

# Exploring LLMs for User Story Extraction from Mockups

Firmenich Diego<sup>1\*</sup>, Antonelli Leandro<sup>2,3</sup>, Pazos Bruno<sup>1,4</sup>, Lozada Fabricio<sup>3,6</sup>,  
and Morales Leonardo<sup>1,4,5</sup>

<sup>1</sup> Departamento de Informática, Facultad de Ingeniería, Universidad Nacional de la Patagonia, Argentina

<sup>2</sup> LIFIA, Facultad de Informática, Universidad Nacional de La Plata, Argentina

<sup>3</sup> CAETI, Facultad de Tecnología Informática - Universidad Abierta Interamericana

<sup>4</sup> Laboratorio de Ciencias de la Imágenes, DIEC, UNS

<sup>5</sup> Instituto Patagónico de Ciencias Sociales y Humanas, CONICET-CENPAT

<sup>6</sup> UNIANDES, Universidad Regional Autónoma de los Andes, Ecuador

\*[dafirmenich@ing.unp.edu.ar](mailto:dafirmenich@ing.unp.edu.ar)

**Abstract.** User stories are one of the most widely used artifacts in the software industry to define functional requirements. In parallel, the use of high-fidelity mockups facilitate end-user participation in defining their needs. In this work, we explore how combining these techniques with large language models (LLMs) enables agile and automated generation of user stories from mockups. To this end, we present a case study that analyzes the ability of LLMs to extract user stories from high-fidelity mockups, both with and without the inclusion of a glossary of the Language Extended Lexicon (LEL) in the prompts. Our results demonstrate that incorporating the LEL significantly enhances the accuracy and suitability of the generated user stories. This approach represents a step forward in the integration of AI into requirements engineering, with the potential to improve communication between users and developers.

**Keywords:** User stories, high-fidelity mockups, language extended lexicon, requirements engineering.

## 1 Introduction

In the dynamic landscape of software development, defining and communicating requirements is paramount to the success of any project [16]. Among the various techniques and artifacts employed for this purpose, user stories have emerged as one of the most widespread. Serving as concise, user-centric descriptions of functionality, user stories facilitate communication between technical teams and stakeholders, fostering collaboration and ensuring that software solutions are truly aligned with user needs and business goals [29]. Their widespread adoption is due to their ability to promote iterative development, facilitate clear communication, and ultimately drive the creation of valuable and user-focused software products throughout the industry [21]. They are also technically simple artifacts

containing text, minimally structured, and easy to manage with various tools. Other ways for end-users to specify requirements include mockups. Through mockups, users can specify their requirements with much greater precision in terms of UI [10]. These techniques, used in web application development, are even more useful and intuitive when, through augmentation techniques, existing web applications are used as a canvas for high-fidelity mockup creation. However, ensuring consistency and correspondence between user stories and mockups is also a challenge [18].

In this way, mockups, as artifacts, complement the definition of user stories, and user stories textually complement mockups. However, they are often insufficient on their own for comprehensive requirements management and developer implementation. This inherent limitation necessitates complementing visual mockups with textual or structured information, even in user-centered processes where users might provide this input directly through annotations or descriptions. Bridging this gap between visual specification and structured textual requirements is a recognized need in the field, explored by various approaches focused on extracting or complementing visual artifacts with textual requirements. [28, 31]

In this work, we show how, based on recent advances in large language models (LLMs), it is possible to derive user stories from images corresponding to mockups. This has great value, as it allows for even more streamlined requirement definition with different artifacts. However, in the same way that traditionally not all users of a particular application specify their user stories with a common language. And this is a source of ambiguity [3]. In this work, we test how, in addition to facilitating the derivation of text for user stories from images of their mockups, LLMs are capable of generating very precise user stories if a specific glossary is provided to them at the time of processing. Particularly, this proposal uses the Language Extended Lexicon (LEL) [26]. This adds another layer, not only facilitating the automatic generation of user stories from mockups, but also facilitating and guaranteeing the definition of these artifacts with the terminology that truly corresponds to each domain, since the LEL is originally built by requirements engineers in each domain, and consists of a series of very well-structured textual descriptions of the application domain.

