

A Requirements Engineering Framework for Human-centered Artificial Intelligence Software Systems

Khlood Ahmad^a, Mohamed Abdelrazek^a, Chetan Arora^b, Arbind Agrahari Baniya^a, Muneera Bano^c and John Grundy^b

^aDeakin University, Geelong, VIC, Australia

^bMonash University, Clayton, VIC, Australia

^cCSIRO's Data 61, Clayton, VIC, Australia

ARTICLE INFO

Keywords:

Requirements Engineering
Software Engineering
Artificial Intelligence
Machine Learning
Human-centered
Conceptual Modeling
Virtual Reality
Empirical Software Engineering


ABSTRACT

[Context] Artificial intelligence (AI) components used in building software solutions have substantially increased in recent years. However, many of these solutions focus on technical aspects and ignore critical human-centered aspects. [Objective] Including human-centered aspects during requirements engineering (RE) when building AI-based software can help achieve more responsible, unbiased, and inclusive AI-based software solutions. [Method] In this paper, we present a new framework developed based on human-centered AI guidelines and a user survey to aid in collecting requirements for human-centered AI-based software. We provide a catalog to elicit these requirements and a conceptual model to present them visually. [Results] The framework is applied to a case study to elicit and model requirements for enhancing the quality of 360° videos intended for virtual reality (VR) users. [Conclusion] We found that our proposed approach helped the project team fully understand the human-centered needs of the project to deliver. Furthermore, the framework helped to understand what requirements need to be captured at the initial stages against later stages in the engineering process of AI-based software.

AI-based software systems are rapidly becoming essential in many organizations [1]. However, the focus on the technical side of building artificial intelligence (AI)-based systems are most common, and many projects, more often than not, fail to address critical *human aspects* during the development phases [2, 3]. These include but are not limited to age, gender, ethnicity, socio-economic status, education, language, culture, emotions, personality, and many others [4]. Ignoring human-centered aspects in AI-based software tends to produce biased and non-inclusive outcomes [5]. Shneiderman [6] emphasizes the dangers of autonomy-first design in AI and the hidden biases that will follow. Misrepresenting human aspects in requirements for model selection and data used in training AI algorithms can lead to discriminative decision procedures even if the underlying computational processes were unbiased [7]. For example, a study by Carnegie Mellon revealed that women were far less likely to receive high-paying job ads from Google than men [8] due to the under-representation of people of color or women in high paying IT jobs.

Studies on human-centered design aim to develop systems that put human needs and values at the center of software development and clearly understand the context of the software system's usage [2, 9]. These human values include security, tradition, achievement, power, etc. [10]. Increasingly more research resources are being invested in developing human-centered AI solutions. Large organizations such as Google, Microsoft, and Apple are moving towards including human-centered values when building AI-based software [11, 12, 13]. These large organizations have now devised explicit guidelines for building human-centered AI systems. While these guidelines provide an excellent platform, they mainly focus on the design phase of software development and not on requirements engineering (RE).

Recent research efforts in RE have focused on including human-centered requirements in building software solutions, such as emotions [14, 15, 16, 17, 18, 19], gender [20, 21], power and politics [22], personality traits [23, 24], age, [25], and mental and physical challenges [26]. For example, Miller et al. [27] argued that "emotions should be considered as first-class citizens in software engineering methodology" and acknowledged the fact that software engineers usually overlook emotion when building software systems. Other studies have stressed that emotions are usually neglected when it comes to RE [28]. Perera et al. [29] emphasized the importance of including human values early on in RE.

 ahmadkhl@deakin.edu.au (K. Ahmad); mohamed.abdelrazek@deakin.edu.au (M. Abdelrazek); chetan.arora@monash.edu (C. Arora); aagraharibaniya@deakin.edu.au (A. Agrahari Baniya); muneera.bano@ (M. Bano); john.grundy@monash.edu (J. Grundy)
ORCID(s):

Although this is a growing research area in RE, there are limited studies that focus on including requirements for human-centered AI-based systems. We examined the studies presented in our systematic literature review (SLR) [30, 31] and a mapping study [32] on RE for AI (RE4AI) and found that most AI-based software lacks human-centered approaches when writing and modeling requirements, and existing research mainly focused on explainability, trust, and ethics with limited empirical evaluations.

We propose a new framework to specify and model requirements for human-centered AI-based software. The framework is based on industrial human-centered guidelines, an analysis of studies obtained from the literature, and a user survey. The survey identified the guidelines and human-centered aspects that should be addressed and prioritized during RE. Results from the survey were used as a baseline to build the RE for human-centered AI (RE4HCAI) framework. The framework is then evaluated in a case study, and requirements are elicited and modeled for a system that uses AI to enhance the quality of 360° videos for virtual reality (VR) users. The purpose of the AI-based software was to improve the quality of experience (QoE) and quality of service (QoS) for systems streaming and rendering 360° video contents. We discovered that some requirements could not be identified during the early stages of the project due to the black-box nature of the AI and the difficulty in explaining how it would respond to the available data. Therefore, a more iterative RE approach should be used to write requirements for AI-based software.

The key contributions of this research can be described as follows:

1. We present a new framework to help elicit and specify requirements for AI-based software. The framework is based on the industry's human-centered AI development, an SLR, and an expert survey. As part of our framework, we provide a catalog to aid in eliciting requirements for AI-based software. The framework further provides a conceptual model to present the requirements visually.
2. We apply and evaluate the framework on a case study that uses AI to enhance the spatial quality of 360° VR videos and report on our findings.

The rest of the paper is structured as following: Section 1 provides a brief background related work. Section 2 presents details of our framework. Section 3 reports on implementing the framework and modeling language on a software system to enhance 360° videos for VR users. Section 4 discusses key results and summarizes emerging theories. Section 5 addresses threats to validity. Section 6 discusses related work, and section 7 concludes.

1. Background and Motivation

The studies that focus on human-centered approaches are selected from an SLR [30, 31] that we conducted on RE4AI. These papers include studies on emotion for a human-centered social robot [33], understating expectations and limitations of AI-based software [34], identifying and mitigating human-centered issues related to fairness and biases [35], ethics [36, 37], explainability [38, 39, 40, 41], and trust [42]. Next, the industrial guidelines on human-centered AI development are analyzed, specifically Google's PAIR guidebook [11], Apple's human interface guidelines for building ML applications [13], Microsoft's eighteen guidelines for human-centered AI interaction [43] and the machine learning (ML) canvas [44]. We note that ML Canvas does not directly address the human-centered aspect. Nevertheless, we chose to include ML Canvas for its relevance to our work as the ML Canvas complemented the industrial guidelines by providing means for capturing relevant information for building ML models, e.g., ML business needs, decisions, data sources, and evaluation, and it facilitates collaboration between different stakeholders.

The gathered requirements from the SLR and the industrial guidelines are summarized, mapped, and categorized into six areas to build human-centered AI-based software. The categorization of six areas aligns with the five categories proposed in Google PAIR, and the sixth area identifies the human-centered approaches used when selecting and training an appropriate AI-based model [45]. Each area is explained in detail below as follows:

1.1. Area#1 Requirements for User Needs

The first area focuses on ensuring that the user needs are captured first when building AI-based software, documenting the limitations and capabilities, what the system can do, how well it can do its assigned tasks, and identifying the users and stakeholders. Other factors include how the user will interact with the system and what approaches should be selected to match the user's needs. Is the system going to be proactive or reactive? For example, in a proactive system, Google search provides a list of results based on the users' entered keywords. The next step is to determine if the user is aware of the AI feature. Visible features are when the user is aware of the AI component,

such as having auto-complete trying to guess what the user wants to input next. Invisible features are when the user is not aware of the AI feature. For example, a map would gather real-time data on traffic to provide the best route to users [11, 12, 13, 44].

Next we need to find if the system is going to automate or augment the user's needs? For instance, any task that does not require human oversight can be automated by AI-based software. Any task that the user enjoys or requires a human-in-the-loop can be augmented by AI to improve the user's experience and efficiency. Finally, evaluating the choice of reward function should be specified. How will the AI choose between right and wrong predictions? For instance, in case of binary classification, a learned AI model can make four different predictions, i.e., true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN). The two reward functions available are precision and recall. Precision is the proportion of correctly picked TP out of all the TP and FP. The system is more confident that the outcome is correct with higher precision. A flight booking app with high precision would predict the cheapest flights for a specific date but might miss out on some flights. In contrast, recall is the proportion of correctly picked TP out of all the TP and FN. Higher recall provides the confidence that the system has included all relevant results [11]. In a search system with high recall, all the results relevant to the search query are retrieved, potentially also including results that might not be relevant. Evaluating the reward function will depend on the specific task and domain of the AI application used [46]. The decision to go for either would be based on the trade-off from selecting precision vs recall. Therefore, when evaluating the reward function, a list of trade-offs should be documented to justify the selected function.

1.2. Area#2 Requirements for Model Needs

The second area includes documenting requirements for model needs. How do we choose an algorithm that will optimize the satisfaction of stakeholders' needs? Do we need a system that is explainable or accurate? For example, shin et al. [47] look for ways to reduce cost by experimenting with different algorithms and found that some algorithms can produce better outcomes with lower costs. However, these algorithms might lack in other aspects, such as explainability. In contrast, different algorithms might provide a better explanation but with predictions with lower confidence [48]. The choice of algorithm will also depend on the user's needs for the system. Thus as explained by Berry [46], as part of the requirements specification for AI, one needs to clearly specify the possible settings of each variable in a given AI algorithm, applied to a particular task in a given context.

A learned model (LM) as defined by Berry [46] "is the result of an instance of ML or deep learning (DL), whether the LM is taught, self-taught, or both with relevant real-world (RW) data". How will incoming data affect the model training? A dynamic system is trained online and improves as the user interacts with the system and often involves using feedback. A static system will improve with system updates and is trained offline. For example, an image recognition model that depends on data that does not change over time, will need updates with new releases, or else the LM might render extremely biased or useless after a while [13].

Model selection depends on the choice of algorithm and can include supervised, unsupervised, and re-enforcement learning. When training the system, we need to specify a threshold for when to avoid over-fitting or under-fitting training the data and select tools that will be used to evaluate the learned model. Also, deciding on how the learned model will improve should be specified. Does the system need a dynamic model that updates and trains online or a static system that improves only with updates? When tuning the learned model, the types of feedback, user behavior, and training data used in model tuning, and adjust the parameters accordingly should be considered [13, 44, 11].

1.3. Area#3 Requirements for Data Needs

The third area focuses on data needs and data collection methods. Data collection includes the type and amount of data needed. Once data collection methods are specified, data requirements, including quality of data used, security, privacy settings, and fairness, should be considered. Data quality includes five components: accuracy, completeness, consistency, credibility, and currentness. 1) Having correct data will ensure accuracy. 2) Completeness refers to the availability of all attributes and events associated with the data. 3) Consistency is having no contradictions in the data used. 4) Ensuring credibility by having truthful data. 5) And finally, currentness means that data must be collected within the correct time frame [49].

Data quantity focuses on data being diverse. And data quality addresses the completeness, consistency, and correctness of data being used [50]. Selecting data should include identifying features, labels, and sampling rates. Labels identify the features needed to train the ML model, such as labeling a scanned image of a tumor as malignant. Explicit labeling is done manually, and implicit labeling is when the model learns the pattern independently. Examples represent a row of data and contain features, and labels represent descriptions given to data. More samples in the dataset

ensure diversity but also increase costs [47]. In this case, a threshold should be identified to set the amount of data needed within the given budget.

Identifying and reporting biases in data must be addressed. Such biases can include automation, selection, group attribution, etc. Automation biases are when preferences are selected based on automated suggestions from the system. Selection bias usually happens when data is not collected randomly from the target population but rather selected based on the stakeholder's requests. Group attribution assumes that an output suitable for an individual will have the same impact on everyone in the group. Identifying key data characteristics should be set early on to avoid discrimination and biases [11, 43].

