Towards Agentic Research Problem Engineering: A Vision for Human-AI Collaboration in Software Engineering Research

Anrafel Fernandes Pereira
PUC-Rio
Rio de Janeiro, Brazil
University of Vassouras
Vassouras / RJ, Brazil
afpereira@inf.puc-rio.br

Rafael da Silva Tomaz
PUC-Rio
Rio de Janeiro, Brazil
rtomaz@inf.puc-rio.br

Paloma Guenes
PUC-Rio
Rio de Janeiro, Brazil
University of Bari
Bari, Italy
pguenes@inf.puc-rio.br

Eduardo Kinder Almentero
UFRRJ
Seropédica, Brazil
almentero@ufrrj.br

Allysson Allex Araújo
Federal University of Cariri (UFCA)
Juazeiro do Norte, Ceará, Brazil
allysson.araujo@ufca.edu.br

Marcos Kalinowski
PUC-Rio
Rio de Janeiro, Brazil
kalinowski@inf.puc-rio.br

Abstract

Artificial Intelligence (AI) is transforming Software Engineering through agentic systems capable of collaborating with humans in complex cognitive activities. However, support for research problem formulation remains limited. This Vision paper introduces *Research Problem Engineering* (RPE) as an emerging engineering discipline and advocates an *Agentic Research Problem Engineering* (ARPE) perspective, where specialized AI agents collaborate with researchers and practitioners throughout the research problem lifecycle. We present a conceptual architecture, position Lean Research Inception (LRI) and LRI Co-Assistant as initial operationalizations, and outline a research agenda for advancing human-agent collaboration in research problem formulation.

Keywords

Research Problem Engineering, Artificial Intelligence Agents, Research Problem Formulation, Agentic Software Engineering, Human-Agent Collaboration

1 Introduction

Research problem formulation is one of the most influential activities in Software Engineering (SE) research because it can shape what will be investigated and directly affects the relevance, rigor, applicability, and potential impact of the resulting knowledge [2, 13]. Previous work has emphasized that high-quality research depends on identifying meaningful, well-defined, and contextually grounded problems [13, 14, 17]. This is particularly important given the long-standing concern with the practical relevance of SE research, where weak industry engagement, unrealistic assumptions, and poorly contextualized problems contribute to the gap between academic research and industrial practice [2, 9]. Despite this, most existing approaches still focus on evaluating research outcomes rather than systematically supporting research problem formulation in the early stages of the research lifecycle [8, 12].

Recent work has started to address this limitation. Pereira *et al.* [10] identified and empirically evaluated seven attributes of practically relevant SE research problems: *practical problem, context, implications / impacts, stakeholders, evidence, objectives,* and *research questions.* Building on these findings, Lean Research Inception (LRI) [11] introduced a collaborative approach for formulating, refining, and initially assessing research problems. Together, these contributions provide initial evidence that research problem formulation can be supported through explicit processes, structured artifacts, and evaluation mechanisms rather than relying solely on researchers' individual experience.

At the same time, Artificial Intelligence (AI) is reshaping both SE and scientific research. Large Language Models (LLMs), AI agents, and agentic systems are increasingly supporting knowledge synthesis, contextual reasoning, hypothesis generation, and scientific discovery [1, 6]. Initiatives such as LLM-Assisted Empirical SE [3], Google's AI Co-Scientist [4], and Agentic SE [5] reflect a broader shift from AI as a productivity tool toward AI as a cognitive teammate in knowledge-intensive activities. This trend is also visible in LRI Co-Assistant [16], which provides an initial AI-assisted operationalization of LRI.

These developments create a promising opportunity for research problem formulation. Activities such as understanding context, identifying stakeholders, gathering evidence, assessing relevance, challenging assumptions, refining problem statements, and formulating research questions are inherently knowledge-intensive and well suited to human-agent collaboration. However, most current agentic SE efforts remain focused on solution-oriented activities such as development, testing, debugging, documentation, and maintenance [15]. Comparatively little attention has been devoted to how AI agents can support the engineering of research problems themselves.

We argue that research problem formulation should therefore be treated as an engineering activity. We introduce *Research Problem Engineering* (RPE) as an emerging discipline concerned with the systematic discovery, formulation, refinement, assessment, and evolution of research problems. In this paper, we advocate an *Agentic Research Problem Engineering* (ARPE) perspective, in which specialized AI agents collaborate with researchers and practitioners throughout the RPE lifecycle. RPE is broader than its agentic realization: human-centered methods, collaborative workshops, AI-assisted tools, and multi-agent ecosystems can all be viewed as alternative operationalizations of the same underlying discipline.

