

# A Catalogue of Concerns for Specifying Machine Learning-Enabled Systems

Hugo Villamizar, Marcos Kalinowski, and H elio Lopes

Pontifical Catholic University of Rio de Janeiro, Rio de Janeiro, Brazil  
{hvillamizar, kalinowski, lopes}@inf.puc-rio.br

**Abstract.** Requirements engineering (RE) activities for machine learning (ML) are not well-established and researched in the literature. Many issues and challenges exist when specifying, designing, and developing ML-enabled systems. Adding more focus on RE for ML can help to develop more reliable ML-enabled systems. Based on insights collected from previous work and industrial experiences, we propose a catalogue of 45 concerns to be considered when specifying ML-enabled systems, covering five different perspectives we identified as relevant for such systems: objectives, user experience, infrastructure, model, and data. Examples of such concerns include the execution engine and telemetry for the infrastructure perspective, and explainability and reproducibility for the model perspective. We conducted a focus group session with eight software professionals with experience developing ML-enabled systems to validate the importance, quality and feasibility of using our catalogue. The feedback allowed us to improve the catalogue and confirmed its practical relevance. The main research contribution of this work consists in providing a validated set of concerns grouped into perspectives that can be used by requirements engineers to support the specification of ML-enabled systems.

**Keywords:** Requirements engineering · requirements specification · perspectives · concerns · machine learning · artificial intelligence.

## 1 Introduction

The incorporation of machine learning (ML) components into the software of companies from all sectors is increasingly common. We refer to these as ML-enabled systems. Ensuring the quality of such systems is essential to understand and evaluate their results. However, this is not an easy task. For instance, a good learning system is one in which the learning evolves and improves over time, the model creation is reproducible and maintainable, the users are aware of how often the predictions are right and wrong, and the customer knows how much ML helps the system to achieve their goals.

Typically, when practitioners evaluate ML models they look only at measures such as accuracy, precision and recall. However, it is important to understand the big picture of the constraints these systems put on the overall development.

Where will the model be executed? What data will it have access to? How fast does it need to be? What is the business impact of a false positive? A false negative? How should the model be tuned to maximize business results? Moreover, the model is just a component of a system as a whole. In fact, there are other components that require attention such as the infrastructure to deploy, update, and serve the model, the integration of the model with the rest of the system functionality and a user interaction design to build better experiences of using ML-enabled systems.

Literature has shown that there are problems and challenges in the development of ML-enabled systems [2][26][33][36]. Requirements engineering (RE) is no stranger to this. The impact of incomplete/hidden requirements on overall system development is considered one of the most critical problems of RE in practice [27], and the difficulty to specify complete requirements increases for ML-enabled systems, where requirements engineers are typically not aware of the concerns that should be considered for such specification. Some recent research papers have drawn the attention of researchers and practitioners on the fact that ML can benefit from RE [9][11][14][15]. However, current research on the intersection between RE and ML mainly focuses on using ML techniques to support RE activities rather than on extending existing requirements techniques or providing new ones to support ML-enabled systems. In line with this, the roadmap for the future of SE proposed by the Carnegie Mellon University Software Engineering Institute [6] emphasizes that existing RE methods will need to be expanded to decouple ML problem and model specification from the system specification. Indeed, recent literature reviews [1][34] show that topics such as identifying quality attributes, specifying them, and understanding how they can be analyzed are not well-established and researched in the context of ML.

RE can improve the development of ML-enabled systems in several ways. For example, by identifying quality metrics beyond accuracy, dealing better with human factors and understanding why models do not fit and for whom they do not fit. With the aim at addressing one of the main problems presented in current RE for ML research, in this work we propose a catalogue of 45 concerns that can be used by requirements engineers to support the specification of ML-enabled systems. Based on our recent systematic mapping study [34], on insights from industrial experiences [17][21] we propose a catalogue organized into five different perspectives: objectives, user experience, infrastructure, model, and data. To validate our catalogue of concerns we conducted a focus group session with eight software professionals experienced with developing ML-enabled systems. The results revealed that the professionals were not explicitly aware of many of the concerns, but that they recognized their relevance and potential impact on the overall system being developed. In general, they agreed with the concerns and the way of grouping them into perspectives. In addition, we received relevant feedback that we used to improve our catalogue.

The remainder of this paper is organized as follows. Section 2 provides the background and presents related work. Section 3 describes the research method-

ology. Section 4 presents our catalogue of concerns for specifying ML-enabled systems. Section 5 presents the focus group session we conducted and how it contributed to our catalogue. We discuss and conclude our work in Section 6 and Section 7, respectively.

