Human-Centered Adoption of LLMs in Industry and Society: Challenges, Opportunities, and Systemic Implications*

Tihana Galinac Grbac, Darko Huljenić

September 5, 2026

Abstract

Large Language Models (LLMs) are increasingly being applied across diverse industrial and societal domains. However, existing research is dominated by domain-specific case studies and technical evaluations, making it difficult to compare findings, identify recurring patterns, and develop generalizable adoption strategies. This paper proposes a human-centered framework for studying LLM adoption as a socio-technical innovation process. This framework is motivated by empirical research on software process improvement in the real industrial context, where successful innovations consistently depended on the interplay between technical solutions, human actors, organizational conditions, and implementation processes. The resulting framework is structured around nine dimensions: Problem, Solution, Human Interaction, Context, Evaluation, Challenges, Benefits, Transferability, and Process. Together, these dimensions provide a common analytical structure for investigating how LLM-based innovations create value, encounter barriers, and become integrated into real-world practice. The framework is proposed as a foundation for workshop discussions and future research contributions, with the goal of fostering cumulative knowledge and enabling cross-domain generalization of LLM adoption experiences.

Keywords: Large Language Models, Human-Centered AI, Socio-Technical Systems, AI Adoption, Responsible AI

1 Introduction

Large Language Models (LLMs) are increasingly being applied across industrial and societal domains, including software engineering, manufacturing, healthcare, education, business operations, and critical infrastructure management.

*Author manuscript version. The final published version is available at IEEE Xplore: https://ieeexplore.ieee.org/document/11592062

Their ability to generate, summarize, reason over, and interact with information has created new opportunities for cognitive automation, decision support, and human–AI collaboration [1]. Consequently, organizations are rapidly exploring how LLM-based solutions can be integrated into products, services, and operational processes [2].

Despite the growing number of applications, research on LLM adoption remains fragmented [3]. Most studies focus on technical performance, benchmark evaluation, and domain-specific use cases. While such studies provide valuable evidence regarding what LLMs can achieve, they offer limited insight into how these technologies are adopted, integrated, governed, and sustained in operational environments. Similar concerns have been raised in recent research on human-centered AI and AI engineering, which emphasizes that successful deployment depends not only on technical capability but also on organizational readiness, user trust, workflow integration, and governance mechanisms [4–6].

The challenge of translating technological capability into organizational value is not unique to LLMs. For decades, software engineering research has investigated similar issues in software process improvement (SPI). A recurring observation across SPI studies [7–15] is that technically promising improvements frequently fail to achieve their intended impact because of organizational, human, or contextual barriers. These experiences suggest that LLM based improvement should be studied not merely as a technical deployment problem, but as a broader socio-technical innovation and improvement process.

To support more systematic and comparable research, this paper proposes a human-centered analytical framework for studying LLM adoption across domains. The framework is organized around nine dimensions—Problem, Solution, Human Interaction, Context, Evaluation, Challenges, Benefits, Transferability, and Process—which emerged from recurring patterns observed in empirical studies of various improvement initiatives. The framework is presented as the conceptual foundation for the workshop "Human-Centered Adoption of LLMs in Industry and Society: Challenges, Opportunities, and Systemic Implications," organized within the Software and Systems Engineering (SSE) track of the MIPRO 2026 Convention. The workshop aims to bring together researchers and practitioners from different domains to discuss experiences with LLM adoption, identify recurring challenges and opportunities, and contribute towards a cumulative body of knowledge on human-centered and sustainable integration of LLM-based technologies. In this context, the proposed framework serves both as an analytical structure for workshop discussions and as a reporting guideline for future empirical studies.

2 Human-Centered Perspective on LLM Adoption

Human-centered adoption extends beyond the traditional evaluation of technical performance and focuses on how technological innovations interact with people, organizations, and their operating environments. From this perspective, LLMs are not viewed as isolated computational artifacts, but as components of broader socio-technical systems whose effectiveness depends on the alignment between technological capabilities, human practices, organizational processes, and contextual constraints.