1.4. Area#4 Requirements for Feedback and User Control

This area deals with identifying and finding which kinds of user feedback need to be established in RE when building AI-based systems. Different types of feedback include implicit, explicit, and calibration. Implicit feedback provides information about the user's interaction, preferences, and behavior of the system. Examples of implicit feedback would include accepting or rejecting a recommendation, times of use, number of hours, when or how many times the user logged on, etc. Explicit feedback is provided by the user when requested by the system, and this might include surveys, forms, ratings, written feedback, likes, or dislikes. Calibration is the initial information the system might need from the user to function, such as scanning your fingerprint for the first time to activate touch ID.

It is essential to identify how and when the feedback will be used in model tuning and what changes it will have on the AI-based system. When asking for feedback, privacy measures should be considered to secure it. Also, give the user a choice to dismiss the feedback if they wish. Allowing the user to feel that they are in control over the system is an essential aspect of human-centered AI and can depend on a number of factors. Providing the user with control can be achieved by either giving them full control of the system in case of a failure or providing them with multiple options to choose from [11, 43, 13].

1.5. Area#5 Requirements for Explainability and Trust

Explaining AI-based systems reasoning can include providing explanations to both the end-user and other stakeholders. From the stakeholders' point of view, explainability requirements should include setting expectations of what the system can do and how it can do it and explaining the system's limitations and capabilities. On the other hand, explaining to the user will involve the data used and informing them about changes to the system that might happen with updates or learned machine improvements. Explaining predictions can be explained by either providing an example or displaying confidence. Confidence could be displayed in many ways depending on the situation and the AI-based system. Each situation should be evaluated to find an appropriate method to show confidence.

Explainable requirements can often conflict with others, such as performance and cost. It might be cheaper to build systems that are not very explainable with better performance [41]. Therefore, it is vital to calculate the trade-off when favoring explainable system or when explaining a system and identify how it might conflict with other requirements. Schoonderwoerd et al., [39] indicate that explanations should be provided based on a specific context, with the need to identify which explanations should be provided and when. However, in some situations, explainability might be important to ensure compliance with the European Union's General Data Protection Regulation (GDPR) [51, 52]. In this case, other measures would need to be considered.

Providing realistic expectations helps users and stakeholders to avoid over-trusting the AI-based software. Explainability requirements might occur during the building process or after the AI-based model is deployed. In [53] explainability requirements are divided into four components to include: Who the explanation is addressed to, what needs to be explained, when should the explanation happen, and who explains? Also, explanations should consist of consequences that might occur due to an action performed by the user. Using confidence can be a way to provide an explanation to users. However, it is essential to determine when and how to display predictions, as sometimes showing confidence could lead to mistrust [11].

1.6. Area#6 Requirements for Errors and Failure

None of the studies in the RE4AI literature mention the need to address errors in RE when building AI-based systems [31]. The focus of most industrial guidelines is on designing AI-based software with a human-centered approach and does not focus on RE. Our survey results [45] show that practitioners working on AI-based systems wanted to know how to deal with errors and specify error sources during RE. The different error types include background errors, context errors, and system limitations. Some errors are more difficult to identify and are invisible to

the end-user, such as background and context errors. Context errors happen due to several reasons and can be avoided by ensuring the user has a good understanding of how the system works and making sure the system is aligned with their needs. They are usually an outcome that is a true positive yet does not provide a prediction that is in line with the user's needs.

Error sources include system, incorrect predictions, data, input, and output errors. Data errors happen with mislabeled data and can be due to either poor training or inaccurate labeling. Prediction errors can occur when an incorrect model is used, the data is not comprehensive, or missing some critical elements. Input errors happen when users input unexpected data and can be due to user's old habits or having an abusive user. Output errors are when the system provides a prediction that is low in confidence or an irrelevant output but high in confidence. And finally, system errors would happen when multiple systems using AI integrate or depend on each other. We need to identify error types and sources in RE and mitigate them by providing an action plan for how they should be addressed and fixed [13, 11, 43].

2. A Framework to Manage RE for Human-centered AI-based systems (RE4HCAI)

This section presents the RE4HCAI framework (as illustrated in figure 1) for eliciting and modeling requirements for human-centred AI-based systems. The framework consists of three layers. The first layer in Section 2.1, presents the human-centered guidelines that should be included in RE. The guidelines are combined and mapped into a reference model. The second layer presents a catalog (provided in the appendix) to help elicit requirements for building human-centered AI-based software, this layer is discussed in Section 2.2. The last layer proposes a modeling language to present the requirements visually as explained in section 2.3.

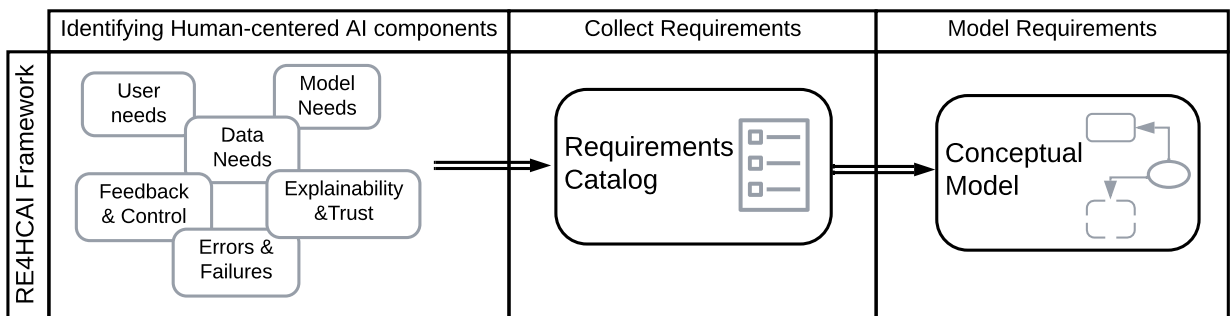


Figure 1: Our proposed framework to elicit and model requirements for human-centered AI

2.1. Human-centered AI Guidelines for RE

The coverage of these six areas in RE are investigated by conducting a user survey with 29 practitioners and researchers working on AI-related projects [45]. The participants included data scientists, machine learning specialists, and software engineers working with AI-based software. We asked each participant to identify which of the mapped human-centered guidelines are, or should be, included in RE based on their experience. The survey results (shown in figure 2) confirmed that all six human-centered areas need to be specified in RE4HCAI. The survey participants responded more favorably towards the inclusion of three areas, namely user needs, model needs and data needs, than the remaining three areas for RE4HCAI. Thus, our evaluation focuses on user needs, model needs and data needs in Section 3.

The six areas of the human-centered AI components that are needed to be considered in RE are mapped to present a reference model that showcases the overall layout of the framework, as shown in figure 3. Each Area presented in the reference model is explained further in the background Section 1.

2.2. Catalog for Collecting Requirements

In this section, a catalog of human-centered AI requirements is presented, along the six areas in our framework. The gathered human-centered requirements and the mapped requirements from the SLR are listed in a tabular format




Figure 2: Human-centered needs for AI and the number of times they were selected by the participants

to form our catalog. The catalog can be used as a checklist to elicit requirements for AI-based software and has six sections, and each section is dedicated to eliciting detailed requirements for each of our six areas explained in the background Section 1. The catalog is provided in the appendix.

2.3. Modeling Requirements for Human-centered AI

In the SLR, we sought any existing modeling language or requirements notations used in RE4AI. UML was the most popular method used, as it was easier for non-software engineers to work with and learn. However, UML has limitations as it does not support the modeling of non-functional requirements (NFR) and business rules. Goal-oriented requirements engineering (GORE) had better support for NFR and business rules but was more challenging to learn and was mainly used by requirements engineers and software engineers. Also, GORE could model requirements at lower levels of abstractions than UML. Silva et al. [54] explained that using GORE to model a requirement or concept could be presented with fewer structural diagrams than UML. However, UML is more widely known and used than GORE, also, GORE was reported to be more difficult to learn among non-software engineers [55, 56].

Looking at the different roles from the survey, we found that around 25% of the people involved in building AI-based software were software engineers and requirements engineers. Data scientists and ML specialists contributed to around 27% of the team building AI-based software. Other roles included system and business analysts, developers, and researchers, as shown in figure 4. Due to the diverse nature of team structure in building AI-based software and the