## 2 Background and Related Work

This section provides a background on RE for ML-enabled systems, presenting particularities that make RE essential in this context. Related work is also presented.

### 2.1 RE for ML-enabled Systems

ML is the study of computer algorithms that explores data to determine the best way to combine the information contained in the representation (training data) into a model that generalizes to data it has not already seen [28]. This type of systems, unlike traditional software systems, base its behavior on data from the external world instead of explicitly programming hard rules. In other words, data, to some extent, replace code. However, data may not be adequate and lead to bad outcomes. The output of a model is a prediction, sometimes surprisingly accurate and sometimes surprisingly inaccurate. This supposes a change in the way of designing, developing and testing ML-enabled systems.

RE and ML have a special connection. We can see a machine-learned model as a specification based on training data since data is a learned description of how the model shall behave [25]. Nevertheless, most ML models lack requirements specifications since current RE practices are not well defined and organized for ML [18][24]. Assuring the quality of the specifications is crucial for project success since misunderstandings and defects in requirements documents can easily lead to design flaws and cause severe and costly problems. Good requirements should be, at a minimum, complete, consistent, correct, unambiguous and testable [32].

There are a considerable number of concerns when designing ML-enabled systems. Given their inherent nature, requirements for these systems must have a stronger focus on data requirements [7]. For instance, Vogelsang [35] based on interviews with data scientists, mentioned that requirements engineer has to identify and specify requirements regarding the quantity, quality, collection, formats and the ranges of data. This with the aim of analyzing the given dataset against requirements and business goals. Furthermore, it is well known that in ML projects, model performance measures must be specified. These measures, to some extent, depend on the quality and pre-processing of the data. However, the importance of specifying these measures in a way that customers can understand and analyze it to make decisions is often overlooked [33]. Another concern in requirements for ML is the definition of quality properties, also known as non-functional requirements (NFRs) [9] [16]. For instance, depending on the context of classification problems, the requirements engineer might have to specify ethical concerns, particularly if it is people who are classified, and define what

characteristics should not be used for classification [3]. Similarly, other quality properties such as explainability, transparency, modularity, testability, security and privacy should be clearly specified [10].

RE practices are not well established for ML-enabled systems. Some secondary studies have recently given attention to this [1][34]. For instance, while there is still limited research, the SE community agrees on the need of extending RE techniques for ML. Furthermore, the new proposed techniques have not yet been applied in practice [31]. On the other hand, there is a misconception about the new role of data scientists that often are taking of responsibilities such as defining and eliciting requirements by themselves [22]. Another particular challenge is the overconfidence of using ML [12]. Customers see ML as magic, in other words, “ML will solve everything”. It’s the job of RE to let stakeholders know the limitations and manage their expectations.

## 2.2 Related Work

The intersection of RE and ML has been studied in recent years by the RE community and discussed in renown SE conferences [11][14][15]. Several studies have investigated what properties or components should be considered by requirements engineers when designing ML-enabled systems. Here, we focus on an overview of work we consider directly related to our research.

Chuprina et al. [8] present an artefact-based RE approach for the development of data-centric systems. The proposal encompasses four layers: context, requirements, system, and data, where each contains a set of concerns. We consider this work strongly related to our research. However, we found some differences. Firstly, their scope concerns data-centric systems, while ours is specifically related to machine learning. Indeed, our intention is to be more specific, including more fine-grained concerns that can be easily considered by practitioners in the ML-enabled system context. In addition, we detail ML-related concerns that we faced in practice that were not considered as part of their proposal, such as concerns related to user experience and infrastructure, which in our context showed being important for the success of ML-enabled systems.

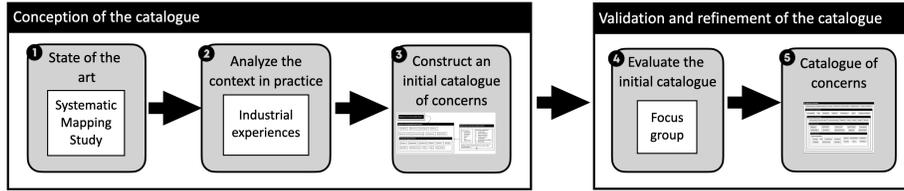
Nakamichi et al. [29] propose a requirements-driven model to determine the quality properties of ML systems. They list a set of issues related to ML. The authors cover perspectives such as environment/user, system/infrastructure, model, data and quality characteristics. In line with [8], the study is a great contribution in the field, however, we believe that our catalogue goes one step further than their described issues, by mapping additional perspectives and relevant concerns, for instance related to user experiences and ML objectives.

