Panel Discussion: Artificial Intelligence meets Enterprise Modelling

Monique Snoeck[1] [0000-0002-3824-3214], Janis Stirna[2] [0000-0002-3669-832X], Hans Weigand[3] [0000-0002-6035-9045], and Hend*erik* A. Proper[4] [0000-0002-7318-2496]

[1] Research Center of Information Systems Engineering, KU Leuven, Leuven, Belgium
[2] Department of Computer and Systems Sciences, Stockholm University, Stockholm, Sweden
[3] Tilburg University, Tilburg, The Netherlands
[4] Luxembourg Institute of Science and Technology, Belval, Luxembourg

Abstract. This paper reports on the panel discussion on "AI meets Enterprise Modelling" held at the PoEM 2019 conference in Luxembourg. The panel was concerned with the interplay between AI and enterprise modelling. Both in terms of how AI may benefit enterprise modelling activities, as well as conversely how the role of AI in modern-day enterprises provides new challenges to be included in enterprise modelling.

1 Introduction

We are moving towards an *Artificial Intelligence* (AI) intensive society, in which AI plays an increasingly prominent role in many, if not all, facets of social and economic life. In our homes, AI enabled thermostats allow us to optimise our energy consumption. When on the road, AI powered apps on our mobile phones help us find the best way to reach our destination. When working in a multi-lingual environment, AI techniques help us translate documents. In healthcare, AI based solutions aid doctors in producing better diagnoses. When mining resources in hard to reach places, AI comes to the aid as well, for example, to drive autonomous vehicles [11]. And, as we venture out to new frontiers, such as deep space, or the bottom of the ocean, AI plays a crucial role by taking over tasks that would be too dangerous for humans to perform [3, 13].

As more and more tasks are entrusted to AI based actors, a rich variety of AI techniques is being deployed, ranging from *machine learning* and *data analytics*, to the more traditional *logic* and *rule-based* approaches.

The topic of AI brings about many opportunities and challenges to enterprises in general, and thus, enterprise modelling in particular. From an entrepreneurial perspective, it offers many new possibilities to optimise existing processes and services, while also offering ample opportunities for new products and services. This is also confirmed by recent survey conducted by EY [6]: "*When asked for the top three desired business outcomes from the application of AI, the answers were: to improve and/or develop new products and services; achieve cost efficiencies and streamlined business operations, and to accelerate decision-making*".

At the same time, AI also raises fundamental concerns regarding different social and economic aspects, including ethics, security, privacy and trust. Workers may fear they will be replaced by AI-based systems, consumers, and lawyers, ponder the question of who is responsible in case a wrong decision is taken by an AI-based system.

The aim of this panel discussion, held at PoEM 2019 in Luxembourg, was to explore the challenges which the emerging AI intensive society brings to the field of enterprise modelling (EM), as well as opportunities that AI brings to enterprise modelling:

- *Are new modelling concepts needed to cater for AI-based actors?*
- *How to capture AI related regulative and / or ethical considerations in enterprise models?*
- *Should enterprise modelling encompass models / concepts for explainable AI?*
- *Can AI be used to automate enterprise modelling activities?*
- *Can enterprise models be used as inputs to AI based solutions?*

The panel session involved Monique Snoeck, Janis Stirna, and Hans Weigand as panellists, while Hend*erik* A. Proper chaired the session. The remainder of this paper provides the position statements of the respective panellists, while the conclusion reflects on the points brought forward in the original position statements, as well as the outcomes of the panel discussion.

2 The AI Actor – Intelligent Enterprise Modelling Needed

Position statement by: Monique Snoeck

For many years, Enterprise Architecture has revolved around a fairly stable set of concerns such as Value, Capability, Organisation structure, Processes, Data. The increased use of AI raises the question whether there is a need for addressing new concerns via separate architectural views. When processing is shifting from "*human intelligent*" actors or simple automation to "*machine intelligent*" actors, we can easily see that the current set of concerns remains relevant in the presence of artificial intelligent actors: data is collected, stored, processed by an actor, and output is generated in the form of information, decisions, new (derived) data. Hence, to a large extent, the current set of addressed concerns seems to provide sufficient coverage. In the end, an artificial intelligent service, is just another kind of actor, besides the human or automated actors.