The contributions of this Vision paper are fourfold. First, we introduce *Research Problem Engineering* (RPE) as an emerging engineering discipline. Second, we propose an *Agentic Research Problem Engineering* perspective based on human-agent collaboration and specialized AI agents supporting the RPE lifecycle. Third, we position Lean Research Inception (LRI) [11] and LRI Co-Assistant [16] as initial human-centered and AI-assisted operationalizations of this vision. Finally, we outline a research agenda for advancing *Agentic Research Problem Engineering* in SE.

2 Towards Agentic Research Problem Engineering

We define *Research Problem Engineering* (RPE) as an emerging engineering discipline concerned with the systematic discovery, understanding, formulation, assessment, and continuous evolution of research problems. In this paper, we advocate an *Agentic Research Problem Engineering* (ARPE) perspective, where these activities are supported through human-agent collaboration based on an orchestrated ecosystem of specialized AI agents. In this discipline, AI agents are not intended to replace human expertise, but to augment it by acting as AI teammates throughout the research problem lifecycle. Its goal is to transform research opportunities, practical challenges, scientific gaps, and initial ideas into research problems that are well-founded, relevant, and investigable.

The rationale behind ARPE is inspired by the Software Requirements Knowledge Area of SWEBOK [7]. SWEBOK defines software requirements as expressions of the needs and constraints imposed on a software solution and recognizes Requirements Engineering (RE) as the discipline responsible for systematically eliciting, analyzing, specifying, validating, and managing those requirements throughout the software lifecycle.

We argue that research problems deserve a similar engineering perspective. Just as the quality of a software system depends on the quality of its requirements, the quality and impact of scientific research largely depend on the quality of the research problems that motivate it. Poorly formulated research problems may propagate throughout the research lifecycle, leading to inadequate research designs, weak industrial relevance, and limited scientific contributions [2, 10, 11].

However, although inspired by RE, ARPE should be understood as a different engineering object. Requirements Engineering aims to define what should be built to solve an identified problem, whereas ARPE seeks to determine what should be investigated before any solution exists. Consequently, ARPE must deal with challenges that are uncommon in RE, including establishing the novelty, practical relevance, and scientific significance of research problems, identifying diffuse stakeholders from both academia and industry, and continuously evolving research problems as new evidence emerges.

Inspired by the development and management activities defined in SWEBOK, Figure 1 presents the envisioned lifecycle of *Research Problem Engineering*, which is discussed in greater detail in Table 1 further details how this lifecycle is augmented with agent support to establish *Agentic Research Problem Engineering*. Rather than viewing problem formulation as an isolated activity, the lifecycle emphasizes a continuous engineering process in which research problems are

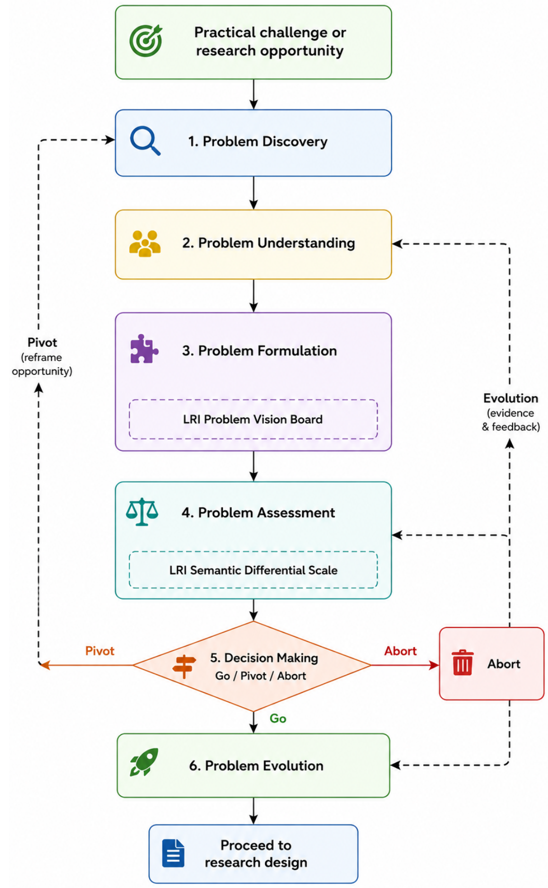

Figure 1: Lifecycle of Research Problem Engineering

discovered, understood, formulated, assessed, and evolved through iterative learning.

