



# A Requirements Engineering Perspective to AI-Based Systems Development: A Vision Paper

Xavier Franch<sup>1,2</sup> , Andreas Jedlitschka<sup>2</sup> , and Silverio Martínez-Fernández<sup>1</sup> 

<sup>1</sup> Universitat Politècnica de Catalunya (UPC), Barcelona, Spain  
{xavier.franch, silverio.martinez}@upc.edu

<sup>2</sup> Fraunhofer IESE, Kaiserslautern, Germany  
andreas.jedlitschka@iese.fraunhofer.de

**Abstract.** *Context and motivation:* AI-based systems (i.e., systems integrating some AI model or component) are becoming pervasive in society. A number of characteristics of AI-based systems challenge classical requirements engineering (RE) and raise questions yet to be answered. *Question:* This vision paper inquires the role that RE should play in the development of AI-based systems with a focus on three areas: roles involved, requirements' scope and non-functional requirements. *Principal Ideas:* The paper builds upon the vision that RE shall become the cornerstone in AI-based system development and proposes some initial ideas and roadmap for these three areas. *Contribution:* Our vision is a step towards clarifying the role of RE in the context of AI-based systems development. The different research lines outlined in the paper call for further research in this area.

**Keywords:** Requirements Engineering · Artificial Intelligence · Machine Learning · AI-based System · Vision Paper · RE · AI · ML

## 1 Introduction

AI-based systems, defined as software systems that integrate artificial intelligence (AI) models and components [22], are becoming increasingly pervasive in society. Being yet-another-type of software system, the development of AI-based systems requires following usual software engineering practices [20] and, in particular, requirements engineering (RE) is expected to be applicable in this context.

Still, RE in the context of AI-based systems (which is sometimes referred to as RE4AI<sup>1</sup>) has been reported as challenging by several authors. Some authors have focused on particular RE issues (e.g., a precise definition of satisfaction of a specification in the presence of AI [3]). Others analyse RE4AI from a wider perspective. For instance, Ishikawa and Yoshioka conducted a questionnaire-based survey with 278 responses and report that “decision making with the customers” is the dominant concern when building

---

<sup>1</sup> Other authors are more specific and talk about RE for machine learning (ML) systems. In this paper, we have adopted the widest AI perspective, which includes ML.

ML-based systems [16]. Several works [1, 12, 22] enumerate a number of challenges related to RE, e.g., importance of context, consideration of data-related requirements and need to define new types of non-functional requirements, this latter aspect also mentioned by Horkoff's seminal paper on the topic [13].

These works, cited as examples, uncover a tension between the current practices of AI-based development and RE. This is partly motivated by the novel and fast emergence of AI in the software arena. The unprecedented evolving pace of new AI solutions and technologies puts the emphasis on creating new models and algorithms to solve all kinds of complex problems, disregarding methodological aspects required by the complexity to integrate these models and algorithms into a large software system [18]. This complexity calls for adopting well-established software engineering practices that have been largely ignored [20], RE being one of them. What are the requirements that apply to these models, to the data needed to build them, and to the algorithms to process them? Who is in charge of formulating these requirements? The answer to this type of questions will shape the form RE4AI will take in the future.

## 2 Background

From a technological stance, a cause of this tension is the data-oriented nature of AI-based systems. Data management has resulted in new **roles** involved in the development of AI-based systems. Besides, data lies at the heart of a major activity in AI-based system development, namely training, which may have its own requirements, different from those for the system-to-be, therefore yielding diverse **requirement scopes**. These new scopes may bring their particular perspectives on requirements, represented by new types of **non-functional requirements**, or redefinition of existing ones. In this paper, to make our vision concrete, we are going to focus on the three aforementioned aspects.

**Roles.** Based on a literature review, Pei et al. present an overview of the different roles involved in RE for ML systems, their RE-related concerns and challenges, and collaboration patterns among them [26]. Starting from the classical RE roles of Business Expert, Requirements Engineer and Software Engineer, they propose adding Domain Expert and Data Scientist. They model the collaboration among these actors using  $i^*$ , although the proposed model does not include the requirements engineer, which makes the responsibilities and dependencies of this role implicit or even hidden. Collaboration among Requirements Engineer and Data Scientist is also stressed as a key factor by Ahmad et al. [1].