This work is organized as follows: Section 2 establishes the background, providing a synthesized review of the pillars that underpin the work: augmentation-based mockups, LLMs, and LEL. Likewise, a brief review of recent related works is made. Section 3 describes the proposed process in more detail. Section 4 shows the evaluation based on a case study, and Section 5 shows the results and discussion. Finally, Section 6 presents our conclusions and proposes future work.

## 2 Background and related works

High-fidelity mockups are interactive and detailed representations of user interfaces that play a crucial role in software development by enhancing com-

munication among stakeholders and optimizing design validation [12, 30]. These mockups, which closely resemble the final product, enable more effective user feedback and ensure that the final product aligns with expected requirements. Advanced tools such as Mockplug [12] and machine learning-based techniques, such as those used in SketchingInterfaces [35], have demonstrated improvements in development efficiency: Mockplug allows for a rapid technical requirement specification directly within the user interface, while SketchingInterfaces enables the automatic conversion of hand-drawn sketches into high-fidelity mockups. Furthermore, automating code generation from these mockups significantly reduces both development costs and time, allowing developers to focus on functionality rather than repetitive design tasks [30].

Mockups complement textual requirement artifacts, such as user stories, by offering a visual representation of the intended functionality. However, maintaining consistency between user stories and mockups remains a challenge [18]. Recent studies have explored the automation of traceability between user stories and graphical artifacts using LLMs [18]. Our work builds on this synergy by exploring how LLMs can derive user stories from high-fidelity mockups while ensuring domain-specific accuracy through the integration of an glossary of the domain: the Language Extended Lexicon (LEL).

The LEL provides a structured representation of linguistic symbols within a specific domain [19]. In Software Engineering, the LEL can be described as an enriched glossary that not only defines terms but also incorporates functions and behaviors, serving as a domain metamodel [33]. Its primary foundation lies in the premise that user and client engagement is strengthened when they share a common language with software engineers, thereby facilitating communication between stakeholders and developers. Its application helps reduce ambiguity, improve alignment between initial requirements and user needs, and expand domain knowledge [4]. Although adequate use of the LEL glossary implies additional work, this extra activity can result in benefit in complex domains or teams without experienced workers, since the LEL glossary optimizes the elicitation of functional requirements from artifacts such as use cases and user stories, providing a more structured and precise representation of system expectations [5]. Moreover, it plays a key role in identifying and addressing non-functional requirements by treating them as first-class elements within the specification process, allowing for a more comprehensive approach to stakeholder demands [23].

By integrating the LEL with LLMs, we leverage this structured linguistic knowledge to enhance the consistency and precision of automatically generated user stories.

LLMs are Artificial Intelligence systems trained on enormous datasets and given their flexibility, with a very large amount of trainable parameters. Thanks to their architecture and multiple features, these models are capable of processing and responding to human-like text inputs by generating outputs with similar characteristics, appearing to understand the meaning of words and sentences given the context in which they are presented.

A big part of these model’s success can be attributed to their architecture known as transformer [32] which basically allow them to “pay attention” to specific portions in the input text and analyze the relationship with the rest of the input data; how much these portions are influenced by other data in the input and conversely how they affect the rest of the data.

With the goal of capturing as much information as possible, LLMs have millions or even billions of trainable parameters. During a pre-training stage [9], these models are trained on huge volumes of unlabeled data in text format [27, 36, 20] in order to create and adjust their knowledge base providing the model with basic notions of the language, its structure, how it is composed and how context is built in it. In further stages, this knowledge base is extended through a process known as fine-tuning, allowing the model to be trained for specific purposes. Nowadays, it is common to find LLMs behind many online services, such as: virtual assistants [6, 2], chatbots [25, 14], automatic translation [8, 15], content generation [13, 24], and sentiment analysis among others.

