

AI applications in requirements engineering: a systematic mapping study

Melih Sah Dogru^{1,✉}, Jan Penkert², Felix Mahr² and Tobias Düser¹

¹ IPEK - Institute of Product Engineering, Karlsruhe Institute of Technology, Germany,

² FAPS - Institute for Factory Automation and Production Systems, Friedrich-Alexander Universität Erlangen-Nürnberg, Germany

✉ melih.dogru@partner.kit.edu

ABSTRACT: Artificial intelligence influences requirements engineering, but it remains unclear which activities benefit and how. This paper reviews 15 studies from the last five years, classifying AI approaches with an established RE framework. Current work focuses on operational tasks: requirements determination, analysis, consolidation, and traceability. About two thirds address single activities rather than integrated solutions. Early-phase tasks like knowledge elicitation receive little support despite being central to practice. The mapping clarifies existing AI support and gaps for future work.

KEYWORDS: artificial intelligence (AI), requirements engineering (RE), systematic literature review, systems engineering (SE)

1. Introduction

Modern product development has become increasingly complex due to emerging technologies, increasing software content, and interconnected product ecosystems, compounded by intensified global competition, especially in the automotive industry (Hertzke et al., 2025; IHK Ulm, 2025; Schmidt, 2025). In the early stages of product development, the influence on the final product is high, as decisions made at this stage affect all subsequent processes (Cooper & Kleinschmidt, 1993). Requirements engineering plays a central role in this phase, as requirements form the basis of product development and extend throughout the entire development process (Siemens Digital Industries Software, 2020). To address this complexity, systems engineering (SE) provides a structured methodology that makes systems manageable through systematic decomposition and consideration of requirements and interactions (Bender et al., 2021). However, managing thousands of interdependent requirements, in complex domains like automotive development, remains resource-intensive, requiring substantial effort for classification, consistency checking, and traceability management. Artificial intelligence (AI) offers promising potential to support requirements engineers in these tasks (Siemens Digital Industries Software, 2020). For instance, a mid-range car already comprises around 50,000 mechanical requirements and legal regulations, while in the E/E and software domain up to 450,000 requirements may be involved. Managing and tracking these requirements is often a manual process that is both time-consuming and highly prone to errors (Siemens Digital Industries Software, 2020). Recent advances in AI, driven by increasing computational power and data availability (Fraunhofer-Institut für Kognitive Systeme IKS, 2025), have enabled sophisticated natural language processing capabilities. Large language models (LLMs) are relevant in this context, as they can understand semantic relationships and support core requirements engineering (RE) activities such as requirements elicitation, classification, and traceability analysis (IBM, 2025). While AI applications in requirements engineering have gained increasing attention, current research remains fragmented across different RE tasks and domains. Although individual studies demonstrate AI's potential for specific activities, such as requirements classification or similarity detection, no systematic overview exists that maps AI applications in

requirements engineering to established RE activity frameworks. This lack of systematic categorization makes it difficult to identify which RE activities are well-supported by current AI approaches and which remain unaddressed, hindering strategic development of AI support for requirements engineering. This study addresses this gap through a systematic analysis of current AI applications in requirements engineering, mapping them to a established RE framework. The findings show where AI support exists and where it is needed, helping direct future development of AI tools in requirements engineering.

2. Theoretical background

This section establishes the foundational concepts necessary for understanding and analyzing AI applications in requirements engineering. [Section 2.1](#) introduces the Systems Engineering context, which defines the application domain for requirements engineering practices. [Section 2.2](#) presents the Requirements Engineering activities framework by [Rupp \(2021\)](#), which provides a comprehensive taxonomy of RE tasks and serves as the basis for systematically categorizing AI applications in the subsequent analysis.

2.1. Systems engineering

Systems Engineering (SE) is an interdisciplinary methodology for developing complex systems through systematic management of requirements, interactions, and trade-offs ([International Council on Systems Engineering \(INCOSE\), 2015](#)). SE emphasizes defining customer needs and required functionality early, then documenting requirements, synthesizing designs, and validating the system ([International Organization for Standardization \(ISO\), 2015](#)). By systematically decomposing systems into manageable domains ([Scheithauer, 2013](#)) while maintaining holistic consideration of stakeholder requirements, SE provides the structured framework in which requirements engineering operates, particularly in complex domains such as automotive and aerospace development.