RE Framework for Human-centered AI-based Software Systems

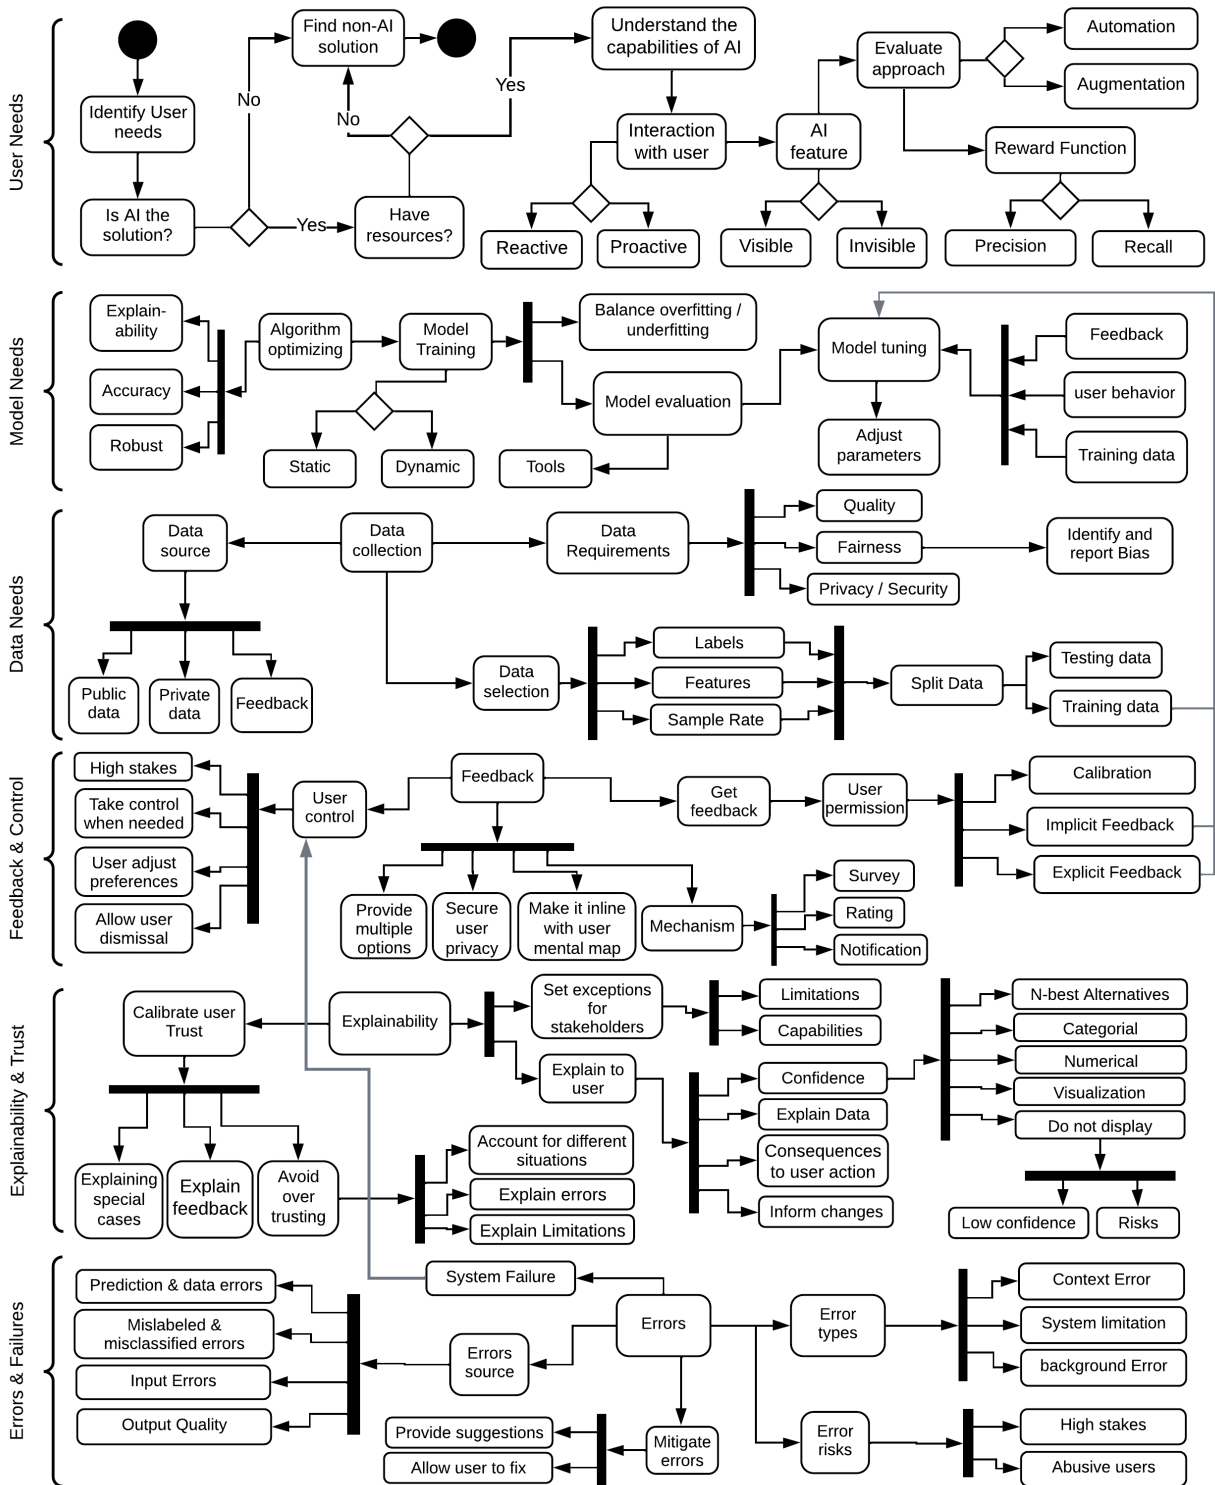


Figure 3: Diagram showcasing the reference model for RE RE4HCAI

need to use more specific modeling concepts to reflect on the process provided in the framework, we decided to create




Figure 4: The different roles of participant in the user survey

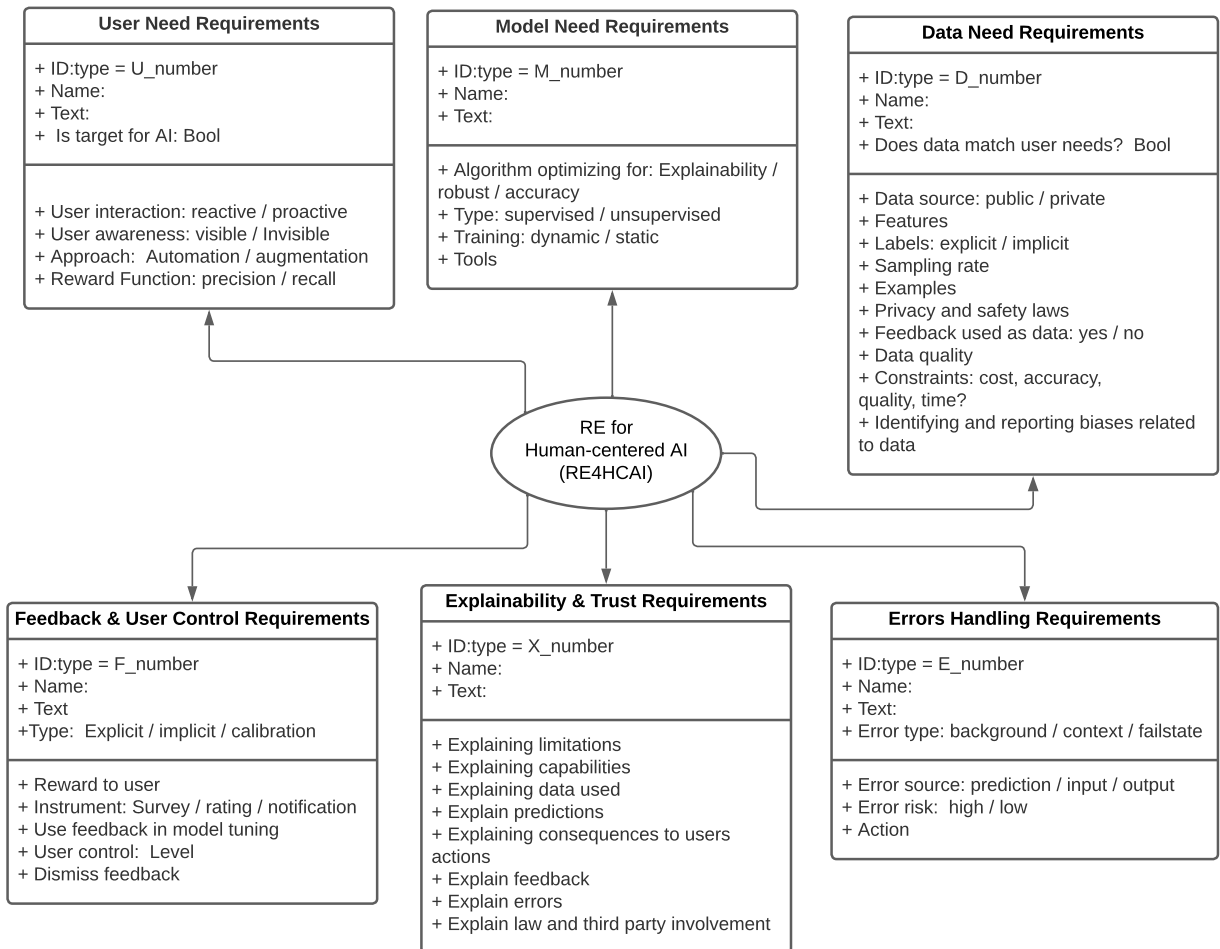




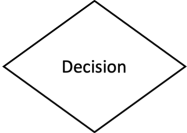

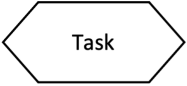



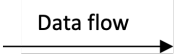
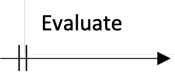


Figure 5: First level of the modeling framework showing the six different area's of the RE4HCAI

a conceptual model for RE4HCAI framework; as conceptual modeling allows for precise concepts to be modeled as well as presenting a holistic view of the application [57].

Our conceptual model consists two levels. The first level presented in figure 5 provides a holistic view of the system requirements. We use an oval shape notation to list the main goal that needs to be addressed (notations in table 1). For example, in the case study we conduct in Section 3 that goal would be enhancing the quality of 360°videos by four times. This goal would connect to the six sub-models to include area#1 user needs, area#2 model needs area#3 data

Table 1
Legend

Notation	Explanation	Notation	Explanation
	Displays a capability of the AI-based system (What the AI component of the system CAN do)		Displays a limitation of the AI-based system (What the AI component of the system can NOT do)
	Display the need for the system		Showing the trade-off between two choices made (e.g. selecting precision vs recall)
	Display a decision or choice to be made		Display the goals needed for the system to be built, goals can be functional requirements
	A process or task that needs to be achieved		An attribute or property that is needed for a process
	Shows the different users involved in building the systems or end-users		Displays databases or data sources used in building the system
	Direction the data flows in		Data flow that needs to be evaluated

needs, area#4 feedback and user control, area#5 explainability and trust, and area#6 errors and failure. For each area we use a UML class diagrams to present it. Each area contains the main attributes listed in the reference model in figure 3. The second level consist of a sub-model to show each area in further detail. We provide the notations to model these requirements in table 1. When building the visual notations for our language we attempt to comply with the nine principles of notations [58] to reduce the complexity and cognitive load for visual notations in SE modeling. We use different shapes, textures, icons, and colors to present our notations.

Goals are presented using an oval shape, and each need includes the components modeled in the sub-models. Since our framework focuses mostly on needs, we use a square with dash-lines to visually emphasize the concept of needs. Elongated hexagons are used to present processes or tasks. Trade-offs are displayed using an equilateral octagon. Limitations and capabilities are modeled using rectangles to show what the AI component can and cannot do with limitations shaded in yellow to differentiate between them visually. Icons were used to present attributes(folder icon), users (stick icon), and data sources (database icon). And finally, the components are joined using connectors to either show the flow of data or evaluate the data.

When presenting the models, we use blue tick icons on notations to represent the parts that are belong to the framework. The notations that do not include a blue tick icon are system specific, as show in figure 6. System requirements will change when modeling a different AI-based system.

3. Case Study Application and Evaluation

This section describes the process we used in designing, selecting, and conducting the case study in order to implement our proposed framework. Our case study answers following research questions (RQs):

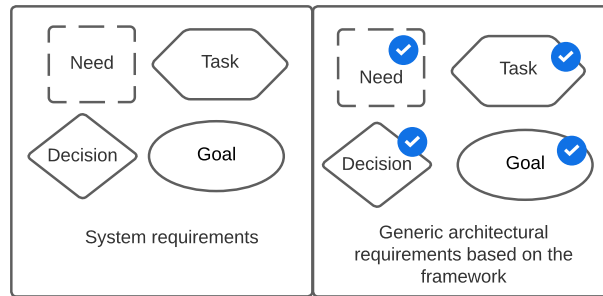


Figure 6: The blue icon used with modeling notations show the difference between architectural and system requirements

- **RQ1. How does our RE4HCAI framework fare for eliciting and modelling requirements for human-centered AI software?** RQ1 aims to assess the usefulness of our framework in eliciting and modelling requirements. It further contrasts the RE process before and after the implementation of our framework in the case study.
- **RQ2. How can the RE practices be aligned in the life cycle of AI projects?** RQ2 aims to provide an understanding and alignment of the RE practices in AI-based software development projects.

The case study is conducted following the guidelines presented in [59] and the steps provided in [60] to investigate how our framework could contribute to engineering AI-based systems with a human-centered perspective. The main reason for conducting this study was to find *how* to extract requirements for building human-centered AI-based software? And if implementing the framework presented benefited the process? The catalog is used to elicit requirements for the AI system, to identify user needs, model needs, and data needs as presented in tables 2, 3, and 4. These requirements are then modeled, as shown in Figs. 8, 9, and 10.

3.1. Study selection

When selecting the case studies we had the following selection criteria in mind:

- **Software project with a major AI component.** Since our framework focuses on RE4HCAI, we required a project that had a major AI component to be able to answer our RQs, and ideally a human-centric project that involved interaction with users at the centre of it.
- **Project Maturity.** The project should be at a reasonably mature stage of development, but ideally still under development. This criterion was motivated by our previous experience where experts tend to miss out on details in projects which have been already completed. Furthermore, in order to address RQ2, we required a project where the requirements had been established using traditional methods.
- **Availability and background of the expert.** We needed to select a project with experts that had the requisite knowledge on the AI component of the system so we could elicit and model the requirements for model and data needs. Also, the case study required substantial time commitment from the expert(s), and hence we sought projects where at least one expert agreed to the time commitment.

Our case study on VR-360° video enhancer matched all the criteria noted above. It builds on a DL model to enhance the quality of 360° videos for VR platforms, and hence has a major AI component - first criteria - and the video enhancement is done for the education and entertainment purposes of the users streaming VR videos. The project was at later stages of development, the requirements of the project had been finalised in the form of a software feature plan, and most of the features had been implemented already - second criteria. Furthermore, we had access to the main ML expert of the project, who developed the AI component and agreed on the time commitment - third criteria.

3.2. Data Collection

We collected data over four sessions with the ML expert. Of these four sessions, three were dedicated to extracting and modeling the requirements for the project and a fourth session to evaluate the modeled human-centered requirements. The ML expert who built the AI project, and ended up being the fourth author, had no prior background knowledge of how our framework or the modeling language worked. They were given a brief overview of our framework in the first session and further explanation on it in the fourth session. In the fourth session, we identified which of the requirements were needed before the start of the project. In other words, these were the requirements that ML expert deemed important for them to have known ideally at the start of the project. However, they were not known at the actual start of the project.