LLMs have been increasingly explored in software engineering applications, including requirements engineering. Recent studies have investigated their ability to generate user stories, extract requirements from textual descriptions, and even refine ambiguous requirements [3, 17]. For example, Kolthoff et al. [18] proposed an LLM-based approach to interlink user stories with graphical user interface prototypes, demonstrating improvements in requirement traceability. Other works have explored the use of LLMs for requirements classification, ambiguity resolution [7], and structured requirements generation [34].

Despite these advancements, relatively few studies have explored the direct derivation of user stories from high-fidelity mockups. Works such as Firmenich et al. [11] and Wimmer et al. [35] have examined how augmented user interactions can facilitate requirement elicitation. Our work aims to address this gap by evaluating how LLMs can process visual representations of requirements and generate structured user stories, particularly when enhanced with an LEL. This approach builds upon prior work in automated requirements engineering and demonstrates how domain-specific lexicons can refine the output of LLMs.

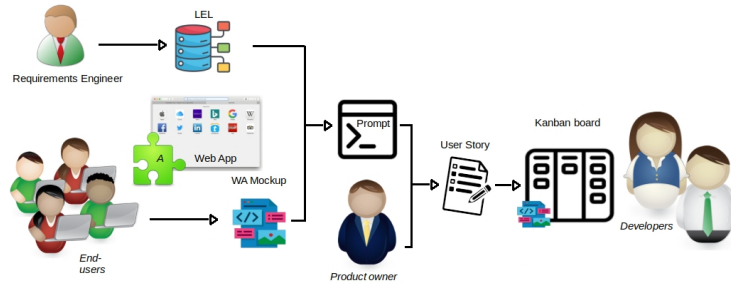
### 3 Contribution

The proposed approach is illustrated in Fig. 1, where the collaborative flow between the various actors involved in the proposal can be appreciated. End-users interact with a tool for high-fidelity mockup creation (identified in the figure as a browser add-on, that has a puzzle piece “A”). This allows them to define their requirements through three key actions: selecting relevant pre-existing elements for their needs, adding new components (widgets, elements from other parts of the application or other applications), and removing unwanted elements.

Using the web application as a canvas, the modified elements (selected or added) adopt a distinctive visual style, employing a characteristic hand-drawn font typical of mockups. This facilitates their identification compared to the elements of the original application.

In parallel, requirements engineers describe the language of the application domain, identifying and describing LEL symbols. The output produced by the requirements engineers is a structured glossary that textually describes all actors (subjects) in the application domain, with all their actions (verbs) and all concepts (objects) involved, along with their situations (states), establishing meaning (notion) for each of them as well as the relationships (behavioral responses) between them.

All this valuable information for each requirement, that is, the screenshot of the original application, the mockup, and the LEL glossary, is used to feed the context of the LLM. Although the LEL glossary is not intuitive and needs some experience, it is not hard to learn and provides the knowledge necessary to complement the visual artifacts. The product owner, using a prompt along with this information, can then complement the definition of the original requirement, obtaining a user story appropriate for the mockup created. Finally, after managing and prioritizing the requirements, the product owner can manage these products, for example, in a Kanban board shared with the developers.



**Fig. 1.** Proposed Collaborative Workflow

Table 1 summarizes the roles of the proposed process, the actions that each actor must perform, as well as the tools involved, and the result and format obtained.

To further streamline the process, once the end-user defines their requirements using the mockup tool and the mockups are finalized, the tool can automatically interact with the LLM via API. This interaction generates the corresponding user story, which can then be seamlessly integrated into requirement management platforms, such as a Kanban board managed by the product owner[22].

## 4 Proof of Concept and Preliminary Validation

This section presents two use cases to illustrate the proposal. The first one is based on a widely known application, while the second one focuses on a niche domain.

**Table 1.** Roles, Actions, Tools, Products, and Formats in the Proposed Approach

Role	Action	Tool	Product	Format
End-user	Defines requirements	High-fidelity mockup tool	WA based Mockup	WA Mockup, Images
Requirements Engineer	Defines the application domain language	LEL edition Tool	LEL	Text
Product Owner	Manages and prioritizes requirements	Kanban, Prompt	Mockup+US	Text, Image
Developer	Implements requirements	Kanban, IDEs	Artifact	Software

For the first case, we selected YouTube, a widely used website where specifying a LEL was not considered necessary. This example showcases the power of LLMs in interpreting mockups independently, relying solely on the mockup itself without additional context.