Unlike traditional AI assistants that primarily respond to isolated user requests, agentic AI systems are capable of autonomously planning tasks, coordinating specialized agents, maintaining contextual memory, invoking external tools, and iteratively refining intermediate artifacts [5]. These capabilities enable AI to move beyond task automation toward collaborative problem solving. In this work, we envision *Agentic Research Problem Engineering* as a discipline inherently designed for this paradigm, where an ecosystem of specialized AI agents collaborates with researchers throughout the lifecycle of ARPE while preserving human oversight and decision authority. Rather than embedding AI into an existing research method, we envision ARPE as one of the first engineering disciplines conceived from the outset for human-agent collaboration.

Our vision complements recent community-driven research agendas on Agents and SE [15]. While these agendas primarily investigate how AI agents can support SE activities such as requirements engineering, software architecture, quality assurance, and software development, we argue that AI agents can also support an even earlier engineering activity: the engineering of research problems that guide future SE research.

Table 1: Examples of specialized AI agents supporting *Agentic Research Problem Engineering*.

ARPE Activity	Specialized Agent	Primary Capability	Human-Agent Collaboration
Problem Discovery	Discovery Agent	Identify research opportunities from industrial challenges, scientific gaps, emerging technologies, and societal needs.	Suggest candidate research opportunities while researchers and practitioners decide which ones deserve further investigation.
Problem Understanding	Context Agent	Analyze domain context, constraints, assumptions, and environmental factors.	Help researchers and practitioners build a shared understanding of the problem context.
Problem Understanding	Stakeholder Agent	Identify stakeholders, their concerns, interests, and potential implications / impacts.	Support stakeholder mapping while researchers and practitioners validate and prioritize stakeholders.
Problem Understanding	Evidence Agent	Retrieve, summarize, and synthesize scientific and industrial evidence.	Collect and organize supporting evidence while researchers and practitioners judge its relevance and credibility.
Problem Formulation	Formulation Agent	Generate and refine problem statements, objectives, expected impacts, and research questions.	Propose alternative formulations while researchers and practitioners iteratively refine the final research problem.
Problem Assessment	Critic Agent	Evaluate practical relevance (clarity, value, feasibility, applicability) and expose weaknesses.	Act as a devil's advocate, challenging assumptions and stimulating critical reflection.
Decision Making	Decision Agent	Compare alternatives and summarize trade-offs.	Support Go/Pivot/Abort decisions while researchers and practitioners remain responsible for the final decision.
Problem Evolution	Evolution Agent	Continuously refine research problems as new evidence, stakeholder feedback, and contextual changes emerge.	Recommend updates throughout the research lifecycle while preserving researcher control.

2.1 Initial Operationalizations

Research Problem Engineering (RPE) has already been partially operationalized through Lean Research Inception (LRI) [11], which provides a structured, human-centered process for collaborative research problem formulation and assessment. LRI Co-Assistant [16] further extends this process by incorporating AI-assisted reasoning into these activities. Although neither approach fully realizes the vision of ARPE, both provide initial operationalizations that demonstrate the feasibility of supporting RPE through structured methods, shared artifacts, and human-AI collaboration.

2.2 ARPE Ecosystem

Recent advances in AI agents enable a shift from tool-assisted problem formulation toward *Agentic Research Problem Engineering*, in which specialized AI agents collaborate with researchers throughout the ARPE lifecycle. Rather than relying on a single conversational assistant, we envision an ecosystem of complementary agents that contribute different reasoning capabilities while keeping researchers responsible for interpretation, trade-off analysis, and final decisions. Figure 2 illustrates the preliminary conceptual architecture of ARPE. The architecture is organized into five layers: (i) a *Human Layer*, where SE researchers and practitioners collaborate and make final decisions; (ii) an *Agent Orchestration Layer*, responsible for task decomposition, agent coordination, context propagation, shared memory management, and synthesis of intermediate artifacts; (iii) a *Specialized AI Agents Layer*, composed of agents that support each activity of the ARPE lifecycle; (iv) a