We note that we could elicit and model the requirements only for the first three areas of the framework collaboratively. This was the case due to the limited availability of the ML expert and the fact that the elicitation and modeling for each area (as noted below) requires major time commitment. While we would have ideally wanted to cover all six areas, however, we had a choice of covering a subset thoroughly or cover all areas superficially. We made the former choice, which is in line with our survey findings that the first three areas are the most focused areas by practitioners when eliciting requirements [45].

In the first session, the ML expert gave an overview of the project, and the expert was given a brief overview of the motivation of the study. The expert was further given a brief overview of the framework. The first session was dedicated to collaboratively first eliciting and then modeling the first area user needs. The first session was the longest and lasted ≈ 4 hours and was divided into two periods. In the first half, the requirements were elicited, and the second half consisted of modeling the requirements. The second session was dedicated to the second area model needs. The session took ≈ 2.5 hours and was done over two periods. Similar to the first session, the requirements were elicited in the first half, followed by modeling the requirements in the second half. In line with the first two sessions, the third session was dedicated to eliciting and modeling the third area data needs. The session took ≈ 3 hours over two periods. The fourth session took two hours, and we evaluated all three models as a whole. We instructed the ML expert that they were free to suggest any changes to the elicited and modeled requirements from the three sessions. The ML expert reflected on the comparison between the RE process followed originally in the project and the sessions using our framework and the holistic view of the system is presented in figure 7.

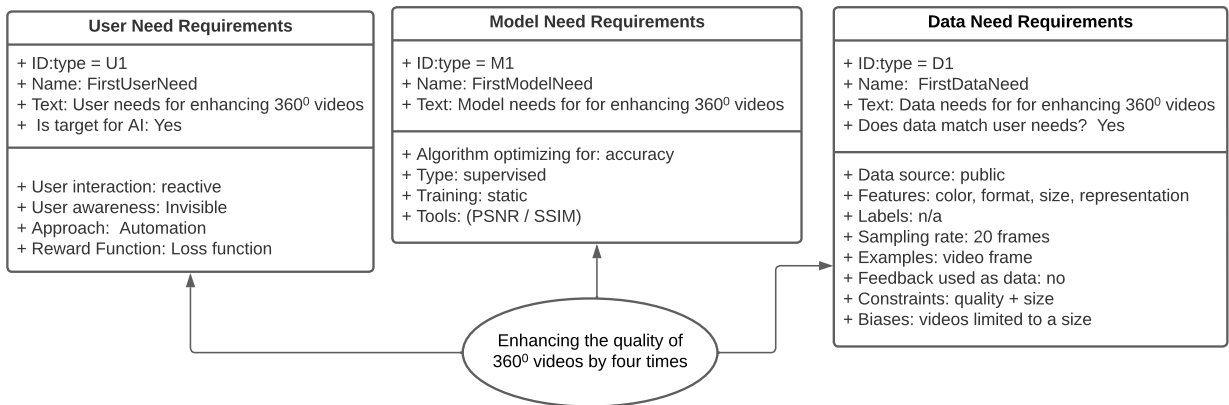


Figure 7: Model presenting holistic view of the requirements elicited for enhancing the quality of 360° video

3.3. Improving the Quality of 360° Videos

The case study was conducted on a project that uses deep learning (DL) to enhance the quality of 360° videos for VR platforms [61]. Current approaches to building 360° videos can capture and facilitate the immersive and interactive viewing experience. However, the quality of the final product can degrade due to the limitations of consumer-grade hardware, bandwidth streaming constraints, and processing that requires stitching videos from multiple cameras [62, 63, 64]. The proposed solution was to enhance the quality of the final product using AI-based software.

Table 2

Identified user needs requirements for the 360 video enhancer. * requirements that could be specified before testing the data on the selected model

Identify need for AI?	Who are the users	* VR users who could watch enhanced 360 videos offline
	Why do we need the system What is the system used for	* Improving the immersive experience for viewers * Entertainment and teaching
Systems capabilities	Limitations	* Limitations to the user: the need for VR equipment with high rendering capabilities Limitations to stakeholders: Hardware resources and processing time
	Capabilities	Improved quality of experience Improved quality of service
	How well can the system do what it does?	* It should improve the resolution of the video by four times
Interaction with user action	Proactive (User requests action) or Reactive (Interacts with the user without requesting)	* Reactive – as it will refine the quality of the video without the users having to ask
Is the user aware of the AI feature?	Visible or Invisible features	* Invisible feature, the user will not be aware of the AI component
Evaluate approach	Augmentation vs Automation	* Automation – not possible for a humans to be directly involved in super-resolving videos
	Reward function	* A loss function List potential pitfalls How do you provide inclusion
		optimise model to predict correct pixel values Loss of visual quality N/A

3.4. Requirements for User Needs

The first session was held to elicit requirements for identifying the user's needs. The session started by identifying the need for the proposed AI-based software using the catalog in the appendix. The need for the system was determined by the quality of the existing 360° videos. Most 360° videos had low resolution, and contain artifacts and noise, thus affecting the quality of the final product [65, 64]. The process of removing unwanted noise and artifacts was not achievable without the AI component. Therefore, establishing the need for AI-based software.

To identify the system's capabilities, we needed to understand what the expected system could and could not do. System capabilities covered both the end-user and stakeholders. From the user's perspective, the end product would improve the resolution four times while using the same hardware equipment. Therefore, improving the QoE in either entertainment or educational settings. The limitations affected both the end-user and the stakeholders. It was important to list these limitations to avoid high expectations from either side. The limitations for the end-user were that users had to own VR equipment with the capability to render high-quality 360° videos. The second limitation to the user was that they could watch these videos only offline. Therefore, online streaming would not be possible. Furthermore, the limitation for the stakeholder was towards building the system, which included hardware resources and time needed to process the videos. The stakeholders needed to establish ways to improve quality without increasing expenses, which will later link to model selection and collected data.

When modeling limitations and challenges, we found that using color helped distinguish between them visually, as shown in figure 8. We modeled the human-centered aspects as goals, and each goal would be mapped to either a need, process, capability, or limitation. For example, the system needed to be automated because it was not possible for a human to do the system's task. Processes involved actions that needed to be implemented, such as building a reactive system that doesn't require user interaction. Also, having different textures for the needs made it easier to identify them visually. The advantage of modeling the first area was that it was easy to recognize limitations, capabilities, and needs when building human-centered AI-based software.

3.5. Requirements for Model Needs

The second session involved collecting requirements for from our catalog, as shown in table 3. Area#2 helped understand the rationale behind the selection of a given AI algorithm, and what to optimize for? The **selected DL** model was optimized for accuracy as the AI model aimed to produce more accurate predictions of better-quality frames. Model selection heavily depended on the type of optimization opted to use. In this case, regression was used as it will always optimize for accuracy. In model needs, we noticed that most of the goals were connected to actions or decisions. For example, the goal involved identifying model tuning techniques involved tasks such as tracking the training losses

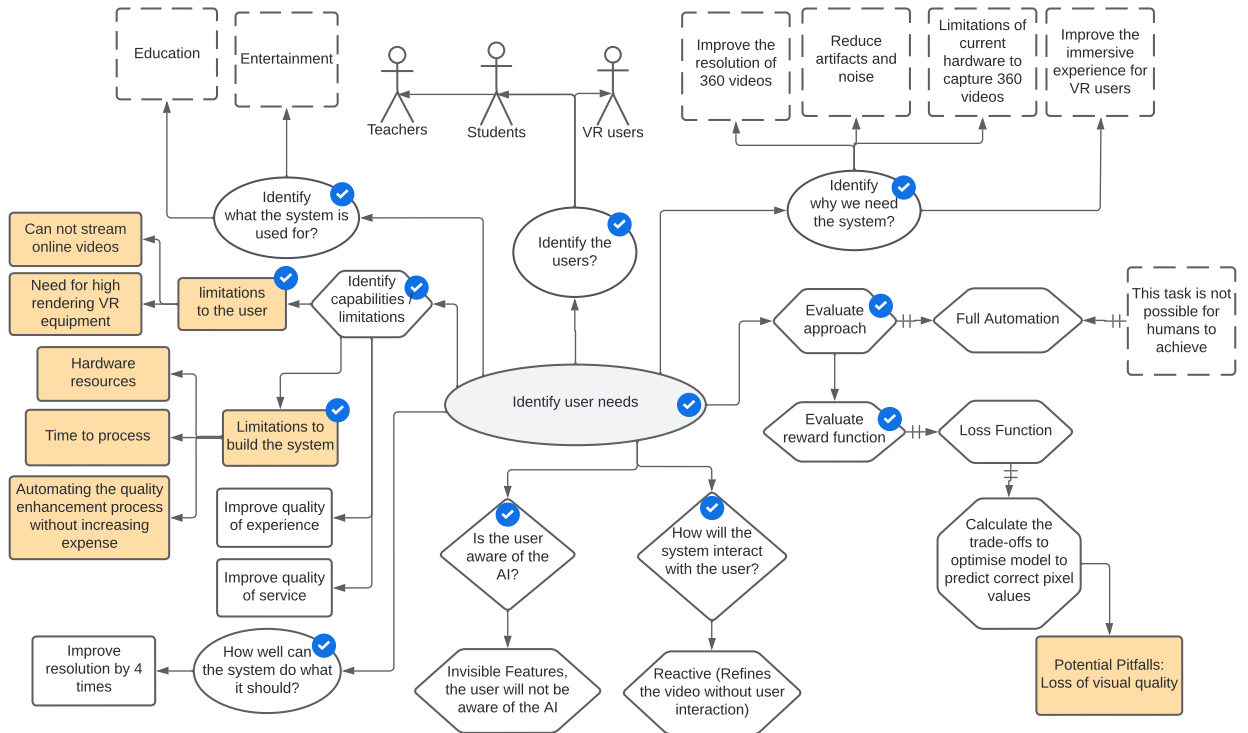


Figure 8: Model presenting the requirements for User Needs

by tuning the different attributes. In cases in which goals or tasks could not be achieved, such as tracing output errors to data, we modeled them as limitations.

Model tuning included adjusting the model to user feedback, user behavior, and output errors. For this particular situation, user feedback would be done with a ranking scale that is given after the user watches the enhanced video. The other form of feedback would be to perform a user study to evaluate the quality of the generated outcome from the model. The study will involve human subjects ranking the quality on a scale of 1 to 100. This will indicate if the human-centered perceptual quality of the predicted outcome is acceptable. If not, the plan would be to change the model training and optimize it to improve perceptual quality.

Trade-offs that needed to be calculated and modeled were connected to scaling the DL model. Scaling was achieved by changing either the width or depth [66]. Increasing the depth meant adding more convolution layers and allowing the model to learn more complex patterns. If the performance was not good enough, the depth would have to increase. However, the trade-off to increasing the depth is that more computational resources are required, resulting in delays in the output generation. On the contrary, changing the width would require that features extracted from the input become wider. This would require more memory allocation for data, as it would store more features for the given input at multiple layers. For the model used, an experiment would be set to identify when to modify width or depth and at what intervals while evaluating the trade-offs for time vs. costs to achieve the most effective results.

