

# RITA: A Tool for Automated Requirements Classification and Specification from Online User Feedback

Manjeshwar Aniruddh Mallya<sup>1</sup>, Alessio Ferrari<sup>2</sup>, Mohammad Amin Zadenoori<sup>3</sup> and Jacek Dąbrowski<sup>1,\*</sup>

<sup>1</sup>Lero, the Research Ireland Centre for Software, University of Limerick, Ireland

<sup>2</sup>University College Dublin (UCD), Ireland

<sup>3</sup>University of Padova, Italy

## Abstract

**Context and motivation.** Online user feedback is a valuable resource for requirements engineering, but its volume and noise make analysis difficult. Existing tools support individual feedback analysis tasks, but their capabilities are rarely integrated into end-to-end support. **Problem.** The lack of end-to-end integration limits the practical adoption of existing RE tools and makes it difficult to assess their real-world usefulness. **Solution.** To address this challenge, we present RITA, a tool that integrates lightweight open-source large language models into a unified workflow for feedback-driven RE. RITA supports automated request classification, non-functional requirement identification, and natural-language requirements specification generation from online feedback via a user-friendly interface, and integrates with JIRA for seamless transfer of requirements specifications to development tools. **Results and conclusions.** RITA exploits previously evaluated LLM-based RE techniques to efficiently transform raw user feedback into requirements artefacts, helping bridge the gap between research and practice. A demonstration is available at: <https://youtu.be/8meCLpwQWV8>.

## Keywords

Requirements Engineering, User Feedback, App Reviews, Requirements Specification, Requirements Classification, Natural Language Processing, Mining Software Repository, Large Language Models, AI4RE, NLP4RE

## 1. Introduction

Online user feedback, such as app store reviews or online discussion forums, contains useful information about what users want, what problems they encounter, and how they perceive software quality [1]. This information is useful for requirements engineering (RE) activities (e.g., requirements specification) [2]. However, the large volume of feedback and its informal and noisy nature make it difficult for requirements engineers to analyse it manually and turn it into structured requirements artefacts [3].

A variety of automated techniques have been proposed to support the analysis of user feedback [1]. These include methods for classifying feedback [4], extracting features, and identifying sentiment or topics [5, 6]. More recently, large language models (LLMs) have attracted attention because they can process natural language inputs, exhibit reasoning-like capabilities, and generate coherent text [7]. This makes them suitable for tasks such as identifying user requests and drafting requirements specifications [8].

In our previous study, we empirically evaluated lightweight open-source LLMs on feedback-driven RE tasks, including user request classification, non-functional requirement identification, and requirements specification generation [9]. We showed that such models can support these tasks with moderate success, while also observing limitations related to consistency, accuracy, and practical use.

Although this prior work provided evidence of what LLMs can do on individual RE tasks, it did not address how these techniques could be combined and used within a single tool that fits practitioners' workflows [2, 7]. Moreover, existing tools typically focus on isolated analysis functions rather than

---

*Joint Proceedings of REFSQ-2026 Workshops, Doctoral Symposium, Posters & Tools Track, and Education and Training Track. Co-located with REFSQ 2026. Poznan, Poland, March 23-26, 2026.*

\*Corresponding author.

✉ mallyaaniruddh@gmail.com (M. A. Mallya); alessio.ferrari@ucd.ie (A. Ferrari); amin.zadenoori@unipd.it (M. A. Zadenoori); jacek.dabrowski@lero.ie (J. Dąbrowski)

ORCID 0000-0002-0636-5663 (A. Ferrari); 0000-0003-4591-153X (M. A. Zadenoori); 0000-0003-3392-0690 (J. Dąbrowski)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

providing support that transforms raw user feedback into artefacts that can be used directly in software development [1].

To address this gap, we developed RITA, a tool that integrates multiple LLM-based analysis tasks into a unified workflow for feedback-driven RE. RITA supports requirements engineers in analysing online user feedback by automatically classifying user requests, identifying non-functional concerns, and generating natural-language requirements specifications. These capabilities are delivered through a graphical interface designed to support user interaction. In addition, the tool is integrated with JIRA, enabling generated requirements artefacts to be transferred directly into widely used development tools. RITA uses lightweight open-source LLMs that we previously evaluated and benchmarked for these tasks; details on the models, datasets, prompts, and performance are reported in our prior study [9].

The goal of RITA is to enable the use of LLM-based techniques in practice during early and ongoing RE activities that rely on continuous user feedback. The tool is intended for practitioners who need lightweight support for analysing feedback and generating draft requirements within existing tools.

