

# Preface: The 4th International Workshop on Requirements Engineering for Artificial Intelligence (RE4AI'23)

Renata Guizzardi<sup>1</sup>, Jennifer Horkoff<sup>2</sup>, Anna Perini<sup>3</sup> and Angelo Susi<sup>3</sup>

<sup>1</sup>University of Twente, The Netherlands

<sup>2</sup>Chalmers | University of Gothenburg, Sweden

<sup>3</sup>Fondazione Bruno Kessler, Trento, Italy

## 1. Preface

It is our pleasure to welcome you to the fourth edition of the International Workshop on Requirements Engineering for Artificial Intelligence (RE4AI'23), held in conjunction with the 29th Working Conference on Requirements Engineering (REFSQ 2023), in the city of Barcelona, Spain.

RE4AI aims to provide a forum for discussing how Requirements Engineering methods, techniques and tools may be used to support the development of Artificial Intelligence systems that are lawful, ethical and robust.

**Workshop Format:** Based on feedback from the previous edition of this workshop, we decided to change the format of the workshop this year, proposing sessions with the presentation of talks by invited speakers. Many of these speakers are the authors of former papers in RE4AI, i.e., researchers who have been busy with the development of work in this area for at least 4 years. Others have been invited due to the recognition that they too develop relevant research in the area, although they haven't yet had the opportunity to publish at this forum.


Specifically we invited speakers to elaborate on their experience with RE4AI, including a consideration of questions like: What is the role of RE in engineering AI-based systems?<sup>1</sup> Are RE foundational concepts and theory still useful?<sup>2</sup> As RE researchers, what can we learn from AI

---

*In: A. Ferrari, B. Penzenstadler, I. Hadar, S. Oyedeji, S. Abualhaija, A. Vogelsang, G. Deshpande, A. Rachmann, J. Gulden, A. Wohlgemuth, A. Hess, S. Fricker, R. Guizzardi, J. Horkoff, A. Perini, A. Susi, O. Karras, A. Moreira, F. Dalpiaz, P. Spoletini, D. Amyot. Joint Proceedings of REFSQ-2023 Workshops, Doctoral Symposium, Posters & Tools Track, and Journal Early Feedback Track. Co-located with REFSQ 2023. Barcelona, Catalunya, Spain, April 17, 2023.*

✉ r.guizzardi@utwente.nl (R. Guizzardi); jennifer.horkoff@gu.se (J. Horkoff); perini@fbk.eu (A. Perini); susi@fbk.eu (A. Susi)

ORCID 0000-0002-2019-5277 (J. Horkoff); 0000-0001-8818-6476 (A. Perini); 0000-0002-5026-7462 (A. Susi)

 © 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

<sup>1</sup>See ideas proposed in IEEE SW - doi 10.1109/MS.2022.3196205 "Can RE Help Better Prepare Industrial AI for Commercial Scale?" By Boris Scharinger, Markus Borg, Andreas Vogelsang, and Thomas Olsson

<sup>2</sup>See "point/counterpoint" discussion provided by Mary Shaw and Liming Zhu - doi 10.1109/MS.2022.3203200 - "Can Software Engineering Harness the Benefits of Advanced AI?"

applications in different domains, such as autonomous vehicles, or social media - e.g. chatGPT, or other real projects you are involved in? How is/will RE integrate with policies and (about to come out) standards?<sup>3</sup>

**Workshop Program:** The workshop held invited talks in two sessions. Each talk was given roughly twenty minutes for both presentation and discussion. We were delighted to attract prominent RE for AI researchers from several countries and universities, with diverse experiences and opinions. In the following, we list the talk titles, authors, and abstracts.

**Presenter:** *Daniel Berry*, University of Waterloo, Canada

**Title:** “RE for AI: A Hot Topic as Viewed by an RE Alter Kaker”

**Abstract:** I explore how not-very-successful attempts to fit requirements for artificial intelligences (AIs) and learned machines (LMs) into the traditional RE mold led to a rethinking about RE for AI. I talk about some implications of the rethinking.