In this straightforward example, the user has simply added a button to enable access to statistics from the channel view, as shown in Fig. 2.

**Fig. 2.** User Creating a High-Fidelity Mockup

It's important to note that this figure highlights the tool the user employed to add a button to an existing channel interface. The new button was added simply by dragging and dropping it onto the web application. Subsequently, the user only needed to edit its properties to set the text to "stats." Through this simple gesture, the user effectively communicated the desire for a new button, its intended location, and its desired label. The mockup visually distinguishes this new button by using a different font style.

A screenshot of the high-fidelity mockup, along with a screenshot of the original version of the application before any modifications, was attached to the

prompt. The prompt used is simply: *'A YouTube user who posts videos, based on the original YouTube version (image 1), has made a high-fidelity mockup to describe a requirement (image 2). Could you make me a user story for that mockup?'*

In the Fig. 3, it can be seen how the LLM provides a very suitable user story for the provided mockup. Although it is a very simple mockup, it can be appreciated how the LLM was able to interpret the image, distinguish that there was a new element, and that the element referred to statistics. The resulting user story explicitly states that a new button is desired, that it should be located next to the existing buttons, and that a consistent design is desired, making the new button look the same as the originals.

**User Story:**

As a user who publishes videos on YouTube,  
I want to have an interface that allows me to easily access the stats of my videos,  
so I can quickly and easily monitor the performance of each video without leaving the main page  
of my channel.

**Acceptance Criteria:**

- The main page of the channel must show a new button or tab labeled "Stats."
- Clicking on "Stats" should display a summary of key metrics (such as views, watch time, likes, and subscribers gained) directly from the main interface.
- This button should be located in an accessible spot at the top of the page, next to existing options like "Manage Videos" and "Customize Channel."
- The functionality must be fast and should not interfere with the user's normal navigation on the main channel page.

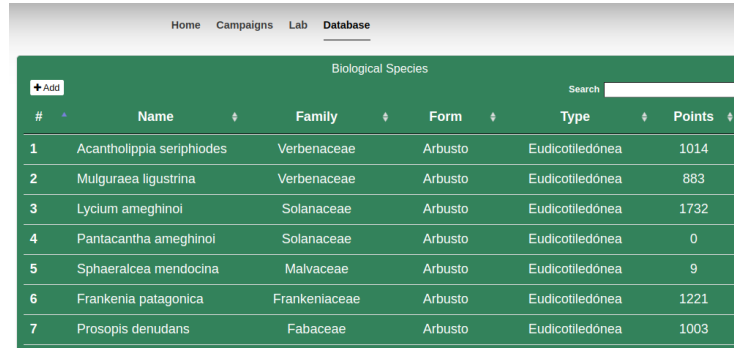
**Technical Notes:**

- The "Stats" button must have a design consistent with the rest of the YouTube interface to maintain visual coherence.
- The stats information should load without redirecting the user to a new page.

**Fig. 3.** User Creating a High-Fidelity Mockup

The second case study corresponds to an application with a very particular domain, such as the LeafLab application[1]. LeafLab is a web application used on an intranet by biological scientists specializing in botany. We have used this application to represent case studies in previous works [22].

In this example, we see how the biologist navigates the list of floristic species, as shown in Fig. 4. In this screenshot, it can be easily seen that the application offers a list of species, ordering them by the *#id* number in the application, and this piece of information is not relevant to the biologist.

