Addressing Challenges in AI-based Systems Development: A Proposal of Adapted Requirements Engineering Process

Kamil JABŁOŃSKI
Gdańsk University of Technology, Gdańsk, Poland
s165165@student.pg.edu.pl

Aleksander JARZĘBOWICZ
Gdańsk University of Technology, Gdańsk, Poland
olek@eti.pg.edu.pl
ORCID: 0000-0003-3181-4210

Abstract

[Context] Present-day IT systems are more and more dependent on artificial intelligence (AI) solutions. Developing AI-based systems means facing new challenges, not known for more conventional systems. Such challenges need to be identified and addressed by properly adapting the existing development and management processes. [Objective] In this paper, we focus on the requirements engineering (RE) area of IT projects and aim to propose the RE process that would be able to address at least some of the reported challenges. No proposal of such process could be found in the existing literature. [Method] We conducted a literature review using a snowballing technique to identify RE-related challenges for AI-based systems. Then, we compared several RE industry guides, selected a well-established RE process and adapted it by introducing additional practices. The additional practices were proposed as result of brainstorming and ideation process. [Results] The contributions of this paper include: a list of identified challenges, a set of additional practices to mitigate challenges and a model of the adapted RE process which integrates such practices. [Conclusions] The proposed process is available for validation activities and can be used by researchers and practitioners as a base for further adaptations of RE approaches to AI solutions.

Keywords: Requirements engineering, Artificial intelligence, RE4AI, challenges

Introduction

In recent years significant advances in the field of Artificial Intelligence (AI) could be observed, as more effective technologies, in particular machine learning and deep learning, enable more effective predictions and decision-making (Lukac, Milic and Nikolic 2018). As a result, there is a growing demand from the business side for AI-based systems that would provide advantage in a competitive market and a large number of industrial projects dedicated to such a purpose is initiated (Dalpiaz and Niu 2020).

Such projects are still IT projects that need to follow software engineering processes. However, the specifics of developing AI-based systems result in a need to adapt the existing processes, practices and techniques used. In this paper, we focus on requirements engineering (RE), which is one of the core activities of every IT project, as it is necessary to capture customers' requirements in order to provide them with a system that matches their needs (Wan et al. 2019; Przybyłek and Zakrzewski 2018; Przybyłek 2014). A substantial body of RE knowledge is available, however it is unrealistic to expect that RE could be conducted exactly the same way as for a more conventional system. For example, conventional systems are in most cases requirements-driven i.e. particular requirements like expected software features are expressed by human stakeholders, while AI-based system is often driven by outcome of AI model i.e. developers are asked to experiment with different solutions to achieve a given metric e.g. prediction rate (Bosch, Olsson and Crnkovic 2018). Another example concerns new key categories of non-functional requirements like transparency or retrainability which were previously either unknown or their significance was negligible (Horkoff 2019). Such and other issues resulted in the emergence of a new research trend (called "RE for AI" or "RE4AI") aimed at developing or adapting RE processes for AI-related projects (Belani, Vukovic and Car 2019).

In our work, we intended to learn about the additional challenges encountered in RE for AI-based systems. Next, we made an attempt to select a base RE process among existing ones and adapt it by introducing additional practices addressing the identified challenges. This paper provides answers to the following research questions:

- **RQ1**: What additional RE-related challenges are reported for AI-based systems, compared to more conventional systems?
- **RQ2**: Which of existing proposals of RE process can be adapted for development of AI-based systems?
- **RQ3**: What practices can be introduced to address the reported challenges?

The paper is organized as follows. First, we outline the background and related work. Next, we describe the literature search on RE-related challenges for AI-based systems and its results. In the subsequent section we present our main proposal - the adapted RE process and the practices introduced to it. The paper ends with conclusions.

Background and Related Work

It is recognized that development of AI-based systems is to some extent different in comparison to more conventional systems and, as such, requires more dedicated methods and processes (Wan et al. 2019; Bosch, Olsson and Crnkovic 2018). This includes, among others, a need for adapted RE processes, adjusted to AI specifics (Dalpiaz and Niu 2020).

RE is a widely recognized discipline, with a documented body of knowledge e.g. (IREB 2017; Wiegers and Beatty 2013), yet still considered as difficult and prone to challenges (Méndez Fernández et al. 2017). A number of recent works on general RE challenges applicable to virtually all IT projects (Méndez Fernández et al. 2017; Jarzębowicz and Ślesiński 2019; Przybyłek 2014) or a specific sub-class of them e.g. Agile projects (Wagner et al. 2018; Przybyłek and Zakrzewski 2018) is available. While it may be interesting to investigate whether all such known general challenges are applicable to AI-related projects, our work focused solely on RE challenges encountered during AI-based systems development and caused by the specific aspects of AI. Such specific challenges are reported by several sources, but they are rather a result of a single study e.g. series of interviews, than a summary of known RE-related challenges for AI-based systems we aimed at. We omit listing and describing such sources here, as the entire next section is dedicated to this purpose.