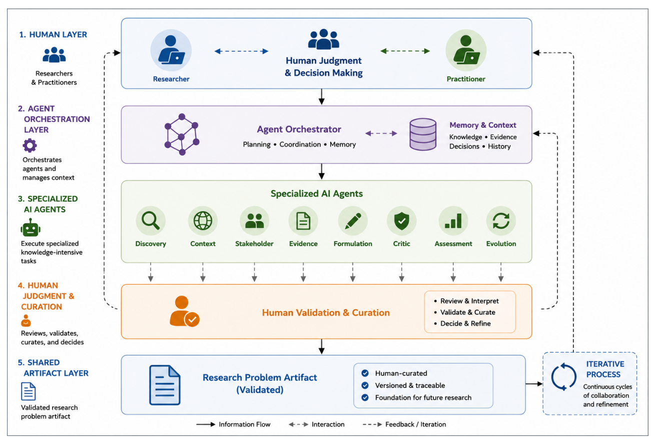

Figure 2: Conceptual Architecture of ARPE

Human Judment & Curation Layer, designed to ensure a human-in-the-loop approach by enabling researchers to review, refine, and validate AI-generated outputs at every step; and (v) a *Shared Artifact Layer*, which maintains the evolving representation of the research problem and enables collaboration among humans and AI agents.

The activities involved in ARPE are heterogeneous, requiring different reasoning capabilities such as contextual analysis, evidence synthesis, critical evaluation, stakeholder reasoning, and iterative refinement. This heterogeneity naturally motivates an ecosystem of specialized agents rather than a single general-purpose assistant.

Table 1 presents examples of specialized AI agents that can support ARPE. These agents form an ecosystem of cognitive teammates rather than isolated assistants. Their purpose is not to automate scientific reasoning, but to augment researchers' and practitioners' ability to explore alternatives, synthesize evidence, expose hidden assumptions, and support reflective decision making. Humans remain responsible for defining research goals, interpreting evidence, negotiating trade-offs, and making final decisions, while AI agents expand the reasoning space and reduce cognitive effort during early-stage research activities.

3 Concluding Remarks and Research Agenda

This Vision paper introduced *Agentic Research Problem Engineering* (ARPE), an emerging engineering discipline inherently designed for human-agent collaboration, where an orchestrated ecosystem of specialized AI agents augments researchers and practitioners throughout the lifecycle of research problem formulation, from problem discovery to continuous evolution. Inspired by Requirements Engineering, we argue that research problem formulation should itself be treated as an engineering activity and that recent advances in AI agents create an unprecedented opportunity to augment this process. Rather than replacing researchers and practitioners, specialized AI agents can act as AI Teammates that expand the exploration of alternatives, synthesize evidence, challenge assumptions, and support reflective decision making throughout the ARPE lifecycle. We further discussed how LRI [11] and LRI Co-Assistant [16] represent initial operationalizations of this vision.

Realizing this vision requires advances across multiple research directions. First, future work should investigate effective models of *human-agent collaboration*, identifying how researchers and specialized agents interact throughout each ARPE activity while preserving human agency and accountability. Second, *multi-agent orchestration* mechanisms are needed to coordinate specialized agents responsible for problem discovery, contextual understanding, evidence gathering, formulation, assessment, decision making, and continuous problem evolution. Third, agents should provide *explainable and evidence-based reasoning*, allowing researchers to understand, validate, and challenge AI-generated suggestions. Finally, the community needs methods for evaluating both the quality of formulated research problems and the effectiveness of ARPE processes, considering dimensions such as clarity, contextualization, practical relevance, stakeholder alignment, and value.

We believe that ARPE opens a promising research direction at the intersection of SE, AI for Science, and Agentic AI, and complements the research agenda [15] by extending the role of AI agents beyond software construction. While AI agents are increasingly transforming software development activities, we envision their next frontier as supporting the engineering of the research problems that shape future scientific discoveries. We hope this vision stimulates new research on AI-supported scientific reasoning and contributes to establishing ARPE as a first-class area within SE.

Artifact Availability

Initial operationalizations of this vision are publicly available. LRI Co-Assistant is available on Zenodo at https://doi.org/10.5281/zenodo.21385837, with an online demo at https://lri-tool.vercel.app using `researcher@example.com` / `researcher123`. Lean Research Inception (LRI) and its collaborative template are available on Miroverse at https://miro.com/templates/lean-research-inception-lri/.

Acknowledgements

We gratefully acknowledge the support of CAPES (Finance Code 001 and PROEX), CNPq (Grants 312275/2023-4 and 420191/2025-9), FAPERJ (Grant E-26/204.256/2024), and Instituto Kunumi for their generous support.

References

[1] Mamdouh Alenezi. 2026. Rethinking software engineering for agentic ai systems. *arXiv preprint arXiv:2604.10599* (2026).