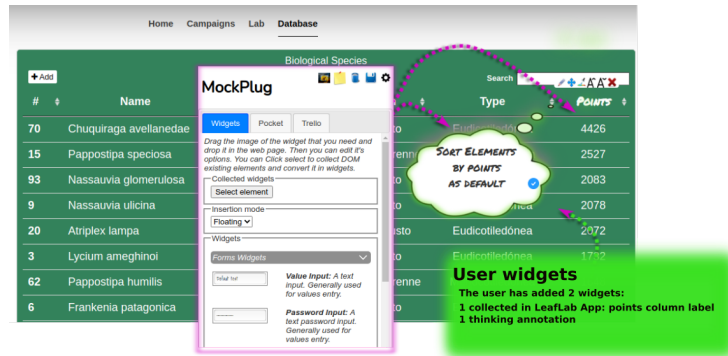


#	Name	Family	Form	Type	Points
1	Acantholippia seriphoides	Verbenaceae	Arbusto	Eudicotiledónea	1014
2	Mulguraea ligustrina	Verbenaceae	Arbusto	Eudicotiledónea	883
3	Lycium ameghinoi	Solanaceae	Arbusto	Eudicotiledónea	1732
4	Pantacantha ameghinoi	Solanaceae	Arbusto	Eudicotiledónea	0
5	Sphaeralcea mendocina	Malvaceae	Arbusto	Eudicotiledónea	9
6	Frankenia patagonica	Frankeniaceae	Arbusto	Eudicotiledónea	1221
7	Prosopis denudans	Fabaceae	Arbusto	Eudicotiledónea	1003

**Fig. 4.** LeafLab web application displaying floristic species list, ordered by system #id.

The biologist wants the species to be listed sorted by relevance in the field, that is, showing on top of the list those species that have been observed most frequently, thus generating, by default, a ranking of predominance.

To specify his need, he created a high-fidelity mockup as shown in Fig. 5. In his mockup, the species are listed by the number of points in which each species appeared in the exploration surveys.



**Fig. 5.** Biologist's mockup creation process for new LeafLab feature.

It is important to note that the word *points* is not just a simple score. For the biologist and within their specific domain, it has a very specific meaning: the listed species are those corresponding to thousands of specimens methodologically surveyed in the field during field campaigns. The word *points* refer to the fact that in each survey, the biologist traverses an imaginary straight line of 300 meters, along which are distributed 100 equidistant sampling points. At each of these points in the field, the biologist takes a sample of a floristic specimen, which must be identified to species, often collecting a sample and using a



stereoscopic microscope in the laboratory to differentiate one from another.

To achieve this, taking the web application as a canvas, the biologist selects the existing title of the *points* column using the mockup tool, converting it into an element of their mockup. Upon performing this action, that element is automatically displayed in a distinctive way from the others in the mockup, through a handwritten-style font. As can be seen in the figure, the user also adds an annotation to the mockup, in the form of a thought bubble, referencing that column, where they textually express *Sort element by points as default*.

In Fig. 6 on the left, we can observe how, when requesting the LLM to generate a user story based on the original version and the mockup, the user story, while surprisingly good, does not make precise use of the domain language. Specifically, when stating that the species should be sorted by the data in the *points* column, it does so in terms of 'score' and not 'relevance', as in the explanatory sentence, it indicates: *So that I can quickly see the most relevant or highest-scoring species without the need for manual searching or sorting*. While on the right, it illustrates how providing the LEL to the prompt enables the LLM to generate a user story that better aligns with the domain. The term 'points' no longer refers to a score but rather to the number of times a species was recorded in the surveyed field. The explanatory sentence is notably more precise and employs terminology specific to the application domain: *So that I can quickly view the most relevant species based on the number of points where they were found, without the need for manual sorting*.

#### User Story:

**Title:** Automatically Sort Biological Species Database Elements by Points by Default.

**As:** A user of the biological species database application,

**I want:** The species table to automatically sort by the points column in descending order upon accessing the page,

**So that:** I can quickly see the most relevant or highest-scoring species without the need for manual searching or sorting.

#### Acceptance Criteria:

- Upon accessing the "Database" section of the biological species, the table should automatically be sorted by the "Points" column in descending order (from highest to lowest score).
- A visual indicator should be displayed in the "Points" column to highlight that it is sorted by this metric.
- The user should have the ability to manually sort by other columns if desired, without losing the functionality of the default sorting by points upon initiation.
- The "Points" column should remain clickable to toggle between ascending and descending order if the user needs it.

