambridge University Press, 2005.

[93] Shengxiang Yang. Genetic algorithms with memory- and elitism-based immigrants in dynamic environments. *Evolutionary Computation*, 16(3):385–416, 2008.

[94] Simon X. Yang and Max Meng. An efficient neural network approach to dynamic robot motion planning. *Neural Networks*, 13(2):143–148, 2000.

[95] Simon X. Yang and Max Meng. An efficient neural network method for real-time motion planning with safety consideration. *Robotics and Autonomous Systems*, 32:115–128, 2000.

[96] B. W. Yap and C. H. Sim. Comparisons of various types of normality tests. *Journal of Statistical Computation and Simulation*, 81(12):2141–2155, 2011.

[97] Xin Yu, Ke Tang, and Xin Yao. An immigrants scheme based on environmental information for genetic algorithms in changing environments. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, pages 1141–1147. IEEE, 2008.