Critical Success Factors for Artificial Intelligence Projects

Victoria Uren (v.uren@aston.ac.uk)
Aston Business School, Aston University Birmingham B4 7ET

Abstract

Artificial intelligence (AI) is being adopted in a wide range of operational settings, but many organizations are still at an early stage of adoption. This work concerns critical success factors of AI projects, and is based on interviews with professionals who have been involved in delivering them. Technology Readiness Levels were used as a surrogate for context. CSFs identified by the study include: address real, carefully scoped business problems, match the AI technology to the problem, build understanding of AI beyond the development team, invest to generate high quality data, develop data capabilities, and embed AI expertise in the organization.

Keywords: Artificial Intelligence, Technology Readiness Level, Critical Success Factors

Introduction

IT innovation can have profound effects on operations. Historically, inventory management, Enterprise Resource Planning systems, scheduling, simulation, CAD and so forth have facilitated more efficient processes, supported decision making and helped to improve the quality of services and products. Artificial Intelligence (AI) is being widely predicted to impact a wide range of industries in the coming years. The UK Industrial Strategy has identified AI as *"one of the great opportunities of our age"* (UK Government, 2018) and the European Commission (2020) is committed to invest in AI with the aim of achieving an *"ecosystem of excellence"* and an *"ecosystem of trust"*.

AI is not a unitary technology. Although many commentators currently conflate AI with machine learning, seen more broadly, AI has several high-level aims. As well as learning, perception, reasoning, communication and knowledge representation also need to be considered.

Machine learning uses a very large number of algorithmic approaches to identify patterns in data. It is suitable for providing decision support in cases where there are large amounts of historical data to learn from. For example, deep learning performance can be as good as or exceed human experts for diagnosis using medical scans. The NHS in the UK is investigating this as a means to alleviate workforce shortages (Joshi and Morley, 2019). Machine Learning also supports process improvement, for example, the prediction of arrival times at Changi Airport, which feeds into improved downstream operations, such as baggage handling (Lee and Miller, 2019).

Perception utilises sensors to gather data from the environment, which can then be used to support autonomous behaviour. An example of technology deployed in logistics operations is cruise control systems in trucks, which can maintain constant distance to the truck in front, but also take in features of the environment, such as gradients, and adjust speed or change gear accordingly (Klumpp and Zijm, 2019, Tambe et al., 2014).

Reasoning in AI is characterized by terms such as probabilistic, symbolic or sub-symbolic reasoning. For example, a Bayesian machine learning algorithm would be considered probabilistic, whereas the neural nets deployed in deep learning are sub-symbolic. Game theory provides an example of an approach to reasoning which can have symbolic components. Work by Tambe and others on multiagent systems has deployed game theory to model the decisions of defender and attacker agents. The models are used to allocate transport security resources in airports, metro systems etc. (Tambe et al., 2014). As legal and ethical requirements push organizations towards explainable AI, the ability to unpack underlying reasoning in ways that can be presented to a human decision maker will become important (European Commission, 2020). This need to communicate the underlying reasoning of an algorithm links to the next aspect of AI.

Communication is the ability for intelligent systems to interact with human users. Natural Language Processing (NLP) is a key communication technology. For example, service agents in the hospitality sector can support multi-lingual communication with guests at reduced cost compared to employing multi-lingual staff 24/7 (Ivanov and Webster, 2017). Complex factors influence human-machine communication, Gursoy et al (2019), for instance, demonstrate the importance of hedonic motivation, social influence and anthropomorphism in users' interactions with service agents.

Knowledge representation is a critical foundation for AI. One example is the pre-processing step and parsing of text to identify parts of speech in NLP (Manning & Schütze 1999). Another is provided by Linked Data, which provides a representation that is used for symbolic reasoning in web-based systems (W3C, 2020).

Complex AI applications integrate multiple AI technologies. Autonomous vehicles provide an example which is of current interest in logistics with the use of robots in picking operations within warehouses (Mahroof, 2019). Such a robot needs sensors to give it information about its surroundings perhaps using a vision system, which would probably be trained using deep learning. It will be able to reason about how to achieve its objectives, and it may apply multiagent systems principles to guide interaction with other robots and human workers.