But even if the high-level concerns to be addressed remain the same, the relative importance of the different views may need to be reconsidered. Recent evolutions of enterprise architecture standards seem to emphasise capability, value and process modelling at the expense of good old school data modelling. Yet, in practice, we see a lot of data-related problems having a negative impact on an organisation's capability to reap benefit from artificial intelligence: Data architecture problems such as a lack of integration, unproperly set authorisation, etc. cause a lot of data preparation work, up to 75% of the project's time. Meta-data problems, such as missing meta-data, cause rework, for example because the predictive features turn out not to correspond to their assumed meaning. Data model problems such as missing or wrongly defined associations may make data unusable or require the need for complex processing to make it usable. Data versioning problems may make data unusable as well: if the historical value of an attribute is required for good predictive algorithms, overwriting of data values makes a data set unusable. In the light of the increasing importance of good data management for reaping the value of data analytics and artificial intelligence, data architecture should more than ever be a primary cornerstone of enterprise architecture.

High quality data modelling also requires a better understanding of the impact of data modelling choices on capability development and service delivery, as well as of the capability debt resulting from data architecture deficiencies. This calls for more intelligent modelling techniques allowing for intelligent analysis and reasoning about data models and their impact on capability development and value creation. An example of such cross model checking for data models versus behavioural models can be found in the MERODE approach [16, 17]. Similar explicit definition and formalisation of correspondences between e.g. a data model and value models or capability models are needed.

Even if the set of concerns and corresponding views does not change drastically, at a more detailed level, one should consider if the models that capture the targeted concern by means of modelling concepts and their relations are still adequate. Modelling always involves reduction of the reality to the aspects deemed necessary to be addressed explicitly, thereby leaving out aspects that can safely be ignored considering the modelling goal. The increased use of artificial intelligence may raise the need for reconsidering what is addressed explicitly and what is left out or addressed only implicitly: a human actor is better able to cater for concerns that have not been considered by service designers than an artificial intelligent actor. As an example, a human actor will more easily adapt decision making to context specific factors, while an automated service will only address context if this was explicitly included in its design. The increased use of automated actors (artificial intelligent or not) therefore calls for more explicit incorporation of contextual, ethical, and other factors that today are either not addressed or only implicitly. Today, data modelling relies solely on the notions of classes, attributes and associations. Besides the data structure, it might be necessary to address more explicitly issues such as ownership, privacy, ethics, bias, etc. Furthermore, automated actors magnify errors, they do not compensate for design errors. Our designs need therefore to be flawless. So, also here, intelligent model checking would be an added value for better enterprise modelling practice.

In summary, the current set of high-level concerns addressed by enterprise modelling seem to provide a largely coverage, but data architecture should more than ever by a primary cornerstone of enterprise architecture. At the detailed level, modelling techniques may need to be enriched to explicitly address aspects that an artificial intelligent agent need to be instructed about. The quality of modelling and in particular making explicit and formalising connections between different architectural views needs to be improved so as to allow for intelligent model analysis and reasoning.

3 Exploring the intersection between AI and Enterprise Modelling

Position statement by: Janis Stirna

Technologies based on Artificial Intelligence (AI) have become more available, and as a result companies, public bodies, and governments have incorporated them in their IT products and services. AI technologies have enabled the creation of entirely new kinds of solutions that were not available even a few years back. Despite the overall enthusiasm and appreciation of these development, there have also been failures and concerns about whether the AI-based solutions are all that helpful and what is their long-term impact. In some cases, especially when it comes to issues of privacy and trust, the AI solutions have raised concerns about their negative societal impact. This position statement ponders on the current situation and seeks to identify areas of potential contribution between AI and enterprise modelling.

Currently, notable advancements of AI-based solutions have been in areas related to detection, recognition, and automatic documentation based on machine learning. These kinds of solutions have also been applied to tasks that require monitoring, adaptation and adjustment as well as gathering feedback loops. According to Pearl & MacKenzie [10] this has been possible already for some time and what the current contributions mostly do is curve fitting to represent irregularities in large sets of data. Some of these efforts create ground breaking results, which has fuelled a lot of enthusiasm even if they have been theoretically possible for some time. According to Chauvet [4], if the past is any tell, the enthusiasm in AI occurs in 30-year cycles – each cycle starting with what is seen as dramatic successes, followed by philosophical critique and difficulties in tackling harder and more novel problems. During the previous two cycles this led to what Chauvet (ibid.) calls stagnation. What distinguishes the current situation from the past is the widespread interest in AI from the business community which is manifested by the availability of venture capital and innovative solutions. While many of these offerings are promising significant impacts there are also concerns that some of them are plain impossible to deliver the advertised value and are, in fact, what Narayanan [12] calls, "*a digital snake oil*". We can, of course, pay no attention to the irrationality of those that think such "*magic tools*" (e.g. [5]) improve their business and regard that the choice to use these tools is at their peril only. However, with the pervasiveness of ICT and its transcending role, such solutions and the decisions to use them have the potential of negatively impacting many individuals and the society. We should also consider that AI is used by public organisations and governments. While goals of such systems are efficiency and, in the case of surveillance systems, public safety, they can also be, and often are, used for actions that contradict with the values of open and democratic societies. They routinely raise questions about trust, transparency, security, privacy, integrity. Hence, a risk is that over time AI solutions become "*the digital asbestos*" – great on the outset but harmful once the content and impact have accumulated. Currently the impact of such solutions has not been studied in depth and the development of AI based ICT falls back on general principles and guidelines, such as the ACM Code of Ethics for software engineers [1] or US Department of Defence guidelines for developing AI solutions [18].

