

# Investigating the Performance of Small Language Models in Detecting Test Smells in Manual Test Cases

Keila Lucas  
Federal Univ. of Campina Grande  
Campina Grande, PB, Brazil  
keila.santos@copin.ufcg.edu.br

Rohit Gheyi  
Federal Univ. of Campina Grande  
Campina Grande, PB, Brazil  
rohit@dsc.ufcg.edu.br

Márcio Ribeiro  
Federal University of Alagoas  
Maceió, AL, Brazil  
marcio@ic.ufal.br

Fabio Palomba  
University of Salerno  
Salerno (SA), Italy  
fpalomba@unisa.it

Luana Martins  
University of Salerno  
Salerno (SA), Italy  
lalmeidamartins@unisa.it

Elvys Soares  
Federal Institute of Alagoas  
Maceió, AL, Brazil  
elvys.soares@ifal.edu.br

## ABSTRACT

Manual testing, in which testers follow natural language instructions to validate system behavior, remains crucial for uncovering issues not easily captured by automation. However, these test cases often suffer from test smells, quality issues such as ambiguity, redundancy, or missing checks that reduce test reliability and maintainability. While detection tools exist, they typically require manual rule definition and lack scalability. This study investigates the potential of Small Language Models (SLMs) for automatically detecting test smells. We evaluate Gemma3, Llama3.2, and Phi-4 on 143 real-world Ubuntu test cases, covering seven types of test smells. Phi-4 achieved the best results, reaching a *pass@2* of 97% in detecting sentences with test smells, while Gemma3 and Llama3.2 reached approximately 91%. Beyond detection, SLMs autonomously explained issues and suggested improvements, even without explicit prompt instructions. They enabled low-cost, concept-driven identification of diverse test smells without relying on extensive rule definitions or syntactic analysis. These findings highlight the potential of SLMs as efficient tools that preserve data privacy and can improve test quality in real-world scenarios.

## KEYWORDS

Test Smells, Small Language Models, Manual Testing.

## 1 Introduction

Manual testing is a widely used software verification technique in which developers execute predefined test steps to detect defects in the system [6, 24]. These steps are typically written in natural language and provide instructions on how to interact with the system and assess its behavior. Manual testing remains essential for identifying issues that are difficult to detect using automated tools [7].

Despite its importance, manual testing often suffers from quality and reliability issues. Test descriptions, written in natural language, are frequently created outside formal software engineering practices [7]. As a result, they may contain test smells—issues that make test cases harder to maintain, ambiguous, redundant, overly complex or insufficient for proper verification [7, 10–12, 24]. A common example of a test smell

is the *Ambiguous Test*, in which vague instructions are given, such as “*Select a keyboard layout and click Continue*”. There is no specification of which layout should be selected – the tester might assume that any layout is acceptable, that the default layout should be chosen, or that a specific layout is implicitly required – which leads to multiple interpretations.

Improving the clarity and specificity of such instructions is a complex task. Furthermore, existing tools aimed at detecting these smells in natural language have limited scope and require significant implementation effort, often relying on manually defined rules or extensive keyword specifications. Recent studies [21, 24] have evaluated issues in manual testing, such as ambiguous descriptions, excessive complexity, actions without checks, or inadequate checks, all of which can severely hinder the effectiveness of the testing process. Soares *et al.* [25] show that test smells can reduce the efficiency of manual test execution. To improve the quality of the testing process, it is essential to detect problems that affect the understanding of the checks, which can negatively impact the maintainability of the software.

In light of the evolution of foundation models, Small Language Models (SLMs) have been gaining prominence in Natural Language Processing tasks [14, 28]. SLMs offer a lightweight, cost-effective, and customizable alternative to Large Language Models (LLMs), which, while powerful, are resource-intensive and less practical for specialized tasks. Moreover, SLMs can be fine-tuned locally, enhancing data privacy and control. From a testing perspective, SLMs offer a promising and accessible solution for automating the detection of test smells in manual test descriptions. However, it remains unclear to what extent SLMs can contribute to the identification of test smells.