The above examples serve to illustrate a few applications of AI which motivate many organizations to consider deploying AI technology in their own operations. The inherent complexity of AI, and its unfamiliarity to many users on the ground, including some working in IT roles, means that IT projects that use AI have additional complexities over and above those of projects using more conventional IT. The current push to deploy the technology means it is important to examine how AI projects differ from more standard IT deployments.

Existing AI innovation research often assesses specific technologies and domains. In contrast, this work looks at success factors at the level of project management, not constrained by particular AI technologies or application domains. By identifying common themes that occur across a range of applications, the aim of the research is to make recommendations that are relevant to a wide range of adoption contexts. The most similar works we are aware of looked at organisational readiness to adopt AI (AlSheibani et al., 2018, Pumplun et al., 2019) and concern factors, such as the technology fit, skills within the organization, management support and the competitive environment. We will give further consideration to those two works and papers on operational application of AI, in the discussion section.

This study provides empirical evidence of critical success factors (CSFs), *"those few things that must go well"* (Boynton and Zmud, 1984), for AI projects. The term critical success factor was coined in 1979 by Rockart (1979) in relation to the information needs of chief executives. It has since been adopted widely, bridging the management and engineering literatures. There are numerous examples in the operations literature of using CSFs to study information systems. Enterprise Resource Planning systems in particular have been thoroughly examined, e.g. (Soliman et al., 2001, Umble et al., 2003, Woo, 2007, Snider et al., 2009). Other examples include B2B eCommerce systems in supply chains (Cullen and Taylor, 2009) and advanced manufacturing technology (Kumar et al., 2018).

Studies of CSFs for information systems projects share some recurring themes, such as the need for senior management support and project management capability (Umble et al., 2003, Snider et al., 2009). However, they do not do not follow a prescribed framework, but seek what is specific to each application. The diversity of CSFs is illustrated by contrasting a couple of examples. Chow & Cao (Chow and Cao, 2008) looked at CSFs in agile software projects and identified the CSFs *"correct delivery strategy", "proper practice of Agile software engineering techniques"* and *"high caliber team"*, which are factors relevant to the IT context. Hong & Kim (Hong and Kim, 2002) looked at ERP system implementation, and identified the organizational fit of the ERP system, coupled with an optimal level of adaptation to the organization, both factors relevant to the business context.

Limitations of CSFs include the fact that participants are constrained by their own experience and so might fail to correctly recognize the factors that led to success (Davis, 1980). This study tries to mitigate that risk by taking a multi-perspective approach, the participants have different roles and backgrounds, so that different factors may be identified from different angles. CSFs have also been criticized on the grounds that they lack consideration of specific contextual circumstances (Axelsson and Melin, 2014). As illustrated above, CSFs have been used in other contexts than Rockart's original (1979) focus on the chief executive's viewpoint, leading to a lack of comparability between studies. As this study examines a wide range of application settings, in order to get multiple perspectives, the well-established, Technology Readiness Level (TRL) scale was used as a context surrogate.

Technology Readiness Levels were developed by NASA to describe the progress of engineering technology development from vision (TRL 1) to proven operational use (TRL 9) (Héder, 2017). The grounds for suggesting that this is an appropriate way of viewing context for AI projects at this time are that (as the findings confirm) much of the work that organizations are currently doing with AI is experimental and exploratory. Our assumption is that the stage of technology readiness for projects might affect both how participants view success and the CSFs they identify. The inclusion of TRL levels in this study is original, and we are not aware of any work which considers this aspect.

Method

Semi structured interviews were conducted with professionals who had been actively involved with AI projects. Participants were sought who had experience of different parts of the AI development journey and projects with different TRLs to provide multiple perspectives. These included (number of participants in brackets):

- UNI - University researchers with experience of knowledge transfer projects involving industry partners (4).

- IT - Consultants in developer and manager roles who had worked on the development of systems for clients (5).
- ENT - Entrepreneur who had built technology businesses using AI systems (1).
- SAF - System testing professional, responsible for validating systems. (1).

The study is work in progress at the time of writing. However, the participants interviewed so far provided insights on AI projects at all TRL levels, from TRL 1, including both University and consultant led ideation sessions, to TRL 9, with applications such as RPA with AI components, and an eCommerce discovery platform. The range of operational contexts discussed was very broad, and encompassed autonomous vehicles, banking, retail and marketing, equipment maintenance, food and agriculture, insurance, health, logistics and smart cites. The full coverage of TRL levels and breadth of contexts provides confidence that initial observations can be made.