This view is consistent with decades of research on software process improvement (SPI) where the success of an innovation was repeatedly shown to depend on more than its intrinsic technical merits [7–13]. Studies of software verification improvement using Six Sigma emphasized systematic problem definition, measurement, root-cause analysis, intervention design, evaluation, and process control as prerequisites for sustainable improvement [16, 17]. Research on agile transformation in globally distributed development highlighted the importance of human collaboration, training, organizational context, and change management in achieving sustainable improvement outcomes [15, 18]. Investigations of software quality assurance and defect detection effectiveness demonstrated the need for clearly defined quality goals, evaluation procedures, risk management practices, measurement systems, and lessons learned that can be transferred across projects and organizations [14, 18]. Various SPI studies have demonstrated that technically sound solutions frequently fail to achieve their intended benefits when insufficient attention is paid to stakeholder involvement, organizational readiness, communication structures, deployment strategies, measurement systems, and long-term integration into operational practice [12, 16, 17].

Many of these challenges reappear in the context of LLM adoption. Organizations introducing LLM-based systems must determine not only whether the technology can perform a task, but also whether it can be trusted, integrated into existing workflows, governed responsibly, accepted by users, and sustained over time. Consequently, successful adoption requires understanding both the technical properties of the solution and the socio-organizational conditions under which it operates. This observation suggests that LLM adoption should be approached as a socio-technical improvement and innovation process. Similar to previous SPI initiatives, researchers and practitioners must understand what problem is being addressed, what solution is proposed, how people interact with the solution, under which contextual conditions it is introduced, what outcomes are achieved, what barriers are encountered, what benefits are created, whether the experience can be transferred to other settings, and how successful adoption can be reproduced. These recurring questions motivate the analytical dimensions of the framework proposed in Section 3.

3 Unified framework for human-centered adoption of LLM - core dimensions for analysis

In this section, we propose a unified framework for reporting a human-centred adoption of LLM. The framework is developed aiming to capture the full lifecycle of LLM adoption in a particular domain, but also aiming to assess its comparability across domains and generalization potential. We identify a set of questions that are minimal to satisfy that aim:

A Problem - What is being solved?

B Solution - How is it solved?

C Human Interaction - How do humans engage?

D Context - Under what conditions?

E Evaluation - What is the effect?

F Challenges - What limits it?

G Benefits - What value is created?

H Transferability - What generalizes?

I Process - How to reproduce it?

Synthesizing recurring concerns explained in Section 2 yielded the nine dimensions outlined above. Across these studies, a recurring pattern was observed: successful improvement efforts consistently required understanding the problem being addressed, the intervention being introduced, the role of human actors, the organizational and technological context, the effects achieved, the barriers encountered, the benefits created, the potential for transfer to other settings, and the mechanisms required for sustained deployment. While originally derived from SPI initiatives, these dimensions are equally applicable to the study of LLM-based innovations, which similarly represent socio-technical interventions requiring both technical effectiveness and organizational acceptance.

We are aware that the list of questions is not unique, and there may be many other ways to systematize this knowledge. That is why we provide this list as part of preliminary communication for the workshop as a structured analytical foundation, where we want to engage the community to move beyond reporting fragmented case studies toward building cumulative knowledge. Therefore, we suggest using this list, which will help systematize knowledge in this direction.

3.1 Problem Definition

Understanding the baseline problem is fundamental. This includes the characterization of workflows, inefficiencies, and cognitive demands.

Importance: Without standardized problem definitions, it is impossible to compare results across domains or identify generalizable patterns. A structured problem typology enables abstraction and synthesis.

3.2 LLM-Based Solution Design

Defines how LLMs are integrated into systems, including interaction modes, prompting strategies, and system boundaries.

Importance: Solution design determines both performance and transferability. Identifying recurring design patterns supports the development of reusable solution archetypes.

3.3 Human-AI Interaction

The focus is on how users interact with LLM systems, including trust, interpretability, and cognitive load.

Importance: Human-AI interaction is critical for adoption success. Poor interaction design can negate technical benefits and hinder scalability.

3.4 Implementation Context

This includes organizational readiness, stakeholder roles, and deployment settings.

Importance: Contextual factors strongly influence outcomes. Capturing them systematically enables understanding of boundary conditions for successful adoption.

3.5 Evaluation Metrics

A standardized evaluation framework should include performance, human factors, risks, and economic impact.

Importance: Current evaluations are inconsistent and domain-specific. Standardization enables comparability and cumulative knowledge building.

3.6 Challenges

Challenges should be categorized into technical, organizational, ethical, and operational dimensions.

Importance: A structured taxonomy allows identification of recurring barriers and supports the development of mitigation strategies.