3.6. Requirements for Data Needs

The third session involved extracting requirements for data needs. While conducting this session, it was found that the data collected originally was not fit to be used in the training of the model. Therefore, the data had to be cleaned and adjusted to become accurate, correct, consistent, and current. Sudden changes in scene, motion, or light and removing rolling and credit information were removed to make the data accurate. Next, to ensure data completeness, a list of all attributes and events needed for the AI product was listed, and made sure each attribute had data to represent. Completeness was established in a technical way to show the diversity of the presented data to ensure that the dataset had diversity in terms of: 1- Motion, 2- Content (objects and scenes), and 3- Light. The initially collected data was

Table 3

Identified model needs requirements for the 360 video enhancer. * Requirements that could be specified before testing the data on the selected model

Optimize for?	Accuracy	* Compare the ground truth of data to the predictions
Choose the ML type	Supervised	* Regression
Model training	Static with off-line training	* Fixed dataset that the model learns from
Balance between overfitting & underfitting	Keep track of training losses	Pinpoint the iteration at which the model when overfitting starts
Model Tuning (Include parameter tuning and architecture changes)	User feedback Adjusting to user behavior Trace output errors to data Parameter tuning Quantitative feedback	Provide a ranking to evaluate the perceptual quality Change model training to optimize losses based on results from a user study to evaluate perceptual quality In case of output errors occur, data will remain constant and architectural/training changes will be made. Observe training losses and make changes to training iterations, learning rates, and data batch sizes based to validation results. Use image quality assessment metric to evaluate model's enhancement capability.
Specify scalability issues	Scaling the model	Change either the width or depth of the model to scale the DL model
Choose tools to use to evaluate the model	* Compute image quality assessment metric (PSNR / SSIM, WS-PSNR / WS-SSIM)	
Evaluate the quality of the model	This is evaluated based on the number of parameters, runtime and evaluation results	

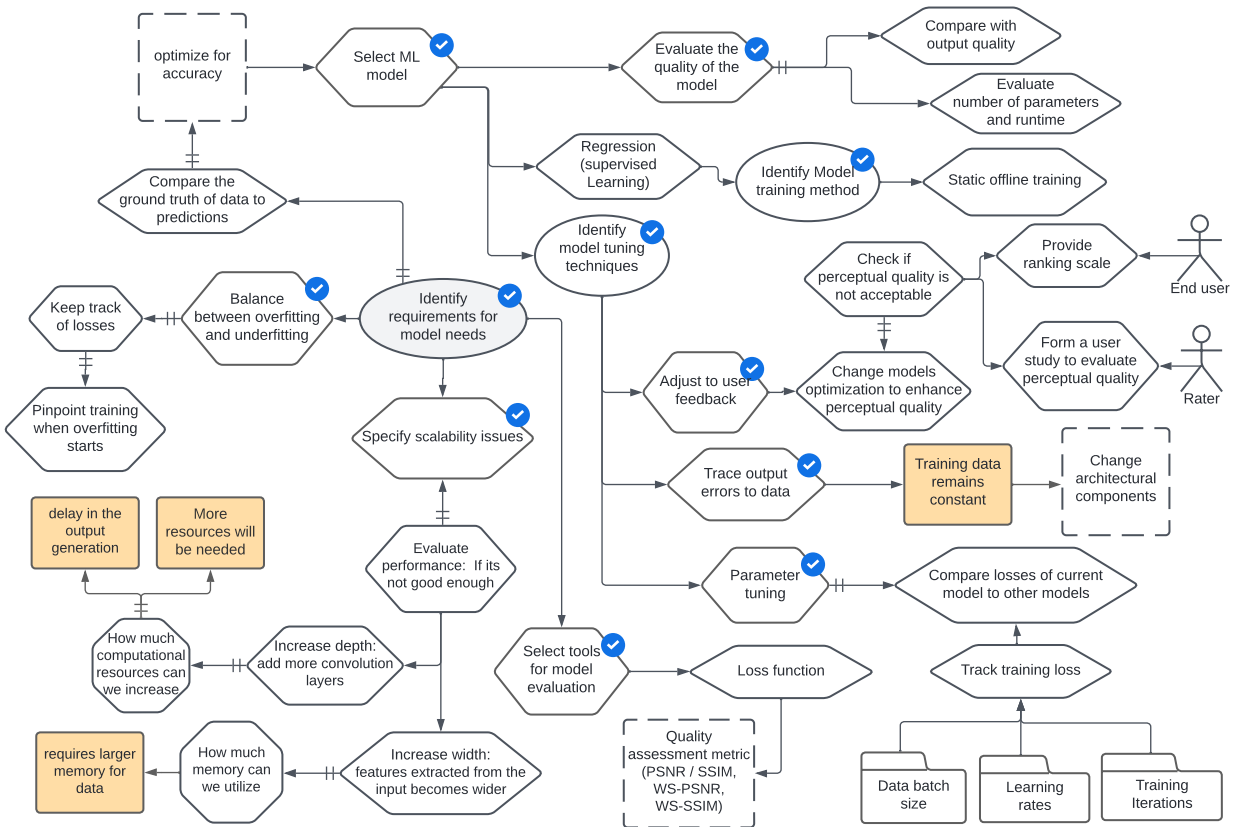


Figure 9: Model presenting the requirements for Model Needs

not consistent in quality. Thus, videos were processed to the same length and resolution to ensure consistency. Finally,

Table 4

Identified data needs requirements for the 360 video enhancer. * Requirements that could be specified before testing the data on the selected model

Data selection	Features	Color: RGB format: 360° Videos Representation: Equirectangular Projection size: 20 frames
	Labels	N/A
	Example	An actual video frame
	Sampling rate	Up-to 20 frames from a given video clip
Data collection	Identify type of data needed	* Diverse contents to include different objects, motions and lightning
	Amount of data needed	Initially the quantity and diversity was not enough so a conventional video dataset was used to initialise the model training
Constraints	Diversity of data	To ensure diversity a complexity analysis was conducted
	cost, accuracy, quality, time?	Quality: Resolution of 480x360 (hardware and memory limitations) Size: up to 20 frames (memory and model limitations)
Data source	Type of data source	* Open-source: from research labs, YouTube, open-source platforms
	Is data responsibly sourced?	* All sources of the data collected are given credit and citations
	Using feedback as data	* N/A
	Charges in obtaining data	* Time taken to gather and clean the data needed
Split Data	What measures are taken to ensure the data is up to date?	* Data was recently collected (less than a month ago)
	Training data (Model tuning)	90% randomly selected for training
	Testing data	10% out of 590 randomly selected for testing
Data quality	Accuracy	* Data is filtered to remove sudden changes in scene, motion, and light. Rolling and credit information is removed. Each video is divided into shots and static shots are removed.
	Completeness	Completeness was insured in a technical way to show the diversity that data presents
	Consistency	Data is processed to become the same length and resolution
	Credibility	N/A
	Currentness	* Creating video shots which represent unique single scene; thereby, ensuring that each clip/shot has current data belonging to continuous period of time
Data requirements	Protect personal information Does data comply with privacy and law	We are not collecting usage/user related information * Videos used for training and testing are not shared with the end users. The user will have their own choice of videos to use!
Fairness	Identify biases:	* Model might not perform equally for videos with larger motion or luminance changes.
	Missing features Under or over representative data	More diverse scenes and objects are needed Possibility of same scene/object represented across multiple videos

currentness was provided by creating video shots that represented a unique single scene. Therefore, ensuring that each clip/shot had current data belonging to a continuous period of time.

Data requirements presented a higher number of needs. The need to specify how the data will be selected and collected for use in model training and testing. For data selection, the needs consisted of setting features, labels, examples, and sampling rates. Features included color, size, location, etc. An essential aspect of data needs was to identify any biases that might happen due to the data used. The proposed system was designed for videos with insignificant motion, content, and light changes across time. However, if the videos had larger motion or luminance changes, the model might not perform equally. Moreover, the diversity of the scenes and objects found were limited to what was available in the open-sourced platforms creating the possibility of the same scene/object represented across multiple videos. Thus, more diverse scenes would be needed to avoid under or over-representing data.

Data collection involved gathering diverse content, including different objects such as trees, buildings, humans, animals, etc., and various motions such as how objects move across the video and lighting situations such as day, night, bright light, and dim. These measures were identified before data collection and obtained from literature based on the International Telecommunication Union's standard guidelines for spatio-temporal complexity measures [67]. Although measures were taken into account, some constraints were found after training the model with the collected dataset. The constraints were related to image quality and size due to hardware and model memory-related limitations. These constraints were modeled as limitations to show that data was limited to the specified size and frame rate.

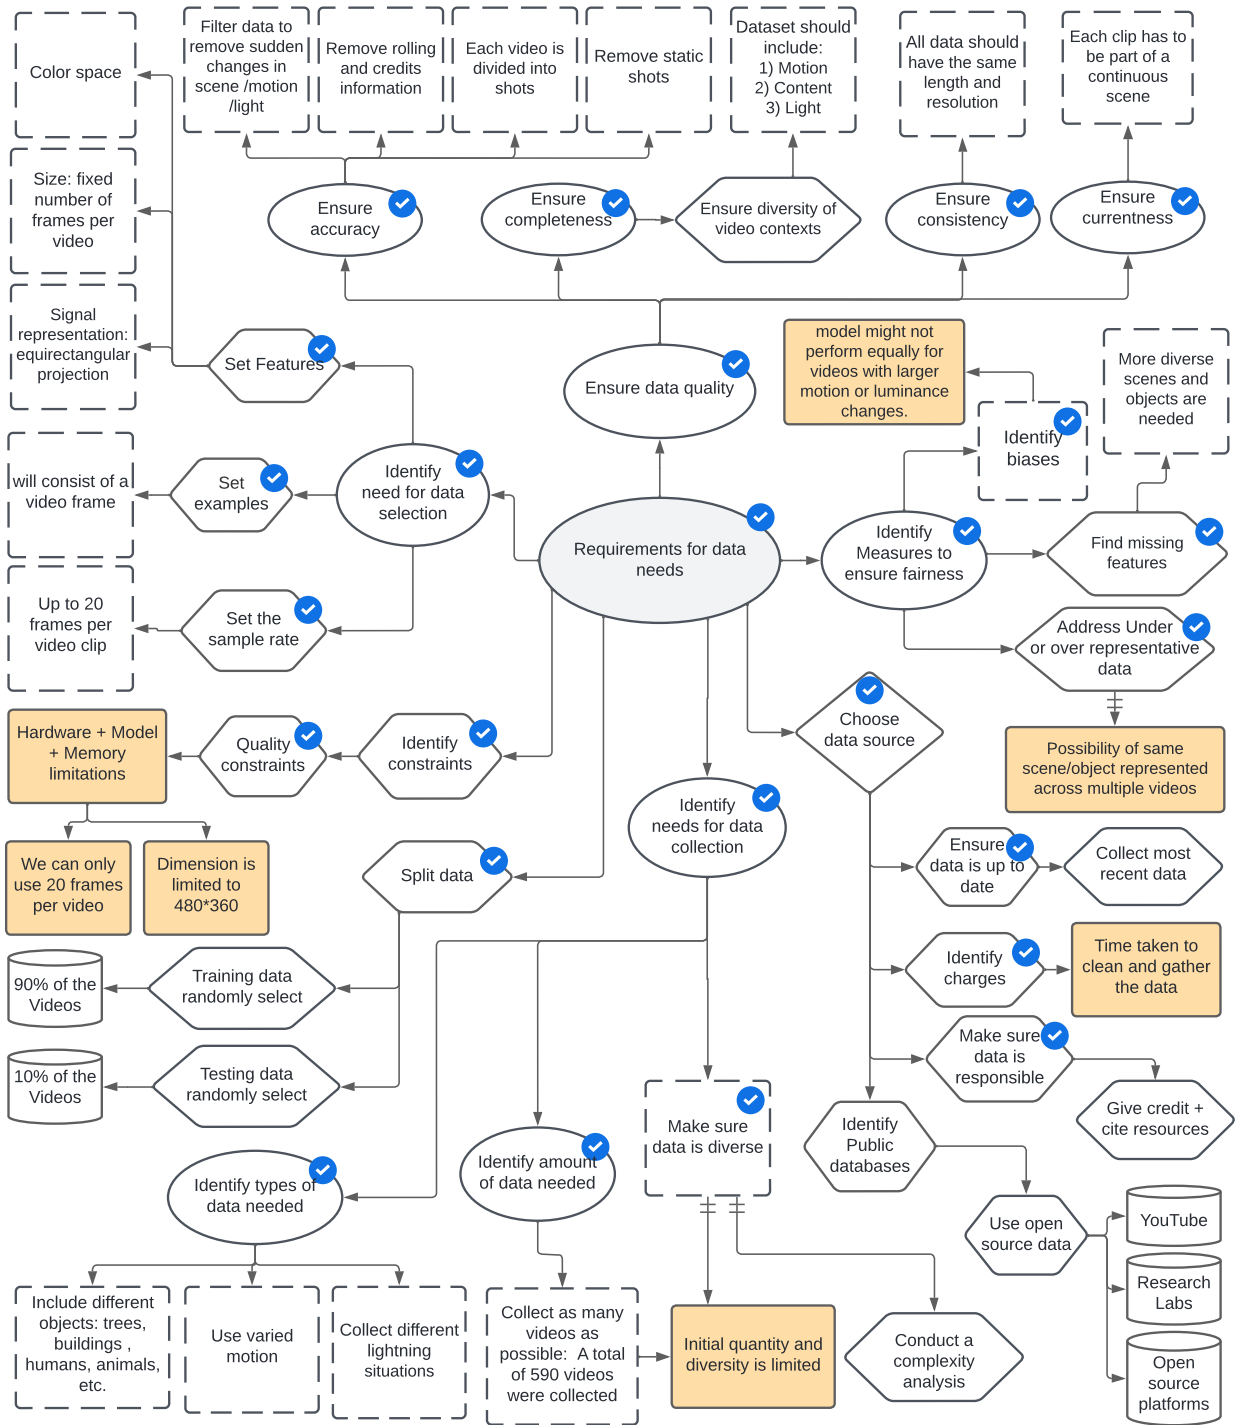


Figure 10: Model presenting the requirements for Data Needs

4. Discussion

In this section, we reflect on how using our framework helped write and model requirements for AI-based systems.