Several ideas on adapting RE practices for AI-based systems were proposed. A general SE process (not focused on RE) was designed by Hesenius et al. (2019). Nalchigar, Yu and Keshavjee (2021) introduced the RE framework, utilizing Goal-Oriented Requirements Engineering models to specify requirements for machine learning systems. Belani, Vukovic and Car (2019) provided a mapping between the challenges concerning AI-related entities and high-level RE activities that should address them. Vogelsang and Borg (2019) proposed additional practices to 4 key RE activities, as a response to challenges communicated by data scientists. There is however, at least to our knowledge, no proposal of a complete RE process for AI-based systems available.

Literature Review on Challenges

To learn about challenges reported about RE for AI-based system we performed a literature search and review using a snowballing approach (Wohlin 2014). Snowballing means starting with an initial set of previously identified papers and investigating the papers related to them through a citation network. The process is iteratively repeated for all papers qualified as relevant to the study's topic. We conducted the search using a number of previously identified papers as a start set and Google Scholar as a means to navigate through the citation network.

The qualification of the candidate sources was based on the following inclusion/exclusion criteria: IC1 – Peer-reviewed sources; IC2 – Sources in English; IC3 – Sources reporting RE-related challenges; EC1 – Non peer-reviewed sources (blogs, white papers etc.); EC2 – Sources in languages other than English; EC3 – Sources not available online; EC4 – Sources that do not contribute to knowledge about RE-related challenges (though instead they could e.g. present other aspects of RE in the context of AI-based systems or challenges that are completely not related to RE).

Data extraction was limited to identifying challenges described in the papers, together with all accompanying information, in order to fully understand a nature of each challenge and its contributing factors. We intentionally tried to be more inclusive i.e. consider the challenges that somehow relate to RE or can be addressed by RE (e.g. a challenge mainly related to testing can partially be mitigated by establishing acceptance criteria during RE). Only the completely unrelated challenges (e.g. the need for more computational power) were excluded. We omitted the explicit quality assessment of the sources found – it would be more appropriate if we e.g. tried to identify some RE/SE methods, tools or controlled experiments, but in case of challenges we did not filter them on the basis of the assessed sources' quality scores. The results of data extraction were further processed to find the same or very similar challenges and group them. The final results of the literature review are presented in Table 1.

Table 1: Challenges related to requirements engineering activities for AI-based systems

ID	Challenge	Description	Sources
Ch1	Difficult decision-making with stakeholders	It is much more difficult to establish stakeholders' expectations, as it requires determining unambiguous criteria regarding AI-based system's operation. Moreover, making key project decisions together with customers is problematic, due to their lack of knowledge and misunderstandings of AI mechanisms and their abilities (what is a realistic expectation, what is possible at all).	(Vogelsang and Borg 2019; Ishikawa and Yoshioka 2019)
Ch2	Effective quality evaluation	It is difficult to evaluate the quality of the proposed solution and select the appropriate metrics. In particular, it may be problematic to determine the performance measures of the solution, as the scope of testing is hardly obvious.	(Vogelsang and Borg 2019; Ishikawa and Yoshioka 2019; Arpteg et al. 2018)
Ch3	Effective configuration and change management	Development of AI-based system requires experimenting with alternative solutions and optimizing the finally selected solution. After some changes/updates are introduced or an alternative is proposed, such new version is evaluated through experimenting on datasets. The management of such experiments is time-consuming, moreover one should expect difficulty in comparing experiment results.	(Ishikawa and Yoshioka 2019; Arpteg et al. 2018)
Ch4	Critical importance of data	Data - especially training data – has a much greater significance for AI-based systems than for conventional ones and to a large extent determines system's future operations. Thus, the data should be prepared and tested as carefully as code for a more conventional system. Both, data quantity and quality should be carefully considered, which may not be possible at the same time, as there is often a choice between small sets of reliable, verified data and large sets of questionable quality, originating from uncertain sources.	(Belani, Vukovic and Car 2019; Vogelsang and Borg 2019; Ishikawa and Yoshioka 2019)
Ch5	The need for unique competencies	Development of an AI-based system and associated experiments/evaluations often require unique competencies, in particular from the mathematical domain (especially statistics).	(Ishikawa and Yoshioka 2019)
Ch6	Complex and diversified testing process	AI-based systems require more extensive testing, including several "objects of interest": the system as a whole, AI training model, other (non-AI) system components and finally the datasets used in training. The testing approaches to each of such objects differ significantly.	(Arpteg et al. 2018)
Ch7	Communication impaired by cultural differences	To succeed, all project team members have to cooperate and communicate, but such communication can be impaired by various differences related to e.g. geographical location, but also to their professional background and related mindsets.	(Arpteg et al. 2018)
Ch8	Restrictions imposed by legal and ethical aspects	Application of the generic legal requirements can result in unexpected consequences. For example, the General Data Protection Regulation states that personal data can only be used in ways specified by an explicit consent of the person involved. It implies that the developers must know what data will be required by their AI model before they start its development. Also, domain regulations (e.g. finances) and ethical guidelines, e.g. by European Commission (EC 2019) may need to be addressed.	(Vogelsang and Borg 2019)