The key contribution of this paper is the move from isolated, offline evaluation of LLM4RE techniques to their use in a fully integrated end-to-end tool. Unlike our prior work [9], which studied individual tasks in isolation, this paper introduces RITA, a working system that combines these tasks into a single interactive workflow and connects them to existing software development tools. By providing a working tool rather than only benchmarks, RITA enables studies of how LLM-based RE techniques are used in practice. A demo<sup>1</sup> and the replication package<sup>2</sup> are publicly available.

## 2. RITA Tool Description

We now describe the features of the RITA tool from the perspective of its end users. We then present an overview of the user flow, the architecture of the tool, and details of its implementation.

### 2.1. Use Cases

This section describes RITA from the perspective of its users. RITA supports requirements analysts, product managers, and software engineers in converting large collections of textual user feedback into structured software requirements. Within the system boundary, RITA supports the following use cases:

- **UC1 – Upload feedback corpus.** The user uploads a dataset of user feedback (e.g., app reviews, survey responses, or issue reports) to be analysed. RITA supports common data formats (txt, csv, and xls) and accepts feedback from multiple sources.
- **UC2 – Configure classification scheme.** The user defines how the feedback is analysed by selecting a classification type (e.g., User Request Types or Non-Functional Requirements), which represent alternative classification schemes, choosing a large language model, and specifying a prompt strategy. This allows the analysis to be adapted to different project or research goals.
- **UC3 – Classify feedback.** RITA processes the uploaded feedback and assigns each entry to one or more requirement-related categories based on the selected configuration, converting raw text into structured requirement labels.
- **UC4 – Inspect classification results.** The user reviews the classification output to see how individual feedback items were interpreted, supporting validation and qualitative analysis.
- **UC5 – Generate Software Requirements Specifications (SRS).** RITA generates formalised SRS documents from the classified feedback, consolidating and structuring the extracted requirements for documentation and communication.
- **UC6 – Generate user stories.** RITA converts classified feedback into agile-style user stories, allowing the user to move from raw feedback to backlog-ready artefacts.

---

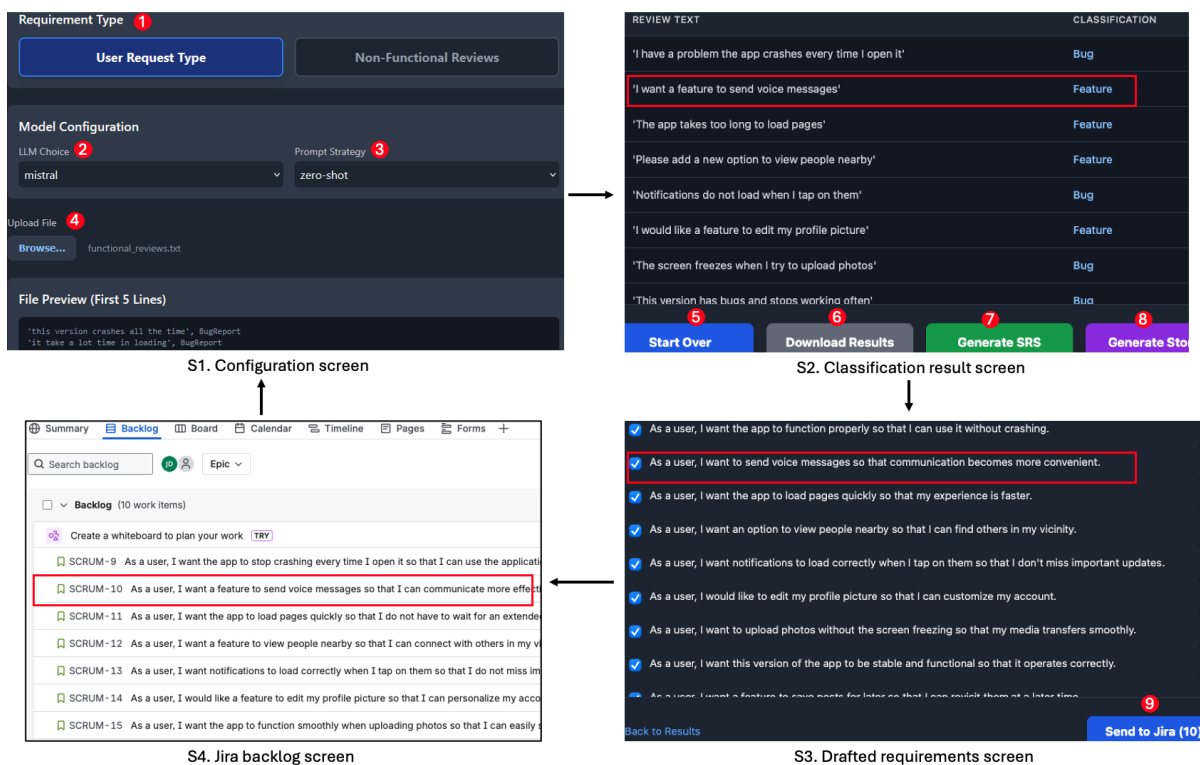
<sup>1</sup><https://youtu.be/8meCLpwQWV8>