Adopting a more specific stance and through an interview-based survey, Vogelsang and Borg take the data scientist perspective, given the importance of this role in ML system development [32]. The paper focuses on the activities done, processes followed and challenges found by data analysts in the context of RE4AI and does not explore connections with other roles. Still, the authors make a clear point that data scientist decisions should be subordinated to the classical job of the requirements engineer.

**Non-Functional Requirements (NFRs).** Several authors have explored which NFRs apply to AI-based systems; in fact, according to a mapping study by Martínez-Fernández et al., this is the hottest topic in the RE4AI-related literature [22].

A good number of papers explore a designated type of NFR in detail, e.g., safety, performance [4][29]. Other authors adopt a holistic perspective and investigate which NFRs apply to AI-based systems. For instance, Habibullah and Horkoff conducted an interview-based survey with ten practitioners [11] to elicit NFR types, their priorities, and most relevant NFR-related challenges. In summary, they state: (1) NFR types can be grouped into thematically-relevant clusters; (2) there are a number of new NFR types specifically related to AI-based systems, or whose relevance excels in this context, such as trust, ethics and explainability [2]; (3) other traditional NFR types, such as usability, are not considered so priority (although as usual, there are conflicting views on the importance of this and other NFR types in AI-based systems [11]).

**Requirements Scope.** Some authors have already considered the need to identify the concrete system part, which is the target of a particular NFR. For instance, performance, as discussed in [29], refers to model performance. More generally, Siebert et al. propose a layered view approach to ML system quality, from Environment to System/Infrastructure and then to ML Components, embracing model and data [28]. This approach is also adopted by Habibullah et al., who argue that requirements (concretely, NFRs) over ML systems may apply to different scopes [10]. They propose as scopes: Training Data, ML Algorithm, ML Model, Results and the whole ML System. Then, they explore which NFRs apply to each scope. In some cases, application requires an adaptation of the standard definition (e.g., from a software system perspective to a data perspective).

At their turn, adopting an intentional viewpoint, Nalchigar et al. identify three perspectives in modelling ML requirements [24]: (1) Business view, expressing stakeholder requirements; (2) Analytics Design view, representing the design of ML solutions for addressing the former requirements; (3) Data Preparation view, conceptualising the design of data preparation tasks. The latter two views are related to some of the scopes identified in [10], although with emphasis on design consequences.

### 3 RE4AI: Vision and Roadmap

In this paper, we envision that **RE shall become the cornerstone that coordinates all roles, activities and artefacts that are involved in the development of AI-based systems.** We support this vision upon the following arguments:

- Requirements engineers possess a number of skills that make them well-suited for this new challenge, especially communication skills [25]. For instance, they know how to talk to people of different profiles and how to bring them together. Therefore, they are in a good position to mediate the communication gap amongst roles.
- “Classical” RE distinguishes different scopes for requirements, e.g. stakeholder requirements, system requirements, etc. [14]. Therefore, considering additional scopes as those mentioned in the Sect. 2, seems to fit naturally in the discipline.
- Lately, new NFRs have been incorporated in the RE body of knowledge, in different types of systems (e.g., mobile games [30]), or due to societal needs (e.g., sustainability [5]). Thus, RE is well-prepared to replicate the process for AI-significant qualities, and help in the processes of which and where apply to every context.

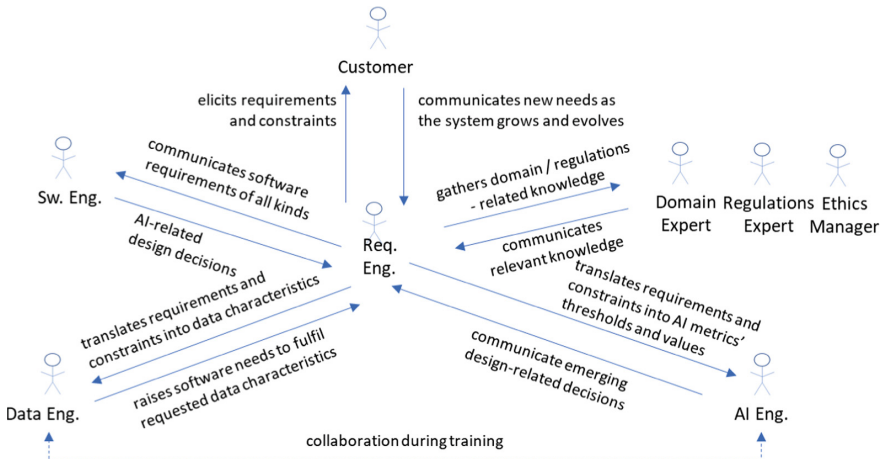
Building upon this vision and the background outlined in Sect. 2, we elaborate a roadmap for each of the three areas, which we are focusing on. The roadmap consists of a baseline research position followed by an enumeration of some research lines.