Another study we consider relevant is one conducted by Nalchigar [30]. They report on an empirical study that evaluates a conceptual modeling framework for ML solution development for the healthcare sector. It consists of three views consumed by business people, data scientists, and data engineers. We know that this contribution can help to address ML model concerns, however, we believe that other views such as infrastructure and user experience should be broken down into concerns to cover the full system perspective.

More recently, Berry [5] discussed some RE related ideas on how to use measures to evaluate AI related solutions from the point of view of recall and precision. Hence, he addresses what we identify as the model performance concern. However, practical experiences show many other relevant concerns that should be considered when designing ML-enabled systems (e.g., ethics, explainability, accountability).

### 3 Methodology for Defining the Concerns

We used the constructionism theory [13] that advocates a person needs to understand how something works before exploring the different ways to construct solutions. Figure 1 illustrates what we did and how we created, validated and improved our catalogue to support the specification of ML-enabled systems.



**Fig. 1.** Overview of the study steps for defining the catalogue

In order to understand what needs to be created, first, we conducted a literature review [34] on how ML could benefit from the RE perspective and what research opportunities could be addressed (step 1). The next step involved getting insights from industrial experiences (step 2). During the last two years, the first author has participated in R&D projects involving the development of ML-enabled systems as part of the ExACTa<sup>1</sup> initiative [21], which is co-coordinated by the second and the third authors. These projects involved several deliveries of solutions involving different types of machine learning problems and algorithms (e.g., decision trees, logistic regression, neural networks) to industrial partners. In experiential learning, this is defined as learning through reflection on doing. We took advantage of these planned synergies between theory and practice to learn more about the context in which the concerns operate. Additionally, we took advice from an industry-oriented publication based on more than a decade of experience building intelligent systems [17], with the aim at positioning the knowledge and insights acquired up to this point.

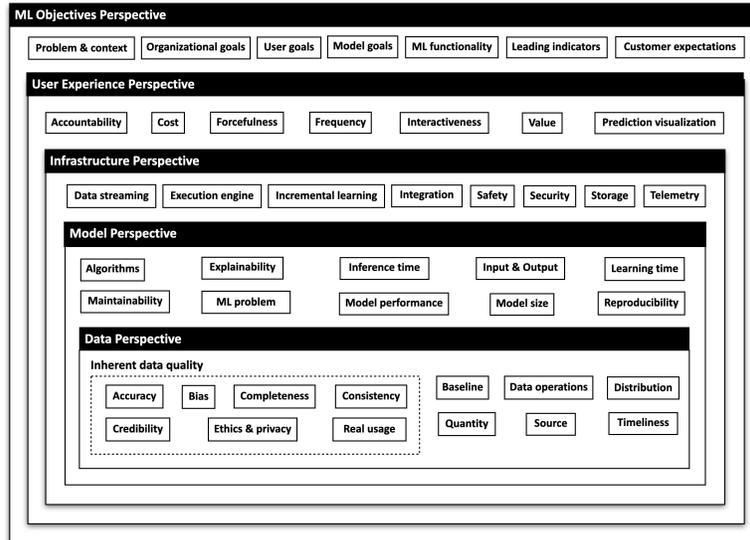
After analyzing the state of the art and the ML-enabled system development context in practice, we created the initial catalogue of concerns (step 3). It is noteworthy that the catalogue was critically reviewed by the second and the third authors, which are scientists active in the areas of software engineering

<sup>1</sup> <http://www.exacta.inf.puc-rio.br>

and data science and that have been exploring the intersection between these two areas. The literature review led us to focus on requirements definition, since this was one of the main identified research challenges, and revealed several quality properties of ML-enabled systems. Our industrial experiences allowed us to validate our findings and revealed complementary perspectives and concerns to be considered for ML-enabled systems. Finally, we conducted a focus group session (step 4) with eight software professionals with large experience developing ML-enabled systems. The results of the focus group allowed us to adjust the initial proposal (step 5).