<sup>2</sup><https://github.com/Aniruddh-Mallya/RITA>

- **UC7 – Export requirements artefacts.** The user exports classification results, SRS documents, and user stories in standard formats (e.g., csv) for further analysis, reporting, or archiving.
- **UC8 – Integrate with Jira.** The user transfers selected user stories into Jira as issues so that the generated requirements can be managed within existing development workflows.

## 2.2. User Flow and Screen Presentation

The RITA tool guides the user through a sequence of screens for configuring the analysis (S1), viewing the classification results (S2), and viewing requirement drafts (S3). The exported output from RITA, in the form of requirements, is then made available in the Jira backlog screen (S4). Figure 1 presents the user flow and screen presentation, focusing on four main screens. The user flow starts on the configuration screen (see S1). Here, the user selects the classification type **1**, the large language model **2**, and the prompt strategy **3** that control how the feedback is processed. These settings can be changed between runs without re-uploading the data. After configuration, the user uploads the feedback file **4**. RITA accepts CSV, Excel, and text files. For tabular inputs, the user selects the column containing the feedback text. Once the file is submitted, the system starts the classification process and displays progress information, such as the number of feedback items processed. The user can cancel the run and restart it with different settings if needed. When processing is complete, the classified feedback is displayed in a structured table on the classification results screen (see S2). Individual entries can be inspected, and the full output can be downloaded. From this screen, users can restart the classification with a different input file or classification type **5**, or download the results **6**. Users can also proceed to generate higher-level artefacts. RITA can generate Software Requirements Specifications **7** and user stories **8** from the classified feedback (see S3). After the user stories are generated, the user can select stories of interest and transfer them to Jira **9**. These user stories are then created as Jira items and appear in the Jira interface (see S4).



**Figure 1:** User Flow and Screen Presentation in RITA Tool and Integrated Jira Project Management Tool.

## 2.3. Overview of the RITA Architecture

We now present a high-level overview of the RITA architecture. This section describes the system's main components, their structural organization, and their deployment in the execution environment. It also explains how these components interact at runtime to realize the intended use cases.

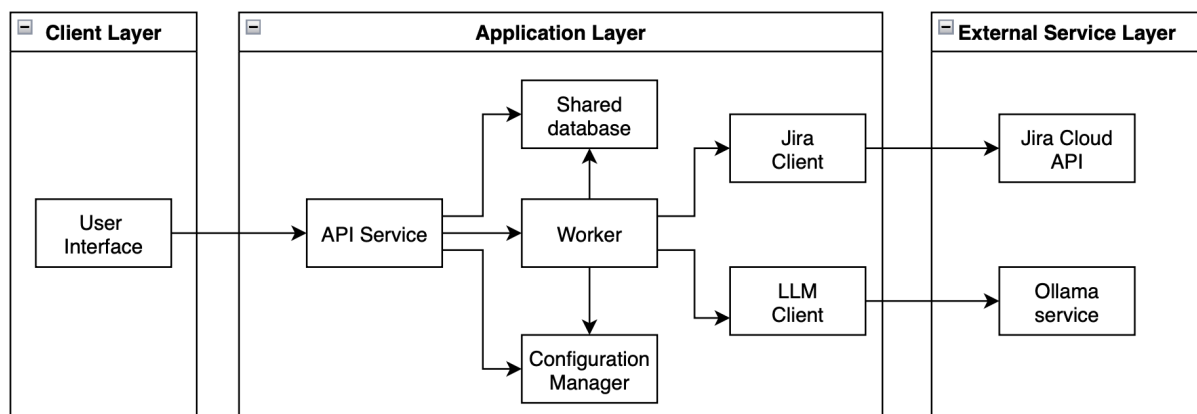
### 2.3.1. Component View

The RITA architecture is organised into three layers: **client**, **application**, and **external service**. It is composed of several components separating user interaction from data analysis and processing. Figure 2 presents the RITA component structure and its interconnections; arrows point from components that initiate interactions to the components they use.