The interview protocol was adapted from Joo's (2011) study on Semantic Web innovation, with additional questions on data and TRLs. The semi-structured approach allowed questions to be modified to suit the different experience and perspective of each participant. A similar order of topics was typically followed for the interviews, unless participants naturally covered some topics earlier. The order was: first the interviewee's role and involvement in AI projects, then project objectives and motivation, TRL levels of projects, the readiness of the (sometimes client) organization to adopt AI, including resources and skills, interactions with partners or clients, data related issues, the success of projects, including what success looks like, and the most important factors for achieving success, and finally the interviewee's view of future trends related to AI. The final question was open to bring in serendipitous insights based on the unique experience of each participant.

This study adapted a version of the TRL scale by Meystel et al., (2003) with diagnostic questions for each level, posed in terms that would be accessible to end users and managers as well as IT specialists. The scale is summarised below.

Table 1 – Technology Readiness Level Scale

TRL 1	Basic Principles and broad vision
TRL 2	Conceptual design
TRL 3	Theoretical & experimental analysis - proof of concept
TRL 4	Component validation in "laboratory" conditions
TRL 5	Component validation in more realistic conditions
TRL 6	Subsystem model or prototype demonstration
TRL 7	System prototype demonstration in operational environment
TRL 8	Actual system completed, tested and demonstrated
TRL 9	Actual system proven through operational use

Interviews were recorded, transcribed by the researcher and the filename prefixed by one of UNI, IT, ENT or SAF, in order to tag the type of participant (as above) who had made each statement. The software used for coding was Nvivo, which supported an approach in which transcripts were first chunked according to the planned topics of the interview. This brought related text from different participants together in themes for more fine-grained analysis, which was recorded as sub-codes. Sub-coding was inductive, allowing themes to emerge from the participants' statements.

Findings

Initial findings are presented here concerning 1) the pathway that organizations follow towards adopting AI, 2) what success looks like at different TRL levels, and 3) critical success factors.

Pathway to AI adoption

Participants identified an innovation pathway, leading to successful AI in the world, that can be described as having three stages: laying the foundations of AI, adoption, and mature AI. Organizations with pre-existing expertise in Big Data, statistics etc. have an advantage in building an understanding of what AI can deliver for them, and may join the pathway at later stages.

Laying the foundations for AI starts with initial interest in the technology, which can be triggered through media coverage and hype, but sometimes by knowledge of a data-related issue in the organization that needs solving. A critical activity at this stage is to improve the organization's internal data practices. Both IT and UNI participants who had interacted with organizations that sought AI collaboration, reported that they had sometimes advised working on internal data quality first, before moving on to AI. At this stage, it was reported that organizations could either be using, or be encouraged to adopt innovative technologies proximal to AI that would solve immediate business problems more easily, while also building foundational knowledge. Typical proximal technologies included Robotic Process Automation, analytics, and automating spreadsheets, what one IT participant described as *"quick win applications, to start off with, to show the value"*. Quick wins also included *"off-the-shelf AI products"*, such as third party APIs, a chat-bot in the case of ENT. These allow organizations to achieve high TRL levels sooner.

During adoption, organizations build capacity, both internally and through collaboration with external partners to access additional skills. As ENT put it *"at every stage as an organization, and especially in the last three to six years when we shifted, or increasingly shifted focus there's been lots of new learning"*. Activities like ideation, i.e. creating suitable business cases to trial AI, experimentation with technologies, and the development of projects which use relatively limited amounts of AI to solve problems, characterize this stage. If adoption progresses well, the ambition of projects increases, higher TRL levels are reached, more complex AI is used and more AI components are built into systems. Participants indicated that, while many of the projects so far had been evolutionary, with small changes made, they could foresee follow-on projects that had potential to be more revolutionary, for example, moving from automation of individual tasks to holistic changes involving entire business processes.

Mature AI was not yet seen as accomplished, although cases of very sophisticated AI applications within tech-oriented private sector companies were noted, and one participant observed how technologies which would once have been considered AI, such as search, come to be regarded as computer science, not as AI, once they are established. Steps that were noted as needed to reach maturity included regulation, development of steps such as ethical review of projects, developing testing processes for systems with evolving behaviour, and a better understanding of societal tolerance of the risk associated with AI, in order to assure the safety of applications.