## 4 Catalogue of Concerns for ML-Enabled Systems

The success of ML-enabled systems is related to taking care of not only models and data, but also business context, user experience and infrastructure. The specification of ML-enabled systems involves concerns that are often not easily identified, resulting in hidden requirements. For instance, it is clear that a model needs good data to be trained and then evaluated, but it is not clear what are the criteria that define the data as good, nor that defining the frequency and forcefulness of the model is important to get better user experiences, nor that the model needs to be integrated with other services to ingest new data. Therefore, herein we propose a catalogue (Figure 2) considering five complementary perspectives that accommodate the findings of our literature review and that showed being relevant in practice, building a big picture for specifying ML-enabled systems that has not previously been formalized.



**Fig. 2.** An overview of our catalogue: Perspectives and Concerns

The perspectives we identified and validated as relevant and complimentary are: objectives, user experience, infrastructure, model, and data. The depicted version already considers the adjustments made based on the focus group feedback. We believe that this catalog can be used by requirements engineers to support the specification of ML-enabled systems, making them aware of the big picture and helping to avoid incomplete/hidden requirements. We suggest the concerns to be analyzed by requirements engineers and discussed with stakeholders to understand the degree to which related requirements should be met. We describe the perspectives and detail their concerns hereafter, focusing on what the concerns mean in terms of RE specifications.

**ML objectives perspective:** Bridging the gap between the high-level goals and the detailed properties of ML models is one of the most common causes of failure to succeed [4]. Table 1 presents the concerns related to ML objectives that could influence the specification of ML-enabled systems.

**Table 1.** ML objective perspective

<b>Concern</b>	<b>Description (Addressing this concerns involves ... )</b>
Problem & Context	Specifying the problem that ML will address and its context before coding. ML must be targeted at the right problem.
Organizational goals	Specifying measurable benefits ML is expected to bring to the organization. E.g., increase the revenue in X%, increase the number of units sold in Y%, number of trees saved.
User goals	Specifying what the users want to achieve by using ML. E.g., for recommendation systems this could involve helping users finding content they will enjoy.
Model goals	Specifying metrics and acceptable measures the model should achieve (e.g., for classification problems this could involve accuracy $\geq$ X%, precision $\geq$ Y%, recall $\geq$ Z%).
ML functionality	Specifying the ML results in terms of functionality that the model will provide (e.g., classify customers, predict probabilities).
Leading indicators	Specifying measures correlating with future success, from the business' perspective. This could include the users' affective states when using the ML-enabled system (e.g., customer sentiment and engagement).
Customer expectations	Specifying expectations of customers and end-user in terms of how the system should behave (e.g., how often they expect predictions to be right or wrong).

**User experience perspective:** Better ML includes building better experiences of using ML. Connecting the predictions with users is critical for achieving success. However, this is not trivial and requires deep analysis. For instance, one mistake in the way of interacting with the user, and all the work spent on data pre-processing and modeling may be wasted. The goal of this perspective is to create effective user experiences (UX) so that users can interact appropriately and understand what is going on with the predictions of the model. Table 2 shows the user experience concerns and a brief description of each one.

**Table 2.** User Experience Perspective.

Concern	Description (Addressing this concerns involves ... )
Accountability	Specifying who is responsible for unexpected model results or actions taken based on unexpected model results.
Cost	Specifying costs involved in executing the inferences and also the user impact of a wrong model prediction.
Forcefulness	Specifying how strongly the system forces the user to do what the model indicates they should (e.g., automatic or assisted actions).
Frequency	Specifying how often the system interacts with users (e.g., interact whenever the user asks for it or whenever the system thinks the user will respond).
Interactiveness	Specifying what interactions the users will have with the ML-enabled system, (e.g., to provide new data for learning, or human-in-the-loop systems where models require human interaction).
Value	Specifying the added value as perceived by users from the predictions to their work.
Prediction visualization	Specifying how the ML outcomes will be presented so that users can understand them (e.g., specifying dashboard and visualization prototypes for validation).

**Infrastructure perspective:** ML models need to be integrated with other services. This includes components such as ingesting and learning from new data. The adoption of ML is growing, but proper implementations are needed to fulfill their promise. Hence, software engineers play an important role to orchestrate these ML components. Table 3 presents the infrastructure concerns.