**Roles.** Our baseline research position aligns with Vogelsang and Borg’ statement on the need of the requirements engineer to act as a bridge among the customer and technical roles as data scientist [32]. For this reason, we place the requirements engineer role in the centre of the scene (see Fig. 1). Surrounding it, we identify several other roles (see definitions in Table 1 and most relevant relationships in Fig. 1):

- We split the concept of Business Expert from [26] into Customer, Domain Expert, Ethics Manager, and Regulation Expert, recognizing the importance of adhering to all kinds of regulations and social demands when developing AI-based systems.
- We introduce the Software Engineer as a multi-facet role embracing all software engineering roles different from RE: software architect, developer, etc.
- We have decided to split the role of data scientist into two: (i) the Data Engineer, who takes care of all data-related aspects in the typical AI/ML pipeline (mining, harvesting, selecting, cleaning, annotating, enriching, augmenting, ...); (ii) the AI Expert, who knows the algorithms and models existing in the AI discipline, when they can be applied and what results do they bring. It is worth remarking that, as usual, a person may play more than one role, therefore our identification of two different roles does not preclude that a single person, who could be labelled as a Data Scientist, ultimately plays both of them together.

**Table 1.** Roles involved in RE4AI.

Stakeholder	Main responsibility
Customer	Has the vision of the AI-based system and provides feedback when requested
Domain Expert	Has knowledge on the domain (including the data in that domain) in which the AI-based system will operate
Ethics Manager	Ensures that the AI-based systems work according to ethical principles
Regulations Expert	Ensures that regulations on trustworthiness, inclusiveness, etc., are fulfilled
Requirements Engineer	Formulates the needs of the customer, collaborating with all other roles
Data Engineer	Gathers, manipulates, and tests data to make it usable by other roles
AI Engineer	Knows the best algorithm to be applied in every situation
Software Engineer	Designs, develops, tests, and deploys software as required



**Fig. 1.** RE4AI: roles and a representative sample of their relationships.

This baseline position opens a research roadmap along the following lines:

- To complete a catalogue of roles and their responsibilities. Concerning responsibilities, goal-oriented (intentional) models as proposed in [26] look as an appropriate approach, also because this type of models is well-suited to include NFRs as discussed below.
- Related to the previous item, it can be argued that the presented figure has a classical flavour, not completely agile. On the one hand, we are not including a role such as Product Owner. On the other hand, all interactions are proposed to go through the Requirements Engineer, who could eventually become a bottleneck. We can envisage more agile micro-interactions, where, e.g., the Data Engineer and the AI Engineer may directly collaborate during the training process to curate the data set to achieve the required values for accuracy (represented with dotted lines in Fig. 1).
- The central position of the Requirements Engineer requires additional knowledge compared to a more traditional setting. For instance, the Requirements Engineer needs to understand what are the data characteristics that matter to Data Engineers (e.g., size, balance, ...) and how requirements relate to them.

**Requirements Scope.** We concur with Habibullah et al.'s vision on the existence of requirements scopes that distinguish software, data and AI algorithms. This baseline position opens a research roadmap along the following lines:

- Determine the full set of relevant scopes. For instance, some scope may be worth adding. Remarkably, we can think of adding a Data Engineering scope from the software perspective. For instance, when new data is needed, it may be necessary to develop some software component to gather this data from the source in appropriate quality, and this component should be developed according to its own requirements. Remarkably, such a Data Engineering scope could be useful in other contexts not

strictly related to AI-based systems where it is still necessary to acquire data from different sources (e.g., from IoT devices).