#### User Story:

**Title:** Automatically Sort Biological Species by Points in the Database.

**As:** A botanist,

**I want:** The biological species discovered in the transects to be automatically sorted by the points column in descending order when accessing the database,

**So that:** I can quickly view the most relevant species based on the number of points where they were found, without the need for manual sorting.

#### Acceptance Criteria:

- Upon accessing the "Database" section of the application, the biological species are automatically sorted by the points column in descending order.
- The points column must display a visual indicator showing that it is being used for the default sort order.
- The botanist must have the option to sort the biological species by family, form, or type if desired, without losing the default sorting by points on the initial view.
- The **points column** must be selectable to toggle between ascending and descending order if the botanist requires it.

#### Involved Behaviors:

- The botanist navigates through the biological species discovered in the transects.
- The botanist sorts the biological species discovered in the transects by family, form, type, or points.

**Fig. 6.** User story for the biologist mockup, comparison without LEL (left) and with LEL (right) in the prompt.

Notably, the distinctive elements introduced by the LEL are highlighted in the figure. It is also worth mentioning that the provided lexicon was incomplete, and only sufficient for this specific use case. As shown in Table 2 and Table 3.

**Table 2.** LEL portion for Subject, Notion, Behavioral responses

Subject	Notion	Behavioral responses
Biologist	User of the application licensed or doctorate in biological sciences and specialized in botany.	The biologist orders the species discovered in the transects by family. The biologist orders the species discovered in the transects by form. The biologist orders the species discovered in the transects by type. The biologist orders the species discovered in the transects by points. The biologist edits the properties of a species.

**Table 3.** LEL portion for Objects and Notions

Object	Notion
Transect	Imaginary straight line of approximately 300 meters, on which biologists sample specimens of biological species. Each transect is composed of 300 points.
Point	At each point of the 300 points of a transect, the biologist takes a sample of a specimen of a biological species found.
Biological Species	Species found in the field campaigns carried out by biologists. Each species may have been found in 0 or more points.
Visit	A traverse of a particular transect. Each traverse of a transect involves walking approximately 300 meters, recording the species of a particular floristic specimen every 3 meters, resulting in 100 points per transect.

To preliminarily validate the proposed approach, we focus on the fundamental types of actions that a user can perform when constructing their high-fidelity mockup on the existing web application. That is, when building their mockup within the proposed approach, the user will specify their requirement based on some combination of the following possibilities:

- A. Add a new element (not previously existing).
- B. Add an existing element to the requirement.
- C. Remove existing elements.
- D. Compose elements from other parts of the application.
- E. Compose elements from other applications.

For each of these options, we conducted tests with and without the LEL, evaluating whether the LLM is capable of generating an appropriate user story for the mockup and with what level of precision it does so. Following the methodology used in the previous examples, a mockup was designed that incorporated different categories of actions, and the adequacy of the generated user story was evaluated. It should be mentioned that both the LEL and the mockups used in the validation were developed by the team responsible for the application, with experience in both the domain and the technologies involved.

Each user story was assigned a score according to the following value scale:

1. Does not distinguish any element of the requirement.
2. Distinguishes only partially the elements that compose the requirement.
3. Distinguishes the elements of the requirement, but does not describe them adequately.
4. Distinguishes the elements of the requirement and describes some of them adequately.
5. Distinguishes the elements of the requirement and describes them adequately in their entirety.

The table summarizes the results, where the score with LEL is greater than or equal to the score without LEL, indicating the preliminary importance of using the glossary LEL.

**Table 4.** Type of Requirement Scores With and Without LEL

Type of Requirement	Score without LEL	Score with LEL
A	3	5
B	3	5
C	4	4
D	3	5
E	3	5

## 5 Conclusion

The rapid advancement of language models in recent years has continuously redefined multiple processes across diverse application domains, and this trend is expected to continue. The capabilities of these new tools continue to expand, driven by rapid advances in the generative AI industry.

Although numerous powerful tools have emerged to assist software engineers, the core processes of software engineering continue to evolve alongside these advancements.