In this study, we explore the use of SLMs to detect seven types of test smells in manual test descriptions [3], aiming to improve the quality and effectiveness of manual testing processes through a more accessible and efficient solution. We evaluate Gemma3 (4.3B), Llama3.2 (3.21B), and Phi-4 (14.7B) models with respect to their ability to detect issues in natural language. Our dataset consists of 261 sentences extracted from 143 manual test cases of the Ubuntu operating system. The performance of each model is assessed based on accuracy and *pass@2* for identifying each type of smell,

focusing on ambiguity, incomplete checks, unnecessary complexity, and redundancy. The results demonstrate that SLMs are capable of effectively detecting test smells in manual test descriptions. Among the evaluated models, Phi-4 achieved the best performance, identifying 216 sentences with test smells, exhibiting the lowest performance variation across different temperatures, and reaching a *pass@2* of 97%. In contrast, the Gemma3 and Llama3.2 models performed significantly worse, with *pass@2* scores of 91% and 90%, respectively.

The results show that SLMs can effectively detect test smells in manual test descriptions and, in some cases, even suggest refactorings without explicit instructions. Their use enables easier detection of various types of smells without extensive syntactic rule definitions, while also offering low computational cost and preserving data privacy. Expanding detection to other types of smells – even in different languages – can be easily done by incorporating their conceptual definitions. SLMs are a promising solution for improving manual testing and demonstrate practical value in real-world scenarios.

## 2 Research Methodology

The *goal* of this study is to evaluate the performance of the models Gemma3, Llama3.2, and Phi-4 in detecting test smells in manual test cases, from the perspective of researchers, and considering smell types previously cataloged in the literature. We investigate the following **research questions (RQs)**:

- RQ<sub>1</sub>** To what extent can Gemma3 detect natural language test smells?
- RQ<sub>2</sub>** To what extent can Llama3.2 detect natural language test smells?
- RQ<sub>3</sub>** To what extent can Phi-4 detect natural language test smells?

To answer our research questions, we selected test case sentences from the Ubuntu manual, each annotated with one of the seven known smells in natural language tests. We applied Directional Stimulus Prompting (DSP) [13] to the SLMs and evaluated each model’s responses in terms of correctness and coherence using the *pass@k* and semantic consistency metrics.

**Models selection.** We selected Llama3.2 because it is more compact than the latest version, Llama3.3. The 3B model is lightweight, and effective in multilingual text generation [18]. The Gemma3 and Phi-4 models are recent models that achieve robust performance relative to their size, especially in reasoning-focused benchmarks [1, 26]. The Phi-4 model provides high performance in NLP tasks, with the advantage of being easily adaptable and implementable on local devices [1].

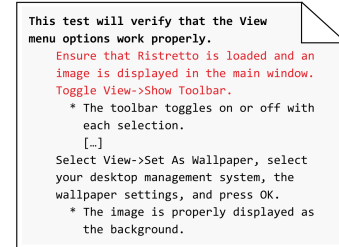
**Model setup.** The SLMs selected for the study, the Gemma3, Llama3.2, and Phi-4 models, were executed using the Ollama framework [9], with all default parameters initialized, except for temperature. The temperature is a hyperparameter that controls the randomness of the model’s output [2, 20], we evaluated the detection accuracy of the models across 11 temperature settings, which were adjusted within the range from 0.0 to 1.0 (with an increase of 0.1).

**Table 1: Sample occurrences per test smell.**

Test Smell	Sample
Ambiguous Test	90
Conditional Test	9
Eager Action	69
Misplaced Action	8
Misplaced Precondition	5
Misplaced Verification	16
Unverified Action	64
<b>TOTAL</b>	<b>261</b>

**Dataset selection.** We selected 261 sentences in natural language extracted from 143 manual test cases of the Ubuntu Operating System [4, 27]. We chose manual test cases from Ubuntu due to its popularity and the large number of open tests that validate various functionalities and behaviors of the system. Each sentence has been manually curated by Aranda *et al.* [3] and represents a set of system verifications and expected results to guide the tester during the execution of the tests. The reference catalog defined by Aranda *et al.* [3] presents seven types of test smells in natural language, as presented in Table 1.