Another possible scope emerges if we consider not just software requirements but system requirements. In this case, we can think of a Hardware scope for which requirements on e.g. the type of processor (for instance, requiring the use of a GPU for efficiency reasons) or additional components (for instance, requiring a wattmeter in order to make energy efficiency measurable) become relevant, given the impact on runtime efficiency and even in accuracy.

- Clarify the workflow among different types of requirements and constraints. While the definition of scopes provides a static view of the types of requirements that apply in AI-based systems, there is a need to put all of them together into a holistic view, clarifying their relationships. See Fig. 2 for an example scenario showing how the Requirements Engineer elicits and documents requirements (R) and constraints (C) from a Customer deploying an app for plant recognition.

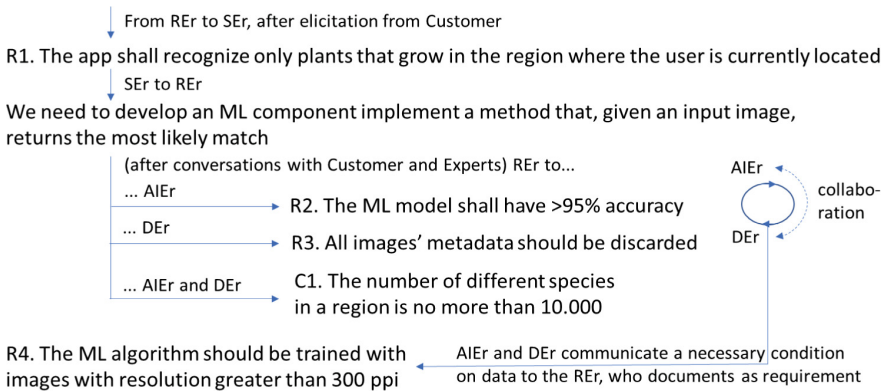


Fig. 2. RE4AI: example scenario showing the flow of requirements (roles identified by initials).


**Non-Functional Requirements.** Current approaches (*cf.* Background) consider all types of NFRs at the same level of abstraction, e.g., Habibullah and Horkoff’s clusters [11]. We envision the convenience of hierarchizing NFR types. In particular, we propose as a baseline position to use the structure proposed in the ISO/IEC 25010 standard [15], which distinguishes quality in use and product quality models, with the former defined in terms of the latter. In addition, because a number of NFR types may not apply to all the requirement scopes, or their definition may vary from scope to scope [11], we propose to replicate this structure for every scope (see Fig. 3).

From this baseline position, we foresee the following research lines:

- The composition and relationships of the different quality models is a significant long-term milestone to achieve by the community. Of course, it may be argued that,

because the use of standards is not widespread in the traditional RE context [7], it can be even harder to push for standards in this lively AI context, but still we believe that the structure that standards provide, entails a benefit *per se* to consolidate what is meant by RE4AI.

- In another vein, as hinted above in Fig. 3 and aligned with the terminology proposed, e.g., by the IREB association [14], we prefer to move from NFRs to quality requirements and constraints. The reason is, on the one hand, to adhere to current terminology promoted by certification bodies and other authors [27], and on the other hand the fact that constraints may play an important role when it comes to understand the limits of data in a particular context: a constraint may well limit the size of data, the period of availability, and other information that can be relevant to Data Engineers and AI Engineers to do their job.
- There are a number of concepts that have arisen in the AI community that relate to NFRs and quality, whose fit to this vision needs to be explored. Examples are: data smells [6], highly related to data requirements; Great Expectations (<https://greatexpectations.io/>), as an open standard for data quality; model cards [23], as an example of description of models which can serve to check whether requirements at the scope of ML Model are satisfied or not.



**Fig. 3.** RE4AI: different quality models.

## 4 Discussion

In this vision paper, we have reflected on the role of RE in the development of AI-based systems (RE4AI) and advocated that RE should articulate all activities and roles around. For space issues, we have focused the vision on three concrete major areas that directly relate to the data-oriented nature of AI-based systems, not considering other that can be equally important [3, 22]. For each area, we have envisaged a roadmap in the form of baseline position and research lines departing from this position.