In this work, we propose the integration of high-fidelity mockups and the glossary LEL to generate user stories for the definition of requirements using an

LLM, and we provide evidence supporting its feasibility. Our results are highly encouraging: deriving user stories from mockups and leveraging a glossary LEL enable agile and precise involvement of end-users in requirement definition, while streamlining communication through existing web applications and appropriate tools. By incorporating the glossary LEL into the prompt before generating user stories, we enhance the value of domain analysis performed by requirements engineers, ensuring terminology consistency, and improving requirement clarity.

This approach has the potential to significantly impact software development workflows by reducing ambiguity in requirement definition and speeding up the initial phases of development. However, further validation is needed to assess its adaptability to diverse application domains and its robustness in handling complex requirements. Addressing potential biases in language models and refining the automation process to ensure alignment with stakeholder expectations remain key challenges.

Future work will focus on refining this integration, expanding its applicability, and conducting empirical studies to measure its effectiveness in real-world software development scenarios. Furthermore, exploring hybrid approaches that combine AI-generated content with human validation could further enhance the reliability of this method.

## References

1. Almonacid, S., Klagges, M.R., Navarro, P., Morales, L., Pazos, B., Puigbó, A.C., Firmenich, D.: Mobile and wearable computing in patagonian wilderness. In: Cloud Computing and Big Data: 7th Conference, JCC&BD 2019, La Plata, Buenos Aires, Argentina, June 24–28, 2019, Revised Selected Papers 7. pp. 137–154. Springer (2019)
2. Amazon.com Inc.: Amazon Alexa. <https://www.alexa.com/>, last access: 2025-02-17
3. Amna, A.R., Poels, G.: Ambiguity in user stories: A systematic literature review. *Information and Software Technology* **145**, 106824 (2022)
4. Antonelli, L., Lezoche, M., Delle Ville, J.: Knowledge extraction from the language extended lexicon glossary using natural language processing. *Tecnológicas* (59), 2 (2024)
5. Antonelli, L., Rossi, G., do Prado Leite, J.C.S., Oliveros, A.: Deriving requirements specifications from the application domain language captured by language extended lexicon. In: WER (2012)
6. Apple Inc.: Siri. <https://www.apple.com/siri/>, last access: 2025-02-17
7. Belzner, L., Gabor, T., Wirsing, M.: Large language model assisted software engineering: prospects, challenges, and a case study. In: International Conference on Bridging the Gap between AI and Reality. pp. 355–374. Springer (2023)
8. DeepL GmbH: DeepL Translator. <https://www.deepl.com/en/translator>, last access: 2025-02-17
9. Devlin, J.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
10. Filipović, M., Vuković, Ž., Dejanović, I., Milosavljević, G.: Rapid requirements elicitation of enterprise applications based on executable mockups. *Applied Sciences* **11**(16), 7684 (2021)