**Presenter:** *Beatriz Cabrero-Daniel*, Chalmers | University of Gothenburg, Sweden

**Title:** How I learned to stop worrying and discuss about AI

**Abstract:** AI-powered tools, such as autonomous vehicles or chatbots, are increasingly relying on foundation models. These models are incredibly complex, and training and operating them require significant computing resources... which can block academic research, key in evaluating the potential harms of such systems and informing about them. There is therefore a need for collaboration between various societal actors comes in. Authorities, engineers, and ethicists must work together to understand the implications of this disruptive AI technology and protect and inform individuals (and society as a whole). The first step to have informed discussions about and governance mechanisms for AI is agreeing on the right terms to use. Nevertheless, one might ask oneself whether it is possible to set solid grounds when the AI field is changing so rapidly: we might as well learn to stop worrying and love the AI ;)

**Presenter:** *Jaelson Castro*, Federal University of Pernambuco, Brazil

**Title:** **Can we trust Socially Assistive Robots? What are their requirements?**

**Abstract:** The development and use of Socially Assistive Robots (SARs) has grown significantly in recent years. Trust is one of the critical aspects for the adoption of robots in a social setting. Trust in Requirements Engineering is considered a non-functional requirement that needs to be properly satisfied. Hence it needs to be elicited, modeled, validated, and managed. The purpose of this research is to investigate what their requirements are. We report on the challenges, candidate solution paths, and research priorities regarding RE4AI for trustable SARs. Our initial goal is to develop a catalog of Non-Functional Requirements (NFRs) that is adequate to support elicitation and specification in projects of socially assistive humanoid robots (SARs), allowing the identification of possible problems from the point of view of Trust.

---

<sup>3</sup>For example: EU AI act <https://artificialintelligenceact.eu/>; ACM Statement on Principles for Responsible Algorithmic Systems [www.acm.org/binaries/content/assets/public-policy/final-joint-ai-statement-update.pdf](http://www.acm.org/binaries/content/assets/public-policy/final-joint-ai-statement-update.pdf); 7000-2021 - IEEE Standard Model Process for Addressing Ethical Concerns during System Design - <https://ieeexplore.ieee.org/document/9536679>

**Presenter:** *Xavier Franch*, Universitat Politècnica de Catalunya, Spain

**Title:** A Requirements Engineering Perspective to AI-based Systems Development

**Abstract:** This talk will reflect on the role that RE should play in the development of AI-based systems with a focus on three areas: roles involved, requirements' scope and non-functional requirements. The position taken is that requirements engineers shall become the cornerstone in AI-based system development in collaboration with other key roles.

**Presenter:** *Hans-Martin Heyn*, Chalmers | University of Gothenburg, Sweden

**Title:** The challenge of creating an architecture framework for Very Efficient Deep Learning in the IoT

**Abstract:** The VEDLIoT (Very Efficient Deep Learning in the Internet-of-Things) project has run since the beginning of 2021 as a cooperation between six academic and five industrial partners from different industry sectors. The aim of the project is to provide tools, methodologies, and experience for supporting the development and deployment of distributed systems with components that rely on deep learning. A key enabler for VEDLIoT is the ability to decompose requirements and architecture based on the needs of the 13 use cases of the project. Therefore, we propose to extend the state of the art on architecture framework by providing a mathematical model for system architectures, which is scalable and supports co-evolution of different aspects of an AI system. In this talk we present our work on Requirement Engineering in VEDLIoT and guidelines based on a mathematical formulation on how a consistent architecture framework can be built up that supports the creation and management of system architectures and requirements for distributed AI systems in VEDLIoT.

**Presenter:** *Eric Knauss*, Chalmers | University of Gothenburg, Sweden

**Title:** The Role of Requirements Engineering in the Development of Automotive Perception Systems

**Abstract:** Software that contains machine learning algorithms is an integral part of automotive perception, for example, in driving automation systems. The development of such software, specifically the training and validation of the machine learning components, require large annotated datasets. An industry of data and annotation services has emerged to serve the development of such data-intensive automotive software components. Wide-spread difficulties to specify data and annotation needs challenge collaborations between OEMs (Original Equipment Manufacturers) and their suppliers of software components, data, and annotations. In this talk, we will describe these challenges from a requirements perspective, including the need to describe driving scenarios in specific contexts, maintaining the scoping of the context descriptions, specifying requirements on data, annotations, and their quality. We describe how this relates to requirements management and traceability, as well as managing value-chains of automotive perception systems across complex customer and supplier relationships. We close this talk with an outlook on how requirements engineering can support this complex domain and which future research directions we foresee.

