



ChatGPT as a Tool for User Story Quality Evaluation: Trustworthy Out of the Box?

Downloaded from: <https://research.chalmers.se>, 2026-05-08 12:19 UTC

Citation for the original published paper (version of record):

Ronanki, K., Cabrero-Daniel, B., Berger, C. (2024). ChatGPT as a Tool for User Story Quality Evaluation: Trustworthy Out of the Box?. *Lecture Notes in Business Information Processing*, 489 LNBIP: 173-181. http://dx.doi.org/10.1007/978-3-031-48550-3_17

N.B. When citing this work, cite the original published paper.



ChatGPT as a Tool for User Story Quality Evaluation: Trustworthy Out of the Box?

Krishna Ronanki^(✉), Beatriz Cabrero-Daniel, and Christian Berger

University of Gothenburg, Gothenburg, Sweden
{krishna.ronanki,beatriz.cabrero-daniel,christian.berger}@gu.se

Abstract. In Agile software development, user stories play a vital role in capturing and conveying end-user needs, prioritizing features, and facilitating communication and collaboration within development teams. However, automated methods for evaluating user stories require training in NLP tools and can be time-consuming to develop and integrate. This study explores using ChatGPT for user story quality evaluation and compares its performance with an existing benchmark. Our study shows that ChatGPT’s evaluation aligns well with human evaluation, and we propose a “best of three” strategy to improve its output stability. We also discuss the concept of trustworthiness in AI and its implications for non-experts using ChatGPT’s unprocessed outputs. Our research contributes to understanding the reliability and applicability of Generative AI in user story evaluation and offers recommendations for future research.

Keywords: ChatGPT · User Stories · Quality · Agile

1 Introduction

In agile software development projects, user stories are one of the most widely used notation to express requirements [1]. They are considered a very granular representation of requirements that developers use to build new features [2] as they help to capture & communicate end-user needs to prioritize & deliver small, working features in each development cycle [3].

The quality of user stories is crucial to the success of a development project as they impact the quality of the system design which, in turn, affects the final product [4]. They provide clear guidance for development efforts, improve communication and collaboration within teams, and help to ensure that development teams have a shared understanding of what needs to be delivered [5].

However, evaluating the quality of user stories manually can be time-consuming. One potential solution for improving agile software development processes is the integration of automated methods. This can be accomplished through modifications to existing workflows and the implementation of evaluation tools [6]. Existing methods for automatically evaluating user stories can be relatively fast and efficient, especially when compared to the time required for human evaluation [7]. Natural Language Processing (NLP) has been identified

as a potential method for evaluating various aspects of user stories. However, the accuracy and effectiveness of this method can be influenced by factors such as the quality of the training data and the complexity of the user stories under evaluation [8]. Unfortunately, the process of developing and incorporating automated methods for evaluating user stories can be a time-intensive endeavour due to the necessity of training NLP tools to accurately construct algorithms [9].

Developers are increasingly exploring the use of standalone general-purpose applications such as ChatGPT to aid in their software development endeavours. ChatGPT, based on the GPT-3.5 language model, is optimized for dialogue and is capable of answering questions in a human-like text [10].

Despite being trained on a large general-purpose corpus and specifically fine-tuned for conversational tasks [11], it has been observed to perform surprisingly well on specific technical tasks [12]. For this study, we investigated how well a general-purpose large language model like ChatGPT performs in evaluating the quality of user stories.

2 Background

The user story technique is a widely used approach for expressing requirements by utilizing a template that consists of the following elements: “As a (role), I want (goal), so that (benefit)” [3]. The primary components of a requirement that are captured by user stories are: who is it for, what it expects from the system, and, optionally, why it is important [3]. We follow this user story structure in our study while using the few-shot prompting technique to evaluate the user story quality using ChatGPT.

Few-shot prompting is a technique where the model is provided with a small number of examples of the task as conditioning in the initial prompt [13]. It refers to the ability of language models to learn a new task with limited training samples provided by the user [14, 15]. We used this prompting technique to provide an example to ChatGPT of what a user story should look like structurally before asking it to evaluate the user story on the defined criteria.