11. Firmenich, D., Firmenich, S., Rivero, J.M., Antonelli, L.: A platform for web augmentation requirements specification. In: Web Engineering: 14th International Conference, ICWE 2014, Toulouse, France, July 1-4, 2014. Proceedings 14. pp. 1–20. Springer (2014)
12. Firmenich, D.A., Morales, L., Mura, G., Calfuquir, N.: Mockplug: A high-fidelity mocking tool for plugging functional requirements into existing web applications. *Computer software and media applications* **7**(1), 6736 (Sep 2024). <https://doi.org/10.24294/csma.v7i1.6736>
13. Geniusee Inc.: Generative AI development services. <https://geniusee.com/generative-ai-development>, last access: 2025-02-17
14. Google LLC.: Gemini. <https://gemini.google.com/>, last access: 2025-02-17
15. Google LLC.: Google Translate. <https://translate.google.com>, last access: 2025-02-17
16. Hussain, A., Mkpojiogu, E.O., Kamal, F.M.: The role of requirements in the success or failure of software projects. *International Review of Management and Marketing* **6**(7), 306–311 (2016)
17. Jin, H., Huang, L., Cai, H., Yan, J., Li, B., Chen, H.: From llms to llm-based agents for software engineering: A survey of current, challenges and future. *arXiv preprint arXiv:2408.02479* (2024)
18. Kolthoff, K., Kretzer, F., Bartelt, C., Maedche, A., Ponzetto, S.P.: Interlinking user stories and gui prototyping: A semi-automatic llm-based approach. *arXiv preprint arXiv:2406.08120* (2024)
19. Leite, J.d.P., Franco, A.P.M.: A strategy for conceptual model acquisition. In: [1993] Proceedings of the IEEE International Symposium on Requirements Engineering. pp. 243–246. IEEE (1993)
20. Li, J., Li, D., Savarese, S., Hoi, S.: Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In: International conference on machine learning. pp. 19730–19742. PMLR (2023)
21. Lucassen, G., Dalpiaz, F., Werf, J.M.E.v.d., Brinkkemper, S.: The use and effectiveness of user stories in practice. In: Requirements Engineering: Foundation for Software Quality: 22nd International Working Conference, REFSQ 2016, Gothenburg, Sweden, March 14-17, 2016, Proceedings 22. pp. 205–222. Springer (2016)
22. Marticorena, L.G., Morales, L.A., Antonelli, L., Rossi, G., Firmenich, D.: Development iterations based on web augmentation and
23. Neto, J.D.M.S., do Prado Leite, J.C.S., Cysneiros Filho, L.M.: Non-functional requirements for object-oriented modeling. In: WER. pp. 109–125 (2000)
24. Onilab Inc.: LLM training and development. <https://onilab.com/services/llm-training-and-development>, last access: 2025-02-17
25. OpenAI: ChatGPT. <https://openai.com/index/chatgpt/>, last access: 2025-02-17
26. do Prado Leite, J.C.S., Franco, A.P.M.: A strategy for conceptual model acquisition. In: Proceedings of IEEE International Symposium on Requirements Engineering, RE 1993, San Diego, California, USA, January 4-6, 1993. pp. 243–246. IEEE Computer Society (1993). <https://doi.org/10.1109/ISRE.1993.324851>, <https://doi.org/10.1109/ISRE.1993.324851>
27. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research* **21** (10 2019), <https://arxiv.org/abs/1910.10683v4>
28. Ravid, A., Berry, D.M.: A method for extracting and stating software requirements that a user interface prototype contains. *Requirements Engineering* **5**, 225–241 (2000)

29. Rehkopf, M.: User stories with examples and a template. Atlassian. URL: <https://www.atlassian.com/agile/project-management/user-stories> [accessed 2021-12-03] (2021)
30. Samir, M., Elsayed, A., Marie, M.I.: A model for automatic code generation from high fidelity graphical user interface mockups using deep learning techniques. *International Journal of Advanced Computer Science & Applications* **15**(3) (2024)
31. Urbieto, M., Torres, N., Rivero, J.M., Rossi, G., Dominguez-Mayo, F.J.: Improving mockup-based requirement specification with end-user annotations. In: *Agile Processes in Software Engineering and Extreme Programming: 19th International Conference, XP 2018, Porto, Portugal, May 21–25, 2018, Proceedings* 19. pp. 19–34. Springer International Publishing (2018)
32. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Łukasz Kaiser, Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems* **2017-December**, 5999–6009 (6 2017), <https://arxiv.org/abs/1706.03762v7>, intentemos citarlo TESINA
33. Wehbe, R.: The language extended lexicon, revisited. In: *XII Argentine Symposium on Software Engineering (ASSE 2011)(XL JAIIO, Córdoba, 1º y 2 de septiembre de 2011)* (2011)
34. Wei, B.: Requirements are all you need: From requirements to code with llms. *arXiv preprint arXiv:2406.10101* (2024)
35. Wimmer, C., Untertrifaller, A., Grechenig, T.: Sketchinginterfaces: a tool for automatically generating high-fidelity user interface mockups from hand-drawn sketches. In: *Proceedings of the 32nd Australian Conference on Human-Computer Interaction*. pp. 538–545 (2020)
36. Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., Fidler, S.: Aligning books and movies: Towards story-like visual explanations by watching movies and reading books (6 2015)