2.2. Requirements engineering activities

Requirements Engineering involves systematic activities for identifying, specifying, and managing requirements throughout the development process. [Rupp and Sophisten \(2021\)](#) offers a comprehensive framework that organizes these RE activities into four main categories. The first category, Knowledge Elicitation, deals with gathering the necessary information, determining what requirements are needed, defining project visions and goals, understanding relevant business processes, clarifying the system context, and identifying where requirements come from. The second category, Deriving Good Requirements, is about ensuring quality by verifying requirements, bringing related requirements together, and continuously analyzing and refining them. The third category, Communicating Requirements, addresses how requirements are shared with different stakeholders, whether through alternative presentation formats, visual models, or natural language documentation. The fourth category, Managing Requirements, covers the ongoing work of editing requirements as they evolve, maintaining traceability throughout the project, and handling the various administrative tasks that arise during development. In this study, we use this framework as the analytical lens to map each identified AI application to specific RE activities (see [Table 3](#)). This reveals which parts of the RE process already receive AI support and which areas remain largely unexplored.

3. Research questions and methodology

3.1. Research questions

Based on the identified gap in systematic understanding of AI applications in requirements engineering, this study poses three research questions:

RQ1: What AI applications for requirements engineering are reported in the systems engineering literature, and which RE activities do they address?

RQ2: What patterns emerge regarding industry domains, implementation status, and AI methods employed in current applications?

RQ3: What implications arise for future development of AI support in requirements engineering?

3.2. Research methodology

This study follows the Design Research Methodology (DRM) proposed by [Blessing and Chakrabarti \(2009\)](#), which structures design research into four phases: Research Clarification, Descriptive Study I, Prescriptive Study, and Descriptive Study II. This paper specifically addresses the Descriptive Study I phase, which aims to build comprehensive understanding of the existing situation through systematic analysis as foundation for subsequent development and validation phases. To answer the research questions, a systematic literature review (SLR) was conducted following the process described by [Xiao and Watson \(2019\)](#). The review proceeded through seven steps: (1) formulating research questions, (2) defining search strategy, (3) executing search, (4) applying selection criteria, (5) analyzing selected publications, (6) synthesizing findings, and (7) presenting results. [Figure 1](#) visualizes the complete SLR process. To answer RQ1, the study first identifies relevant AI applications through this systematic literature review. Subsequently, these applications are mapped to [Rupp's \(2021\)](#) Requirements Engineering framework, which organizes RE activities into four main groups comprising specific activities. This mapping reveals which RE activities currently receive AI support and which remain largely unaddressed. Additional analyses of industry domains, implementation status, and AI methods address RQ2, while RQ3 is explored through synthesis of these findings in the discussion. Subsequent DRM phases (Prescriptive Study and Descriptive Study II) are beyond the scope of this paper.




Figure 1. Process of SLR showing the seven steps according to [Xiao and Watson \(2019\)](#)

3.3. Search strategy

To identify relevant publications, three scientific databases were selected: Scopus, Web of Science, and Google Scholar. These databases were chosen to capture both established academic publications and emerging research in the field. The search string was developed based on the main research areas: requirements engineering, artificial intelligence, and systems engineering. Since requirements engineering is sometimes referred to as requirements management in practice, both terms were included. Similarly, various abbreviations and full forms of artificial intelligence terms were incorporated to ensure comprehensive coverage:

("Requirements Engineering" OR "Requirements Management") AND ("AI" OR "Artificial Intelligence" OR "LLM" OR "Large Language Model") AND ("Systems Engineering")

After initial testing, the search string was also translated into German to capture relevant German-language publications, particularly from German-speaking academic and industrial contexts. Given the rapid development of AI technologies, particularly large language models, in recent years, the search was restricted to publications from 2020 to 2025 to focus on current approaches and technologies. Only peer-reviewed journal articles and conference proceedings in German or English were considered to ensure academic quality.

3.4. Selection process and criteria

The English search returned results across all databases: Scopus (331), Web of Science (42), and Google Scholar (around 5,800). The German search yielded 17 additional results, all from Google Scholar. For Scopus and Web of Science, all results were included in the screening process, while for Google Scholar, the first 100 results from both English and German searches were considered to focus on the most relevant documents. After removing duplicates, 490 publications remained for initial screening.

