On the Interplay between Requirements, Engineering, and Artificial Intelligence

Blagovesta Kostova
EPFL
blagovesta.pirelli@epfl.ch

Seda Gürses
TU Delft / KU Leuven
f.s.gurses@tudelft.nl

Alain Wegmann
EPFL
alain.wegmann@epfl.ch

Abstract

With this paper, we present our reflections on the issues that are faced by the Requirements Engineering academic discipline and practice. The current reality of the Artificial Intelligence and Machine Learning hype penetrating from research into all industry sectors and all phases of system design and development is a transformative shift that influences the way Requirements Engineering is conducted and the nature of the systems that are engineered. We identify two sides of this transformation with regards to the Requirements Engineering discipline: (1) Artificial Intelligence tools are used more and more during the Requirements Engineering process, (2) the Requirements Engineering process for systems that include Artificial Intelligence is different. By identifying and framing these changes, we pose questions about what it means to engineer requirements. Our analysis asks more questions than it answers. We hope to engage the Requirements Engineering academic community in a larger conversation about the role of Requirements Engineering in the changing world and about a possible new vision of engineering becoming secondary to requirements in Requirements Engineering.

1 Introduction

To better understand the new situation that Requirements Engineering (RE) practitioners and academics face, we begin by inquiring how these actors construct their worldview, regarding the RE process. According to Jureta *et al.* [JMF09]: “Because requirements-related information is elicited through the communication between the stakeholders and the software engineers, it is via communication that stakeholders reveal their understanding of the system-to-be and its relevant surroundings. To ground an RE ontology in how stakeholders reveal their understanding of the system and its environment is to ground the ontology in a conceptualization of the communication process. [...] Namely, the scope of RE conceptualizations should be determined by **how and what can be communicated by the stakeholders during requirements elicitation**.”[1] This description represents the dominant assumption in the RE field, also inline with Zave and Jackson [ZJ97]: Reality is observed and communicated by human stakeholders.

In: M. Sabetzadeh, A. Vogelsang, S. Abualhaija, M. Borg, F. Dalpiaz, M. Daneva, N. Fernández, X. Franch, D. Fucci, V. Gervasi, E. Groen, R. Guizzardi, A. Herrmann, J. Horkoff, L. Mich, A. Perini, A. Susi (eds.): Joint Proceedings of REFSQ-2020 Workshops, Doctoral Symposium, Live Studies Track, and Poster Track, Pisa, Italy, 24-03-2020, published at http://ceur-ws.org

[1] Emphasis added.

Even though this assumption is not entirely outdated, as humans are still very much involved in the RE process, the introduction of data-driven RE augments the way stakeholders obtain and communicate the information about their reality. Let us take requirements elicitation as an example and reason about these changes. Traditionally, requirements are thought to be elicited with the help of various methods: interviews, workshops, questionnaires, ethnographic studies (cf. [oBA15]). Recently, these mostly human-to-human methods are supplemented and, in some cases, completely replaced by data that is gathered from preexisting sources: gathering existing data from social media, online forums [WM17], automated tools to elicit requirements [FBG18, RM19], prioritizing requirements ([PSA13, GIG17]), refining requirements into specifications [MDM$^+$19], interpreting requirements and classifying them [DDAÇ19, KM17, WV16].

In this position paper, we outline certain open questions regarding the use of Artificial Intelligence (AI) in RE in Section 3 and the changes in the RE process for AI-based systems in Section 4. To explain what we consider AI, we begin by a short terminology section in Section 2. We conclude with an outlook for the future in Section 5.

2 What is AI?

We begin with a note on AI terminology. AI is a term used throughout the industry to signify a system that is self-intelligent and capable of learning on its own and of taking decisions without the need of human supervision. The AI components of systems usually require a resource-intensive (resources being data, time, computing power, expertise) training phase that *could* output a useful algorithmic model. The correctness of a model is measured in terms of accuracy over new data (unseen during the training phase). Recently, other metrics have attempted to address concerns about the "ethical accuracy" of algorithms, e.g., fairness [Nar18].