The user story quality criteria we used in our study were presented by Lucassen et al. [16] in their work which focuses on proposing a holistic approach for ensuring the quality of agile requirements expressed as user stories. The approach is comprised of two components: (i) the QUS framework, which is a collection of 13 criteria that can be applied to a set of user stories to assess the quality of individual stories and the set, and (ii) the AQUASA software tool, which utilizes state-of-the-art NLP techniques to automatically detect violations of a selection of the quality criteria in the QUS framework. Tõemets’ work investigates whether it is feasible to predict the quality of user stories for monitoring purposes and to determine the correlation between user story quality and other aspects of software development [17]. The user stories we chose to evaluate as part of our study and the benchmark evaluation scores of the selected user stories using the AQUASA tool were also included in this work.

3 Method

In our study, we performed a comparative analysis of manual and automated evaluation of user stories. Firstly, we assessed the quality of user stories manually, and then we employed ChatGPT for the same task. Our aim was to determine the effectiveness of ChatGPT in evaluating user stories and to compare its performance with human evaluation (Fig. 1).

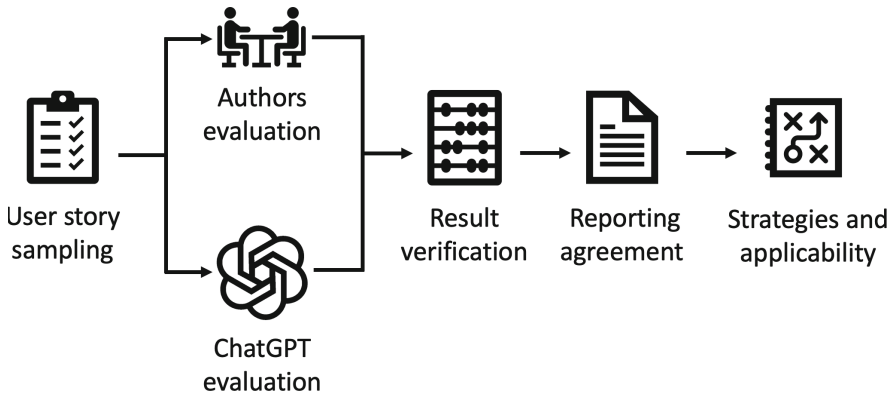


Fig. 1. Methodology and verification plan (during workshop)

To assess the ability of ChatGPT to replicate human evaluations of user stories, an open-source database presented in Töemets [17]¹ was selected as it came with benchmark evaluation of the user stories using the AQUSA tool, which also allows us to refer to an accepted baseline. After retrieving the benchmark, we conducted a double-blinded manual evaluation of the randomly selected set of user stories to assess their quality in terms of atomicity, well-formedness, minimality, conceptual soundness, unambiguity, completeness (full sentence or not), and estimability. The sole criterion for selection was the presence of a benchmark evaluation established using the AQUSA tool. However, the evaluation done by AQUSA presented in their study focuses on appraising only the following aspects: atomicity, well-formedness, and minimality.

To evaluate the performance of ChatGPT (March 23 version) for user story quality evaluation, we conducted a series of tests using the one-shot prompting method [15], a variation of the few-shot prompting method. Specifically, we presented ChatGPT with a set of criteria, user story pairs and recorded its responses. To ensure the reliability and consistency of ChatGPT's performance, we repeated this process three times. The evaluation was carried out based on seven criteria presented by Lucassen et al. [16].

¹ Visit <https://github.com/TanelToemets/Analysing-The-Quality-Of-User-Stories-In-Agile-Software-Projects>.

Finally, we compare the data from the human evaluation, the AQUUSA tool benchmark evaluations and the ChatGPT evaluation and present our findings as tables. The comparison was done for each of the seven criteria and for the overall precision, recall, specificity, and F1 score. The comparison was made to identify any significant differences between the two tools and to ascertain the accuracy of ChatGPT in replicating human evaluation.

The results of our experiments raise important issues related to the usability and transparency of ChatGPT's outputs, particularly for non-expert users. In this regard, the discussion section of our paper highlights the need to carefully consider the trustworthiness of ChatGPT's raw outputs and the importance of ensuring that users have the necessary tools to understand and interpret them correctly. By addressing these concerns, we can enhance the usability and effectiveness of ChatGPT as a tool for supporting decision-making in a variety of contexts.

3.1 Threats to Validity

Validity threats can arise in the benchmark creation process due to the limited scope of evaluation criteria used. The authors of the AQUUSA tool evaluated only the atomic, well-formed, and minimal criteria for the sampled user stories. Furthermore, there are concerns about the reliability and accuracy of the evaluation data since it was not provided by the AQUUSA authors themselves, but from a master thesis based on Lucasen et al. [16]. Another potential threat to validity is the use of human raters who may not have been experienced practitioners, thus leading to concerns about their reliability. To mitigate this, an independent rating of user stories was conducted, and in case of disagreement, a consensus was reached through a meeting. Moreover, ChatGPT was tested only three times, and further testing could yield different results. Nonetheless, we argue that this is sufficient to assert that developers cannot blindly trust ChatGPT's outputs and integrate them into their agile software development process. This finding emphasizes the need for cautious and careful consideration when incorporating natural language processing (NLP) tools like ChatGPT into agile development practices.