The screening followed two stages. First, all 490 titles were reviewed for alignment with these criteria, which resulted in 149 publications warranting filtering abstracts. After conducting the abstract analysis, 45 documents remained for full-text assessment. In the second stage, these 45 documents were read in full to verify their relevance. During this process, six publications could not be accessed despite attempts through institutional access and interlibrary services. After thorough assessment, 24 publications were

excluded as they either lacked sufficient detail on AI implementation, focused primarily on AI engineering aspects rather than RE support, or did not provide transferable insights for systems engineering contexts. Forward and backward citation searches were conducted on the remaining 15 publications. While several additional documents were identified through these searches, none met all inclusion criteria when assessed, particularly regarding the timeframe (2020-2025) and the focus on practical AI application in requirements engineering. The final dataset thus consisted of 15 publications. Figure 2 illustrates this screening process, showing the systematic reduction from 490 initial search results through various filtering stages to the final 15 publications.

Table 1. Overview of inclusion and exclusion criteria

Publications were included if they:	Publications were excluded if they:
-Addressed concrete processes, methods, or tools in requirements management -Demonstrated the application of AI techniques with a focus on natural language processing or semantic analysis -Provided transferable insights for technical system development	-Focused solely on theoretical AI development without practical RE application -Came from domains with no transferability to technical system development -Lacked sufficient detail on AI implementation




Figure 2. Flow diagram showing literature screening process

4. Results: AI applications in requirements engineering

4.1. Overview of identified publications

Table 2 provides an overview of the identified publications, showing their industry domain, implementation status, specific RE tasks addressed, and AI methods employed. The column ‘State’ indicates whether the referenced work describes a framework (F) or applies an implementation (I) of an AI application. No distinction is made regarding whether the reference uses test data or real field data from industry.

Table 2. Overview of identified publications on AI in RE

References	Industry	State	RE & SE Context	AI-Method
(Timperley et al., 2025)	Aerospace	I	- Creation of a list of functions based on functional requirement - Architecture development according to MBSE from a list of functions	LLM (Chat-GPT-4)
(Zhang et al., 2024)	Automotive	I	- Identification of semantic similarities between scenarios and product features	BERT
(Jiang et al., 2025)	Manufacturing	I	- Identify requirement from error characteristics and customer requirement to derive relevant requirement	BERT
(Bashir, 2024)	Railway	I	- Identification of requirement from documents - Allocation of requirements to a system	BERT

(Continued)

Table 2. Continued

References	Industry	State	RE & SE Context	AI-Method
(Uygun & Momodu, 2024)	Automotive	I	- Information gathering through analysis of requirement documents using a chatbot	LLM (Nous-Hermes-13B-GPTQ)
(Riesener et al., 2024)	Automotive	F	- Structuring and deriving use cases from requirements	BERT
(Schleifer et al., 2024)	Automotive	I	- Analysis and classification of requirements	BERT
(Gräßler, Oleff, Hiebb, et al., 2022)	Automotive	I	- Dependency analysis of requirements	BERT
(Gräßler, Oleff, & Preuß, 2022)	Robotic	I	- Identification of requirement dependencies - Classification of dependency types	BERT
(Gräßler et al., 2023)	Automotive	I	- Requirements revised from the original set	LLM (GPT-J)
(Dehn et al., 2023)	Automotive	F	- No RE activity	/
(Van Remmen et al., 2023)	Overarching	F	- No RE activity	/
(Bourdon et al., 2025)	Nuclear	I	- INCOSE-compliant extraction, classification, and reformulation of requirements - Functional architecture creation and system allocation - Regulatory validation (ISO 15288) - Report generation and change impact analysis	LLM (Mistral-7B-Instructv0.2)
(Wang et al., 2025)	Overarching	I	- Extraction of requirements - Tracking and management of requirements - Structuring of requirements	LLM (Chat-GPT-4)
(Hovemann et al., 2025)	Overarching	I	- System context analysis - Development of alternative architectures - Derivation of test cases	LLM

4.2. Description of identified publications

The identified publications are summarized below in order of appearance in Table 2.