4.1. RQ1. How does our RE4HCAI framework fare for eliciting and modelling requirements for human-centered AI software?

Most traditional RE methods are not equipped to manage AI-based software, and requirements are sometimes difficult to write, especially during the early stages of systems development. The problem with specifying requirements for a AI-based software project is that it is difficult to explain the black-box nature of AI and how the system would work or what outcome it will predict [68]. In our case study, we found that the process of eliciting and specifying requirements for AI is vastly different from traditional approaches, as some aspects of the system are unpredictable and difficult to explain.

To evaluate the framework and modeling tool, we compared the building process of the 360° VR model before and after the implementation of our framework. Before implementing the framework, many human-centered aspects were missing from the final product, and most of the development was based on the technical aspects of building the AI-based system. For example, classifying what data was needed was based on technical aspects (e.g., model and hardware needs) rather than focusing on the human-perspective. Whereas, after applying the framework we were able to identify possible biases that might lead to issues, e.g., biases related to the lack of diversity of data sources and the biases related to the lack of diversity of type of videos.

The use of the catalog helped in eliciting most of the requirements to build the AI-based software project from a human-centered perspective. The models presented in figures 8, 9, and 10 helped the ML expert visualize what needed to be considered when implementing the AI-based software. The conceptual models for our RE4HCAI framework provided a visual view of all the components in the system. Our participant said “It helped me visualize all the components needed to be done in my project”. Also, we found that limitations and system needs were easier to spot with the visual representation via the conceptual model. Another observation we found missing was using feedback in the learned model. The initial configuration of the learned model did not consider the users’ view of how good the perceptual quality of the improved 360° videos was. However, after applying the framework, the decision was made to include user feedback to account for model tuning in future work of the 360° VR project.

We also noticed that some aspects of the framework would change depending on the application domain and the AI type used. For example, in this case study regression was used. Therefore, the optimized reward function would be the loss function. Therefore, we need to consider other possible requirements that might need to be elicited and specified in such situations, in our framework. This relates to the Berry’s finding [46] that different AI algorithms and their individual settings in different application contexts would require different requirements specifications. However, had the system be designed using a classification algorithm and had to choose between precision or recall, one would need to calculate the trade-offs. When calculating the trade-off between recall vs. precision, one would have to list the implications of a FP vs that of a FN. In this particular case if precision is favored, i.e., one opts to reduce the number of FPs at the cost of tolerating more FNs, then the 360° VR system will use the correctly selected TP frames, but remove some frames from the final outcome FN. Thus, the final outcome will have visual gaps, as some relevant frames will be missing. On the other hand, favoring recall meant that most frames would be included, i.e., more FPs. Thus, the final output will also include frames that are not enhanced FP, thus, compromising the quality of experience as the visual quality will not always be the same in the final product and some of the frames will still have the same quality as the original video.

The downside to favoring precision would be the possibility of having lags when rendering frames, as the system might choose to exclude frames that should have been enhanced from the frame sequence. These lags would increase the chance of latency which is a potential cause of VR sickness. Decreasing this delay or latency would result in reduced sickness [69]. Recall would still include these frames when rendering, even if these frames might not have the expected quality. In this case, the trade-off is to compromise the visual quality instead of causing lags or delays in the visual display.

4.2. RQ2. How can the RE practices be aligned in the life cycle of AI projects?

When applying the framework to the case study, not all requirements could be elicited at the start of the project. We found that it was important to identify some requirements at the start of RE. For example, when identifying user needs, we found that it was necessary to have initial specifications before building AI-based software. Also, some of the capabilities could be identified only after getting initial results from the model training, such as how much it would improve the quality of experience and if it would affect the quality of service provided. Furthermore, limitations such

as hardware resources and processing time could not be identified at the start of the project. The extent of the limitation was found after training the model with the existing dataset. The reason was that the ML expert could not know how much scaling the model needed. Adding more features and layers resulted in higher processing time and required more resources to accommodate the change. They could establish how many features or layers were needed to improve model performance only after several training iterations.

Similar patterns were found in both requirements for model needs and data needs. We observed that these requirements would change over time as they learned how the model would interact with the available data. For example, information regarding how many frames are needed could be specified only after initial testing. The datasets were limited to using up to 20 frames per clip, which was unknown at the beginning of the project. However, after initial testing, it was found that this was a constraint due to the model and hardware limitations. The same applied to identifying the quantity and diversity of the data. After experimenting on the first round of training and testing, they found that the existing 360 dataset was inadequate. Thus, the dataset had to be modified accordingly.

Meanwhile, it was important to specify some requirements for data needs, such as data sources, charges, diversity of collected data, data quality, and if data was up to date. Other requirements involved cleaning the data to provide accurate and current results. Some of these data requirements were obtained from literature and standard guidelines. Therefore, we found that not all requirements could be specified at the start of the project, and some might change or appear while model testing is in progress. We highlight the requirements that we could identify prior to starting the project in tables 2, 3, and 4 by adding a * character before the requirement.

5. Threats to Validity

In all the phases of our research method and case study design we attempted to mitigate and reduce any threat to validity as follow:

5.1. Internal Validity:

Case studies are easier to perform than experiments. However, they are more challenging to interpret [60]. The main disadvantage to conducting case studies is that the outcomes are more susceptible to researcher bias and evaluation usually depends on how the results are interpreted by the researchers [59]. To reduce potential threats and biases, we made sure when selecting the case study that the person leading the AI project had no prior knowledge of our framework. Also, we collected requirements based on our framework and compared the results to how it was built without using the proposed framework.

5.2. Construct Validity:

Another threat was the selection of studies to build the catalog. The risks of building a catalog for requirements are that it might miss out on some requirements and not be comprehensive [70]. We established our requirements for the framework based on the guidelines and the literature presented on human-centered AI. Although we might have missed out on some requirements, we did base our findings on our SLR, which covered a comprehensive list of studies on RE4AI research. Also, we used industrial guidelines such as Google and Microsoft, which have already done extensive research to provide their guidelines.

5.3. External Validity:

The framework has been applied to one case study only, which is a threat to external validity. We will address this threat by conducting more case studies in future work, and investigate how the framework applies to projects from different application domains and multiple stages in the software development lifecycle. Also, due to time limitations we could only assess the first three area's of the framework and plan to conduct further evaluation to assess the last three areas in future work.

6. Related Work

AI-based software should be carefully assessed to not replace people's abilities but rather augment their capabilities and allow people to make the final choice [71]. Shneiderman [72] explained the importance of enabling the human to be in control of the AI-based system. Including human needs and values in building AI-based software needs to be researched and examined carefully. For example, AI bots such as Alexia and Siri require extensive training to obtain

the right personality [73] and some companies are investing in creating algorithms for chatbots that can detect sarcasm or respond with empathy [74].

Recent studies have shown that many AI-based systems lack requirements specifications [75, 76, 77], which is mainly due to the difference in the building process between traditional systems vs. AI-based software [78]. We conducted an SLR [30] to identify literature that focused on using RE4AI. From the results, we were able to identify existing frameworks in RE4AI to include [79, 80, 37, 81]. Nalchigar et al. [80, 82] offer a framework to manage RE in building ML systems. The presented GR4ML framework covered three views: the business view, analytic design view, and data preparation view. A conceptual modeling language was used to present each view visually. The framework was applied to a case study that investigated the use of ML in the medical domain.

Bosch et al. [79] proposed a framework Holistic DevOps framework for building AI-based software. The framework combined three practices: requirements-driven, data-driven, and AI-driven software systems. Requirements-driven approaches are used for systems that don't require frequent changes. This method was found to use fewer resources when testing the system [83]. On the other hand, data-driven approaches are used for systems that are constantly changing, and updated [79] and are designed based on the analysis of available data [84]. Most AI-based systems are guided by large amounts of data and limited given resources. Companies that usually use automation, such as speech and image recognition, tend to use AI-driven approaches [79].

Aydemir and Dalpiaz [37] proposed a framework to aid requirements engineers and stakeholders in analyzing ethical requirements. The framework assisted in extracting, managing, and evaluating ethical requirements. However, the focus was on ethical requirements only and did not include other aspects of human-centered AI requirements. The last framework [81] provides a virtual framework to test specified requirements for a self-driving car. Also, there are frameworks built to manage aspects of AI-based systems such as Khalajzadeh et al. [85, 86] who created a domain-specific modeling tool to support data analytic solutions for ML systems (BiDaML). The toolset provides five diagrams to support a different aspect of big data analytics.

Shneiderman [87] presented a framework that proposes 15 recommendations for engineering trustworthy, reliable, and safe human-centered AI-based systems. To ensure reliable human-centered AI, software engineering teams need to apply technically sound practices such as using appropriate analysis tools, updating workflows for each task and domain, using new forms of validation and verification, testing for bias detection, and providing explainable user interfaces. For reliable human-centered AI, there should be more commitment and training to ensure safety measures and reporting and addressing errors and failures. And lastly, promoting trustworthy systems by providing independent oversight reviews to support legal and ethical codes of conduct. Another method used to aid in creating awareness to ethics when building AI software is the ECCOLA [88] method which is based on the ethical AI guidelines and involves using a deck of 21 cards, where each card targets an aspect in ethics.

Although, several frameworks are proposed to provide methods for managing RE4AI. None of these frameworks have focused on delivering an overall solution to building human-centered AI for RE. An exception to Shneiderman's [87] framework that relies entirely on promoting human-centered AI-based software. However, it focuses on the entire Software Engineering process and not on RE. This calls for the need to research and study appropriate human-centered AI methods before including them in software systems. We argue that the proposed solutions in RE should address human needs, as having a very well-engineered product would not be useful if it does not satisfy the user's needs.

7. Conclusion & Future Work

In this paper, we offer a framework to extract and model requirements for human-centered AI. The proposed framework consists of three layers. The first layer provides a reference map to the six Areas of human-centered AI guidelines. The second layer presents a catalog to elicit these requirements. The third layer provides a modeling tool to show the elicited requirements from the second layer visually. The framework was applied to a case study, and we extracted requirements for the first three layers user, model, and data needs. The case included specifying and modeling requirements for an AI system that enhanced the quality of 360° VR videos. We found that some requirements were more difficult to specify at the start and could be identified only after testing the model with the existing data.

We plan to further evaluate our framework in the future. We plan to evaluate it in a number of workshops and compare it with existing frameworks and modeling platforms. Also, we do not provide details of how attributes are connected to each other in sub-level models. For example, in model needs, when modeling the goal for balancing

between overfitting and underfitting we did not model the actual process, and this would have to be presented in a separate sub-model. We plan to extend our model to present this level of detail in future work.