What Success Looks Like

Participants were asked what success in an AI project looked like to them. It was noteworthy that participants saw success as having a number of facets that went beyond whether a piece of AI technology worked as intended. These can be categorized as commercial success, contractual success, ethical success, learning as success, and measurable success.

Commercial success is usually only be seen once a project has reached TRL 9 and a product is released, which has value for the organization. For example, in ENT's eCommerce platform, success looked like *"highly accurate predictions that are ultimately recommendations"*. Although this has a measurable component (*"highly accurate"*) it also goes right to what ENT's clients buy, which are trusted recommendations. For IT respondents that worked for multiple clients, commercial success was defined more generally as achieving desired business benefits or "value".

Contractual success was defined by both the IT and UNI participants as *"achieving intended outcomes"* [IT] which would be agreed with project stakeholders in the early phases of the project. These may be business outcomes, as above, but could also be related to commissioned activities. This kind of success could be achieved at all TRL levels in well managed projects.

Ethical success can be about avoiding using AI for bad ends, *"just because you can do something doesn't mean you should"* [IT], but can also include positive outcomes. An example given by one UNI participant was using optimization algorithms in logistics planning to help reduce air pollution. Avoiding bad outcomes produced by the AI itself was also raised. Examples were, concerns over biases in machine learning algorithms, concerns about lack of explainability in black-box machine learning algorithms, the issue of physical safety of autonomous vehicles, as well as ethical issues in application areas such as health and justice.

Participants who had been involved in projects at lower TRL levels (approximately TRL 1 to 5) often identified learning as a kind of success, which might even occur in projects that were technical failures: *"in some senses, from a project point of view, that was a failure, ... [but] they recognized the benefits, that we'd explored certain scenarios, certain options and narrowed them down to where they might focus the next phase of the business."* [UNI]. Furthermore, one UNI participant noted that success could come as an unintended consequence of learning that might only be seen in retrospect, giving the example of a researcher who had left their University to start an AI based business. This maybe a view peculiar to the higher education sector as there was no direct benefit either to the respondent's organization or to one of its clients, but the success addressed the high-level goals of an educational institution.

Measurable success perhaps comes closest to the traditional computer science paradigm of assessing the technical performance of AI algorithms. Some of the success measures that were discussed were related to commercial success. Specific examples included logistical efficiency, the ability of a business to do better prediction, and quality improvement e.g. *"The benefit could be that once this automation is introduced the quality score on this process, which when humans do it is 65%, goes up to 95%."* [IT]. However, for the safety engineer, the ability to measure confidence in the behaviour of an AI system, particularly one which might evolve with time, was a fundamental problem that still needs to be resolved: *"Artificial intelligent part perhaps isn't deterministic ... we think we know what it's going to do, and it should do that every time, but it's a learning thing and it can learn different behaviours over time... So, with unanticipated emergent behaviours how do you know that's still going to be a successful thing?"* [SAF]. This implies that continued success once TRL 9 is achieved will have a maintenance cost.

Critical Success Factors

Two kinds of success factors were identified by participants. The first kind was defined by one IT participant as *"general good development practices"* and are recognizable as critical success factors that have been identified in an IT context in many CSF studies. These included, getting engagement from stakeholders, both end users and IT stakeholders in client organizations who would have to support the technology once it went live, getting senior management support, change management, including cultural change, and properly understanding the requirements. As these have been discussed extensively elsewhere, this paper focusses on the second kind which were either specific to AI or had an AI twist.

The importance of addressing real business problems was emphasized by most participants, as opposed to the fallacy of seeking problems that fit available technology. As one UNI respondent who had been working on small scale projects with SMEs put it *"you can't just do AI for the sake of doing AI"*. Linked to this is the need to define appropriate project scope and to communicate well with users, so that their expectations of what AI can achieve are realistic. Thus avoiding disappointment and loss of credibility further down the line.