Timperley et al. (2025) used the LLM ChatGPT-4 in the aerospace sector to first create a list of functions from functional requirements. In a further step, an architecture was generated according to MBSE. For this purpose, different prompts were transmitted to the language model via a Python script, with the most detailed prompt with the most accurate task description leading to the highest traceability and the best result. Zhang et al. (2024) used various Bidirectional Encoder Representations for Transformers (BERT) models to match usage scenarios for Volkswagen vehicles with product features. They succeeded in identifying relevant requirements and uncovering undiscovered customer needs. The comparison was made by determining the semantic similarity between the textual content, which was calculated using cosine similarity. Jiang et al. (2025) examined customer requirements and fault characteristics of used products in the context of remanufacturing and translated them into design requirements. To identify the requirements, a similarity calculation based on cosine similarity was performed using a BERT model. This method was used to identify and merge requirements that were formulated differently but were semantically similar. In the railway industry, Bashir (2024) uses BERT to identify requirements from extensive tender documents and allocate them to subsystems. The process model is domain-specific due to the large-scale nature of the industry. For requirement allocation, 15 correctly assigned requirements served as examples for BERT. Tests on real industry data showed significant reductions in processing time, fewer assignment errors, and improved traceability. To retrieve information from a large document repository within requirements management more quickly and efficiently, Uygun and Momodu (2024) are developing a Q&A chatbot. For this purpose, a RAG is used, which enriches the LLM Nous-Hermes-

13B-GPTQ with domain-specific information from a knowledge database, thereby improving requirements analysis. The chatbot can answer questions about documents from the automotive research sector. Due to the authors' data protection concerns, the LLM is operated locally. [Schleifer et al. \(2024\)](#) developed a framework for structuring requirements and subsequently deriving use case diagrams from system requirements. Due to the lack of homogeneity of requirement sets in industry, rule-based approaches are only of limited suitability for the task. A potential analysis shows that NLP-supported assistance systems can achieve efficiency gains in requirements analysis and mapping. The usability of the developed framework has been verified with industry experts from the automotive industry, but no evaluation with real data has been carried out. [Riesener et al. \(2024\)](#) require a labelled data set for fine-tuning a named entity recognition (NER) model in the field of requirements management. To reduce the high manual effort involved in creation by requirements engineers, BERT is used for the automated analysis and classification of requirements. The requirements are processed sentence by sentence and the categories entity, action, attributes, relative operator and quantity are used for classification, as this structure reflects the typical structure of technical requirements. Using the automatic approach, the authors have succeeded in reducing the manual effort by more than 70%. [Gräßler et al. \(2022\)](#) propose an AI-supported method for proactive management of requirement changes in complex technical systems. It combines NLP with graph-based approaches, using BERT to uncover semantic similarities between requirements and analyze dependencies. Based on these dependencies, a PageRank algorithm evaluates change impact and calculates change probability. The probability of a change is visually represented by node coloring in the graph. The authors also developed a prototype and tested it with real requirements from the automotive industry. Since the PageRank algorithm is not an AI but a graph-based algorithm, the task of change impact analysis and change probability assessment is not listed in [Table 3](#) among the tasks of AI in the SE context. In order to identify requirement dependencies and more accurately recognize the classification of dependency types such as Requires, Refines, Refined by and Required by, [Gräßler et al. \(2022\)](#) use a BERT model with active learning, known as active-BERT. The authors describe that both the quality and quantity of data are crucial success factors for establishing NLP solutions in requirements management. When merging requirements from different areas, there is a particular risk that semantically similar requirements will be formulated in different ways, making classification more difficult. [Gräßler et al. \(2023\)](#) develop a method for automatic extraction and classification of requirements through data augmentation. Requirements from specification documents across three domains are manually classified into functional, non-functional, and unclassifiable. The LLM GPT-J then generates artificial requirements in these categories, which are manually checked and reclassified. However, combining real and augmented data negatively affected classification performance due to domain discrepancies. The authors conclude that artificially generated data is particularly helpful when domain-specific data is insufficient. Using a framework consisting of 18 standardised steps, [Dehn et al. \(2023\)](#) systematically describe the possible use of artificial intelligence in the requirements engineering process. They define RE flow types, which function as data objects, and RE operations, which are applied to the data objects. The combination results in a total of 18 Elementary Process Steps (EPS), which can be represented as an input-output model. In addition, the authors provide a list of AI technologies for each process step. The list of recommended AI algorithms is not included in [Dehn et al. \(2023\)](#) and no real-world application of the framework is described. [Van Remmen et al. \(2023\)](#) developed a framework to guide AI system design in requirements management. The framework identifies potential NLP applications by deriving NLP building blocks from literature and mapping them to [Rupp and Sophisten's \(2021\)](#) RE activities. It also integrates company-specific factors such as product typology and reference systems to provide implementation recommendations. The framework has been demonstrated only through an academic example, not yet validated in practice. [Van Remmen et al. \(2023\)](#) note that NLP potential increases with the number and complexity of requirements, thus growing as development progresses and requirements accumulate. [Bourdon et al. \(2025\)](#) fine-tuned CurieLM, based on Mistral-7B-Instruct-v0.2, for nuclear engineering to automate requirements management tasks. Safety standards were used to generate specification reports and extract, classify, and structure requirements. Via a human-machine interface, engineers could apply INCOSE rules for reformulation, which the LLM executed and proposed for manual validation. While CurieLM occasionally misinterpreted or overapplied rules, human oversight ensured compliance. Requirement allocation followed a top-down subsystem approach, achieving the best results with low model temperature. In a case study, CurieLM reduced processing time by 88% for extraction, 87% for reformulation, and 66% for the Product Breakdown Structure, confirming the value of human-in-the-loop collaboration ([Bourdon et al., 2025](#)). [Wang et al. \(2025\)](#) introduced an information extraction framework