**Presenter:** *Alejandro Maté*, University of Alicante, Spain

**Title:** Towards a more systematic design of Artificial Intelligence

**Abstract:** Thanks to its many advantages, Artificial Intelligence (AI) has become a common

practice for companies and academia. However, AI is far from a simple process where an algorithm is trained and the project is complete. AI requires many steps such as domain understanding, data preprocessing, model selection, training and evaluation as well as deep knowledge of a wide variety of traditional and novel algorithms that behave differently. To make things worse, most AI projects rarely define clear-cut objectives, often producing failures or larger than expected costs when stakeholders notice that the model is less-than-optimal for the real-world use. To improve AI practice, we propose to exploit the power of Requirements Engineering. By clearly describing the objectives pursued together with the qualities of the desired solution such as explainability or performance with large datasets, we can discard from the initial moment inadequate AI solutions. Even more, an adequate specification of AI Requirements provides the necessary framework to evaluate the results of the project in a comprehensive way rather than just looking at model performance and questioning whether the achieved accuracy was high enough or not. In order to create the ideal AI Requirements language, we consider that we must bridge the gap between the stakeholders and the developers, being understandable enough to specify the high-level objectives while at the same time sufficiently specific to link these objectives to the underlying solution.

**Presenter:** *Kurt Schneider*, Leibniz Universität Hannover, Germany

**Title:** Requirements Engineering and Artificial Intelligence in One

**Abstract:** Many contributions either address RE for AI or AI for RE. In this talk, I will present a case in which both directions occur at the same time.

About 10 years ago, machine learning techniques were used to classify bugs, feedback, and even requirements. We identified security-related requirements using and optimizing Bayesian classifiers. This was AI4RE, since security requirements could be identified and considered by applying ML. During the last years, app reviews were classified and interpreted in order to derive change requests and requirements automatically; AI4RE again.

The reverse direction (RE4AI) provides performance requirements and constraints for the AI part of a software. For example, bounds for precision, recall, or accuracy are required. Users, stakeholders, and developers use concepts applicable to the statistical nature of AI and ML to express robustness and correctness requirements, e.g. via False Negatives and False Positives. The ultimate goal is to define and phrase requirements on AI in a way that is comprehensible to humans – and that provides guidance for AI engineers.

Explainability is a quality aspect that gains importance as software gets more complex. It has two facets: (1) In the artificial intelligence developer community, interpretability is the most relevant variant of explainability. If an ML algorithm can explain how it reached a result, ML engineers can use that information to improve the algorithms, the training set, or hyperparameters. (2) Software users and stakeholders require certain explanations to be given under specific conditions. Often, an explanation should indicate why the software behaved as it did. However, what are meaningful and realistic requirements on explanations?

In the softXplain project, we try to identify reasonable requirements for explainability – no matter whether AI is involved or not. Many people do not know – or do not even care – whether a piece of software contains AI. While using the software, they just need to know what happened, and why, if software behaves in an unexpected way. Explanations may refer to traditional software, or it may contain AI/ML components. For example, we need to de-

termine when and under which conditions (“triggers”) we want an explanation to be given. One option could be to use ML for triggering explanations as required. AI and RE merge into one.

**Presenter:** *Renata Guizzardi*, University of Twente, the Netherlands

**Title:** Dealing with Ethical Requirements for AI Systems: A Requirements Engineering Perspective

**Abstract:** Ethics is a central concern in any civilized society, and we have to take ethical decisions everyday. For example, in face of a Pandemic and the consequent shortages in hospital resources, who should be the first patients to be treated? How to guarantee the fair grading of students exams when you have a huge class of 300 students and thus require support from teaching assistants in grading? And so on. Although ethical issues are hard, taking decisions when they come up is important and thus, we shouldn't run from it. Nowadays, more and more, we are relying on systems (and especially on AI systems) to take decisions on our behalf. So a question that comes to mind is: are these systems being developed to effectively handle ethical issues? A look at the literature already suggests they aren't. Thus, our work has been focused on proposing the development of ethical systems by design, claiming this can only be done if Requirements Engineering (RE) activities are realized in such a way that ethical concerns are focused since the start, and also throughout the whole system's life cycle. Such an RE method is an open issue, and our first results show that ontologies can be supportive in this context. Using ontologies can help overcome misunderstandings and support the focus on the things (i.e., domain concepts and relations) that matter. There is still much to be learned by the RE foundations, for instance, on selecting proper stakeholders, representing and formalizing requirements in different ways etc. Moreover, the important recent works on legal and ethical frameworks (created by governmental and standard organizations) should be taken into account. And finally, promoting relevant forums like RE4AI, where people working on related topics can gather and share ideas and experiences is paramount to make consistent steps towards ethical systems' development.