[2] Vahid Garousi, Markus Borg, and Markku Oivo. 2020. Practical relevance of software engineering research: synthesizing the community's voice. *Empirical Software Engineering* 25 (2020), 1687–1754.

[3] Victoria Gomes, Delaney Selb, Fabio Palomba, Rodrigo Spinola, and David Lo. 2026. LLM-Assisted Empirical Software Engineering: Systematic Literature Review and Research Agenda. *arXiv preprint arXiv:2604.26192* (2026).

[4] Juraj Gottweis, Wei-Hung Weng, Alexander Daryin, Tao Tu, Anil Palepu, Petar Sirkovic, Artiom Myaskovsky, Felix Weissenberger, Keran Rong, Ryutaro Tanno, et al. 2025. Towards an AI co-scientist. *preprint arXiv:2502.18864* (2025).

[5] Ahmed E Hassan, Hao Li, Dayi Lin, Bram Adams, Tse-Hsun Chen, Yutaro Kashiwa, and Dong Qiu. 2025. Agentic software engineering: Foundational pillars and a research roadmap. *arXiv preprint arXiv:2509.06216* (2025).

[6] Junda He, Christoph Treude, and David Lo. 2025. Llm-based multi-agent systems for software engineering: Literature review, vision, and the road ahead. *ACM Transactions on Software Engineering and Methodology* 34, 5 (2025), 1–30.

[7] IEEE Computer Society. 2026. *Guide to the Software Engineering Body of Knowledge (SWEBOK Guide), Version 4.0.* IEEE Computer Society Press, Los Alamitos, CA. https://computer.org

[8] Martin Ivarsson and Tony Gorschek. 2011. A method for evaluating rigor and industrial relevance of technology evaluations. *Empirical Software Engineering* 16 (2011), 365–395.

[9] Dusica Marijan and Arnaud Gotlieb. 2020. Lessons learned on research co-creation: making industry-academia collaboration work. In *2020 46th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, 272–275.

[10] Anrafel Fernandes Pereira, Maria Teresa Baldassarre, Daniel Mendez, Jürgen Börstler, Nauman bin Ali, Rahul Mohanani, Darja Smite, Stefan Biffl, Rogardt Heldal, Davide Falessi, Daniel Graziotin, and Marcos Kalinowski. 2026. Attributes to Support the Formulation of Practically Relevant Research Problems in Software Engineering. In *3rd International Workshop on Methodological Issues with Empirical Studies in Software Engineering (WSESE@ICESE'26)*.

[11] Anrafel Fernandes Pereira, Marcos Kalinowski, Maria Teresa Baldassarre, Jürgen Börstler, Nauman bin Ali, and Daniel Mendez. 2025. Towards Lean Research Inception: Assessing Practical Relevance of Formulated Research Problems. In *Evaluation and Assessment in Software Engineering (EASE '25)*. ACM, Istanbul, Turkiye.

[12] Kai Petersen, Jürgen Börstler, Nauman Bin Ali, and Emelie Engström. 2024. Revisiting the construct and assessment of industrial relevance in software engineering research. In *Proceedings of the 1st IEEE/ACM International Workshop on Methodological Issues with Empirical Studies in Software Engineering*. 17–20.

[13] Mary Shaw. 2002. What makes good research in software engineering? *International Journal on Software Tools for Technology Transfer* 4, 1 (2002), 1–7.

[14] Klaas-Jan Stol and Brian Fitzgerald. 2018. The ABC of software engineering research. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 27, 3 (2018), 1–51.

[15] Davide Taibi, Henry Muccini, Karthik Vaidhyanathan, Marcos Kalinowski, Michele Albano, Antonio Pedro Santos Alves, Renato Cerqueira, Mateus Devino, Matteo Esposito, Rodrigo Falcão, et al. 2026. A research agenda on agents and software engineering: Outcomes from the rio A2SE seminar. *arXiv preprint arXiv:2605.11720* (2026).

[16] João Pedro Theodoro, Victoria Cruz de Figueiredo, Anrafel Fernandes Pereira, Wallace Albertini Boquimpani, Mariana Crisostomo Martins, and Marcos Kalinowski. 2026. An AI-Assisted Tool for Practically Relevant Research Problem Formulation in Software Engineering. In *40th Brazilian Symposium on Software Engineering (SBES 2026)*.

[17] Roel J Wieringa. 2014. *Design science methodology for information systems and software engineering.* Springer.