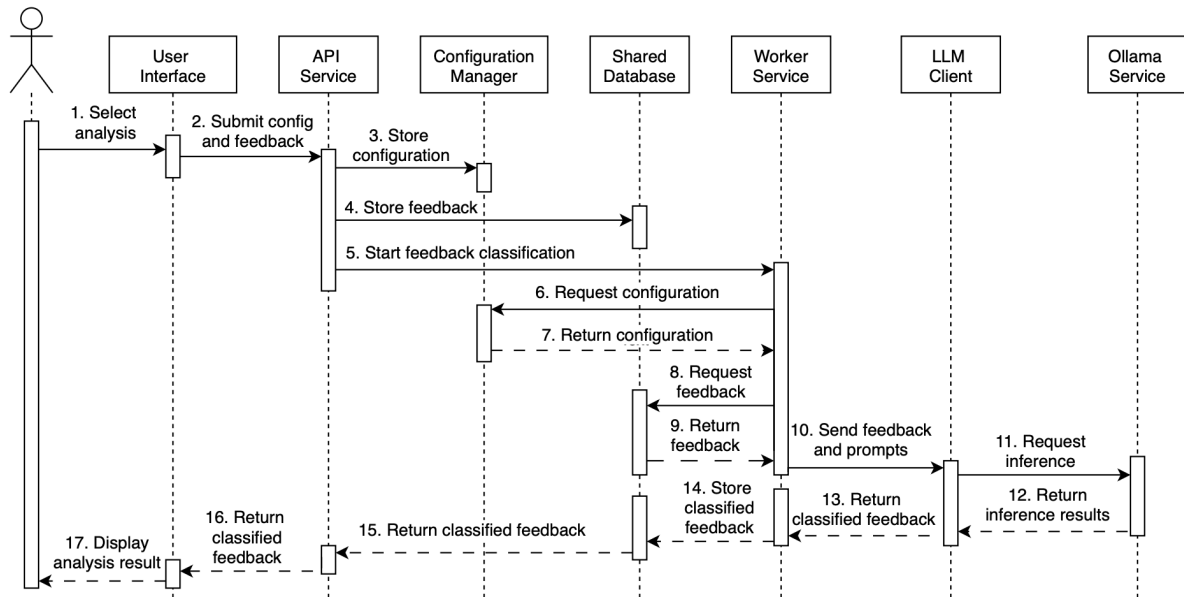
The **client layer** consists of a user interface (UI) that runs in the user's web browser. This interface provides a Web interface for users to upload review files and configure analysis options. It also allows users to initiate classification and requirements generation and to view the generated results. All user interactions are translated into structured HTTP requests, sent to the backend *API service* for processing.

The **application layer** acts as the central coordination point of the system and consists of six components. The *API service* serves as the entry point to the backend. It is directly connected to the client layer, the *configuration manager*, the *worker service*, and the *shared database*. The API service receives requests from the client layer and validates the inputs. It stores uploaded data, starts background processing, and retrieves results for display in the UI. The *configuration manager* is accessed by both the API service and the worker service. It maintains analysis parameters, model settings, and project-specific options used during execution. The *worker service* is responsible for executing computationally intensive and long-running tasks. These tasks include analysing large collections of review text and generating requirements artefacts. The worker service is invoked by the API service and retrieves the configuration from the configuration manager. During execution, it reads input data from the *shared database* and stores intermediate and final results there. For language model inference, the worker service uses the *LLM client* to communicate with the language model service. It also uses the *JIRA client* to transfer generated requirements artefacts to external Jira systems. The *shared database* is accessed by both the API service and the worker service. It stores uploaded reviews, intermediate processing states, and generated artefacts. This enables persistent storage and retrieval across user sessions.

The **external service layer** consists of two components. The *Ollama local service* hosts large language models. It is accessed only through the LLM client to perform automated analysis and text generation. The *Jira Cloud applications* are accessed through the JIRA client. They receive exported user stories and requirements artefacts generated by the RITA tool.



**Figure 2:** Structural overview of RITA. Arrows indicate usage dependencies from invoking to used components.



**Figure 3:** Sequence Diagram of RITA (Review Classification Workflow)

### 2.3.2. Use Case Realisation

We now describe how RITA implements its two main use cases: feedback classification and requirements specification, and Jira integration with requirements export. The first is explained in more detail and supported by a supplementary sequence diagram showing system interactions. The second is described more briefly, as it follows the same interaction structure.