**Table 3.** Infrastructure Perspective.

Concern	Description (Addressing this concerns involves ... )
Data streaming	Specifying what data steaming strategy will be used (e.g., real time data transportation or in batches).
Execution engine	Specifying how the model of the ML-enabled system will be executed and consumed (e.g., client-side, back-end, cloud-based, web service end-point).
Incremental learning	Specifying the need for ML-enabled system abilities to continuously learn from new data, extending the existing model's knowledge.
Integration	Specifying the integration that the model will have with the rest of the system functionality.
Safety	Specifying how the system deals with risks to prevent dangerous failures. Critical systems that incorporate ML should analyze the probability of the occurrence of harm and its severity.
Security	Specifying how the system deals with security issues (e.g., vulnerabilities) to protect the data. ML systems often contain sensitive data that should be protected.
Storage	Specifying where the ML artifacts (e.g., models, data, scripts) will be stored.
Telemetry	Specifying what ML-enabled system data needs to be collected. Telemetry involves collecting data such as clicks on particular buttons and could involve other usage and performance monitoring data.

**Model perspective:** Building a model implies not only training an algorithm with data to predict or classify well some phenomenon. Many other aspects determine its success. The aim of this perspective is to provide a set of model concerns a requirements engineer should analyze. Table 4 presents these concerns.

**Table 4.** Model Perspective.

Concern	Description (Addressing this concerns involves ... )
Algorithms	Specifying the set of algorithms that could be used/investigated, based on the ML problem and other concerns to be considered (e.g., constraints regarding explainability or model performance, for instance, can limit the solution options).
Explainability	Specifying the need to understand reasons of the model inferences. The model might need to be able to summarize the reasons of its decisions. Other related concerns, such as transparency and interpretability, may apply.
Inference time	Specifying the acceptable time to execute the model and return the predictions.
Input & Output	Specifying the expected inputs (features) and outcomes of the model. Of course, the set of meaningful inputs can be refined/improved during pre-processing activities, such as feature selection.
Learning time	Specifying the acceptable time to train the model.
Maintainability	Specifying the need for preparing the model to go through changes with reasonable effort (e.g., refactoring, documentation, automated redeployment).
ML problem type	Specifying the problem type tackled by the ML algorithm (e.g., classification, regression, clustering, extract information from text).
Model performance	Specifying the metrics used to evaluate the model (e.g., precision, recall, F1-score, mean square error) and measurable performance expectations.
Model size	Specifying the size of the model in terms of storage and its complexity (e.g., for decision trees there might be needs for pruning).
Reproducibility	Specifying the need for replicating the model creation process and its experiments.

**Data perspective:** Data is essential for ML-enabled systems. Poor data will result in inaccurate predictions, which is referred to in the ML context as “garbage in, garbage out”. Hence, ML requires high-quality input data. From the viewpoint of RE, it is clear that data constitutes a new type of requirements [7], [35]. Based on the Data Quality model defined in the standard ISO/IEC 25012 [20] and our own experience, we elaborate on the data perspective. Table 5 presents the data concerns and a brief description of each one.

## 5 Focus Group

### 5.1 Research Questions

Based on our goal, we defined the following research question: Is our catalogue promising and could it support the requirements specification of ML-enabled systems? To answer this question, we evaluate this work from three angles. First, the perception of importance to know if the catalogue of concerns is addressing a relevant problem. Second, the perception of quality to establish if the catalogue of concerns is complete, consistent, correct and unambiguous, and third, the perception of feasibility to have an idea to what extent the catalogue of concerns can be applied in practice. For this purpose, we designed a focus group session for promoting in-depth discussion about the catalogue and its suitability. Focus group is a qualitative research method based on gathering data through the conduction of group interviews and it has been conducted in SE for revealing arguments and feedback from practitioners [23].

**Table 5.** Data Perspective.

Concern	Description (Addressing this concerns involves ... )
Accuracy	Specifying the need to get correct data.
Baseline	Specifying the need for a baseline dataset approved by a domain expert that reflects the problem. It is employed to monitor other data acquired afterwards.
Bias	Specifying the need to get data fair samples and representative distributions. Eventually this concern may also apply to the model perspective.
Completeness	Specifying the need to get data containing sufficient observations of all situations where the model will operate.
Consistency	Specifying the need to get consistent data in a specific context.
Credibility	Specifying the need to get true data that is believable and understandable by users.
Data operations	Specifying what operations must be applied on the data (e.g., data cleaning and labeling).
Distribution mismatches	Specifying expected data distributions and how data will be split into training and testing data.
Ethics and privacy	Specifying the need to get data to prevent adversely impacting society (e.g., listing potential adverse impacts to be avoided). Eventually this concern may also apply to the model perspective.
Quantity	Specifying the expected amount of data according to the type of the problem and the complexity of the algorithm.
Real usage	Specifying the need to get real data representing the real problem.
Source	Specifying from where the data will be obtained.
Timeliness	Specifying the time between when data is expected and when it is readily available for use.