Ch9	Negative side-effects of profiling	Many AI-based systems use profiling to determine user's characteristics and consequently provide him/her with the more fitting information. However, it results in filtering information, as such solutions omit everything that does not adhere to a given level of computed similarity. Therefore a user potentially loses a lot of valuable content. An appropriate balance should be found for profiling mechanisms and the corresponding requirements should be agreed between stakeholders.	(Kostova, Gürses and Wegmann 2020)
Ch10	Lack of oracle	It is difficult or even impossible to clearly define the correctness criteria for AI-based system outputs as well as the right outputs for each individual input.	(Ishikawa and Yoshioka 2019)
Ch11	Imperfection	It is intrinsically impossible to make adequate outputs for any of various possible inputs (i.e., 100% accuracy). It is unlikely that the same accuracy can be achieved for any input, for example neural networks are known to be prone to so called "adversarial examples", where a small modification of input (e.g. a few pixels of an image) results in completely different response of a network.	(Ishikawa and Yoshioka 2019)
Ch12	Explainability	AI-based systems are not necessarily transparent in their operation. It can be very difficult to explain the model (what has been learned) and even harder to explain particular decisions/predictions of the model.	(Vogelsang and Borg 2019; Ishikawa and Yoshioka 2019) (Arpteg et al. 2018)
Ch13	Interdependencies between system components	The components responsible for AI models and algorithms are just a part of a larger system. The data flow and dependencies between components are not easy to track (thus so called "unintended feedback loops" can be introduced). It can also be hard to determine which component or particular code fragment is responsible for implementing a given requirement. Moreover, a significant effort is required to update components, adjust interfaces etc. without introducing CACE (Changing Anything Changes Everything) effect.	(Belani, Vukovic and Car 2019; Arpteg et al. 2018)
Ch14	Effort estimation	The estimation of time and resources required is problematic due to the specifics of AI components' training and operation. A goal of the system can be well defined, but it does not directly translate into e.g. the number of iterations and experiments necessary before the acceptable results can be achieved.	(Arpteg et al. 2018)
Ch15	Ensuring data privacy	Ensuring the adequate level of users' privacy is not easy, especially regarding users' data included in training datasets and the data used during system's operation. External regulations may restrict the way the data can be used e.g. request that only aggregated and/or anonymized data is used as input. Another issue is to prevent users' data retrieval - although the information in an AI model is obscured and not easy to transform back to humanly readable form, it is not impossible to do so. Finally, there is also a need to efficiently perform data exploration, develop models, and troubleshoot problems, thus a proper balance between that and privacy has to be found.	(Arpteg et al. 2018)
Ch16	Freedom from discrimination	AI-based systems are designed to "discover" patterns in the training data and apply them to make decisions/predictions during their operation. However, some possible patterns would be clearly unacceptable according to law regulations and/or social standards e.g. filtering job candidates according to race or gender. In case of AI-based systems, the discrimination is not easy to determine because it is not reflected in explicit encoded rules, but can e.g. be a result of an unbalanced training set, where some groups are underrepresented.	(Vogelsang and Borg 2019)

A Proposal of the Requirements Engineering Process for AI-based Systems

Selection of the base process

We planned to address the identified challenges by introducing dedicated mitigating practices to the RE process. This brought us to RQ2 and the decision which RE process should be used as a base. As we wanted it to be a process used in practice, we reviewed the available industrial standards and guidelines dedicated to RE and to business analysis (BA). BA is a more general domain that focuses on identifying the needs of an organization and introducing changes that will deliver value to stakeholders. Such change can have nothing to do with developing or modifying IT system(s), but in case it does, sources on BA provide a good guidance on RE activities, thus they were considered as well

We selected four well-established sources: Certified Professional for Requirements Engineering syllabus by International Requirements Engineering Board (IREB 2017), Business Analysis Body of Knowledge (BABOK v3) by International Institute of Business Analysis (IIBA 2015), and two sources published by Project Management Institute: Business Analysis for Practitioners: A Practice Guide (PMI 2015) and Requirements Management: A Practice Guide (PMI 2016). All the documents were reviewed and compared with respect to scope and detailed contents (the results of scope comparison are shown in Fig. 1, where corresponding areas described in particular sources are juxtaposed). The results were thoroughly discussed by both authors and finally a decision to select (PMI 2016) was reached, because this source covers a wide spectrum of areas/processes (see Fig. 1) and thus enables RE activities to have impact on other areas of development and management in an IT project. Moreover, the source is dedicated to RE, without dealing with some issues covered by BA sources, but not applicable to IT projects.