A word of caution we would like to raise here is that AI technologies are predominantly used to optimize business processes in enterprises and not as fully autonomous systems [KOTG20]. The objective of the optimizations is to remove any inefficiencies in both the delivery of a system (inside the enterprises) and the consumption of a system (behavior of the consumer, see popular literature on behavior engineering [Eya14]) in order to maximize profit. Therefore, we focus on the current role of AI as a means to optimize perceived inefficiencies rather than as a technology that is autonomous and salient.

In the field of AI, there is a concept called technological singularity. It signifies the moment when the technology will become independent of its designers: Artificial General Intelligence (AGI) [Cha09]. However, this is a much disputed topic and even if AGI is possible, which is still questionable, singularity is a futuristic idea that we leave out of the scope of our analysis. Therefore, we focus on the impact of level of automation provided by non-general AI (which we label this automated decision making) on the RE process rather than the potential impact by hypothetical autonomous AGI.

3 AI for RE: What Changed?

3.1 What AI-inspired methods for RE are in use?

The brief glance at the literature, as well as other events such as workshops (see AIRE [air]), panels, and altogether themes of the RE flagship conferences ("RE and Collective Intelligence in the Days of AI" of RE'19 [DPL19]), shows that the RE academic community is making efforts to produce and to use automated decision-making tools during the RE process. However, this new trend of increasing the level of automation during the RE process changes the nature of the RE practice, as well as the research in the area. We believe that the automation tools change the essence of RE and that we, as a community, have not yet systematically studied this transformation in levels of automation. We believe that this area merits research to better understand, in the depth and the breadth, the transition towards automated decision-making methods during the RE process. The inquiry concerns both the state of the art and the state of practice.

3.2 How we know what we know in the RE process?

Based on the changed nature of the RE tools, there is a also a change in the way the RE discipline views the world. The assumption, as described by Zave and Jackson [ZJ97] and Jureta *et al.* [JMF09], is that the RE practitioner will communicate with stakeholders to elicit requirements from the "messy" world and to bring some order into these observations, while collecting data from a mix of qualitative and quantitative methods and refining the requirements into a technical specification. The existing RE ontologies do not consider the automation of feedback and data collection, because this current level of automation did not exist when the ontologies were introduced. Currently, the RE process is being fragmented and transformed based on the environment of the systems that are studied through data that are generated by systems that we have engineered. Our observations about the context of the system are increasingly based only on secondary data. The core RE ontology is changing because of the inclusion of AI, hence, studying this change is an avenue for future research inquiries.

4 RE for AI: What Will Change?

4.1 How can RE bring reason to buzzword concepts?

Understanding the ontological change also raises a question about the role of RE. RE is traditionally a crossroad of disciplines and is uniquely positioned among other domains to address questions of not only what and how to engineer a system but also of whether we should engineer these systems, why, to what extent, and *if at all*. These questions might seem philosophical, yet, AI has become a buzzword. RE could be a place where we counter buzzwords and evaluate systematically what is of actual use in our systems by answering these seemingly philosophical questions.

4.2 Is there a place for an RE human in the loop?

There is a historical distrust in humans that has gone deep into technical designs and engineering methods (starting with the Cold War) [Bai83]. The question is whether there is a possibility to recenter the human (and at times the non-human, such as the environment) in computer science. Certain methods try to mimic human-like or human-friendly characteristics while still eliminating human intervention and interaction with the systems, i.e., chaos techniques (introduce “real” randomness) or to introduce so-called interpretable models whose understandability is debatable [Lip18]. RE is traditionally a discipline related to software engineering. Yet, RE has two components: requirements and engineering. If we assume that the engineering part is where we activate and use our analytical thinking, we can manage the engineering of the requirements. But to devise the requirements, we might need to engage in an activity of an intuitive nature (artistic practices, design of socio-technical systems in a broad sense and not only software engineering.) Not all requirements lead to software systems and this point of departure is far-fetched in the current reality where software is everywhere and the RE tradition is to engineer requirements for software systems. Yet, we can imagine a future that does not require us to engineer but rather to understand and to design social solutions that are, possibly but not necessarily, supported by technology. All these creative problem-solving activities are uniquely human.