## 5.2 The Participants

We invited eight practitioners who have been actively working with the development of ML-enabled systems in industry. Before conducting the focus group session, we applied a characterization form. We asked them about the role they perform within their company, and their experience in years and number of ML projects they participated in. Table 6 shows an overview of the participants.

**Table 6.** Overview of the participants

Id	Role	# years	# ML projects
P1	Data scientist	13	12
P2	Data scientist	9	7
P3	Data scientist	1	3
P4	Developer	1	1
P5	Developer	3	2
P6	Project lead	1	2
P7	Project lead	2	5
P8	Project lead	2	2

## 5.3 Execution

Before starting the focus group, we introduced to the participants the main challenges when engineering ML-enabled systems and presented how RE may address some of them. The focus group was conducted in a Zoom meeting recorded for the study. The study was planned to be executed in two phases. In the first phase that took 20 minutes,

we explained our catalogue by decomposing each requirement perspective in its related concerns. The second phase, that took 45 minutes, was a question and answer session about the importance, quality and feasibility of our catalogue of concerns. The focus group was moderated by the first and second author and was recorded and transcribed.

#### 5.4 Results

**Perception of Importance:** We asked the participants how they define, document and organize requirements for ML-enabled systems and if they think it is important. P7 stressed the lack of formal methods to support their definition, modeling and documentation: “I have constant difficulties to find tools and methods to help my team and customers understand ML requirements”. On the other hand, P2 stated: “In my opinion, the requirements process for ML is ad-hoc, which makes it highly dependent on people’s knowledge” and P3 manifested: “I noticed that requirements have constant rework in ML projects”. Considering the overall discussion, we understood that creating new methods in this direction is absolutely important to address the problems practitioners are facing.

**Perception of Quality:** The participants evaluated our catalogue by (i) analyzing the concerns and perspectives in terms of completeness, consistency, correctness and ambiguity, and by (ii) measuring the capacity of our catalogue to support the specification of ML-enabled systems. Overall, there was a clear consensus that practitioners are unaware of the big picture. They did not know about many of the concerns, while judging them as relevant and helpful to support more precise specifications. P1 mentioned that “I wish I had such concerns specified upfront in my ML projects, decisions regarding these concerns should not be taken without appropriately involving stakeholders or when coding”. When analyzing the perspectives, P2 emphasized the importance of the ML objective perspective: “I understand the need to consider data, model, user and infrastructure, but in my opinion, the functional behaviour of ML models, which is reflected in the objectives, is crucial”. Regarding the concerns, P5 stated that from the technical view, the concerns and their grouping make sense: “When seeing the concerns I was able to relate them to problems and tasks I faced in the past”. P1 manifested the importance to evaluate in depth the completeness of the concerns. For instance: “In the data perspective, a common concern is the definition of a baseline that helps in the acquisition of new data. I would definitely consider it”. We improved the catalog based on the session feedback.

**Perception of Feasibility:** The participants found the catalogue of concerns useful to support the requirements specification of ML-enabled systems. P6 stressed the concerns organization into perspectives: “I think the catalogue can help us to analyze ML requirements since it covers several perspectives for different situations”. In addition, P8 stated: “We need to use this type of proposals in practice due to the overview that it provides and the concerns that may apply in our context”. Hence, we understood that it is feasible to further evaluate our catalogue practical contexts.

## 6 Discussion

We are aware that not every ML-enabled system needs to address all the concerns we proposed and not every ML-enabled system needs to implement them to the same degree. Our intention is to provide an overview of concerns so that requirements engineers can analyze the needs of their ML-enabled systems with stakeholders.

It is noteworthy that this overview focuses on concerns related to ML-enabled part of ML-enabled systems and the integration of this part with the remainder of the system. However, when considering the overall system, general quality characteristics of software products such the ones mentioned in the ISO/IEC 25010 standard [19], should also be analyzed.

The main contribution of our work is to provide a set of concerns grouped into perspectives that can be used by requirements engineers to support the specification of ML-enabled systems. Nevertheless, we believe our work may eventually be useful in various situations. First, to validate an already specified system. In this case, our concerns would be a reference since they come from a literature review and different industrial experiences on building ML-enabled systems. Second, our work helps to understand ML modularity since it provides several components and details at different levels about functional and non-functional aspects. Third, it is applicable to the most common ML approaches. Our work is focused on supervised and unsupervised ML problems. Both of them learn from data. Therefore, ML-enabled system would benefit from, at least, analyzing the perspectives and concerns we proposed.