These three areas have been presented as independent, but they are clearly interrelated. For instance, some NFR types will not apply to all scopes, or some scopes will not be of interest for all roles. In order to integrate these areas (and others that we are not addressing, e.g. verification and validation), we think of constructing conceptual models such as ontologies for knowledge representation [17] which can integrate all these concepts into a holistic model, as we have done in the field of architectures for AI-based systems [9]. Going further, we can think of linking requirements with design

decisions (e.g., which algorithms work better for the elicited requirements) and apply situational method engineering with this purpose, as we have done in previous works related to data-driven methods for RE [8].

We think that the vision presented in this paper may impact future research and practice in RE4AI: concerning research, we have delineated a number of research lines, which may trigger investigation in the community; concerning practice, this vision may contribute to clarify practical aspects that arise in every AI project, by identifying responsibilities of different roles, defining scopes that are different than in traditional systems, and helping to understand quality requirements and constraints in the context of AI-based systems. We acknowledge that practical impact needs to be considered in the long-term, once research progresses more in the short- and mid-term through new results in the suggested research lines. To make this impact possible, we foresee different actions that the community can take. Some are low-hanging fruits, such as continuing the series of workshops related to the topic, notably AIRE and RE4AI, associated with conferences as REFSQ and IEEE RE, and to educational programs in software and systems engineering curricula. Others can be more ambitious, e.g. promoting a new RE certification program in the IREB association, which could have a high practical impact.

**Acknowledgments.** This paper is part of the project TED2021-130923B-I00, funded by MCIN/AEI/<https://doi.org/10.13039/501100011033> and the European Union “NextGenerationEU”/PRTR.

## References

1. Ahmad, K., Bano, M., Abdelrazek, M., Arora, C., Grundy, J.: What’s up with requirements engineering for artificial intelligence systems? RE 1–12 (2021)
2. Balasubramaniam, N., Kauppinen, M., Hiekkänen, K., Kujala, S.: Transparency and explainability of AI systems: ethical guidelines in practice. In: Gervasi, V., Vogelsang, A. (eds.) REFSQ 2022. LNCS, vol. 13216, pp. 3–18. Springer, Cham (2022). [https://doi.org/10.1007/978-3-030-98464-9\\_1](https://doi.org/10.1007/978-3-030-98464-9_1)
3. Berry, D.M.: Requirements engineering for artificial intelligence: what is a requirements specification for an artificial intelligence? In: Gervasi, V., Vogelsang, A. (eds.) REFSQ 2022. LNCS, vol. 13216, pp. 19–25. Springer, Cham (2022). [https://doi.org/10.1007/978-3-030-98464-9\\_2](https://doi.org/10.1007/978-3-030-98464-9_2)
4. Burton, S., Gauerhof, L., Heinzemann, C.: Making making the case for safety of machine learning in highly automated driving. In: SAFECOMP, pp. 5–16 (2017)
5. Duboc, L., Penzenstadler, B., Porras, J., et al.: Requirements engineering for sustainability: an awareness framework for designing software systems for a better tomorrow. Requirements Eng. **25**, 469–492 (2020)
6. Foidl, H., Felderer, M., Ramler, R.: Data smells: categories, causes and consequences, and detection of suspicious data in AI-based systems. In: CAIN, pp. 229–239 (2022)
7. Franch, X., Glinz, M., Méndez, D., Seyff, N.: A Study about the knowledge and use of requirements engineering standards in industry. IEEE Trans. Software Eng. **48**(9), 3310–3325 (2022)
8. Franch, X., Henriksson, A., Ralyté, J., Zdravkovic, J.: Data-driven agile requirements elicitation through the lenses of situational method engineering. In: RE@Next, pp. 402–407 (2020)