3.7 Benefits

Benefits should be analyzed across operational, cognitive, and strategic levels.

Importance: Multi-dimensional benefit analysis reveals the full value of LLM adoption and supports decision-making.

3.8 Transferability and Generalization

Authors must distinguish between domain-specific and transferable aspects of their solutions.

Importance: This dimension directly supports the development of generalizable frameworks and cross-domain insights.

3.9 Adoption Process Mapping

Each study should map its findings to a structured adoption lifecycle, including stages such as problem identification, design, implementation, and scaling.

Importance: This enables the construction of systematic procedures for LLM adoption and identifies critical transition points.

4 Illustrative Framework Application

To provide an initial assessment of the applicability of the proposed framework, we selected a representative sample of research studies from domains where LLM adoption is actively being investigated. The purpose of this initial evaluation is not to establish statistical validity, but to demonstrate content coverage and practical applicability. A more comprehensive evaluation involving systematic coding of a larger corpus of studies and inter-rater reliability assessment is left for future work. The selection criteria were as follows:

- The study explicitly addressed the application, adoption, evaluation, or organizational implications of LLMs,
- The study represented a distinct application domain in order to support cross-domain comparison,
- The study was published in a reputable venue and attracted attention within its research community,
- The study contained sufficient information regarding the problem, solution, and at least some discussion of practical implications, user interaction, organizational considerations, or deployment challenges.

Following these criteria, eight representative studies were selected from software engineering, healthcare, education, organizational adoption, and human-AI interaction research, listed in the first collumn of the Table 1. The sample includes empirical studies, domain-specific investigations, managerial perspectives, and multidisciplinary analyses, providing a diverse set of adoption scenarios. Each paper was independently analyzed against the nine dimensions of the proposed framework: Problem (P), Solution (S), Human Interaction (HI), Context (C), Evaluation (E), Challenges (Ch), Benefits (B), Transferability (T), and Process (Pr). The objective of the analysis was to determine (i) whether the dimensions could be identified within the selected studies and (ii) whether

Table 1: Illustrative application of the proposed framework to representative LLM adoption studies

Study	P	S	HI	C	E	Ch	B	T	Pr
Vaithilingam et al. [19] (Code Generation Usability)	H	H	H	M	H	H	H	L	L
Nguyen-Duc et al. [3] (SE Research Agenda)	H	H	M	H	M	H	H	H	M
Rico and Öberg [2] (Managerial Perspective)	H	H	H	H	M	H	H	H	M
Kung et al. [20] (Medical Education)	H	H	M	H	H	H	H	L	L
Singhal et al. [21] (Clinical Knowledge)	H	H	M	H	H	H	H	M	L
Kasneci et al. [22] (Education)	H	H	H	H	M	H	H	M	L
Dwivedi et al. [23] (Organizational Adoption)	H	H	H	H	M	H	H	H	M
Amershi et al. [4] (Human-AI Interaction)	M	H	H	H	M	H	H	H	M
Average coverage	H	H	M/H	H	M	H	H	M	L/M

P = Problem, S = Solution, HI = Human Interaction, C = Context, E = Evaluation, Ch = Challenges, B = Benefits, T = Transferability, Pr = Process.
H = High coverage, M = Medium coverage, L = Low coverage.

the framework reveals differences in reporting emphasis and potential gaps in current research practice.

The illustrative analysis demonstrates that the proposed framework can be applied consistently across diverse domains including software engineering, healthcare, education, organizational adoption, and human–AI interaction research. The selected studies generally provide strong coverage of the Problem, Solution, Challenges, and Benefits dimensions, reflecting the current emphasis on demonstrating technical capabilities and application potential. Context-related aspects are also relatively well represented, particularly in healthcare, education, and organizational studies where environmental constraints and stakeholder characteristics are frequently discussed. In contrast, the dimensions of Transferability and Process receive substantially less attention. Most studies provide limited evidence regarding how adoption experiences can be generalized across organizations, how deployment strategies should be reproduced, or how successful implementations can be sustained over time. Human Interaction is addressed unevenly, with usability and trust receiving attention in some domains while being largely absent in technically oriented studies. More precisely, studies on LLM-assisted code generation [19] primarily demonstrate technical capabilities and usability benefits. Extending such studies with explicit analysis of organizational context, deployment process, and transferability could reveal under which development practices, team structures, and governance mechanisms the observed benefits are most likely to be sustained. Similarly, healthcare studies [20, 21] frequently discuss model performance and clinical knowledge but provide limited insight into how trust, accountability, workflow integration, and regulatory constraints influence practical adoption. Incorporating the Human Interaction, Context, and Process dimensions would enable a deeper understanding of deployment barriers and success factors. Educational applications [22] often emphasize opportunities and challenges for learners and teachers. However, extending these studies with systematic analysis of Transferability and Process dimensions could reveal how educational outcomes vary across institutions, curricula, and organizational settings.