<table>
<tr><th>BABOK</th><th>IREB</th><th>PMI BA Guide</th><th>PMI RM Guide</th></tr>
<tr><td>Strategy analysis</td><td></td><td>Needs assessment</td><td>Needs assessment</td></tr>
<tr><td>Business analysis planning and monitoring</td><td></td><td>Business analysis planning</td><td>Requirements management planning</td></tr>
<tr><td>Elicitation and collaboration</td><td>Requirements elicitation</td><td rowspan="3">Requirements elicitation and analysis</td><td>Requirements elicitation</td></tr>
<tr><td rowspan="2">Requirements analysis and design definition</td><td>Requirements documentation</td><td rowspan="2">Requirements analysis</td></tr>
<tr><td>Requirements negotiation and validation</td></tr>
<tr><td>Requirements life cycle management</td><td>Requirements management</td><td>Traceability and monitoring</td><td>Requirements monitoring and controlling</td></tr>
<tr><td>Solution evaluation</td><td></td><td>Solution evaluation</td><td>Solution evaluation</td></tr>
<tr><td></td><td></td><td></td><td>Project or phase closure</td></tr>
</table>

Fig 1. The scope of RE and BA industrial guides – a comparison

Adaptation of the process to address challenges

The next step was to adapt the selected process, so it would mitigate the challenges described in the previous section). We planned to achieve it by introducing additional practices, rather than by significantly modifying the base process. The reason was that the base process covers the RE activities to be done for probably any IT project – it is hard to expect that in case of an AI-based system e.g. requirements elicitation or solution evaluation could be omitted. During our work, we thoroughly reviewed the guide and identified (sub)activities where remedies to particular challenges could be added. We also relied on brainstorming and exchange of ideas to establish the challenge-mitigating practices. The existing literature was used as one of the inputs to brainstorming e.g. some of the papers found in our literature review mentioned potential remedies. We built upon such proposals, however due to a very dynamic and adaptive nature of our ideation process, it is impossible to document full traceability of each idea's origins. Our work took several iterations as both authors exchanged and discussed ideas, before deciding about the final adapted RE process.

The description of the resulting process

The RE process according to (PMI 2016) is shown in Fig. 2. The figure depicts the activities and their main sub-activities (except **Needs assessment**, which is not decomposed). The arrows indicate the ordering the activities take place in the RE process. The loop between **Requirements elicitation**, **Requirements analysis** and **Solution evaluation** indicates that several iterations including these activities can take place. **Requirements monitoring and controlling** is not connected using arrows, because it is conducted continuously "in the background", like most of management activities in IT projects. The red circles with identifiers like P1 symbolize additional practices we propose to include. In the remainder of this section we briefly outline the activities of the RE process, based on (PMI 2016). Due to space limitations and possible copyright violations we are not able to provide all the details and the interested reader is referred to the source document. Our descriptions will focus on the additional practices that can be included in the process to address the challenges identified for AI-based systems. The key areas and activities of the RE process as well as challenges and practices discussed are distinguished using bold fonts.

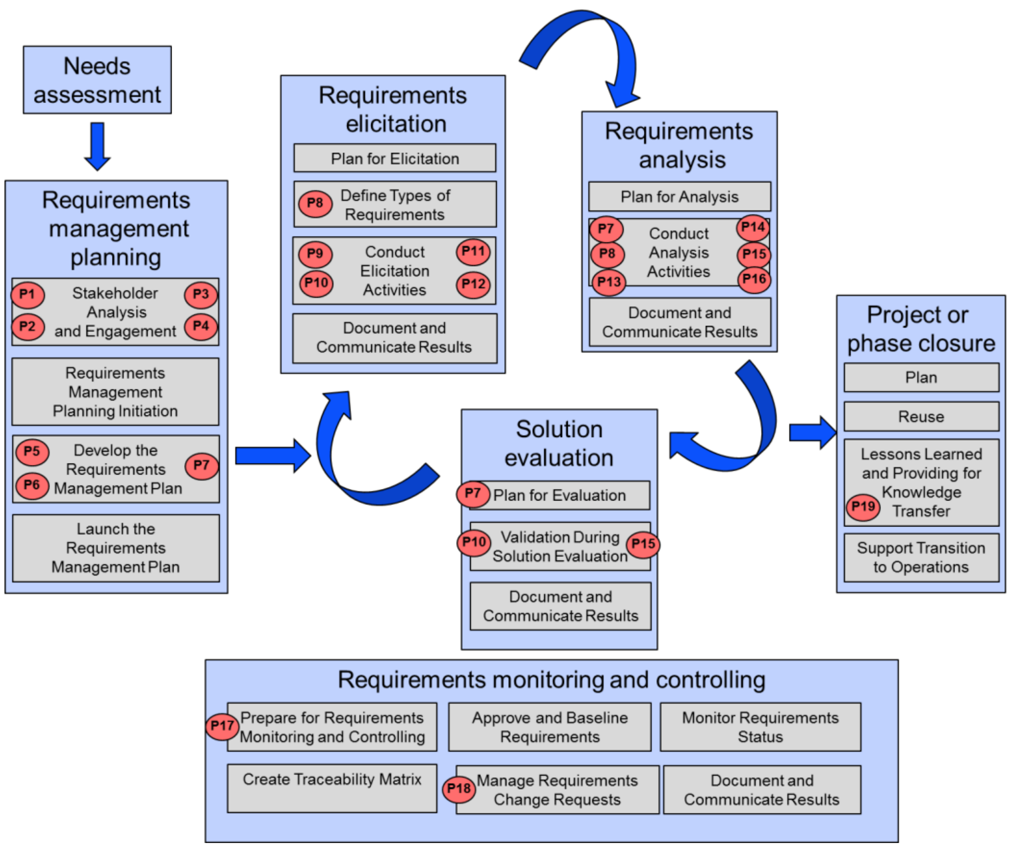

Fig 2. Adapted requirements engineering process model with additional practices

The RE process starts with **Needs assessment**, which often precedes the project itself. It is supposed to identify and analyze a business problem or strategic organizational need in order to determine high-level needs definition that will ultimately be used to determine viable solution options. As the business problem is usually quite abstract, our analysis did not identify any additional practices for this activity. For this reason we also omitted visualizing its sub-activities in Fig. 2.