References

- [1] S. Amershi, A. Begel, C. Bird, R. DeLine, H. Gall, E. Kamar, N. Nagappan, B. Nushi, T. Zimmermann, Software engineering for machine learning: a case study, in: Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice, IEEE Press, 2019, pp. 291–300.
- [2] M. Maguire, Methods to support human-centred design, *International journal of human-computer studies* 55 (4) (2001) 587–634.
- [3] A. Schmidt, Interactive human centered artificial intelligence: a definition and research challenges, in: Proceedings of the International Conference on Advanced Visual Interfaces, 2020, pp. 1–4.
- [4] J. C. Grundy, Impact of end user human aspects on software engineering., in: ENASE, 2021, pp. 9–20.
- [5] S. Amershi, M. Cakmak, W. B. Knox, T. Kulesza, Power to the people: The role of humans in interactive machine learning, *Ai Magazine* 35 (4) (2014) 105–120.
- [6] B. Shneiderman, Human-centered ai, *Issues in Science and Technology* 37 (2) (2021) 56–61.
- [7] T. Calders, I. Zliobaitė, Why unbiased computational processes can lead to discriminative decision procedures, in: *Discrimination and privacy in the information society*, Springer, 2013, pp. 43–57.
- [8] K. Lloyd, Bias amplification in artificial intelligence systems, arXiv preprint arXiv:1809.07842.
- [9] B. Shneiderman, *Human-Centered AI*, Oxford University Press, 2022.
- [10] S. H. Schwartz, An overview of the schwartz theory of basic values, *Online readings in Psychology and Culture* 2 (1) (2012) 2307–0919.
- [11] Google Research, The people + AI guidebook, [online] Available at: <https://research.google/teams/brain/pair/> [Accessed 1 April 2020] (2019).
- [12] Microsoft, Guidelines for human-ai interaction, [online] <https://www.microsoft.com/en-us/research/project/guidelines-for-human-ai-interaction/> [Accessed 1 Feb 2022].
- [13] Apple Developer, Human interface guidelines, [online] <https://developer.apple.com/design/human-interface-guidelines/technologies/machine-learning/introduction> [Accessed 1 May 2020].
- [14] M. Kissoon Curumsing, N. Fernando, M. Abdelrazek, R. Vasa, K. Mouzakis, J. Grundy, Emotion-oriented requirements engineering: a case study in developing a smart home system for the elderly, *Journal of systems and software* 147 (2019) 215–229.
- [15] I. Ramos, D. M. Berry, Is emotion relevant to requirements engineering?, *Requirements Engineering* 10 (2005) 238–242.
- [16] N. Maiden, Spocks and kirks in the requirements universe, *IEEE software* 29 (04) (2012) 84–85.
- [17] C. Thompson, N. Maiden, M. Nouri, K. Zachos, Evoking emotion through stories in creative dementia care, in: *Looking into the Future of Creativity and Decision Support Systems: Proceedings of the 8th International Conference on Knowledge, Information and Creativity Support Systems*, 2013.
- [18] T. Miller, S. Pedell, L. Sterling, F. Vetere, S. Howard, Understanding socially oriented roles and goals through motivational modelling, *Journal of Systems and Software* 85 (9) (2012) 2160–2170.
- [19] D. Callele, E. Neufeld, K. Schneider, Emotional requirements, *IEEE software* 25 (1) (2008) 43–45.
- [20] M. Burnett, A. Peters, C. Hill, N. Elarief, Finding gender-inclusiveness software issues with GenderMag: a field investigation, in: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, ACM, 2016, pp. 2586–2598.
- [21] M. Vorvoreanu, L. Zhang, Y.-H. Huang, C. Hilderbrand, Z. Steine-Hanson, M. Burnett, from gender biases to gender-inclusive design: An empirical investigation, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI'19). ACM, New York, NY, USA. <https://doi.org/10.1145/3290605.3300283>, 2019.
- [22] A. Milne, N. Maiden, Power and politics in requirements engineering: embracing the dark side?, *Requirements Engineering* 17 (2012) 83–98.
- [23] T. Kanij, R. Merkel, J. Grundy, An empirical investigation of personality traits of software testers, in: Proceedings of the Eighth International Workshop on Cooperative and Human Aspects of Software Engineering, IEEE Press, 2015, pp. 1–7.
- [24] A. B. Soomro, N. Salleh, E. Mendes, J. Grundy, G. Burch, A. Nordin, The effect of software engineers' personality traits on team climate and performance: A systematic literature review, *Information and Software Technology* 73 (2016) 52–65.
- [25] J. McIntosh, X. Du, Z. Wu, G. Truong, Q. Ly, R. How, S. Viswanathan, T. Kanij, Evaluating age bias in e-commerce, in: 2021 IEEE/ACM 13th International Conference on Cooperative and Human Aspects of Software Engineering (CHASE), IEEE, 2021, pp. 31–40.
- [26] J. Grundy, K. Mouzakis, R. Vasa, A. Cain, M. Curumsing, M. Abdelrazek, N. Fernando, Supporting diverse challenges of ageing with digital enhanced living solutions., *Studies in health technology and informatics* 246 (2018) 75–90.
- [27] T. Miller, S. Pedell, A. A. Lopez-Lorca, A. Mendoza, L. Sterling, A. Keirnan, Emotion-led modelling for people-oriented requirements engineering: The case study of emergency systems, *Journal of Systems and Software* 105 (2015) 54–71.
- [28] R. Colomo-Palacios, A. Hernández-López, Á. García-Crespo, P. Soto-Acosta, A study of emotions in requirements engineering, in: *World Summit on Knowledge Society*, Springer, 2010, pp. 1–7.
- [29] H. Perera, W. Hussain, J. Whittle, A. Nurwidyantoro, D. Mougouei, R. A. Shams, G. Oliver, A study on the prevalence of human values in software engineering publications, 2015–2018, in: 2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE), IEEE, 2020, pp. 409–420.
- [30] K. Ahmad, M. Bano, M. Abdelrazek, C. Arora, J. Grundy, What's up with requirements engineering for artificial intelligence systems?, in: 2021 IEEE 29th International Requirements Engineering Conference (RE), IEEE, 2021, pp. 1–12.
- [31] K. Ahmad, M. Abdelrazek, C. Arora, M. Bano, J. Grundy, Requirements engineering for artificial intelligence systems: A systematic mapping study, arXiv preprint arXiv:2212.10693.
- [32] H. Villamizar, T. Escovedo, M. Kalinowski, Requirements engineering for machine learning: A systematic mapping study, in: 2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA), IEEE, 2021, pp. 29–36.

- [33] B. Bruno, F. Mastrogiovanni, A. Sgorbissa, Functional requirements and design issues for a socially assistive robot for elderly people with mild cognitive impairments, in: 2013 IEEE RO-MAN, IEEE, 2013, pp. 768–773.
- [34] K. Sandkuhl, Putting AI into context-method support for the introduction of artificial intelligence into organizations, in: 2019 IEEE 21st Conference on Business Informatics (CBI), Vol. 1, IEEE, 2019, pp. 157–164.
- [35] T. M. Fagbola, S. C. Thakur, Towards the development of artificial intelligence-based systems: Human-centered functional requirements and open problems, in: 2019 International Conference on Intelligent Informatics and Biomedical Sciences (ICIIBMS), IEEE, 2019, pp. 200–204.
- [36] H. Kuwajima, F. Ishikawa, Adapting SQuaRE for quality assessment of artificial intelligence systems, in: 2019 IEEE International Symposium on Software Reliability Engineering Workshops (ISSREW), IEEE, 2019, pp. 13–18.
- [37] F. B. Aydemir, F. Dalpiaz, A roadmap for ethics-aware software engineering, in: 2018 IEEE/ACM International Workshop on Software Fairness (FairWare), IEEE, 2018, pp. 15–21.
- [38] M. Hall, D. Harborne, R. Tomsett, V. Galetic, S. Quintana-Amate, A. Nottle, A. Preece, A systematic method to understand requirements for explainable ai (xai) systems, in: Proceedings of the IJCAI Workshop on eXplainable Artificial Intelligence (XAI 2019), Macau, China, Vol. 11, 2019.
- [39] T. A. Schoonderwoerd, W. Jorritsma, M. A. Neerinx, K. van den Bosch, Human-centered xai: Developing design patterns for explanations of clinical decision support systems, International Journal of Human-Computer Studies (2021) 102684.
- [40] D. Cirqueira, D. Nedbal, M. Helfert, M. Bezbradica, Scenario-based requirements elicitation for user-centric explainable ai, in: International Cross-Domain Conference for Machine Learning and Knowledge Extraction, Springer, 2020, pp. 321–341.
- [41] M. A. Köhl, K. Baum, M. Langer, D. Oster, T. Speith, D. Bohlender, Explainability as a non-functional requirement, in: 2019 IEEE 27th International Requirements Engineering Conference (RE), IEEE, 2019, pp. 363–368.
- [42] G. Amaral, R. Guizzardi, G. Guizzardi, J. Mylopoulos, Ontology-based modeling and analysis of trustworthiness requirements: Preliminary results, in: International Conference on Conceptual Modeling, Springer, 2020, pp. 342–352.
- [43] S. Amershi, D. Weld, M. Vorvoreanu, A. Fournery, B. Nushi, P. Collisson, J. Suh, S. Iqbal, P. N. Bennett, K. Inkpen, et al., Guidelines for human-AI interaction, in: Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019, pp. 1–13.
- [44] Louis Dorard, The machine learning canvas, <https://www.louisdorard.com/machine-learning-canvas>. Accessed [March, 2020].
- [45] K. Ahmad, M. Abdelrazek, C. Arora, M. Bano, J. Grundy, Requirements practices and gaps when engineering human-centered artificial intelligence systems, arXiv preprint arXiv:2301.10404.
- [46] D. M. Berry, Requirements engineering for artificial intelligence: What is a requirements specification for an artificial intelligence?, in: International Working Conference on Requirements Engineering: Foundation for Software Quality, Springer, 2022, pp. 19–25.
- [47] C. Shin, S. Rho, H. Lee, W. Rhee, Data requirements for applying machine learning to energy disaggregation, Energies 12 (9) (2019) 1696.
- [48] J. Krause, A. Perer, K. Ng, Interacting with predictions: Visual inspection of black-box machine learning models, in: Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems, 2016, pp. 5686–5697.
- [49] H. Challa, N. Niu, R. Johnson, Faulty requirements made valuable: On the role of data quality in deep learning, in: 2020 IEEE Seventh International Workshop on Artificial Intelligence for Requirements Engineering (AIRE), IEEE, 2020, pp. 61–69.
- [50] A. Vogelsang, M. Borg, Requirements engineering for machine learning: Perspectives from data scientists, IEEE International Requirements Engineering Conference Workshops.
- [51] B. Goodman, S. Flaxman, European union regulations on algorithmic decision-making and a “right to explanation”, AI magazine 38 (3) (2017) 50–57.
- [52] S. Wachter, B. Mittelstadt, C. Russell, Counterfactual explanations without opening the black box: Automated decisions and the gdpr, Harv. JL & Tech. 31 (2017) 841.
- [53] N. Balasubramaniam, M. Kauppinen, K. Hiekkänen, S. Kujala, Transparency and explainability of ai systems: Ethical guidelines in practice, in: International Working Conference on Requirements Engineering: Foundation for Software Quality, Springer, 2022, pp. 3–18.
- [54] J. Silva, A. Salmon, P. del Foyo, J. Silva, D. d. E. Mecânica, Requirements engineering at a glance: Comparing gore and uml methods in the design of automated systems, in: proc. of the XXII Brazilian congress on automation. Brazilian Society of Automation, 2018.
- [55] K. Neace, R. Roncace, P. Fomin, Goal model analysis of autonomy requirements for unmanned aircraft systems, Requirements Engineering 23 (4) (2018) 509–555.
- [56] G. Dimitrakopoulos, E. Kavakli, P. Loucopoulos, D. Anagnostopoulos, T. Zographos, A capability-oriented modelling and simulation approach for autonomous vehicle management, Simulation Modelling Practice and Theory 91 (2019) 28–47.
- [57] D. W. Embley, B. Thalheim, Handbook of conceptual modeling: theory, practice, and research challenges, Springer, 2012.
- [58] D. Moody, The “physics” of notations: toward a scientific basis for constructing visual notations in software engineering, IEEE Transactions on software engineering 35 (6) (2009) 756–779.
- [59] S. Easterbrook, J. Singer, M.-A. Storey, D. Damian, Selecting empirical methods for software engineering research, in: Guide to advanced empirical software engineering, Springer, 2008, pp. 285–311.
- [60] B. Kitchenham, L. Pickard, S. L. Pfleeger, Case studies for method and tool evaluation, IEEE software 12 (4) (1995) 52–62.
- [61] A. Agrahari Baniya, G. Lee, P. Eklund, S. Aryal, Omnidirectional video super-resolution using deep learning.
- [62] C.-L. Fan, W.-C. Lo, Y.-T. Pai, C.-H. Hsu, A survey on 360 video streaming: Acquisition, transmission, and display, ACM Computing Surveys (CSUR) 52 (4) (2019) 1–36.
- [63] M. S. Anwar, J. Wang, W. Khan, A. Ullah, S. Ahmad, Z. Fei, Subjective qoe of 360-degree virtual reality videos and machine learning predictions, IEEE Access 8 (2020) 148084–148099.
- [64] G. d. A. Roberto, N. Birkbeck, F. De Simone, I. Janatra, B. Adsumilli, P. Frossard, Visual distortions in 360° videos, IEEE Transactions on Circuits and Systems for Video Technology 30 (8) (2019) 2524–2537.
- [65] J. Huang, Z. Chen, D. Ceylan, H. Jin, 6-dof vr videos with a single 360-camera, in: 2017 IEEE Virtual Reality (VR), IEEE, 2017, pp. 37–44.
- [66] K. Kawaguchi, J. Huang, L. P. Kaelbling, Effect of depth and width on local minima in deep learning, Neural computation 31 (7) (2019) 1462–1498.