based on the five SysML requirement types: functional, interface, performance, physical, and design constraints. Unstructured requirement texts are pre-processed and analysed with ChatGPT-4, using few-shot learning to enhance accuracy. The extracted data form a knowledge graph revealing semantic relations between requirements. In a use case comparison, GPT-4 achieved 25% higher accuracy than BERT and 5% higher than GPT-3.5-turbo, improving overall traceability and requirement management. [Hovemann et al. \(2025\)](#) used the V-model to recommend advanced prompt engineering techniques for systems engineering activities. Four exemplary tasks were defined for each phase, and three use cases were implemented: AI-assisted correction of context diagrams, development of alternative system architectures, and derivation of test cases from requirements. The results, evaluated by experienced engineers, showed higher-than-expected quality. The study demonstrates that prompt engineering can add significant value to RE tasks with minimal effort, while further research is needed on task-specific adaptations.

4.3. Mapping to requirements engineering activities

To systematically analyze which RE activities receive AI support, the identified applications were mapped to [Rupp's \(2021\)](#) Requirements Engineering framework, which organizes RE activities into four main categories comprising specific sub-activities. [Table 3](#) illustrates this mapping.

Table 3. Mapping of identified publications to RE activities

Reference	Identify knowledge				Derive good requirements			Communicate requirements				Manage requirements			
	Determine requirements	Define visions and goals	Determine business processes	Define context	Identify sources of requirements	Review requirements	Consolidate requirements	Analyze and improve requirements	Communicate reqs. alternatively	Communicate reqs. in natural language	Model requirements	Edit requirements	Establish traceability	Other administrative tasks	Supplement administrative information
(Timperley et al., 2025)	○	○	○	○	○	○	●	○	○	○	○	○	○	○	○
(Zhang et al., 2024)	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○
(Jiang et al., 2025)	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○
(Bashir, 2024)	○	○	○	○	●	○	○	○	○	○	○	○	○	○	○
(Uygun & Momodu, 2024)	●	○	○	○	○	○	○	○	○	○	○	○	○	○	○
(Riesener et al., 2024)	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○
(Schleifer et al., 2024)	○	○	○	○	○	○	●	○	○	○	●	○	○	○	○
(Gräßler, Oleff, Hieb, et al., 2022)	○	○	○	○	○	○	○	○	○	○	○	○	●	○	○
(Gräßler, Oleff, & Preuß, 2022)	○	○	○	○	○	○	○	○	○	○	○	○	●	●	○
(Gräßler et al., 2023)	○	○	○	○	○	○	○	●	○	○	○	○	○	○	○
(Dehn et al., 2023)	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
(Van Remmen et al., 2023)	○	○	○	○	○	○	○	○	○	○	○	○	○	○	○
(Bourdon et al., 2025)	●	○	○	○	○	●	○	●	○	●	○	○	○	●	○
(Wang et al., 2025)	●	○	○	○	○	○	●	○	○	○	○	○	●	○	○
(Hovemann et al., 2025)	○	○	○	●	○	○	○	○	○	○	○	○	○	○	○

The analysis reveals a clear concentration on specific RE activities, with four activities appearing in at least three references: determining requirements (4 references), analyzing and improving requirements (4 references), consolidating requirements (3 references), and establishing traceability (3 references). While [Bourdon et al. \(2025\)](#) address five RE activities and [Wang et al. \(2025\)](#) cover three activities, two-thirds of the references focus on only a single RE activity. This pattern indicates that current AI solutions primarily target isolated RE tasks rather than providing comprehensive support across the RE process.