The **Requirements management planning** activity includes the sub-activities that take place at the beginning of the project. One of them is **Stakeholder analysis and engagement**, which should result in creating a stakeholder register, analyzing stakeholders' characteristics and initiating contact with them to foster their engagement. The RE process for AI-based system should involve specific additional stakeholders. The critical importance of data (**Ch4**) and the expectations to ensure freedom from discrimination (**Ch16**) suggest a need for adequate competencies in preparing and processing data i.e. involvement of a data scientist (**P1**). The impact of legal and ethical requirements (**Ch8**) translates into a need for a legal expert as an additional stakeholder (**P2**). The adequate data privacy (**Ch15**) would probably require competencies in both legal regulations and AI mechanisms, thus both stakeholders (**P1**&**P2**) should cooperate on that throughout the project. The need for unique competencies in mathematics and statistics (**Ch5**) has also to be considered and, in case project team lacks them, additional people should be hired or at least some consultancy provided (**P3**). The analysis of stakeholders' characteristics should consider the potential cultural differences (**Ch7**) and ensure that the overall group of stakeholders covers all the relevant viewpoints with respect to nationalities, mindsets etc. (**P4**).

Another sub-activity (**Development of requirements management plan**) is supposed to define how requirement activities of the project will be planned and managed. The plan should foresee the need to consider the issue of profiling (**Ch9**) - such task has to be conducted jointly by all relevant stakeholders in order to reach a consensus about the degree to which profiling will be used by the system (**P5**). The anticipated communication problems with customer representatives (**Ch1**) should result in planning additional resources and time dedicated for educating stakeholders and for eliciting, analyzing and validating their requirements (**P6**). Another issue is the difficulty to evaluate the quality of the proposed solution (**Ch2**). Handling this challenge requires allocating necessary time and resources to think over and define the metrics and acceptance criteria as well as to ensure that all interested parties understand and accept such metrics/criteria (**P7**).

Requirements elicitation activity covers the discovery process by gathering information from stakeholders and other sources. Effective elicitation is warranted by active involvement of stakeholders and communication with them. One of its initial sub-activities is to **Define types of requirements**, which should result in designing the classification of requirements to be elicited with respect to the level of abstraction/detail (e.g. business goals, user requirements, system requirements) and addressed aspect (e.g. functional, non-functional, constraints). Different classifications can be found in renown sources - e.g. (Wiegers and Beatty 2013) vs. (IIBA 2015) - in each project it is possible to develop a classification most suitable for it. However, requirements related to data are rarely distinguished, but rather included in e.g. functional requirements. Given the importance of data (**Ch4**), we propose to explicitly establish a category of "data requirements", define the attributes that have to be specified for such requirements as well as possible interdependencies between them and other categories of requirements (**P8**). It will allow to group such key requirements and consider them jointly, instead of distributing them between various categories. The most elaborate sub-activity is to **Conduct elicitation activities** and a number of practices can be introduced here. One of the elicitation techniques mentioned in (PMI 2016) is document analysis, which usually means a review of documents related to the client organization. We advocate referring in addition to the sources which summarize the state-of-the-art AI solutions, e.g. (Stateoftheart AI 2023)(**P9**). Such solutions are quite often a result of cooperation of scientific/industrial teams and significant effort put into such enterprise. The challenges of decision-making with stakeholders (**Ch1**) and coping with imperfection (**Ch11**) caused by stakeholders' attitude, can be mitigated by confronting stakeholders with the best known solutions. If such solutions are not able to achieve a given result or ensure 100% correctness for any input, then probably stakeholders can be persuaded to adjust their expectations.

The use of other elicitation techniques can be found helpful in the context of AI-based system. The needs to acquire unique competencies (**Ch5**) and to facilitate consensus about key issues (like: data to be used (**Ch4**), quality metrics/criteria (**Ch2**), legal and ethical aspects (**Ch8**)) can be addressed by group-based elicitation techniques like workshops or focus groups (**P10**). These techniques allow to confront different viewpoints, share knowledge, discuss and reconcile differences. As tracking dataflow and dependencies between the components of an AI-based system is challenging (**Ch13**), a special attention should be paid to the technique called interface analysis. It can help to establish interdependencies and boundaries by determining the input and output needs of each interfacing component. This technique can be supported by document analysis focused on technical documentation of components and external cooperating systems (**P11**).

Prototyping is a RE technique often used to learn more about expectations about user interface and human-computer interaction. It is also however possible to design technical prototypes aimed at developing a working solution to test whether a given idea or requirement is technically feasible. Given difficulty in effort estimation (**Ch14**) and large costs (AI model training can take weeks), it seems worthy to build a technical prototype with a reduced training dataset (in cases it is possible) and demonstrate it to stakeholders in order to obtain their feedback (**P12**).