In the test case *1600\_Ristretto* (Figure 1), we highlight one example sentence that contains the *Eager Action* test smell, as multiple actions are combined into a single instruction without ensuring that each action is completed before proceeding to the next. If Ristretto is not loaded or an image is not displayed, attempting to toggle the toolbar may fail, and the test might not correctly identify the cause of the error.



**Figure 1: Ubuntu test case 1600\_Ristretto.**

**Prompt design.** The queries to the models were guided by DSP. It is a prompting technique designed to optimize the LLM’s response in generating the desired outcome. In this approach, we provide a stimulus or clue related to the specific concept of test smell that we aim to detect in the test sentences. In DSP, we structure the prompt by presenting the concept of the test smell, a guiding instruction, and the test sentence for analysis. For each type of test smell, we adjust the concept and the question, as shown in Figure 2. To guide the models in detecting test smells, we included the conceptual definition of each type in the prompt, in English, as presented in Table 2. All responses generated by the models are available in the study’s artifacts [15].

The models were implemented and executed via Google Colab, using a T4 GPU and Python. The dataset of sentences

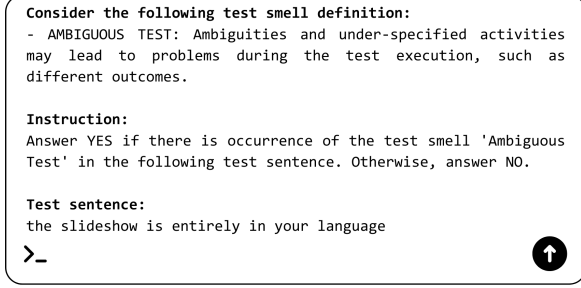


Figure 2: The Directional Stimulus Prompting.

Table 2: Test Smell Definitions.

Test Smell	Definition
Ambiguous Test	The smell occurs when ambiguities and under-specified activities may lead to problems during the test execution, such as different outcomes.
Conditional Test	The smell occurs when conditionals introduce uncertainty and variability to the test.
Eager Action	The smell occurs when a single step groups multiple actions.
Misplaced Action	The smell occurs when an action is placed in the verification list.
Misplaced Precondition	This smell occurs when a precondition is placed as the first action of the test.
Misplaced Verification	The smell occurs when there is a verification in the actions list.
Unverified Action	The smell occurs when there is no verification for a given action.

was loaded from a CSV file with the help of the Pandas library. All queries to the SLMs were performed in March 2025, and the models were downloaded between March 11<sup>th</sup> and 14<sup>th</sup>, as detailed in Table 3.

Table 3: SLMs Setup.

	Gemma3:4B	Llama3.2:3B	Phi-4:14B
Arch	gemma3	llama	phi3
Parameters	4.3B	3.21B	14.7B
Quantization	Q4_K_M	Q4_K_M	Q4_K_M
Updated	Mar/2025	Out/2024	Jan/2025
Download	Mar/14/2025	Mar/11/2025	Mar/13/2025

**Test smell detection workflow.** The test smell detection process was developed in four phases, as illustrated in Figure 3. First, we selected 261 natural language sentences extracted from manual tests of the Ubuntu operating system, which were manually reviewed [3]. Next, the test sentences were submitted to the SLMs using DSP. Each model processed the prompt and generated a textual response indicating whether or not it detected the specific test smell in the given sentence. The models justified their answers based on the definition provided in the prompt. For each sentence, we collected 11 outputs (one for each adjusted temperature setting), across two execution attempts.