4 Results and Analysis

4.1 Comparing the Evaluations to the AQUUSA Benchmark

The AQUUSA benchmark comprises three key criteria for assessing the quality of a story. The first criterion is whether the story is well-formed, which means it includes a role and the expected functionality, commonly referred to as the means. The second criterion is whether the story is atomic, which implies that it addresses only one feature. The third criterion is whether the story is minimal, which requires that it contains a role, a means, and one or more ends [16].

Of all pairs {criteria, user story}, results showed that human evaluators and AQUUSA agreed on only **55%** of the pairs {criteria, user story} as reported in

Table 1, indicating a moderate level of agreement between the two methods. Human evaluators and AQUASA were in agreement in identifying well-formed and atomic user stories in a majority of cases (81.82% and 63.64%, respectively). However, the agreement rate between the two parties dropped significantly when it came to identifying minimal user stories, with only 18.18% agreement observed. The findings of the study indicate that the AQUASA tool exhibits a moderate level of concurrence with human evaluators when it comes to detecting user stories that are well-constructed and atomic in nature, but it currently falls short in identifying minimal user stories.

To enable a fair comparison of results, we conducted evaluations of the same user stories using ChatGPT. The evaluations were performed using two distinct accounts with the history log being cleared between each evaluation. We repeated this process thrice to account for any instability in the results. Table 1 displays the results of three evaluations conducted to assess the agreement rate and F1 scores of ChatGPT. The findings reveal that ChatGPT demonstrated a consistent agreement rate throughout the evaluations. Furthermore, the F1 scores recorded during the assessments ranged from 81% to 86%.

Table 1. Percentage of agreement, rounded to 2 decimals, between human evaluations and two tools, AQUASA and ChatGPT, across 11 randomly sampled user stories

Type	Metric	AQUASA	ChatGPT		
		Benchmark	Run 1	Run 2	Run 3
Criteria	Well-formed	81.82%	81.82%	81.82%	81.82%
	Atomic	63.64%	63.64%	90.91%	90.91%
	Minimal	18.18%	81.82%	72.73%	54.55%
Aggrega.	Agreement rate	54.55%	75.76%	81.82%	75.76%
	Precision	62.50%	80.95%	85.71%	74.07%
	Recall	71.43%	80.95%	85.71%	95.24%
	Specificity	25.00%	66.67%	75.00%	41.67%
	F1 score	66.67%	80.95%	85.71%	83.33%

4.2 ChatGPT-Human Agreement Rate

While performing the evaluation using ChatGPT, measures were taken to cover the rest of the metrics described in Dalpiaz et al. [16]. In terms of agreement rate with human evaluators, ChatGPT’s performance was relatively stable across the three assessments, as reported in Table 2, with agreement rates ranging from 73% to 75%. However, this suggests that there is room for improvement in the agreement rate between ChatGPT and human evaluators. A 25% error rate may be problematic in certain situations. As a result, enhancing ChatGPT’s performance could increase its reliability and effectiveness in various applications.

Table 2. Agreement with human evaluations of ChatGPT, across 11 randomly sampled user stories, using different strategies to interpret the output

Type	Metric	ChatGPT			Interpretation strategy		
		Run 1	Run 2	Run 3	AL1	BO3	PA3
Criteria	Well-formed	81.82%	81.82%	81.82%	81.82%	81.82%	88.89%
	Atomic	63.64%	90.91%	90.91%	72.73%	90.91%	100.00%
	Minimal	81.82%	72.73%	54.55%	45.45%	81.82%	100.00%
	Conceptually sound	90.91%	90.91%	63.64%	63.64%	81.82%	100.00%
	Unambiguous	54.55%	54.55%	63.64%	72.73%	54.55%	62.50%
	Full sentence	45.45%	63.64%	81.82%	81.82%	63.64%	80.00%
	Estimable	90.91%	72.73%	81.82%	72.73%	81.82%	88.89%
Aggrega.	Agreement rate	72.73%	75.32%	74.03%	70.13%	76.62%	87.23%
	Precision	83.33%	82.69%	77.05%	72.06%	81.82%	94.74%
	Recall	75.47%	81.13%	88.68%	92.45%	84.91%	90.00%
	Specificity	66.67%	62.50%	41.67%	20.83%	58.33%	71.43%
	F1 score	79.21%	81.90%	82.46%	80.99%	83.33%	92.31%
	Coverage	100.00%	100.00%	100.00%	100.00%	100.00%	61.04%