Application of the proposed framework would facilitate the identification of recurring socio-technical patterns across domains. For example, studies from software engineering, and healthcare may reveal that successful adoption depends less on marginal improvements in model accuracy and more on factors such as explainability, workflow compatibility, trust, governance mechanisms, and organizational readiness. Such findings would support the development of generalized adoption strategies and reusable implementation patterns that extend beyond individual domains. Consequently, the primary value of the framework is not only to improve reporting completeness, but also to enable cumulative knowledge building. By systematically capturing technical, human, organizational, and implementation aspects of LLM adoption, the framework provides a foundation for cross-domain synthesis, comparative analysis, and the development of evidence-based guidelines for responsible and sustainable deployment of LLM-based innovations.

5 Expected benefits of using defined framework and implications for research and practice

The proposed framework aims to transform fragmented and domain-specific investigations of LLM adoption into a structured and cumulative body of knowledge. Such structured communication would further enable systematic comparison between studies, cross-domain synthesis, and the identification of transferable patterns. Therefore, our framework supports both methodological advancement and practical guidance for the responsible integration of LLMs in industry and society. The key results of such an act are elaborated below.

5.1 Toward Generalizable Knowledge

The proposed framework enables the aggregation of heterogeneous studies into a coherent body of knowledge. By enforcing common dimensions, it becomes possible to identify patterns across domains and develop generalized models.

5.2 Advancing Methodological Rigor

A shared structure improves research quality by encouraging comprehensive and comparable analyses. It also facilitates meta-analysis and synthesis.

5.3 Supporting Industrial Adoption

For practitioners, the framework provides guidance on critical factors influencing LLM adoption, reducing uncertainty and improving decision-making.

5.4 Enabling Community Building

A common framework fosters collaboration among researchers and supports the emergence of a cohesive research community focused on LLM adoption.

6 Conclusion

We are witnessing an explosion of research in the LLM field, and LLM applications have to shift from isolated case studies to structured, human-centered, and generalizable research approaches. Current LLM studies consistently report technical solutions and performance measures, but rarely discuss transferability, implementation processes, or organizational conditions. The proposed framework exposes these gaps and provides a common structure for future comparative studies. This paper proposes a common framework to guide such efforts, emphasizing key analytical dimensions that support comparability and synthesis. Human-centered LLM adoption is not an entirely new problem. It is the latest manifestation of a broader class of organizational improvement and sociotechnical transformation initiatives that have been studied for decades. The proposed framework, built on 9 analytical dimensions, is therefore grounded both in established theory and in empirical evidence accumulated through industrial improvement programs.

By aligning research efforts around shared structures and objectives, the community can move toward the development of standardized procedures for LLM adoption, ultimately enhancing both scientific understanding and societal impact.

Acknowledgment

The authors would like to thank the research community contributing to the emerging field of LLM adoption and human-centered AI systems. This work is funded by European Union NextGenerationEU under the Juraj Dobrila University of Pula institutional research project number IIP_UNIPU_010162.

References

[1] X. Houet al., "Large Language Models for Software Engineering: A Systematic Literature Review", *ACM Trans. Softw. Eng. Methodol*, vol. 33, no. 8, Article 220 (December 2024), 79 pages, 2024.

[2] S. Rico and L.-M. Öberg, "Challenges and Opportunities for Generative AI in Software Engineering: A Managerial View," in *Companion Proceedings of the 33rd ACM International Conference on the Foundations of Software Engineering (FSE Companion '25)*, ACM, New York, NY, USA, pp. 1338–1344, 2025.

[3] A. Nguyen-Duc et al., "Generative Artificial Intelligence for Software Engineering—A Research Agenda,"*Software: Practice and Experience*, vol. 55, no. 11, pp. 1806–1843, 2025, doi: 10.1002/spe.70005.