## 7 Concluding Remarks

The development of ML-enabled systems involves understanding business context and problems, translating them into ML tasks, designing and experimenting with algorithms, evaluating models, designing pipelines, among other tasks. This needs to be considered from early stages of ML software development. Based on the literature and on practical experiences, we proposed a catalogue of 45 concerns to support the specification of ML-enabled systems that covers five perspectives: objectives, user experience, infrastructure, model, and data. This is the first effort aiming at providing the big picture of concerns for specifying ML-enabled systems. With this catalogue, we seek to empower requirements engineers with an ML overview of concerns that should be analyzed together with business owners, data scientists, software engineers and designers.

We evaluated the catalogue by conducting a focus group session with eight ML practitioners involved in developing ML-enabled systems. The purpose was to gain insights about the relevance of the problem we are addressing, the benefits of using it and its feasibility. The results indicated that practitioners consider the identified concerns relevant and the catalogue useful. They stated that grouping perspectives and organizing concerns can help them to identify constraints upfront with other practitioners. Therefore, we believe that the conceptual perspectives and concerns we herein proposed can be helpful to support specifying ML-enabled systems. Future work includes conducting additional evaluations in industry settings, including case studies.

## References

1. Ahmad, K., Bano, M., Abdelrazek, M., Arora, C., Grundy, J.: What's up with requirements engineering for artificial intelligence systems? In: 2021 IEEE 29th International Requirements Engineering Conference (RE). pp. 1–12 (2021)
2. Arpteg, A., Brinne, B., Crnkovic-Friis, L., Bosch, J.: Software engineering challenges of deep learning. In: 2018 44th Euromicro Conference on Software Engineering and Advanced Applications (SEAA). pp. 50–59 (2018)
3. Aydemir, F.B., Dalpiaz, F.: A roadmap for ethics-aware software engineering. In: International Workshop on Software Fairness (FairWare). pp. 15–21 (2018)
4. Barash, G., Farchi, E., Jayaraman, I., Raz, O., Tzoref-Brill, R., Zalmanovici, M.: Bridging the gap between ml solutions and their business requirements using feature interactions. In: ACM Symposium on the Foundations of Software Engineering (FSE). pp. 1048–1058 (2019)
5. Berry, D.M.: Requirements engineering for artificial intelligence: What is a requirements specification for an artificial intelligence? In: International Working Conference on RE: Foundation for Software Quality. pp. 19–25 (2022)
6. Carleton, A.D., Harper, E., Robert, J.E., Klein, M.H., De Niz, D., Desautels, E., Goodenough, J.B., Holland, C., Ozkaya, I., Schmidt, D., et al.: Architecting the future of software engineering: A national agenda for software engineering research and development. Tech. rep., Carnegie-Mellon Univ Pittsburgh PA (2021)
7. Challa, H., Niu, N., Johnson, R.: Faulty requirements made valuable: on the role of data quality in deep learning. In: International Workshop on Artificial Intelligence for Requirements Engineering (AIRE). pp. 61–69 (2020)
8. Chuprina, T., Mendez, D., Wnuk, K.: Towards artefact-based requirements engineering for data-centric systems. arXiv preprint arXiv:2103.05233 (2021)
9. Cysneiros, L.M., do Prado Leite, J.C.S.: Non-functional requirements orienting the development of socially responsible software. In: Enterprise, Business-Process and Information Systems Modeling, pp. 335–342 (2020)
10. Cysneiros, L.M., Raffi, M., do Prado Leite, J.C.S.: Software transparency as a key requirement for self-driving cars. In: International Requirements Engineering Conference (RE). pp. 382–387 (2018)
11. Dalpiaz, F., Niu, N.: Requirements engineering in the days of artificial intelligence. *IEEE Software* **37**(4), 7–10 (2020)
12. DiMatteo, J., Berry, D.M., Czarnecki, K.: Requirements for monitoring inattention of the responsible human in an autonomous vehicle: The recall and precision tradeoff. In: REFSQ Workshops (2020)
13. Fosnot, C.T.: Constructivism: Theory, perspectives, and practice. Teachers College Press (2013)
14. Habibullah, K.M., Horkoff, J.: Non-functional requirements for machine learning: Understanding current use and challenges in industry. In: 2021 IEEE 29th International Requirements Engineering Conference. pp. 13–23 (2021)
15. Heyn, H.M., Knauss, E., Muhammad, A.P., Eriksson, O., Linder, J., Subbiah, P., Pradhan, S.K., Tungal, S.: Requirement engineering challenges for ai-intense systems development. In: 1st Workshop on AI Engineering – Software Engineering for AI (WAIN2021) (2021)
16. Horkoff, J.: Non-functional requirements for ml: Challenges and new directions. In: International Requirements Engineering Conference. pp. 386–391 (2019)
17. Hulten, G.: Building Intelligent Systems. Springer (2019)