With respect to the current state of the art in the development of AI based solutions, EM is able to contribute with capturing and analysing the business motivation and requirements for them. Sandkuhl et al. [15] report that one of the motivation of EM use in organisations is the need to elaborate the details of business and IT alignment. In this case enterprise models serve as blueprints or requirements specifications for the information system. Using enterprise models in AI projects would allow addressing the two issues that Howard [9] suggests for organisations, namely, (i) being more considerate about their data and (ii) asking the right questions and not being fixated on designing solutions around the data that is most easily available. Modelling the business dimension of AI solutions would mitigate the problem reported by Sandkuhl in [14] that many companies consider developing the solutions first before a clear business problem is identified. The latter paper (ibid.) also takes a stance that AI projects are similar to most information system development projects and hence they need to follow steps that ensure their fit to the business and its surroundings. A particular attention should be paid to the role an AI based solution has on the environment, considering all its facets – individual, business, and societal. In essence, capturing and documenting these issues requires modelling of digital business ecosystems, including interdependence of all actors in the business environment, who co-evolve their capabilities and roles [7].

Contemporary EM methods and tools have been developed for modelling organisations from various perspectives such as goals, processes, concepts, information, actors, etc. The predominant approach of analysis has been from a position of a single organisation or solution, including its customers and suppliers. This needs to be extended for the purpose of modelling the interdependence of all actors in the business ecosystem, including their evolution, self-organisation, and self-optimisation. From the point of view of EM this raises a number of interesting research questions, such as, should enterprise models be discovered from the existing solutions; what data sources can be credibly used in this process; can such models be used to analyse the compliance and impact; can they be used to determine the mismatch between the design and the actual implementation; should the EM languages be more formal, which would be contradicting the overall principle of offering a light-weight language to support discussions with stakeholders; and what new perspectives on the organisation and ecosystem should be integrated in EM methods and tools.

4 What's in a model? Enterprise Modelling meets AI

Position statement by: Hans Weigand

The basic idea of enterprise modelling is to offer different (abstract) views on an enterprise. The two main goals are (a) making up the design of an information system (b) providing understanding, first between system analysts and the system developers, but also between different business functions. The underlying worldview is that of automated systems that can be specified completely on a formal declarative level on the one hand and human agents in control over these systems who need to coordinate their actions in order to enable integration of the systems, on the other hand. This worldview is being challenged nowadays with the arrival of AI and robotics in the enterprise. Although so far it is more a vision than reality, clear steps into this direction are made, e.g. under the umbrella of Industry 4.0. The alternative worldview is one of intelligent systems working and evolving largely autonomously in cooperation with or completely replacing human agents. Therefore, the systems are not under complete control of the human agents. Coordination and integration will still be needed, but in a network of both intelligent systems and human agents. What is the role of enterprise models in this alternative world, and are the current modelling languages still adequate?

First, my claim is that the autonomy of intelligent systems should not be overstated. Even if we ignore the ethical choices about human responsibility, we must acknowledge that also future intelligent systems will be embedded in an environment. Autonomy should not be seen as an assumed requirement as but as a required assumption. It is not that the intelligent system must be autonomous, but we must assume autonomy (even if we know that it is not autonomous). Or, to be more specific: autonomy should be the guiding principle for how agents interact with each other.

Second, although autonomy is not that absolute, still the old worldview of a completely controlled system seems to become obsolete. In my view, this requires a shift in the role of "*models*". One shift is that instead of design specifications, we should think of models as approximate ways to understand complex systems. The other shift is that the models should focus on the *interactions* within the complex. The latter shift is not that big – interactions are already in the scope of enterprise models, although perhaps not always respecting the assumed autonomy of the elements. Both shifts have been explored already for years in the field of multi-agent systems (MAS).