- [67] P. ITU-T RECOMMENDATION, Subjective video quality assessment methods for multimedia applications.
- [68] H. Belani, M. Vukovic, Ž. Car, Requirements engineering challenges in building AI-based complex systems, in: 2019 IEEE 27th International Requirements Engineering Conference Workshops (REW), IEEE, 2019, pp. 252–255.
- [69] M. Mihelj, D. Novak, S. Beguš, Virtual reality technology and applications.
- [70] U. Frank, Domain-specific modeling languages: requirements analysis and design guidelines, in: Domain engineering, Springer, 2013, pp. 133–157.
- [71] M.-L. How, S.-M. Cheah, Y.-J. Chan, A. C. Khor, E. M. P. Say, Artificial intelligence-enhanced decision support for informing global sustainable development: A human-centric AI-thinking approach, *Information* 11 (1) (2020) 39.
- [72] B. Shneiderman, Human-centered artificial intelligence: three fresh ideas, *AIS Transactions on Human-Computer Interaction* 12 (3) (2020) 109–124.
- [73] H. J. Wilson, P. R. Daugherty, Collaborative intelligence: humans and AI are joining forces, *Harvard Business Review* 96 (4) (2018) 114–123.
- [74] P. R. Daugherty, H. J. Wilson, *Human+ machine: Reimagining work in the age of AI*, Harvard Business Press, 2018.
- [75] L. E. Lwakatare, A. Raj, J. Bosch, H. H. Olsson, I. Crnkovic, A taxonomy of software engineering challenges for machine learning systems: An empirical investigation, in: *International Conference on Agile Software Development*, Springer, 2019, pp. 227–243.
- [76] H. Kuwajima, H. Yasuoka, T. Nakae, Engineering problems in machine learning systems, *Machine Learning* (2020) 1–24.
- [77] K. Zamani, D. Zowghi, C. Arora, Machine learning in requirements engineering: A mapping study, in: 2021 IEEE 29th International Requirements Engineering Conference Workshops (REW), 2021, pp. 116–125. doi:10.1109/REW53955.2021.00023.
- [78] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, D. Dennison, Hidden technical debt in machine learning systems, in: *Advances in neural information processing systems*, 2015, pp. 2503–2511.
- [79] J. Bosch, H. H. Olsson, I. Crnkovic, It takes three to tango: Requirement, outcome/data, and AI driven development., in: *SiBW*, 2018, pp. 177–192.
- [80] S. Nalchigar, E. Yu, K. Keshavjee, Modeling machine learning requirements from three perspectives: a case report from the healthcare domain, *Requirements Engineering* 26 (2) (2021) 237–254.
- [81] J. S. Becker, Partial consistency for requirement engineering with traffic sequence charts., in: *Software Engineering (Workshops)*, 2020.
- [82] S. Nalchigar, E. Yu, Business-driven data analytics: a conceptual modeling framework, *Data & Knowledge Engineering* 117 (2018) 359–372.
- [83] C. E. Tuncali, G. Fainekos, D. Prokhorov, H. Ito, J. Kapinski, Requirements-driven test generation for autonomous vehicles with machine learning components, *IEEE Transactions on Intelligent Vehicles* 5 (2) (2019) 265–280.
- [84] J. Bach, J. Langner, S. Otten, M. Holzäpfel, E. Sax, Data-driven development, a complementing approach for automotive systems engineering, in: 2017 IEEE International Systems Engineering Symposium (ISSE), IEEE, 2017, pp. 1–6.
- [85] H. Khalajzadeh, M. Abdelrazek, J. Grundy, J. Hosking, Q. He, Bidaml: A suite of visual languages for supporting end-user data analytics, in: 2019 IEEE International Congress on Big Data (BigDataCongress), IEEE, 2019, pp. 93–97.
- [86] H. Khalajzadeh, A. J. Simmons, M. Abdelrazek, J. Grundy, J. G. Hosking, Q. He, Visual languages for supporting big data analytics development., in: *ENASE*, 2020, pp. 15–26.
- [87] B. Shneiderman, Bridging the gap between ethics and practice: guidelines for reliable, safe, and trustworthy human-centered ai systems, *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10 (4) (2020) 1–31.
- [88] V. Vakkuri, K.-K. Kemell, M. Jantunen, E. Halme, P. Abrahamsson, Eccola—a method for implementing ethically aligned ai systems, *Journal of Systems and Software* 182 (2021) 111067.

Appendix

1.2

Table 5: Catalog presenting requirements needed for Human-centered AI

User Needs		
Can AI add value or solve the problem to user's needs?	Do you have the expertise to manage AI? (Data scientist, ML specialist, SE) Is data available?	IF no – use a non-AI solution (format – amount – diverse)
Identify need for AI?	Who are the users Why do we need the system What is the system used for	
What are the system capabilities?	Limitations Capabilities (What can the system do)	Limitations to the user: Limitations to stakeholders: How well can the system do what it does?
Interaction with user action	Proactive vs. Reactive	(Provides results when people request them) or Reactive (Interacts with the user without requesting)
Is the user aware of AI feature?	Visible or Invisible features	
Evaluate approach	Augmentation vs Automation	(user enjoyment, control, high stakes) vs (boring, dangerous, cannot be performed by humans)
Reward function	Impact of having a false positive Impact of having a false negative Trade off - Precision vs Recall Choice of reward function Impact on user List potential pitfalls How do you provide inclusion Predictions	Precision (Excludes relevant results, removes all false positives, and misses some of the true positive predictions) vs Recall (Includes irrelevant results as it includes all true positive but might also include some false positive.) Precision or Recall (How are predictions used in making decisions)
Predictions	How are predictions used in making decisions Impact of new input on predictions Impact of feedback on predictions	
Model Needs		
Optimize algorithm for?	Optimize for accuracy, explainability, robust, etc.	
Choose the ML type	Supervised , Unsupervised, Reinforcement learning	
Model training	Static vs Dynamic	(Offline evaluation / Training improves with updates) vs (Learns from user behavior, Improves from feedback)
Balance between overfitting & underfitting		
Model Tuning (Include parameter tuning and architecture changes)	User feedback Adjusting to user behavior Trace output errors to data Parameter tuning Quantitative feedback	
Specify scalability issues		
Choose tools to use to evaluate the model		
Evaluate the quality of the model		
Data Needs		

Data selection	Features	(a characteristics of an input variable. E.g. Feature would include color, size, location, etc)
	Labels	descriptions given to data (Explicit -manually vs Implicit – model learns)
	Example	(a row of data and contains features and labels)
	Sampling rate	amount of data in a given dataset
Data collection (Match data to user needs)	Identify type of data needed Amount of data needed Diversity of data	
Constraints	cost, accuracy, quality, time?	
Data source	Type of data source Is data responsibly sourced? Using feedback as data Charges in obtaining data Measures taken to ensure data is up to date?	(public, private, mixed)
Split Data	Training data (Model tuning) Testing data	
Data quality	Accuracy Completeness Consistency Credibility Currentness	(The correctness of the data collected) (all attributes and events should have data to associate them). - List all attributes and events needed for the AI product - Make sure each attributes has data to represent (Collected data has to be free from contradictions) (Credible data had to be authentic and truthful) (corresponded to having data collected within the correct time frame; for example, when collecting images of people for a facial recognition application, a picture of an adult person when they were a baby will not be the correct data)
Data requirements	Protect personal information Does data comply with privacy and law	
Fairness	Identify biases Missing features Under or over representative data Other	
Reporting biases	Automation bias Selection bias Impact bias	Automation biases are when preferences are selected based on automated suggestions from the system Selection bias usually happens when data is not collected randomly from the target population but rather selected based on the stakeholder's requests Making an assumption based on the persons judgment (e.g. all wedding dresses are white)
	Group attribution Other	Group attribution assumes that an output suitable for an individual will have the same impact on everyone in the group
Address biases	What methods will be used?	
Feedback & Control		
Explicit feedback (Only ask when needed)	What reward will be provided to the end-user or raters to provide explicit feedback If the reward is symbolic How is it going to be reviewed and evaluated?	(Material / symbolic) What impact will it have on the user? What benefit will the user gain from it?

	What changes will it make to the AI model? How will it be used in model tuning? How is feedback going to be conducted	(surveys, notifications, ratings, etc.)
Implicit feedback	User interactions / behaviours to use: Review feedback from user interaction to how it will make changes to AI How will it be used in model tuning	Frequency of use, time of use, amount of time spent interacting with the system, accept / reject recommendations, etc.
Calibration	Reason for using calibration When should the user use calibration How is it going to affect the AI	
When asking for feedback	Measures taken to secure user privacy Provide multiple options to feedback Is feedback in line with user mental map or users understanding of the system Allow the user to dismiss feedback or opt-out	(what options will you provide?) (How are you going to provide this option?)
User control	When to give the user control Allow user to adjust preferences Allow user dismissal Level of control	(High stakes, legal, safety) (high, medium, low or no control)
Explainability & Trust		
Explain to stakeholders	Limitations of AI Capabilities of AI	
Explain to user	Inform user when changes happen Explain consequences to users action	
Explain data:	Who can access or use the user's personal data How data is shared between apps Explain how predictions are based on data	
Explain predictions	Explain with examples Explain with confidence: Do not display Confidence:	Data visualization (Expert users). Numerical (Might cause confusion). Categorical (e.g., Low - Med - High). N-best Alternatives (Low confidence) Risks (misleading, no impact to user decisions when explaining output, distracting). Low confidence (Can provide partial explanation)
Calculate Trade-off	Identify conflicts with other requirements	
Calibrate user trust		
Explain special cases	Law & rules Third party involvement	
Explain feedback	Impact of feedback on AI How feedback is used to improve model	
Avoid over trusting	Only show relevant info Account for different situations Explain errors	- Low stakes. High stakes
Errors & Failures		
User perceived errors (might change over time)	Failstates Context errors (these are true positive but)	True negative (Not included in system training). System limitation AI making incorrect assumptions. User has poor mental model. System not aligned with users needs
Background errors	Invisible errors the user cannot perceive	
	Prediction & data errors Mislabelled or misclassified results Incorrect Model Incomplete data	Poor training. Incorrect labeling

Input Errors

Unexpected input
 Old habits
 Mis-calibrated input

Output Quality

Low confidence

Lack of data. Uncertain prediction accuracy.
 Unstable information

systems hierarchy error

Multiple systems
 Crashing signals

(Allow user to give one of the system priorities)
 (Allow signals from the primary system only)

Error risks	High stakes	Health, safety and financial decisions - Sensitive / private data
Action	Abusive users Mitigate Errors Allow users to fix mistakes Provide suggestions when in doubt	