The agreement rates reported in Tables 1 and 2 include both true positives, where human raters and tools agreed on an overall positive evaluation of a user story, and true negatives, where humans and the tools agreed on a negative evaluation. Future work, though, might look into database entries where human raters and ChatGPT do not agree on the evaluations.

4.3 How to Select an Answer Based on ChatGPT’s Output

In our study, we evaluated the consistency and reliability of ChatGPT in evaluating user stories against a set of predetermined criteria. Our results indicate that ChatGPT was consistent with itself in 61% of the evaluations, meaning it gave the same response for a given pair {criteria, user story} in three separate runs (PA3). Furthermore, ChatGPT agreed with itself in at least two out of three runs in 83.11% of the evaluations. These findings suggest that ChatGPT’s evaluations are relatively stable when it comes to evaluating user stories. Moreover, we observed that in the subset of evaluations where ChatGPT was consistent across all three runs, the agreement rate with humans was 87%, with precision and recall scores of 95% and 90%, respectively.

The higher rate of agreement between humans and ChatGPT has encouraged us to explore stricter criteria for identifying positive responses. The use of a “best of three” approach in which ChatGPT is required to give a positive response in at least two out of three attempts resulted in a slightly higher agreement rate between humans and ChatGPT, reaching (77%). However, when responses from ChatGPT were classified as positive if at least one positive response was given (AL1), there was more variability in ChatGPT’s responses, leading to a decrease in the agreement rate to 70%.

5 Discussion

The agreement rate remained constant across the first three rounds, as evidenced by Tables 1 and 2. Thus, we opted to conclude our testing after these three runs. However, to establish the reliability of these initial findings, additional evaluations of new user stories and more ChatGPT runs are necessary.

Table 2 shows that the criteria with the highest and lowest agreement rates differed in the three runs, which suggests that certain criteria may have unclear definitions or abstract qualities that made it difficult for ChatGPT to consistently agree with human evaluators. To better understand this discrepancy, further research could examine the specific criteria that posed challenges for ChatGPT or where it exhibited inconsistencies.

Although ChatGPT's consistency in generating responses may correlate with the level of agreement from human evaluators, it is important to note that consistency does not necessarily equate to the accuracy or appropriateness of the output. Additionally, the consistency could be due to potential bias present in the training data.

However, integrating ChatGPT into agile software development requires a thorough assessment of its capabilities, strengths, and limitations. While ChatGPT has shown promise in this task, its performance is not flawless, and it remains an emerging technology that is susceptible to potential biases and limitations. Therefore, careful consideration of ChatGPT's applicability and limitations is necessary before its integration into agile software development processes [18]. On the other hand, GPT-4's expanded architectural model size might play a pivotal role in enhancing its proficiency in NLP, which could lead to increased accuracy and relevance in the generated responses [19].

However, a major obstacle to using ChatGPT in the requirements elicitation process is the issue of extrinsic hallucinations [20]. Non-experts who rely on AI systems might not possess the technical knowledge to evaluate the accuracy and reliability of the generated outputs in some cases. This issue highlights the importance of ensuring the trustworthiness of the AI systems being implemented. Ensuring trustworthiness in AI, particularly in the context of non-experts using ChatGPT for user story evaluation, requires careful consideration of several factors including transparency, explainability, bias mitigation, and continuous improvement through user feedback to be incorporated into the development and implemented process of these AI systems in such human-centric processes.

6 Conclusion and Future Work

The study examines the effectiveness of ChatGPT in assessing the quality of user stories, particularly when it produces consistent results across multiple evaluations. The research focuses on the agreement rate between humans and ChatGPT in evaluating user stories based on said criteria. The results indicate that ChatGPT is more capable of replicating human evaluation (approximately 75%) than the AQUASA tool as demonstrated in Tõemets [17].

While the model performs sufficiently well in independent runs, it exhibits inconsistency in its Boolean outputs when tested multiple times. This suggests caution in interpreting its evaluations and underlines the need for further research into the factors affecting ChatGPT's consistency and reliability.