In case of **Requirements analysis** activity, its main sub-activity **Conduct analysis activities** is of primary interest. This sub-activity has a substantial scope and includes: developing requirements' attributes, selecting requirement models, deriving additional requirements, assigning priorities and conducting verification and validation. The significance of data to be used in AI solution's training and operation (**Ch4**) implies that a particularly careful analysis is necessary for all items from "data requirements" category. It is a follow-up of the work done in **Define types of requirements**. Data requirements must ensure that the training dataset is representative to the target operating environment. Both the quantity and quality of training data have to be considered and a decision balancing them has to be made (larger datasets with more examples vs. completeness, consistency, quality of annotations etc.) (**P8**).

In addition, data requirements should tend to be assigned with higher priorities, as their importance to the final project's outcome cannot be overstated (**P13**). In case of more iterative/adaptive development process, the higher priorities will also cause such requirements to be implemented sooner and thus allow to minimize the related risk. Similarly, the requirements about the explainability of the system's operation (**Ch12**) should be assigned with higher priorities and carefully analyzed, verified and validated with stakeholders (**P14**).

Requirements analysis should also focus on metrics and acceptance criteria for quality evaluation (**Ch2**). Similarly to requirements elicitation, additional resources should be allocated to ensure that such metrics and criteria are commonly understood and acceptable to all interested stakeholders (**P7**). Another issue is the selection of a testing dataset that will be used to check if the developed system fulfills its requirements, as testing is always limited and cannot demonstrate 100% correctness (**Ch11**) is A question arises: who should be responsible for such selection? The developers may not be the best choice, as they could focus on optimizations for this dataset only, while system's effectiveness in target environment will turn out to be worse. The customer and non-technical stakeholders may in turn lack sufficient competencies. A possible solution is to involve an independent third party assigned with such responsibility (**P15**).

Considering the challenges of tracking relationships between system components (**Ch13**) and problematic configuration and change management (**Ch3**), the selection of requirement models should take such issues into account. It can be beneficial to rely on modeling techniques that allow to capture the interfaces between components or systems and the interaction between them expressed as events, dataflows etc. (**P16**). A number of techniques is potentially useful here, including scope models (context diagram, ecosystem map), data models (data flow diagram, data dictionary) and interface models (system interface table, N2 diagram).

The activity of **Solution evaluation** is performed to validate the solution (system), to determine how well it meets the expressed needs. No new practices were proposed here, but some of those described earlier have impact on **Solution evaluation** as well. In case additional effort was made to explain the metrics and acceptance criteria to the stakeholders (**Ch2**), all agreements made then (metrics, evaluation techniques) should be incorporated in **Plan for evaluation** sub-activity (**P7**). The group-based techniques (**P10**), already suggested for requirements elicitation, will also be suitable for **Validation during solution evaluation** sub-activity, as they enable knowledge sharing and discussions. This sub-activity may also involve an independent third party (**P15**), if a decision was previously made to make them responsible for conducting validation.

The continuous process of **Requirements monitoring and controlling** covers issues like configuration management, change management, maintaining traceability and monitoring the current state of requirements. Such tasks are performed according to procedures defined during **Prepare for requirements monitoring and controlling** sub-activity. One should expect difficulties in configuration and change management caused by specifics of AI-related projects – introducing changes in datasets and/or AI models, tracking such "versions", comparing them etc. differs from the established practices of software engineering (**Ch3**). For this purpose, a dedicated change management procedure that incorporates specifics of AI context should be explicitly defined (**P17**). The already introduced challenges about change management (**Ch3**) and establishing training data (**Ch4**) suggest that **Manage requirements change requests** can benefit from more advanced techniques of dependency analysis and impact analysis, as well as traceability matrices and change control boards (**P18**).

Project or phase closure is the last activity with the purpose of finalizing all works to formally complete the project or phase. It includes **Lessons learned and providing for knowledge transfer** sub-activity. In the context of AI-related projects, the knowledge about comparing different AI models and managing changes introduced to them (**Ch3**), as well as effort estimations and their accuracy known on hindsight (**Ch14**) are especially valuable. Thus, project closure needs to ensure that essential information (especially about datasets, AI models and metrics used in the project) is properly documented and can be transferred to other projects (**P19**).

In addition to Fig. 2 and the descriptions provided in this section, Table 2 shows the mapping between challenges and practices that address them.

