

Preface: The 7th International Workshop on Requirements Engineering for Artificial Intelligence (RE4AI'26)

Angelo Susi^{2,*}, Jennifer Horkoff^{2,**} and Anna Perini^{3,*}

¹Fondazione Bruno Kessler, Trento, Italy

²Chalmers University of Technology and University of Gothenburg, Sweden

1. Preface

The RE4AI workshop held its 7th edition with REFSQ 2026 on Monday the 23rd of March in Poznań (Poland). The workshop is devoted to exploring and highlighting ongoing efforts shaping and applying RE research for systems with AI components, or AI systems. Based on the success of past year's events, we had a program of invited speakers and an interactive session shared with the NLP4AI workshop.

Our program is as follows:

Presentation Session 14:00 - 15:30

Jacek Dąbrowski

Title: Are We Advancing? Requirements Engineering for Machine Learning: From Models to Agentic Systems

Abstract: In this invited talk, I offer a reflective and critical perspective on the evolution and maturity of Requirements Engineering (RE) for Machine Learning (ML) systems. Drawing on nearly a decade of work in the area, I highlight key shifts in focus, recent advances, and persistent limitations. I use current research trends to question how effectively the community builds on prior knowledge in Requirements Engineering, and whether we are genuinely advancing the field or, in some cases, rediscovering well-known problems under new labels. To situate RE for ML within the broader RE landscape, I briefly contrast its trajectory with related directions such as NLP4RE and LLM4RE. I conclude by outlining a small number of forward-looking research challenges, with particular attention to emerging agentic technologies, and by reflecting on how RE can contribute to the responsible development of trustworthy systems with long-term societal impact.

Adam Dąbrowski

Title: The Human Factor: How Bias Shapes Requirements Engineering for AI. An industry perspective

Abstract: In the rapidly evolving field of artificial intelligence, defining requirements is often influenced by human bias. Through real-life industry examples, we will explore how biases such as Hofstadter's Law, the Curse of Knowledge, Confirmation Bias, and Anchoring Bias subtly yet profoundly impact decision-making, task prioritization, and requirement definition. The examples, drawn from actual projects and challenges faced by AI teams, will highlight the tangible consequences of bias and provide actionable insights for mitigating its effects in real-world AI development.

Quim Motger

Title: An (AI) Agent Joined Your Requirements Engineering Team: Now What?

Joint Proceedings of REFSQ-2026 Workshops, Doctoral Symposium, Posters Tools Track, and Education and Training Track. Co-located with REFSQ 2026. Poznan, Poland, March 23-26, 2026

✉ susi@fbk.eu (A. Susi); jennifer.horkoff@gu.se (J. Horkoff); perini@fbk.eu (A. Perini)

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



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Abstract: Agentic AI and generative large language models are rapidly reshaping software engineering practices. In requirements engineering, a discipline grounded in language, negotiation, interpretation, and context sensitivity, their impact is particularly profound. Requirements are textual, social, and iterative artifacts. They emerge from conversations, trade-offs, and evolving stakeholder needs. As AI systems become capable of participating in interviews, drafting specifications, analysing feedback from stakeholders, or supporting strategic analyses, the boundaries of traditional RE activities begin to blur. However, much current research focuses on adapting isolated RE tasks to LLMs as black-box components. Are we optimizing the wrong unit of analysis? Is the paradigm shifting from task automation to lifecycle redesign? What does it mean to evaluate not just outputs, but processes, collaboration dynamics, scalability, and sustainability? The talk highlights the need to move beyond single-task benchmarks toward holistic, end-to-end perspectives on the RE lifecycle, to reconsider evaluation and dataset design, and to reflect on how AI agents should be positioned within RE teams.

Başak Aydemir

Title: Keeping the Human Sharp: A New Quality for AI-Tools

Abstract: The rise of AI-powered tools in software engineering has introduced a quiet paradox: tools that make requirements activities faster may simultaneously requirements engineers less capable. As large language models take on tasks once central to RE practice we risk optimizing for output while eroding the human judgment, domain intuition, and professional autonomy that give those outputs their value. This talk argues that the field needs a new efficiency criterion for AI in RE: one that measures not only what a tool produces, but what it preserves in the human using it. We propose that human-AI collaboration in RE is not a single phenomenon but a spectrum of interaction modes from AI as scribe to AI as peer to AI as decision-maker each carrying distinct implications for skill retention, consent, and cognitive agency.

Joint discussion

Joint Interactive Session with NLP4RE 16:00 - 17:30

We deeply thank the presenters for speaking at RE4AI'26. We finally thank the REFSQ'26 workshop chairs, PC chairs and the organizing committee for their trust and support. And please stay tuned for the next editions of RE4AI.

RE4AI 2026 Organizers

Angelo Susi, Jennifer Horkoff, Anna Perini