To address the issue of unstable outputs, the paper suggests strategies such as selecting the “best of three” approach. However, the question of whether ChatGPT's raw outputs can be used directly by non-expert users raises important concerns about the trustworthiness of AI systems. As a result, high-level trustworthiness requirements must be established to ensure that ChatGPT and other AI tools are integrated into Agile software development processes following trustworthy AI principles. The integration of ChatGPT, into agile software development processes, requires careful consideration of its limitations and strengths and the potential impact on the development process. Further research is needed to explore ways to ensure that ChatGPT and other AI systems can be used reliably and effectively in Agile development environments.

References

1. Lucassen, G., Dalpiaz, F., Van Der Werf, J.M.E., Brinkkemper, S.: Forging high-quality user stories: towards a discipline for agile requirements. In: IEEE International Requirements Engineering Conference (RE), pp. 126–135. IEEE (2015)
2. Lucassen, G., Dalpiaz, F., Werf, J.M.E.M., Brinkkemper, S.: The use and effectiveness of user stories in practice. In: Daneva, M., Pastor, O. (eds.) REFSQ 2016. LNCS, vol. 9619, pp. 205–222. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30282-9_14
3. Cohn, M.: User Stories Applied: For Agile Software Development. Addison-Wesley Professional, Boston (2004)
4. Amna, A.R., Poels, G.: Systematic literature mapping of user story research. IEEE Access **10**, 51723–51746 (2022)
5. Mustaffa, S.N.F.N.B., Sallim, J.B., Mohamed, R.B.: Enhancing high-quality user stories with AQUASA: an overview study of data cleaning process. In: 2021 International Conference on Software Engineering & Computer Systems and 4th International Conference on Computational Science and Information Management (ICSECS-ICOCSIM), pp. 295–300 (2021)
6. Humayoun, S.R., Dubinsky, Y., Catarci, T.: User evaluation support through development environment for agile software teams. In: Caporarello, L., Di Martino, B., Martinez, M. (eds.) Smart Organizations and Smart Artifacts. LNISO, vol. 7, pp. 183–191. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-07040-7_18
7. Jurisch, M., Lusky, M., Iglar, B., Böhm, S.: Evaluating a recommendation system for user stories in mobile enterprise application development. Int. J. Adv. Intell. Syst. (2017)
8. Peña, F.J., Roldán, L., Vegetti, M.: User stories identification in software's issues records using natural language processing. In: 2020 IEEE Congreso Bional de Argentina (ARGENCON), pp. 1–7 (2020)
9. Sharir, O., Peleg, B., Shoham, Y.: The cost of training NLP models: a concise overview. [arXiv:2004.08900](https://arxiv.org/abs/2004.08900) (2020)
10. Zhang, B., Ding, D., Jing, L.: How would stance detection techniques evolve after the launch of ChatGPT? arXiv preprint: [arXiv:2212.14548](https://arxiv.org/abs/2212.14548) (2022)

11. Shen, Y., et al.: ChatGPT and other large language models are double-edged swords (2023)
12. Choi, J.H., Hickman, K.E., Monahan, A., Schwarcz, D.: ChatGPT goes to law school. Available at SSRN (2023)
13. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al.: Language models are unsupervised multitask learners. *OpenAI Blog* **1**(8), 9 (2019)
14. Perez, E., Kiela, D., Cho, K.: True few-shot learning with language models. In: *Advances in Neural Information Processing Systems*, vol. 34, pp. 11054–11070 (2021)
15. Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G.: Pre-train, prompt, and predict: a systematic survey of prompting methods in natural language processing. *ACM Comput. Surv.* **55**(9), 1–35 (2023)
16. Lucassen, G., Dalpiaz, F., van der Werf, J.M.E., Brinkkemper, S.: improving agile requirements: the quality user story framework and tool. *Requirements Eng.* **21**, 383–403 (2015)
17. Tõemets, T.: Analysing the quality of user stories in open source projects. PhD thesis, University of Tartu (2020)
18. Borji, A.: A categorical archive of ChatGPT failures (2023)
19. Koubaa, A.: GPT-4 vs. GPT-3.5: a concise showdown. Available in <https://doi.org/10.36227/techrxiv.22312330> (2023)
20. Bang, Y., et al.: A multitask, multilingual, multimodal evaluation of ChatGPT on reasoning, hallucination, and interactivity. arXiv preprint: [arXiv:2302.04023](https://arxiv.org/abs/2302.04023) (2023)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