[4] S. Amershi et al., "Guidelines for Human-AI Interaction," in *Proc. 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*, New York, NY, USA: ACM, 2019, Paper 3, pp. 1–13, doi: 10.1145/3290605.3300233.

[5] B. Shneiderman, "Human-Centered Artificial Intelligence: Reliable, Safe and Trustworthy," *International Journal of Human–Computer Interaction*, vol. 36, no. 6, pp. 495–504, 2020.

[6] National Institute of Standards and Technology (NIST), "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100–1, Gaithersburg, MD, USA: U.S. Department of Commerce, 2023. Available: https://doi.org/10.6028/NIST.AI.100-1

[7] W. S. Humphrey, *Managing the Software Process*, Addison-Wesley, Reading, MA, USA, 1989.

[8] V. R. Basili and G. Caldiera, "Improve Software Quality by Reusing Knowledge and Experience," *Sloan Management Review*, vol. 37, no. 1, pp. 55–64, 1995.

[9] K. El Emam, B. Smith, and P. Fusaro, "Success Factors and Barriers in Software Process Improvement," in R. Messnarz and C. Tully (Eds.), *Better Software Practice for Business Benefit: Principles and Experience*, IEEE Computer Society Press, 1999, pp. 355–371.

[10] T. Dybå, "An Empirical Investigation of the Key Factors for Success in Software Process Improvement," *IEEE Transactions on Software Engineering*, vol. 31, no. 5, pp. 410–424, 2005.

[11] M. Niazi, D. Wilson, and D. Zowghi, "A Framework for Assisting the Design of Effective Software Process Improvement Implementation Strategies," *Journal of Systems and Software*, vol. 78, no. 2, pp. 204–222, 2005.

[12] T. Galinac Grbac, "Empirical Evaluation of Selected Best Practices in Implementation of Software Process Improvement," *Information and Software Technology*, vol. 51, no. 9, pp. 1351–1364, 2009.

[13] J. D. Herbsleb, "Global Software Engineering: The Future of Socio-Technical Coordination," in *Future of Software Engineering (FOSE 2007)*, IEEE Computer Society, pp. 188–198, 2007.

[14] T. Galinac Grbac and D. Huljenic, "Defect Detection Effectiveness and Product Quality in Global Software Development," in *Lecture Notes in Computer Science*, vol. 6759, pp. 113–127, Springer, 2011.

[15] M. Ivcek and T. Galinac, "Adapting Agile Practices in Globally Distributed Large Scale Software Development," in *Proceedings of the International Conference on Telecommunications*, 2009.

[16] T. Galinac and Z. Car, "Software Verification Process Improvement Proposal Using Six Sigma," in *Lecture Notes in Computer Science*, vol. 4589, pp. 51–64, Springer, 2007.

[17] T. Galinac, "Six Sigma Methodology for Software Development Process Improvement," in *Proceedings of the IEEE MELECON*, pp. 773–776, 2006.

[18] M. Ivcek and T. Galinac, "Aspects of Quality Assurance in Global Software Development Organization," in *Proceedings of the International Conference on Telecommunications*, 2009.

[19] P. Vaithilingam, T. Zhang, and E. Glassman, "Expectation vs. Experience: Evaluating the Usability of Code Generation Tools Powered by Large Language Models," in *Proc. CHI Conference on Human Factors in Computing Systems*, 2022, pp. 1–20.

[20] T. H. Kung et al., "Performance of ChatGPT on USMLE: Potential for AI-Assisted Medical Education Using Large Language Models," *PLOS Digital Health*, vol. 2, no. 2, 2023.

[21] K. Singhal et al., "Large Language Models Encode Clinical Knowledge," *Nature*, vol. 620, pp. 172–180, 2023.

[22] E. Kasneci et al., "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education," *Learning and Individual Differences*, vol. 103, 2023, 102274, ISSN 1041-6080, https://doi.org/10.1016/j.lindif.2023.102274.

[23] Y. K. Dwivedi et al., "So What if ChatGPT Wrote It? Multidisciplinary Perspectives on Opportunities, Challenges and Implications of Generative Conversational AI for Research, Practice and Policy," *International Journal of Information Management*, vol. 71, 2023, 102642, ISSN 0268-4012, https://doi.org/10.1016/j.ijinfomgt.2023.102642.