5. Discussion

The classification of the identified publications according to [Rupp's \(2021\)](#) framework shows a clear focus on operational RE activities. Four activities appear in at least three publications: determining requirements (4x), analyzing and improving (4x), consolidating (3x), and establishing traceability (3x). These belong to the categories “Deriving good requirements” and “Managing requirements”, areas with repetitive tasks and measurable results. The categories “Determining knowledge” and “Communicating requirements,” on the other hand, are hardly addressed. Two-thirds of the papers focus on individual activities, with only [Bourdon et al. \(2025\)](#) covering five and [Wang et al. \(2025\)](#) covering three activities. Wang et al. combine classification, extraction, and semantic representation because only this integrated approach can effectively resolve ambiguity, lack of structure, and missing traceability in text-based requirements. Classification provides the typological basis for precise extraction, while the knowledge graph links extracted elements into a coherent semantic model. This yields a consistent, machine-interpretable requirements representation suitable for complex technical systems. [Bourdon et al. \(2025\)](#) examine numerous areas of requirements management and systems engineering because they aim for a holistic transformation of the entire engineering process in nuclear plant development. The GenSE project is intended not only to automate individual tasks but to improve the entire system lifecycle, including requirements elicitation, quality assurance, architecture modeling, interface management, compliance, and change impact analysis. This broad perspective is necessary because nuclear megaprojects are highly complex, heavily regulated, and safety-critical; as a result, the individual disciplines influence each other and cannot be optimized in isolation. One important finding concerns the identification of requirements: All of the approaches examined work with existing requirements, they extract them from existing documents, reformulate them in accordance with standards, or assign them to categories. The actual creation of new requirements based on stakeholder discussions or workshop results is not addressed. However, it is precisely this step, from often unstructured stakeholder statements to precise technical requirements, that represents a core challenge in RE. Even LLMs, which demonstrate creative text generation in other areas, are not used for this task. This may be due to the required combination of different skills: subject-specific knowledge, knowledge of technical conditions, and the balancing of different stakeholder perspectives must come together here. The AI methods used show a balanced distribution: half of the implementations are based on BERT, the other half on large language models. However, this even distribution reflects different degrees of maturity. BERT-based approaches, which have been in use since 2018, achieve reliable results for defined tasks, for example, over 90% accuracy in requirement classification or effective duplicate detection. LLM-based work appears to be concentrated from 2022 onwards, correlating with the widespread availability of ChatGPT at the end of 2022. Easy access to these models, without needing expensive computing infrastructure, allowed more researchers to test them. However, the publications analyzed show that LLMs are primarily used for extraction, reformulation, and analysis in the studies examined. Generative applications, such as the creation of user stories, test cases, or complete sets of requirements from keywords, are not found in the publications. The automotive industry leads with 40% of publications, likely due to the high complexity in this sector. Other industries appear only sporadically: the nuclear industry in the management of regulatory requirements ([Bourdon et al., 2025](#)), the rail sector in the processing of extensive tender documents ([Bashir, 2024](#)), and aviation in MBSE integration ([Timperley et al., 2025](#)). Industries such as medical technology and finance, which also work with complex requirements and regulations, are not represented in the analyzed literature. The reasons for this uneven distribution cannot be deduced from the available data. The study is subject to methodological limitations that are typical for systematic literature reviews. The search terminology selected focused on established terms, which means that works using alternative terminology may not have been included. The focus on current developments (2020-2025) enables the analysis of state-of-the-art AI approaches, but excludes earlier pioneering work. The concentration on English and German publications reflects the main publication languages in the RE,

but may leave out contributions from other language areas. Beyond technical aspects, the adoption of AI in requirements engineering faces practical and ethical barriers. Sensitive stakeholder information such as the origin, intention, and context of statements often contains confidential business data, raising concerns about data protection and intellectual property. In safety-critical domains, it must be ensured that product safety is not compromised by incorrect or misleading AI outputs. A human-in-the-loop approach, in which domain experts validate and approve AI-generated suggestions before baselining, remains essential to ensure accountability, traceability, and reliability.

6. Summary and outlook

This systematic literature review examined 15 publications on AI in requirements engineering, mapping them to Rupp's RE framework. Current implementations concentrate on four operational activities: determining, analyzing, consolidating requirements, and establishing traceability. Early-phase work and generating requirements from stakeholder input remain unaddressed. Two-thirds of approaches target single activities rather than integrated solutions, focusing on structured tasks like classification and traceability. These gaps, particularly in creative and early-phase RE processes, represent opportunities for future work. This mapping reveals where AI support exists and where it is needed most, helping direct research toward integrated solutions and early-phase activities, especially exploring how Large Language Models can support stakeholder-driven requirements generation. By reducing the time required for requirements creation and enabling stakeholders to collaboratively generate and refine requirements directly with the assistance of LLMs, such approaches open the potential for more immediate feedback, accelerated iteration, and significantly improved alignment between stakeholder intent and documented requirements.