9. Franch, X., Martínez-Fernández, S., Ayala, C., Gómez, C.: Architectural decisions in ai-based systems: an ontological view. In: QUATIC, pp. 18–27 (2022)
10. Mohammad Habibullah, K., Gay, G., Horkoff, J.: Non-Functional Requirements for Machine Learning: An Exploration of System Scope and Interest. CoRR abs/2203.11063 (2022)
11. Mohammad Habibullah, K., Horkoff, J.: Non-functional requirements for machine learning: understanding current use and challenges in industry. In: RE:, pp. 13–23 (2021)
12. Heyn, H.-M., Knauss, E., Pir Muhammad, A., et al.: Requirement engineering challenges for AI-intense systems development. In: WAIN, pp. 89–96 (2021)
13. Horkoff, J.: Non-functional requirements for machine learning: challenges and new directions. In: RE, pp. 386–391 (2019)
14. The International Requirements Engineering Board: IREB Certified Professional for Requirements Engineering – Foundation Level – Syllabus, v. 3.1.0 (2022)
15. ISO/IEC 25010:2011. Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — System and software quality models
16. Ishikawa, F., Yoshioka, N.: How do Engineers perceive difficulties in engineering of machine-learning systems' questionnaire survey. In: CESSER-IP, pp. 2–9 (2019)
17. Jurisica, I., Mylopoulos, J., Yu, E.: Ontologies for knowledge management: an information systems perspective. *Knowl. Inf. Syst.* **6**(4), 380–401 (2004). <https://doi.org/10.1007/s10115-003-0135-4>
18. Khomh, F., Adams, B., Cheng, J., Fokaefs, M., Antoniol, G.: Software engineering for machine-learning applications: the road ahead. *IEEE Softw.* **35**(5), 81–84 (2018)
19. Kuwajima, H., Yasuoka, H., Nakae, T.: Engineering problems in machine learning systems. *Mach. Learn.* **109**(5), 1103–1126 (2020)
20. Lwakatere, L.E., Raj, A., Crnkovic, I., Bosch, J., Holmström Olsson, H.: Large-large-scale machine learning systems in real-world industrial settings: a review of challenges and solutions. *Inf. Software Technol.* **127**, 106368 (2020)
21. Martínez-Fernández, S., Franch, X., Jedlitschka, A., Oriol, M., Trendowicz, A.: Developing and operating artificial intelligence models in trustworthy autonomous systems. In: Cherfi, S., Perini, A., Nurcan, S. (eds.) RCIS 2021. LNBIP, vol. 415, pp. 221–229. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-75018-3\\_14](https://doi.org/10.1007/978-3-030-75018-3_14)
22. Martínez-Fernández, S., Bogner, J., Franch, X., et al.: Software engineering for AI-based systems: a survey. *ACM Trans. Software Eng. Methodol.* **31**(2), 37e:1–37e:59 (2022)
23. Mitchell, M., Wu, S., et al.: Model cards for model reporting. In: FAT\*, pp. 220–229 (2019)
24. Nalchigar, S., Yu, E., Keshavjee, K.: Modeling machine learning requirements from three perspectives: a case report from the healthcare domain. *Requirements Eng.* **26**(2), 237–254 (2021)
25. Paech, B.: What is a requirements engineer? *IEEE Softw.* **25**(4), 16–17 (2008)
26. Pei, Z., Liu, L., Wang, C., Wang, J.: Requirements engineering for machine learning: a review and reflection. In: REW, pp. 166–175 (2022)
27. Pohl, K.: *Requirements Engineering - Fundamentals, Principles, and Techniques*. Springer (2010). <https://doi.org/10.1007/978-3-642-12578-2>
28. Siebert, J., et al.: Towards towards guidelines for assessing qualities of machine learning systems. In: Shepperd, M., Brito e Abreu, F., Rodrigues da Silva, A., Pérez-Castillo, R. (eds.) QUATIC 2020. CCIS, vol. 1266, pp. 17–31. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58793-2\\_2](https://doi.org/10.1007/978-3-030-58793-2_2)
29. Tuncali, C.E., Fainekos, G., Prokhorov, D., Ito, H., Kapinski, J.: Requirements-driven test generation for autonomous vehicles with machine learning components. *IEEE Transactions on Intelligent Vehicles* **5**(2), 265–280 (2020)
30. Valente, L., Feijó, B., Leite, J.C.S.P.: Mapping quality requirements for pervasive mobile games. *Requirements Eng.* **22**(1), 137–165 (2017)

31. Villamizar, H., Escovedo, T., Kalinowski, M.: Requirements engineering for machine learning: a systematic mapping study. In: SEAA, pp. 29–36 (2021)
32. Vogelsang, A., Borg, M.: Requirements engineering for machine learning: perspectives from data scientists. In: REW, pp. 245–251 (2019)