**Feedback Classification and Requirements Specification** Figure 3 illustrates the interaction sequence for the feedback classification workflow. The process begins when the user selects the classification configuration and uploads a feedback dataset through the user interface. The user interface forwards the configuration and data to the backend *API service*. The API service stores the configuration via the *configuration manager*, persists the feedback corpus in the *shared database*, and invokes the *worker service* to start the classification job. The worker service retrieves the active configuration from the configuration manager and requests feedback items from the shared database. Each feedback item is submitted to the language model through the *LLM client*, which forwards the inference request to the *Ollama local service*. The classification results returned by the language model are passed back through the LLM client to the worker service and stored in the shared database. Once available, the API service retrieves the stored classification results from the shared database. The results are then returned or streamed to the user interface, where they are displayed to the user.

After reviewing the classification results, the user may initiate the generation of higher-level requirements artefacts. This process follows the same interaction pattern as the feedback classification workflow; therefore, a separate sequence diagram is omitted for the sake of brevity. The user interface forwards the generation request to the *API service*, which triggers the *worker service* to retrieve the classified feedback from the *shared database* and the active configuration from the *configuration manager*. The worker service constructs generation prompts and submits them to the language model via the *LLM client*. The generated artefacts, such as SRS and user stories, are stored in the *shared database*. They are later retrieved by the API and displayed in the UI for inspection and export.

**Jira Integration and Requirements Export** After reviewing the generated requirements artefacts, the user may choose to export user stories through the user interface. The user interface forwards this request to the backend *API service*, which triggers the *worker service* to initiate the export process. The

worker service uses the *JIRA client* to transfer the selected user stories to the *Jira Cloud applications*. The exported user stories are created as Jira issues and become part of the target project's backlog.

### 2.3.3. Deployment View

RITA is deployed as a self-contained system on a single host machine and is accessed through a web browser. The **client layer** runs in the browser, while the **application layer** is deployed on the host machine. The **application layer** includes locally deployed components, including the *LLM client* and the *JIRA client*. These components run on the host machine and act as connectors to external services. Language model inference is performed by the *Ollama local service* in the **external service layer**. It is accessed through the *LLM client* and runs on the same host machine, ensuring that review data remains local. The *JIRA client* connects from the host machine to the *Jira Cloud applications* in the **external service layer** over the network. This connection is used to export generated requirements and artefacts into selected Jira projects. These deployment relationships are illustrated in the deployment diagram.

## 2.4. Implementation

RITA is a Docker-containerised system supporting transparent and reproducible experimentation. The React<sup>3</sup> user interface provides a stateless workflow for uploading datasets, configuring parameters, and inspecting results. The FastAPI service<sup>4</sup> forms the Python-based backend of the system and implements the core orchestration logic. It exposes REST endpoints for submitting analysis requests and retrieving results. It provides streaming endpoints to deliver incremental output to the user interface. The backend orchestrates execution and manages data flow by invoking a background processing component. This component performs computationally intensive tasks, including feedback classification and requirements generation. A local SQLite database<sup>5</sup> is used to store uploaded data, intermediate results, and generated artefacts during execution. Language model inference is performed by a locally hosted Ollama service<sup>6</sup>. This service is accessed through a dedicated LLM client. The client encapsulates model-specific communication and enables interchangeable inference configurations. Optional Jira Cloud integration<sup>7</sup> is provided by a client that exports selected user stories as Jira issues.

## 3. Challenges, Next Steps and Evaluation

**Challenges.** As RITA is still under development, some functions remain unstable. Consequently, the generated user stories and SRS do not always conform to the required structure or level of detail. We will address this by improving prompt design, strengthening post-processing, and systematically debugging the implementation. The lightweight open-source LLMs limit accuracy and consistency, sometimes misclassifying feedback or generating unclear requirements. Future work will customise models, improve prompts, and fine-tune them on domain-specific data. Integration into real workflows remains challenging. While RITA supports Jira, its artefacts do not always align with industrial standards.