References

- Bashir, S. (2024). Towards AI-centric Requirements Engineering for Industrial Systems. *Proceedings of the 2024 IEEE/ACM 46th International Conference on Software Engineering: Companion Proceedings*, 242–246. <https://doi.org/10.1145/3639478.3639811>
- Bender, B., Gericke, K., Pahl, G., & Beitz, W. (Hrsg.). (2021). *Pahl/Beitz Konstruktionslehre: Methoden und Anwendung erfolgreicher Produktentwicklung* (9. Auflage). Springer Vieweg. <https://doi.org/10.1007/978-3-662-57303-7>
- Blessing, L. T. M., & Chakrabarti, A. (2009). *DRM, a Design Research Methodology*. Springer London. <https://doi.org/10.1007/978-1-84882-587-1>
- Bourdon, J., Rodriguez, J., Lesigne, Q., Suchet, P., Fister, B., Montagne, L., Malhomme, O., Benmiloud-Bechet, L., & Plana, R. (2025). Transforming Systems Engineering in Nuclear Projects with Generative AI: A Path to Efficiency and Compliance: Proceedings of the 13th International Conference on Model-Based Software and Systems Engineering, 395–406. <https://doi.org/10.5220/0013443400003896>
- Cooper, R. G., & Kleinschmidt, E. J. (1993). Screening new products for potential winners. *Long Range Planning*, 26(6), 74–81. [https://doi.org/10.1016/0024-6301\(93\)90208-W](https://doi.org/10.1016/0024-6301(93)90208-W)
- Dehn, S., Jacobs, G., Zerwas, T., Berroth, J., Hötter, M., Korten, M., Müller, M., Gossen, N., Striegel, S., & Fleischer, D. (2023). On identifying possible artificial intelligence applications in requirements engineering processes. *Forschung Im Ingenieurwesen*, 87(1), 497–506. <https://doi.org/10.1007/s10010-023-00657-8>
- Fraunhofer-Institut für Kognitive Systeme IKS. (2025, Juli 12). Künstliche Intelligenz (KI) und maschinelles Lernen. <https://www.iks.fraunhofer.de/de/themen/kuenstliche-intelligenz.html>
- Gräßler, I., Oleff, C., Hieb, M., & Preuß, D. (2022). Automated Requirement Dependency Analysis for Complex Technical Systems. *Proceedings of the Design Society*, 2, 1865–1874. <https://doi.org/10.1017/pds.2022.189>
- Gräßler, I., Oleff, C., & Preuß, D. (2022). Proactive Management of Requirement Changes in the Development of Complex Technical Systems. *Applied Sciences*, 12(4), 1874. <https://doi.org/10.3390/app12041874>
- Gräßler, I., Preuß, D., Brandt, L., & Mohr, M. (2023). EFFICIENT FORMALISATION OF TECHNICAL REQUIREMENTS FOR GENERATIVE ENGINEERING. *Proceedings of the Design Society*, 3, 1595–1604. <https://doi.org/10.1017/pds.2023.160>
- Hertzke, P., Schaufuss, P., & Kampshoff, P. (2025). *New twists in electric-vehicle transition: A consumer perspective*. McKinsey Center for Future Mobility. <https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/new-twists-in-the-electric-vehicle-transition-a-consumer-perspective>
- Hovemann, A., Bitá, I. M., Aldade, A. A., Von Heißen, O., & Dumitrescu, R. (2025). Prompt Engineering im Systems Engineering: Potenziale und Grenzen moderner großer Sprachmodelle im V-Modell. *Zeitschrift für wirtschaftlichen Fabrikbetrieb*, 120(s1), 101–106. <https://doi.org/10.1515/zwf-2024-0139>
- IBM. (2025, Juli 12). Was sind große Sprachmodelle (Large Language Models, LLMs)? <https://www.ibm.com/de-think/topics/large-language-models>