Table 2: Requirements engineering additional practices introduced to address challenges

Challenge	Practice	Practice description
Ch1: Difficult decision-making with stakeholders	P6	Plan additional resources and time dedicated for eliciting, analyzing and validating stakeholders' requirements
	P9	Conduct document analysis and review of sources summarizing state-of-the-art solutions and trends in the AI domain
	P14	Assign higher priorities to requirements on explainability
Ch2: Effective quality evaluation	P6	Plan additional resources and time dedicated for eliciting, analyzing and validating stakeholders' requirements
	P7	Plan additional resources and time for defining adequate metrics/acceptance criteria and for educating customer representatives and other stakeholders in this matter
	P10	Use group-based techniques like workshop or focus group for the purpose of requirements elicitation and solution evaluation
Ch3: Effective configuration and change management	P16	During requirements analysis, use modeling techniques that enable capturing interdependencies and interactions between components/systems (e.g. context diagram, ecosystem map, data flow diagram, data dictionary, state diagram, system interface table, N2 diagram)
	P17	Define (and follow) a dedicated change management procedure that incorporates specifics of AI context
	P18	Conduct dedicated dependency analysis and impact analysis for each change to requirements from Data Requirements category
	P19	Document lessons learned including all datasets, AI models and metrics used in the project in a suitable form that enables easy reuse and information retrieval
Ch4: Critical importance of data	P1	Include data scientists as stakeholders (or, if already present, extend their responsibilities in the project)
	P8	Define a new category of requirements: Data Requirements and pay special attention to eliciting, analyzing and verifying such requirements
	P10	Use group-based techniques like workshop or focus group for the purpose of requirements elicitation and solution evaluation
	P13	Assign higher priorities to requirements from Data Requirements category
	P18	Conduct dedicated dependency analysis and impact analysis for each change to requirements from Data Requirements category
Ch5: The need for unique competencies	P3	Involve experts from mathematics and statistics domains as stakeholders
	P10	Use group-based techniques like workshop or focus group for the purpose of requirements elicitation and solution evaluation
Ch6: Complex and diversified testing process	P8	Define a new category of requirements: Data Requirements and pay special attention to eliciting, analyzing and verifying such requirements
Ch7: Communication impaired by cultural differences	P4	Ensure that the overall group of stakeholders includes the representatives who cover all the relevant viewpoints with respect to different cultures and mindsets
Ch8: Restrictions imposed by legal and ethical aspects	P2	Include legal experts as stakeholders (or, if already present, extend their responsibilities in the project)
	P10	Use group-based techniques like workshop or focus group for the purpose of requirements elicitation and solution evaluation
Ch9: Negative side-effects of profiling	P5	Guide the stakeholders to reach a consensus and explicitly state the degree to which user profiling is desirable and how the information gathered this way is to be processed
Ch10: Lack of oracle	P7	Plan additional resources and time for defining adequate metrics/acceptance criteria and for educating customer representatives and other stakeholders in this matter
Ch11: Imperfection	P9	Conduct document analysis and review of sources summarizing state-of-the-art solutions and trends in AI domain
	P15	Involve an independent third party responsible for preparing testing dataset and for solution evaluation

Ch12: Explainability	P14	Assign higher priorities to requirements on explainability
Ch13: Interdependencies between system components	P11	Use interface analysis and document analysis in requirements elicitation
	P16	During requirements analysis, use modeling techniques that enable capturing interdependencies and interactions between components/systems (e.g. context diagram, ecosystem map, data flow diagram, data dictionary, state diagram, system interface table, N2 diagram)
Ch14: Effort estimation	P12	Use technical prototypes in requirements elicitation
	P19	Document lessons learned including all datasets, AI models and metrics used in the project in a suitable form that enables easy reuse and information retrieval
Ch15: Ensuring data privacy	P1	Include data scientists as stakeholders (or, if already present, extend their responsibilities in the project)
	P2	Include legal experts as stakeholders (or, if already present, extend their responsibilities in the project)
Ch16: Ensuring freedom from discrimination	P1	Include data scientists as stakeholders (or, if already present, extend their responsibilities in the project)

Conclusions

In this paper we reported a research study aimed at identifying RE-related challenges for AI-based systems and addressing them with a tailored RE process which includes additional dedicated practices. The answers to the research questions posed were obtained through literature search (RQ1), analysis and comparison of industrial RE and BA guides (RQ2) and ideation process (RQ3). We were able to identify 16 unique challenges specific for RE activities applied to AI-based systems (RQ1). We reviewed 4 well-established industry guides with respect to their scope, structure and detailed contents and decided to use one of them – “Requirements management: A practice guide” published by Project Management Institute (PMI 2016) as a base for designing a dedicated RE4AI process (RQ2). We identified 19 additional RE practices addressing the challenges and positioned them within the specific activities of the base RE process (RQ3). We also started validating such results through interviews with industry experts and the feedback obtained up to now is encouraging (validation activities are not described here due to paper size restrictions).

Our study has several limitations which could affect its validity. The literature search is prone to omitting some relevant sources due to a non-representative start set or wrong decisions about papers’ qualification. We tried to minimize these threats by following the guidelines on snowballing, but they could not be entirely eliminated. The ideation process could fail to produce the most optimal result as it is dependent on its contributors and their creativity.

As (according to our knowledge) no complete RE process for AI-based systems was published, our proposal’s implications for research include the possibility for other researchers to modify the described process by adding/changing practices or to design a completely alternative one (e.g. using other standard/guide as a basis). A comparison and validation of such RE processes is also a way our contributions can be built upon. As for implications for practice, our work can be used by practitioners, especially business and system analysts, working on AI-based systems to improve their RE processes.

The promising directions of future research include further validation and improvements of the process model according to feedback obtained from the industry. It is also possible (and desirable) to keep track of new reports on challenges and, accordingly, adjust the RE process to address them as well.