In the response analysis phase, we evaluated the performance of each SLM using the *pass@k* and *semantic consistency* metrics, aiming to assess both the correctness and coherence of the generated solutions. Responses were considered correct when the model’s justification matched the test smell annotation in the manual baseline. Based on the initial analysis, we identified an opportunity to improve the prompting technique to obtain more accurate, and reliable responses. This enhancement process is discussed in Section 4.4. Finally, we conducted a final review, considering the effectiveness of the changes made to the prompt formulation.

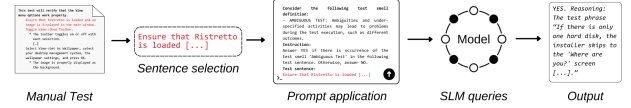


Figure 3: Approach to Manual Test Smell Detection.

**Pass@k and Semantic Consistency metrics.** We evaluated the performance of the models by observing the *pass@k* metric, which is an evaluation used to measure the probability that, by generating *k* samples, at least one of them is correct [5]. This metric provides the probability of success when generating multiple samples, reflecting the model’s ability to generate correct solutions within a limited number of attempts.

### 3 Analysis of the Results

Phi-4 model achieved a *pass@2* of 97%, while the Gemma3 and Llama3.2 models reached a *pass@2* of 91% and 90%, respectively. Regarding the outputs generated by the models, we observed Semantic Consistency (SC), which evaluates how reliably a model reproduces correct outputs in repeated submissions [19]. As shown in Table 4, all models achieved SC above 65%, with Phi-4 standing out with high consistency (90%). Therefore, the model’s reasoning is stable across repeated attempts, surpassing random correctness.

Table 4: Evaluation of SLMs across *pass@k*, Test Smells, and Semantic Consistency.

	Gemma3	Llama3.2	Phi-4
<b>Accuracy (pass@2)</b>	91.1%	91.0%	97.0%
Ambiguous Test	73.3%	77.0%	93.3%
Conditional Test	55.5%	100.0%	78.0%
Eager Action	97.1%	68.1%	85.5%
Misplaced Action	68.1%	62.5%	62.5%
Misplaced Precondition	100.0%	100.0%	20.0%
Misplaced Verification	69.0%	75.0%	69.0%
Unverified Action	39.0%	69.0%	76.5%
<b>Semantic Consistency</b>	65.0%	66.1%	90.0%

Despite the lower performance compared to Phi-4, the Llama3.2 model achieved a *pass@2* of 90%, with fewer parameters (3B) in its architecture, and being lighter than the Phi-4 model (14B). In scenarios with strict resource constraints, the Llama model can be combined with a second model for test smell detection.

We varied the temperature parameter to assess its impact on test smell detection. Phi-4 consistently outperformed the other models across all settings (77–81%), with its best result at 0.2. Gemma3 and Llama3.2 showed stable performance overall, peaking at 0.4 and 1.0, with detection rates ranging from 49–54% and 45–53%, respectively.

In terms of accuracy variation across temperature settings, Gemma3 and Phi-4 showed differences of 13 and 12 detections, respectively, between the maximum and minimum accuracy values. Additionally, the Llama3.2 model showed the highest variation, with a 25-point difference in accuracy between temperatures 0.6 and 1.0. This suggests that the temperature increase had a stronger influence on the accurate detection of test smells.

In the overall results, considering the best performance of each model, 253 sentences affected by test smells were detected by at least one of the models. Meanwhile, in 78 sentences, all three models detected the specific test smells. In only 8 sentences of the dataset, none of the models identified the presence of any of the targeted test smells

## 4 Discussion & Further Analysis

### 4.1 Explanations

SLMs generate responses for the detection of test smells, offering detailed explanations for each identified test smell. These justifications, grounded in the models' reasoning, enhances the review process, enabling developers to gain a deeper understanding and identification of the issues. This represents a significant advantage over other detection approaches that focus solely on syntactic analysis.