18. Ishikawa, F., Yoshioka, N.: How do engineers perceive difficulties in engineering of ml systems?-questionnaire survey. In: 6th International Workshop on Software Engineering Research and Industrial Practice (SER&IP). pp. 2–9 (2019)
19. ISO/IEC: Iso/iec 25010:systems and software quality requirements and evaluation (square) — system and software quality models (2011)
20. ISO/IEC: Iso/iec 25012: Software engineering – software product quality requirements and evaluation (square) – data quality model (2012)
21. Kalinowski, M., Lopes, H., Teixeira, A.F., da Silva Cardoso, G., Kuramoto, A., Itagyba, B., Batista, S.T., Pereira, J.A., Silva, T., Warrak, J.A., et al.: Lean r&d: An agile research and development approach for digital transformation. In: Product-Focused Software Process Improvement (PROFES). pp. 106–124 (2020)
22. Kim, M., Zimmermann, T., DeLine, R., Begel, A.: Data scientists in software teams: State of the art and challenges. *IEEE Trans. on Soft. Eng.* **44**(11), 1024–1038 (2017)
23. Kontio, J., Lehtola, L., Bragge, J.: Using the focus group method in software engineering: obtaining practitioner and user experiences. In: International Symposium on Empirical Software Engineering, 2004. ISESE'04. pp. 271–280 (2004)
24. Kuwajima, H., Yasuoka, H., Nakae, T.: Engineering problems in machine learning systems. *Machine Learning* **109**(5), 1103–1126 (2020)
25. Kästner, C.: Machine learning is requirements engineering—on the role of bugs, verification, and validation. Medium post, Accessed January **25** (2022)
26. Lwakatare, L.E., Raj, A., Bosch, J., Olsson, H.H., Crnkovic, I.: A taxonomy of se challenges for ml systems: An empirical investigation. In: International Conference on Agile Software Development. pp. 227–243 (2019)
27. Mendez, D., Wagner, S., Kalinowski, M., Felderer, M., Mafra, P., Vetrò, A., Conte, T., Christiansson, M.T., Greer, D., Lassenius, C., et al.: Naming the pain in requirements engineering. *Empirical software engineering* **22**(5), 2298–2338 (2017)
28. Mitchell, T.M., et al.: Machine learning. 1997. Burr Ridge, IL: McGraw Hill **45**(37), 870–877 (1997)
29. Nakamichi, K., Ohashi, K., Namba, I., Yamamoto, R., Aoyama, M., Joeckel, L., Siebert, J., Heidrich, J.: Requirements-driven method to determine quality characteristics and measurements for machine learning software and its evaluation. In: International Requirements Engineering Conference (RE). pp. 260–270 (2020)
30. Nalchigar, S., Yu, E., Keshavjee, K.: Modeling machine learning requirements from three perspectives: a case report from the healthcare domain. *Requirements Engineering* **26**(2), 237–254 (2021)
31. Shin, C., Rho, S., Lee, H., Rhee, W.: Data requirements for applying machine learning to energy disaggregation. *Energies* **12**(9), 1696 (2019)
32. Shull, F.J.: Developing techniques for using software documents: a series of empirical studies. University of Maryland (1998)
33. de Souza Nascimento, E., Ahmed, I., Oliveira, E., Palheta, M.P., Steinmacher, I., Conte, T.: Understanding development process of machine learning systems: Challenges and solutions. In: International Symposium on Empirical Software Engineering and Measurement (ESEM). pp. 1–6 (2019)
34. Villamizar, H., Escovedo, T., Kalinowski, M.: Requirements engineering for machine learning: A systematic mapping study. In: 2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA). pp. 29–36 (2021)
35. Vogelsang, A., Borg, M.: Requirements engineering for machine learning: Perspectives from data scientists. In: International Requirements Engineering Conference Workshops (REW). pp. 245–251 (2019)
36. Wan, Z., Xia, X., Lo, D., Murphy, G.C.: How does machine learning change software development practices? *IEEE Transactions on Software Engineering* (2019)