References

- Arpteg, A., Brinne, B., Crnkovic-Friis, L. and Bosch, J. (2018), ‘Software Engineering Challenges of Deep Learning,’ In Proc. of 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), 50-59.

- Belani, H., Vukovic M., and Car, Z. (2019), 'Requirements Engineering Challenges in Building AI-Based Complex Systems,' In Proc. of 27th International Requirements Engineering Conference Workshops, 252-255.
- Bosch, J., Olsson, H. and Crnkovic, I. (2018), 'It takes three to tango: Requirement, outcome/data, and AI driven development,' In Proc. of the International Workshop on Software-intensive Business: Start-ups, Ecosystems and Platforms (SiBW 2018), Espoo, Finland, 3 December 2018, 177-192.
- Dalpiaz, F. and Niu, N. (2020), 'Requirements engineering in the days of artificial intelligence,' *IEEE Software*, 37(4), 7–10.
- European Commission's Independent High-level Expert Group on Artificial Intelligence (2019), 'Ethics Guidelines for Trustworthy AI,' Available at: https://ec.europa.eu/newsroom/dae/document.cfm?doc_id=60419 (Accessed: 29 September 2023).
- Hesenius, M., Schwenzfeier, N., Meyer, O., Koop, W. and Gruhn, V. (2019), 'Towards a software engineering process for developing data-driven applications,' In Proc. of 7th International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering, 35-41.
- Horkoff, J. (2019), 'Non-functional requirements for machine learning: Challenges and new directions,' In Proc. of 27th International Requirements Engineering Conference (RE), 386-391.
- IIBA International Institute of Business Analysis (2015), 'A guide to the business analysis body of knowledge (BABOK Guide),' ver. 3.
- IREB International Requirements Engineering Board (2017), 'IREB Certified Professional for Requirements Engineering,' ver. 2.2
- Ishikawa, F. and Yoshioka, N. (2019), 'How Do Engineers Perceive Difficulties in Engineering of Machine-Learning Systems? - Questionnaire Survey,' In Proc. of IEEE/ACM Joint 7th International Workshop on Conducting Empirical Studies in Industry (CESI) and 6th International Workshop on Software Engineering Research and Industrial Practice (SER&IP), 2-9.
- Jarzębowicz, A. and Ślesiński, W. (2019), 'What is troubling IT analysts? A survey report from Poland on requirements-related problems,' In Proc. of 20th KKIO Software Engineering Conference, 3-19.
- Kostova, B., Gürses, S. and Wegmann, A. (2020), 'On the Interplay between Requirements, Engineering, and Artificial Intelligence,' In Proc. of 26th International Conference on Requirements Engineering: Foundation for Software Quality (REFSQ) Workshops.
- Lukac, D., Milic, M. and Nikolic, J. (2018), 'From Artificial Intelligence to Augmented Age An Overview,' In Proc. of 2018 Zooming Innovation in Consumer Technologies Conference, 100-103.
- Méndez Fernández, D., Wagner, S., Kalinowski, M., Felderer, M., Mafra, P., Vetrò, A. and Conte, T. (2017), 'Naming the pain in requirements engineering,' *Empirical Software Engineering*, 22(5), 2298-2338.
- Nalchigar, S., Yu, E. and Keshavjee, K. (2021), 'Modeling machine learning requirements from three perspectives: a case report from the healthcare domain,' *Requirements Engineering*, 26, 237-254.
- PMI Project Management Institute (2015), 'Business analysis for practitioners. A practice guide'.
- PMI Project Management Institute (2016), 'Requirements management: A practice guide'.
- Przybyłek, A. (2014), 'A Business-Oriented Approach to Requirements Elicitation,' In Proc. of 9th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE), 152-163.
- Przybyłek, A. and Zakrzewski, M. (2018), 'Adopting Collaborative Games into Agile Requirements Engineering,' In Proc. of 13th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE), 54-64.
- Stateoftheart AI (2023), 'Stateoftheart AI community website,' Available at: https://www.stateoftheart.ai (Accessed: 2 October 2023).
- Vogelsang, A. and Borg, M. (2019), 'Requirements Engineering for Machine Learning: Perspectives from Data Scientists,' In Proc. of 27th International Requirements Engineering Conference Workshops, 245-251.

- Wagner, S., Méndez Fernández, D., Kalinowski, M. and Felderer, M. (2018) ‘Agile requirements engineering in practice: Status quo and critical problems,’ *CLEI Electronic Journal*, 21(1), 6:1-6:15.
- Wan, Z., Xia, X., Lo, D. and Murphy G. (2019), ‘How does machine learning change software development practices?,’ *IEEE Transactions on Software Engineering*, 47(9), 1857-1871.
- Wiegers, K. and Beatty, J. (2013), ‘Software Requirements,’ 3rd ed., Microsoft Press, ISBN: 978-0735679665.
- Wohlin, C. (2014), ‘Guidelines for snowballing in systematic literature studies and a replication in software engineering,’ In Proc. of 18th Conference on Evaluation and Assessment in Software Engineering (EASE), 1–10.