Matching the right technology to the problem follows on from addressing the right problem. One UNI machine learning expert described this as *"Applying a wide variety of different models without prejudice"* [UNI], while another observed, as a criticism, that *"I know there are people who are very keen on, I have this hammer and I'll use this no matter what"* [UNI]. It was also emphasized by several participants that AI was not always the best solution, even if the client came asking for it. For example, one IT participant observed *"a Robotic Process Automation or an analytics solution might give you the same benefit, but at a cheaper cost or a quicker implementation"*.

Building understanding of AI technology beyond the development team was considered important. At the lower TRL levels, engagement with projects was often seen as helping to build organizational resources which could support AI adoption in the future. This might be about understanding the capabilities of the technology *"it laid the foundation, it started to change perceptions, it started to open eyes about what might be possible and what's more difficult"* [UNI], or it might involve preparatory implementation, like starting to build data resources for future use *"if the data's a mess, you know, garbage in garbage out, the AI won't give you good data, so it's trying to maybe give a roadmap towards it"* [IT].

Indeed, data was reported to be critical to success. Three broad themes emerged: data access, data quality issues and data capabilities. Data access was reported as essential for projects to proceed, and had two aspects. The first was establishing trust between project partners to allow data to be shared through legal agreements and so forth. The second involved the technical means of sharing data, for example, through APIs or Cloud. Data quality issues were so common that some participants perceived them as the norm. This view was prevalent in IT participants. A selection of issues that were cited includes data volume (too much or too little), biases, dispersed data sources, the challenges of fusing data from multiple sources, the need to produce labelled training data, duplicate data and missing data.

It follows that, building data capabilities is a CSF for AI projects. The first, most fundamental capability, is awareness, as one UNI participant put it *"the more sceptical they are about their own data quality the better their understanding of the problem is"* [IT]. Awareness of data management issues was sufficiently important in itself that improving it was cited as a success outcome of lower TRL projects. Second, the capability of organizations to generate data of the right kind and quality from their processes was identified. Helping organizations get the infrastructure and data management practices in place to start collecting data that could be used for AI in the future was, like generating awareness, perceived as a kind of success which would lead on to the ability to succeed at AI in future. Third, data governance is a capability needed to get AI projects to the higher operational levels. With its focus on addressing privacy, security and the ethical use of data, this contributes to the ability to make AI legally compliant and achieve societal acceptance.

To reach the highest TRL levels, it was seen as necessary to start to embed AI specific skills in the organization. Organizations which had related skills on which to build, for example, statistical knowledge or data science expertise, were considered likely to achieve success more easily. The need to be able to support AI in-house and continue to innovate was key for ENT *"we did engage an external consultant to help us construct it initially and then since that initial work we've had our engineers train and set up an in house online certification and training house so that we've been able to work through the iterations and latest releases"* and also *"We have an orientation to constant self-learning"*. This was also observed by other participants working in a consultancy capacity (both IT and UNI) who reported organizations recruiting data scientists and AI developers as a route to AI success.

Discussion

Although they are still preliminary, the findings of this study contribute to the broader debate around adoption of AI.

Examples of the kind of pathways to adoption identified by our participants can be seen in the literature. Lee & Miller (2019) discuss a roadmap approach to growing the portfolio of AI projects at Changi Airport in Singapore. Mahroof (2019) provides a case of a logistics company using AI for planning first, whilst putting technical infrastructure into the warehouse in preparation for later AI adoption there.

Data and the capability to handle it well are widely acknowledged to be crucial. Pumplun et al.'s (2019) study of organizational readiness factors for AI highlighted the issue of the *"availability, quality and protection of data"*. Authors from IBM discuss the importance of a data pipeline as part of mature machine learning within enterprises (Akkiraju et al., 2018). The finding that data capabilities are a CSF supports AlSheibani's (2018) hypothesis that *"Human, enterprise and technology resources positively influence AI readiness"*.

Our findings point to the importance of building understanding of AI beyond the developer community. Work by Klumpp and Zjim (2019) points to the risk of creating an *"artificial divide"* between workers in the logistics industry who do and do not know how to work with AI. Mahroof (2019) observes different attitudes towards the readiness to adopt AI between operational managers and the implementation team, pointing to different understandings of the importance of preparatory work on infrastructure, as well as fear of the technology from those who had not experienced it before. Developers' skills are discussed from a planning perspective by Lee and Miller (2019) in their report on lessons learned from AI projects at Changi Airport. They observe that AI requires a rethink of the IT skills they have in-house and those they need to out-source, not least to protect the proprietary knowledge they see as a source of competitive advantage.