MAS are a convenient way of understanding, modelling, designing and implementing different kind of (distributed) systems. Firstly, MAS is a programing paradigm to develop operational software system with its own languages (like JADE) and constructs (like goals, beliefs, intentions, contracts). MAS are particularly suited to deploy distributed software systems that run in computational contexts wherein a global control is hard or not possible to achieve. In the course of time, MAS have been used more and more for simulation purposes. MAS simulation environments (testbeds) can be small or complex. According to [8], quite a few simulation systems for multi-agent systems have been built, but, due to the diversity of agents and objectives of simulation studies, a widely accepted modelling and simulation approach has remained elusive. I do not have a simple solution. However, I do want to suggest that enterprise modelling (a) would seek closer collaboration with the MAS field, and see how it can incorporate proven concepts (taking advantage of what MAS has learned over the years), and (b) should shift its attention from design models towards simulation models. The latter means, among others, that the goal is not the specification of a running system, but finding out what is an optimized system, given certain stochastic assumptions about inputs and environmental constraints, or what-if scenarios. The resulting model can be used for realizing an operational system, but this realisation is better described as intervention than as implementation.

If autonomy is taken as the guiding interaction principle, this has an impact on the interaction modelling. In the MAS field, contracts have been introduced as a flexible way in which agents can reduce some uncertainty and avoid some risks, without prescribing each other actions in all details. Regulative and ethical considerations can be included explicitly in these contracts, as part of the design. Transparency (as responsibility to account for some action or non-action) can also be part of the contractual requirements (see [2] for a literature review of explainable agency in MAS).

5 Conclusions

In this section, we summarise the key findings from the panel discussion.

Data remains a key resource – A first conclusion from the panel discussion is that with the advent of AI the need for traditional *data management* increases rather than decreases. With data management we refer to the usual mix of collecting, storing, and processing of data by (a hybrid mix of human and digital) actors, while the resulting data may provide other (human or digital) actors with the information to take decisions. The dependence of AI for data, both for training statistical AI, as well as operational use of AI solutions, makes it more important than ever to include data related aspects in enterprise modelling activities. In practice, we see a lot of data-related problems having a negative impact on an organisation's capability to reap the benefits from AI and / or to avoid risks of failing AI due to incorrect handling of data.

Data comes at a price – Modern day AI applications are "addicted" to data. This not only makes it important to consider data from a traditional data management perspective, but rather also include additional concerns, such as ownership of the data, ownership of "the original" (i.e. privacy), ethics, biases, etc. As such, data indeed comes at a price.

Design for AI – Next to providing foundations for proper data management, enterprise modelling can contribute to the design of enterprise-ready AI based solutions, by capturing and analysing the business motivation and needs for them.

More study of impact needed – AI solutions run the risk that over time they turn out to be "*digital asbestos*" – initially seen as efficient for the intended purpose, but harmful once the impact and side effects of the application have accumulated. The health hazards of asbestos were discovered only years later and currently the impact of AI-based solutions has not been studied in depth. Such impact studies should also address breadth, i.e. the whole ecosystem in which they operate. Enterprise modelling techniques can be used to better chart out the (potential and materialised) impact of AI, thus also gathering, and sharing, experiences.

Compliant AI – Whether or not AI can achieve real autonomy can be debated. Both from a practical and a philosophical stance. What is more important, however, is the fact that AI based systems are part-of / embedded-in a context. This implies that, regardless of the level of autonomy an AI based system achieves, it must comply to the regulations (and ethics) of the socio-technical environment in which it operates.

Data for enterprise modelling – The availability of large amounts of data does not only benefit AI applications. Enterprise modelling itself may also clearly benefit from the data as it can be used (1) as a base to derive / mine models (in other words, "*enterprise model mining*" in general, and not just e.g. process model mining), (2) as validation / evidence of the correctness of enterprise models, and (3) provide real-time insights in the performance of an enterprise, using enterprise models to provide context / structure to the data. As such, the panel also discussed the need to make a distinction between "design models" that portray (parts or aspects of) a possible future / desired state of affairs of an enterprise, and "observational models" that portray (parts / aspects of) the current / past affairs of an enterprise, including models representing runtime of an enterprise.