4.3 What is the role of RE in the design of automated decision-making systems?

Today, the designers of algorithms and systems are humans. By acknowledging and deliberately owning the design of systems, we can re-introduce humans in the loop. And if there is a place for a human translator, we have to answer more questions regarding our role in the RE process such as, Are we re-affirming existing structures, are we instilling our biases, and are we aware of it?

We believe we could be designing systems that enable people to be more than or different from what the data predict about them, in order for them to become what they could. In AI, the automated decision-making components are self-fulfilling prophesies – the models are tuned to find more of the same and to discard the unique. Uniqueness is outlawed, it is seen as a mistake in the pattern and is reported as an error. Yet, originality and ingeniousness are what enable humankind to question the current status and to progress. If our systems are designed to ensure that we are only what we are in the current moment, based on historical data, then we, as designers of these systems, are hindering ourselves from becoming different. Dreams cannot be codified.

RE could be a solution to counter the trend of engineering systems that push their users to regress to the mean. We believe there is a place for us to rethink what it means to be a deviation in the pattern and to allow for errors in the AI model. It might turn out that the mistake is correct. Currently, people are being profiled and shown a version of reality that is optimized for certain behavior (e.g., recommendations for what to watch, buy, like, read). This profiling specializes the reality that we observe from our screens, thus providing a unique user experience for each one of us. Yet, the set of actions that we take within these systems is limited. Consequently, we are both influenced by the data generated by others that look like us data-wise and isolated in such a fine-tuned single reality that prompts us to act in certain limited ways and still, we cannot reconcile with others what we see. RE is the only place to address this problem due to its interdisciplinary nature, with a strong technical emphasis.

4.4 What is a requirement for AI?

The reality is that AI components are being developed and used throughout many applications. We still have to answer questions about the pipeline of requirements and the RE process for integrating AI components in our systems. What would be a requirement for these AI components? How do we evaluate a system composed of AI components? The state of the use of data-driven services by industry has advanced quickly, and the industry has established their own practices. Hence we might have to study more closely the recent developments, yet distance ourselves in order to propose theoretically grounded methodologies for a sensible use of AI in systems, thus combat buzzwords. For example, AI could be designed for and evaluated through the lens of socio-technical systems and the collective good [CV18].

Only a limited number of studies have looked into how the introduction of AI, and mostly ML, in software systems changes the RE and the software engineering processes [ABB+19, VB19, Hor19]. All of these studies identify one of the starkest differences between data-driven software engineering and traditional engineering: the management of the data. There is even an emerging field, called ML Ops, for the study of the operationalization and quality assurance/testing of ML[2].

5 Outlook for the Future

We cannot ignore the current reality that dictates that AI is going to be integrated into our systems in the short term. We can change, however, the narrative to see through the buzz of the term AI and to look at current AI components as a machinery for automated decision-making. We believe that we need to study the domain further and to propose theories, methodologies, and tools to answer the questions we outlined here. (1) AI for RE: *What AI-inspired methods for RE are in use? How we know what we know in the RE process?* And (2) RE for AI: *How can RE bring reason to buzzword concepts? Is there a place for an RE human in the loop? What is the role of RE in the design of automated decision-making systems? What is a requirement for AI?* If the current trend of incorporating more and more AI tools in the RE process continues (AI for RE), while also having this RE process engineer AI tools (RE for AI), then soon we will combine the two implications into a single conjecture that RE will become automated and that there will be only AI for AI. We also do not have to subscribe to the worldview of today and can imagine that there is another way for our systems to be engineered and that the role of RE will be pivotal in bringing the human back in the loop and in making that human an integral part of the process.

References

[ABB+19] Saleema Amershi, Andrew Begel, Christian Bird, Robert DeLine, Harald C. Gall, Ece Kamar, Nachiappan Nagappan, Besmira Nushi, and Thomas Zimmermann. Software engineering for machine learning: a case study. In *International Conference on Software Engineering: Software Engineering in Practice, ICSE SEIP*, 2019.

[air] *International Workshop on Artificial Intelligence for Requirements Engineering.* IEEE.