The finding that AI needs to meet a strategic business need, and be a good fit for the problem is also a theme of the broader literature. Pumplun (2019) discusses the need to fit AI to specific business needs, and advises using conventional systems where they are more appropriate. This supports AlSheibani's (2018) hypothesis that *"Compatibility between the AI business case and an organisation's existing strategies positively influences AI readiness"*.

Operational safety, data security and ethical success are considered by authors such as Hengstler et al. (2016), who discuss the importance of trust in technology for getting user acceptance of smart medical technology and autonomous vehicles, and van Esch et al. (2019), who discuss the use of AI in recruitment processes and how it might impact the anxiety of applicants. There is apparent disagreement in the literature on whether the effect of regulation on AI success is positive or negative: AlSheibani et al. (2018) hypothesise that *"Government regulations can have a positive influence on AI readiness"*, while Pumplun et al. (2019) propose that *"Strict laws regarding the processing of personal data will hamper the training of intelligent machines"*. Although participants were alert to their obligations, the evidence collected so far in this study is insufficient to contribute on that point. Future interviews can target application areas such as healthcare, where ethical issues are important in order to probe this.

Examples of themes which occur in the literature but were not found in the preliminary findings of this study are competitive pressure to adopt AI and demand from customers for more intelligent products and services. For example, Pumplun et al. (2019) propose that *"demanding customers will nudge the companies to design individualized, intelligent products"* and AlSheibani et al. (2018) hypothesise that *"competitive pressure has a positive influence on AI readiness"*. This absence may reflect the bias of the sampling of participants interviewed so far towards implementation, with nine out of the eleven interviews coming from IT professionals working either in academia or IT companies. This needs to be corrected as the study proceeds to include more participants like ENT who are using AI strategically in their operations.

Conclusion

This work contributes to the information systems adoption literature. It is relevant to organizations currently innovating with AI, consultants supporting innovation, and Universities in technology transfer partnerships with organizations. What successful AI innovation looks like depends on the TRL level of the project and an organization's stage on the innovation pathway. Organizations with strengths in areas like statistics and analytics have a head start in building an understanding of AI, and "quick win" solutions can accelerate learning towards mature AI.

CSFs identified by the study include: address real business problems which have been carefully scoped, match the AI technology to the problem, build understanding of AI beyond the development team, invest in processes and systems which ensure high quality data, develop data capabilities, and embed AI expertise in the organization. Building understanding beyond the IT function is particularly important to managers and operations experts: to achieve success, AI shouldn't just be left to the technologists.

The study is limited by the perspectives of the participants interviewed so far. The continuing data collection work aims to recruit more participants from organizations which are using AI. These will include SMEs and large organizations, public and private sector, and more entrepreneurs, as well as established firms. The use of TRL as a contextual surrogate intentionally prompted participants to think about their projects in terms of an innovation pathway. This is useful while AI is not mature, but future CSF studies will likely need to look at alternative approaches to context.

References

Akkiraju, R., Sinha, V., X, A., Mahmud, J., Gundecha, P., Liu, X. & Schumacher, J. 2018. Characterizing machine learning process: A maturity framework. *ArXiv*, abs/1811.04871.

AlSheibani, S., Cheung, Y. & Messom, C. Artificial Intelligence adoption: AI-readiness at the firm level. *22nd Pacific Asia Conference on Information Systems*, 2018 Japan.

Axelsson, K. & Melin, U. Contextual Factors Influencing Health Information Systems Implementation in Public Sector – Investigating the Explanatory Power of Critical Success Factors. *EGOV 2014, LNCS* 8653, 2014. 59-71.

Boynton, A. C. & Zmud, R. W. 1984. An Assessment of Critical Success Factors. *Sloan Management Review*, 25, 17-27.