Return to multi-agent systems – With more and more "AI-driven autonomous actors" that collaborate with humans, it becomes relevant for enterprise modelling to embrace lessons learned in the multi-agent systems (MAS) community. As modern-day enterprises increasingly involve a hybrid mix of human and digital agents, it is relevant to consider these as large-scale (and hybrid) multi-agent systems.

References

[1] The Software Engineering Code of Ethics and Professional Practice. Technical report, The Association for Computing Machinery, Inc. and the Institute for Electrical and Electronics Engineers, Inc., 1999. https://ethics.acm.org/code-of-ethics/software-engineering-code/

[2] S. Anjomshoae, A. Najjar, D. Calvaresi, and K. Främling. Explainable Agents and Robots: Results from a Systematic Literature Review. In *Proceedings AAMAS '19*, pages 1078-1088, 2019.

[3] AZoRobotics. Using artificial intelligence and autonomous robotics for rapid exploration of deep-sea ecosystems, September 2018. https://www.azorobot-ics.com/News.aspx?newsID=10037

[4] J.-M. Chauvet. The 30-Year Cycle In The AI Debate. Technical report, 2018. https://arxiv.org/abs/1810.04053v1

[5] Harwell D. A face-scanning algorithm increasingly decides whether you deserve the job. *The Washington Post*, November, 6th 2019. https://www.washingtonpost.com/technol-ogy/2019/10/22/ai-hiring-face-scanning-algorithm-increasingly-decides-whether-you-de-serve-job/

[6] EY. The Growing Impact of AI on Business. MIT Technology Review. https://www.tech-nologyreview.com/s/611013/the-growing-impact-of-ai-on-business/

[7] Nachira F., Nicolai A., Dini P., Le Louarn M., and Rivera Le'on L., editors. *Digital Business Ecosystems*. European Commission, Bruxelles, 2007. ISBN: 92-79-01817-5

[8] M. Fabien, J. Ferber, and A. Drogoul. Multi-Agent Systems and Simulation: A Survey from the Agent Community's Perspective. In *Multi-Agent Systems*, number 17-66. CRC Press, 2018.

[9] A. Howard. Demystifying the Intelligence of AI. *MIT Sloan Management Review*, Nov 5 2019. https://sloanreview.mit.edu/article/demystifying-the-intelligence-of-ai/?utm_source=twitter&utm_medium=social&utm_campaign=sm-direct

[10] Pearl J. and MacKenzie D. *The Book of Why: The New Science of Cause and Effect*. Ingram. ISBN: 9780465097609

[11] B. Marr. The 4th Industrial Revolution: How Mining Companies Are Using AI, Machine Learning And Robots, September 2018. https://tinyurl.com/y2evvodz

[12] A. Narayanan. How to recognize AI snake oil, presentation. Presentation, Princeton University, 2019. https://www.cs.princeton.edu/ arvindn/talks/MIT-STS-AI-snakeoil.pdf ~

[13] S. O'Kane. NASA's Mars rover is really good at laser-blasting rocks without human input - Curiosity has been studying Mars on its own for a year, 6 2016. https://ti-nyurl.com/yaodb7p5

[14] K. Sandkuhl. Putting AI into Context - Method Support for the Introduction of Ar-tificial Intelligence into Organizations. In *IEEE 21st Conference on Business Informatics (CBI)*, 2019. doi:10.1109/CBI.2019.00025

[15] K. Sandkuhl, J. Stirna, A. Persson, and M. Wißotzki. *Enterprise Modelling: Tackling Business Challenges with the 4EM Method*. Springer, Heidelberg, Germany, 2014. ISBN: 978-3-662-43724-7S

[16] M. Snoeck and G. Dedene. Existence dependency: The key to semantic integrity between structural and behavioral aspects of object types. *IEEE Transactions on software engineering*, 23(4):233-251, 1998.

[17] M. Snoeck, C. Michiels, and G. Dedene. Consistency by construction: The case of merode. In M. A. Jeusfeld and O. Pastor, editors, *Conceptual Modelling for Novel Application Domains*, number 2814 in Lecture Notes in Computer Science, pages 105-117. Springer, Heidelberg, Germany, 2003. ISBN: 978-3-540-39597-3 doi:10.1007/978-3-540-39597-3_11

[18] AI Principles: Recommendations for the Ethical Use of Artificial Intelligence by the Department of Defense. Technical report, US Department of Defense, United States of America. https://admin.govexec.com/media/dib_ai_principles_-_supporting_document_-_embargoed_copy_(oct_2019).pdf