[Bai83] Lisanne Bainbridge. Ironies of automation. *Autom.*, 19(6), 1983.

[Cha09] D Chalmers. The singularity: A philosophical analysis. *Science fiction and philosophy: From time travel to superintelligence*, 2009.

[CV18] Federico Cabitza and Francesco Varanini. Going beyond the system in systems thinking: The cybork. In Cecilia Rossignoli, Francesco Virili, and Stefano Za, editors, *Digital Technology and Organizational Change*, 2018.

[DDAÇ19] Fabiano Dalpiaz, Davide Dell'Anna, Fatma Basak Aydemir, and Sercan Çevikol. Requirements classification with interpretable machine learning and dependency parsing. In *International Requirements Engineering Conference, RE*, 2019.

[DPL19] Daniela E. Damian, Anna Perini, and Seok-Won Lee, editors. *International Requirements Engineering Conference, RE*, 2019.

[2]https://www.forbes.com/sites/cognitiveworld/2020/01/03/how-do-you-test-ai-systems/

[Eya14] Nir Eyal. *Hooked: How to build habit-forming products.* Penguin, 2014.

[FBG18] Edwin Friesen, Frederik Simon Bäumer, and Michaela Geierhos. CORDULA: software requirements extraction utilizing chatbot as communication interface. In *International Conference on Requirements Engineering: Foundation for Software Quality Workshops, REFSQ Workshops*, 2018.

[GIG17] Emitza Guzman, Mohamed Ibrahim, and Martin Glinz. Prioritizing user feedback from Twitter: A survey report. In *International Workshop on Crowd Sourcing in Software Engineering*, 2017.

[Hor19] Jennifer Horkoff. Non-functional requirements for machine learning: Challenges and new directions. In *International Requirements Engineering Conference, RE*, 2019.

[JMF09] Ivan Jureta, John Mylopoulos, and Stéphane Faulkner. A core ontology for requirements. *Applied Ontology*, 2009.

[KM17] Zijad Kurtanovic and Walid Maalej. Automatically classifying functional and non-functional requirements using supervised machine learning. In *International Requirements Engineering Conference, RE*, 2017.

[KOTG20] Bogdan Kulynych, Rebekah Overdorf, Carmela Troncoso, and Seda F. Gürses. POTs: protective optimization technologies. In *Conference on Fairness, Accountability, and Transparency*, 2020.

[Lip18] Zachary C. Lipton. The mythos of model interpretability. *Commununication of the ACM*, 2018.

[MDM+19] Laurens Müter, Tejaswini Deoskar, Max Mathijssen, Sjaak Brinkkemper, and Fabiano Dalpiaz. Refinement of user stories into backlog items: Linguistic structure and action verbs - research preview. In *Requirements Engineering: Foundation for Software Quality, REFSQ*, 2019.

[Nar18] Arvind Narayanan. Translation tutorial: 21 fairness definitions and their politics. In *Conference on Fairness, Accountability, and Transparency*, 2018.

[oBA15] International Institute of Business Analysis. *A guide to the business analysis body of knowledge (BABOK Guide).* 2015.

[PSA13] Anna Perini, Angelo Susi, and Paolo Avesani. A machine learning approach to software requirements prioritization. *IEEE Trans. Software Eng.*, 2013.

[RM19] Tim Rietz and Alexander Maedche. Ladderbot: A requirements self-elicitation system. In *International Requirements Engineering Conference, RE*, 2019.

[VB19] Andreas Vogelsang and Markus Borg. Requirements engineering for machine learning: Perspectives from data scientists. In *International Requirements Engineering Conference Workshops, REW*, 2019.

[WM17] Grant Williams and Anas Mahmoud. Mining Twitter feeds for software user requirements. In *International Requirements Engineering Conference, RE*, 2017.

[WV16] Jonas Winkler and Andreas Vogelsang. Automatic classification of requirements based on convolutional neural networks. In *International Requirements Engineering Conference, RE*, 2016.

[ZJ97] Pamela Zave and Michael Jackson. Four dark corners of requirements engineering. *ACM Trans. Softw. Eng. Methodol.*, 1997.