- IHK Ulm. (2025, Juni 21). Herausforderungen und Chancen der Automobilzulieferindustrie. <https://www.ihk.de/ulm/hauptnavigation/online-magazin/im-fokus/herausforderungen-und-chancen-der-automobilzulieferindustrie-6649638>
- International Council on Systems Engineering (INCOSE). (2015). INCOSE systems engineering handbook: A guide for system life cycle processes and activities (4. Aufl.). John Wiley & Sons. <https://doi.org/10.1002/9781118999400>
- International Organization for Standardization (ISO). (2015). ISO/IEC/IEEE 15288:2015—Systems and software engineering—System life cycle processes (No. ISO/IEC/IEEE 15288:2015). International Organization for Standardization. <https://www.iso.org/standard/63711.html>
- Jiang, Z., Sun, B., Zhu, S., Yan, W., Wang, Y., & Zhang, H. (2025). A knowledge graph – based requirement identification model for products remanufacturing design. *Journal of Engineering Design*, 36(7–9), 1401–1424. <https://doi.org/10.1080/09544828.2024.2339163>
- Riesener, M., Kuhn, M., Schümmelfeder, S., Xiao, D., Norheim, J., Rebentisch, E., & Schuh, G. (2024). Active Learning with Pre-trained Language Models for Named Entity Recognition in Requirements Engineering. *Procedia CIRP*, 128, 339–344. <https://doi.org/10.1016/j.procir.2024.06.027>
- Rupp, C. & SOPHIST-Gesellschaft für Innovatives Software-Engineering. (2021). Requirements-Engineering und -Management: Das Handbuch für Anforderungen in jeder Situation (7., aktualisierte und erweiterte Auflage). Hanser. <https://doi.org/10.3139/9783446464308>
- Scheithauer, D. (2013). Analyse und Synthese im Systems Engineering. In M. Maurer & S.-O. Schulze (Hrsg.), *Tag des Systems Engineering* (S. 1–12). Carl Hanser Verlag GmbH & Co. KG. <https://doi.org/10.3139/9783446436039.001>
- Schleifer, S., Lungu, A., Kruse, B., Van Putten, S., Goetz, S., & Wartzack, S. (2024). Automatic derivation of use case diagrams from interrelated natural language requirements. *Proceedings of the Design Society*, 4, 2725–2734. <https://doi.org/10.1017/pds.2024.275>
- Schmidt, S. (2025). Automobilindustrie im Umbruch—Handlungsmöglichkeiten und Perspektiven. *ATZelextronik*, 20(1–2), 58–58. <https://doi.org/10.1007/s35658-024-1977-7>
- Siemens Digital Industries Software. (2020). Orchestrieren der Entwicklung von eingebetteten Anwendungen in der Automobilbranche: Definition und Planung von Anwendungen (White paper). Siemens Digital Industries Software. <https://www.sw.siemens.com/de-DE/>
- Timperley, L. R., Berthoud, L., Snider, C., & Tryfonas, T. (2025). Assessment of large language models for use in generative design of model based spacecraft system architectures. *Journal of Engineering Design*, 36(4), 550–570. <https://doi.org/10.1080/09544828.2025.2453401>
- Uygun, Y., & Momodu, V. (2024). Local large language models to simplify requirement engineering documents in the automotive industry. *Production & Manufacturing Research*, 12(1), 2375296. <https://doi.org/10.1080/21693277.2024.2375296>
- Van Remmen, J., Horber, D., Pickel, J., Goetz, S., & Wartzack, S. (2023). Potentials of NLP-based Assistance System in Requirements Management. *Tagungsband zum Stuttgarter Symposium für Produktentwicklung SSP 2023*, 91–102. <https://doi.org/10.18419/opus-13131>
- Wang, L., Wang, M.-C., Zhang, Y.-R., Ma, J., Shao, H.-Y., & Chang, Z.-X. (2025). Automated Identification and Representation of System Requirements Based on Large Language Models and Knowledge Graphs. *Applied Sciences*, 15(7), 3502. <https://doi.org/10.3390/app15073502>
- Xiao, Y., & Watson, M. (2019). Guidance on Conducting a Systematic Literature Review. *Journal of Planning Education and Research*, 39(1), 93–112. <https://doi.org/10.1177/0739456X17723971>
- Zhang, Y., Larsson, A., Larsson, T., Tian, W., Zhang, L., & Wang, W. (2024). Utilizing Natural Language Processing for Enhancing Collaborative Value-Driven Design of Smart Product Service System: Smart E-Vehicle Application. In L. M. Camarinha-Matos, A. Ortiz, X. Boucher, & A.-M. Barthe-Delanoë (Hrsg.), *Navigating Unpredictability: Collaborative Networks in Non-linear Worlds* (Bd. 726, S. 305–318). Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-71739-0_20