Chow, T. & Cao, D.-B. 2008. A survey study of critical success factors in agile software projects. *The Journal of Systems and Software*, 81, 961-971.
European Commission. 2020. *On Artificial Intelligence - A European approach to excellence and trust.*
Cullen, A. J. & Taylor, M. 2009. Critical success factors for B2B e-commerce use within the UK NHS pharmaceutical supply chain. *International Journal of Operations and Production Management*, 29, 1159-1185.
Davis, G. B. 1980. Letter to the Editor. *MIS Quarterly*, 4, 69-70.
Gursoy, D., Chi, O. H., Lu, L. & Nunkoo, R. 2019. Consumers' acceptance of artificially intelligent (AI) device use in service delivery. *International Journal of Information Management*, 49, 157-169.
Héder, M. 2017. From NASA to EU: the evolution of the TRL scale in Public Sector Innovation. *The Innovation Journal*, 22.
Hengstler, M., Enkel, E. & Duelli, S. 2016. Applied artificial intelligence and trust - The case of autonomous vehicles and medical assistance devices. *Technological Forecasting and Social Change*, 105, 105-120.
Hong, K.-K. & Kim, Y.-G. 2002. The critical success factors for ERP implementation: an organizational fit perspective. *Information & Management*, 40, 25-40.
Ivanov, S. & Webster, C. 2017. Adoption of robots, artificial intelligence and service automation by travel, tourism and hospitality companies – a cost-benefit analysis. *International Scientific Conference "Contemporary tourism – traditions and innovations"*. Sofia.
Joo, J. 2011. Adoption of Semantic Web from the perspective of technology innovation: a grounded theory approach. *International Journal of Human Computer Studies*, 69, 139-154.
Joshi, I. & Morley, J. 2019. *Artificial Intelligence: How to get it right. Putting policy into practice for safe data-driven innovation in health and care*. London, UK: NHSX.
Klumpp, M. & Zijm, H. 2019. Logistics innovation and social sustainability: how to prevent an artificial divide in human–computer interaction. *Journal of Business Logistics*, 40, 265-278.
Kumar, R., Singh, H. & Chandel, R. 2018. Exploring the key success factors of advanced manufacturing technology implementation in Indian manufacturing industry. *Journal of Manufacturing Technology Management*, 29, 25-40.
Lee, S. & Miller, S. 2019. Distilling managerial insights and lessons from AI projects at Singapore's Changi Airport (Part 2). *Asian Management Insights*, 6, 1-6.
Mahroof, K. 2019. A human-centric perspective exploring the readiness towards smart warehousing: The case of a large retail distribution warehouse. *International Journal of Information Management*, 45, 176-190.
Manning, C. D. & Schütze, H. 1999. *Foundations of Statistical Natural Language Processing*, MIT Press.
Meystel, A., Albus, J., Messina, E. & Leedom, D. 2003. Performance Measures for Intelligent Systems: Measures of Technology Readiness. *Performance Metrics for Intelligent Systems (PERMIS'03).*
Pumplun, L., Tauchert, C. & Heidt, M. A new organizational chassis for artificial intelligence - exploring organizational readiness factors. *Proceedings of the 27th European Conference on Information Systems (ECIS)*, 2019 Sweden.
Rockart, J. F. 1979. Chief executives define their own data needs. *Harvard Business Review*, 57, 81-93.
Snider, B., Da Silveira, J. C. & Balakrishnan, J. 2009. ERP implementation at SMEs: analysis of five Canadian cases. *International Journal of Operations and Production Management*, 29, 4-29.
Soliman, F., Clegg, S. & Tantoush, T. 2001. Critical success factors for integration of CAD/CAM systems with ERP systems. *International Journal of Operations and Production Management*, 21, 609-629.
Tambe, M., Jiang, A. X., An, B. & Jain, M. 2014. Computational Game Theory for Security: Progress and Challenges. *AAAI spring symposium on applied computational game theory*.
UK Government 2018. *Industrial Strategy Artificial Intelligence Sector Deal*.
Umble, E. J., Haft, R. R. & Umble, M. M. 2003. Enterprise resource planning: Implementation procedures and critical success factors. *European Journal of Operational Research*, 146, 241-257.
Van Esch, P., Black, J. S. & Ferolie, J. 2019. Marketing AI recruitment: The next phase in job application and selection. *Computers in Human Behaviour*, 90, 215-222.
W3C 2020. *JSON-LD 1.1: A JSON-based Serialization for Linked Data*.
Woo, H. S. 2007. Critical success factors for implementing ERP: the case of a Chinese electronics manufacturer. *Journal of Manufacturing Technology Management*, 18, 431-442.