**Next Steps and Evaluation.** The next phase of this work focuses on improving the stability of RITA and the quality of its outputs. We will fix known implementation issues, improve the user interface, and refine the formatting of SRS and user stories. We will also experiment with additional open-source models, improve prompt design, and fine-tune or retrain models using requirements and user feedback data. In addition, we will integrate commercial LLMs to increase output accuracy and consistency. We plan to evaluate RITA with requirements engineers, product managers, and software developers using both quantitative and qualitative methods. Quantitatively, we will compare the time required to analyse the same feedback with and without RITA, measure the number of identified requirements, and

---

<sup>3</sup><https://react.dev/>

<sup>4</sup><https://fastapi.tiangolo.com/>

<sup>5</sup><https://www.sqlite.org/>

<sup>6</sup><https://github.com/ollama/ollama-python>

<sup>7</sup><https://www.atlassian.com/software/jira>

assess the effort needed to revise generated artefacts. Qualitatively, we will collect practitioner feedback through interviews and questionnaires to evaluate usability, perceived usefulness, and confidence in the results. This evaluation will determine RITA's practical value for feedback-driven RE.

## 4. Conclusion

This paper presented RITA, a tool that integrates lightweight open-source LLMs into a unified workflow for feedback-driven RE. Building on our prior empirical work [9], RITA applies previously evaluated LLMs to analyse online user feedback. It supports user request classification, non-functional requirement identification, and requirements specification generation using a graphical interface; and integration with the JIRA project management tool. The main contribution of RITA is its end-to-end support for bridging data-driven RE research with practical RE workflows. By combining multiple feedback analysis tasks in a single tool and transferring generated artefacts directly to JIRA, RITA may enable the systematic and structured use of LLM-based techniques in industrial settings. This paper focuses on the tool's design and capabilities. Future work will evaluate its real-world usefulness and effectiveness. It will also extend the tool with additional feedback analysis features and deeper process integration.

## Acknowledgments

A major part of this work was conducted as part of the MSc thesis of M. A. Mallya, supervised by J. Dąbrowski [10]. The results contribute to the Prompt Me project [11]. This publication has emanated from research jointly funded by Taighde Éireann – Research Ireland under Grant Number 13/RC/2094\_2, and co-funded by the European Union under the Systems, Methods, Context (SyMeCo) programme Grant Agreement Number 101081459. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

## Declaration on Generative AI

During the preparation of this work, the authors used ChatGPT-5.2 [12] for sentence polishing, rephrasing, and improving readability. All generated content was reviewed and edited by the authors, who remain fully responsible for the final text.

## References

- [1] J. Dąbrowski, E. Letier, A. Perini, A. Susi, Analysing app reviews for software engineering: a systematic literature review, *Empirical Softw. Engg.* 27 (2022).
- [2] J. Dąbrowski, E. Letier, A. Perini, A. Susi, Mining user feedback for software engineering: Use cases and reference architecture, in: *RE, IEEE, 2022*, pp. 114–126.
- [3] A. A. Al-Subaihin, F. Sarro, S. Black, L. Capra, M. Harman, App store effects on software engineering practices, *IEEE Trans. Softw. Eng* 47 (2021) 300–319.
- [4] W. Maalej, Z. Kurtanovic, H. Nabil, C. Stanik, On the automatic classification of app reviews, *Requir. Eng.* (2016).
- [5] Q. Motger, M. Oriol, M. Tiessler, X. Franch, J. Marco, What about emotions? guiding fine-grained emotion extraction from mobile app reviews, in: *Proc. IEEE Int. Requirements Engineering, 2025*.
- [6] M. Tiessler, Q. Motger, Re-miner 2.0: A holistic framework for mining mobile application reviews, in: *REFSQ 2025 Posters & Tools, 2025*.
- [7] M. A. Zadenoori, J. Dąbrowski, W. Alhoshan, L. Zhao, A. Ferrari, Large language models (llms) for requirements engineering (re): A systematic literature review, *arXiv abs/2509.11446* (2025).

- [8] L. Pasquale, A. Ragone, E. Piemontese, A. Amiri Darban, Exploring the use of llms for requirements specification in an it consulting company, in: Proc. IEEE RE 2025, 2025.
- [9] M. A. Mallya, A. Ferrari, M. A. Zadenoori, J. Dąbrowski, From online user feedback to requirements: Evaluating large language models for classification and specification tasks, in: Proc. REFSQ, 2026.
- [10] M. A. Mallya, Developing Software Requirements Using LLMs: A Preliminary Empirical Study, Master's thesis, University of Limerick, Ireland, 2025. Supervised by Dr. Jacek Dabrowski.
- [11] J. Dąbrowski, Prompt Me: Intelligent Software Agent for Requirements Engineering, <https://prompt-me.github.io>, 2026.
- [12] OpenAI, Chatgpt-5: Large language model, <https://chat.openai.com/>, 2026. Version GPT-5, accessed February 2026.