In response to our **RQ<sub>1</sub>**, we observed that some of the Gemma's responses lacked accompanying explanations, consisted only of binary outputs such as 'YES' or 'NO'. The explanations generated by model are generally brief, both for positive and negative results. In the example of the sentence "*The slideshow is entirely in your language*", the model points out that the sentence is vague: "*It doesn't specify 'which' language, or what constitutes **entirely***", but the model's response did not offer deeper insight on the test smell issue, nor does it provide suggestions to adjust the problem. It is essential to evaluate not only the model's ability to classify sentences as containing test smells or not, but also how it justifies these decisions. The absence of justification limits our ability to assess the reasoning process behind the final decision.

Regarding the Llama3.2 model (**RQ<sub>2</sub>**), the model provided explanations for all the evaluated sentences, at each adjusted temperature. In the *Llama3.2 Explanations* box, the negative explanations generated by the Llama model are generally shorter (i.e., contain fewer tokens) than the positive ones. However, some longer explanations were produced, even though they were *False Negatives*.

#### Examples of Llama3.2 explanations for Eager Action

##### True Positive

**Sentence:** "Allow the machine to reboot, select the first option at the grub menu."

**Output:** Yes. The test sentence can be broken into multiple steps: (1) Allow the machine to reboot, (2) Select the first option at the grub menu. The action of selecting is grouped with rebooting, violating the single-step rule.

##### False Positive

**Sentence:** "Select a non-English language and press Enter"

**Output:** No. This does not exhibit Eager Action since selecting a non-English language and pressing Enter is a single action group, not multiple. Hence, the answer is NO.

The results of the Phi-4 model confirm the best performance for the task of test smell detection in natural language (**RQ<sub>3</sub>**). Phi-4 generates standardized explanations, typically showing consistency, whether in positive or negative responses. Even with shorter responses, the model generated explanations for all evaluated sentences across all adjusted temperatures. The model demonstrated significant potential beyond the detection of test smells. It also proposed improvements—even without direct instruction in the prompt—that help developers understand and address test smell issues in test sentences. We highlight one response to an ambiguity issue in the *Example of Phi-4 explanations with DSP* box.

#### Example of Phi-4 explanation with DSP

**Sentence:** The machine logs in successfully.

**Output:** Indicate which credentials should be used and what the expected outcomes are (e.g., access to the desktop).

### 4.2 False Negatives

All sentences in our dataset are affected by test smells; however, all models produced negative responses for a subset of them. These responses were often accompanied by brief justifications. In some cases, the models identified more general issues which, although potentially relevant, did not align with the specific categories of test smells targeted in our study.

Llama3.2 generated 80 negative responses regarding the presence of test smells in the sentences. We highlight the sentence "*Verify that the full drive space is allocated*" which is recorded in the dataset with the test smell Misplaced Verification, however, the Llama3.2 was unable to identify the issue, providing the following response: "*Since the test sentence only contains one action ('Verify'), and it does not contain any additional actions that would include more than just the 'Verify', there is not enough information about other instructions being impacted.*".

The Phi-4 model was able to detect the Misplaced Verification smell in the previously highlighted sentence, but it produced *False Negative* responses for this type in five other cases. In the sentence "*note the state of each check mark*", the model pointed out that *the sentence simply instructs to observe and record the state without verifying it against the expected results*. This reasoning reveals a limitation of the model in identifying test smells when they result in *False Negatives*, especially for those that require context beyond a single sentence, as is the case with the Misplaced Action and Misplaced Verification types.

### 4.3 Limitations on Test Sentences

The *Misplaced Action* and *Misplaced Verification* smell types require analysis of the affected sentence concerning other

correlated sentences within the test instruction. When the model receives an isolated sentence, it becomes challenging to identify an action or verification that is out of sequence.

In the sentence “*Open some windows*”, the Gemma3 and Phi-4 models indicate that it is a simple action not presented within a list of verifications, as defined by the *Misplaced Action* test smell. The Llama3.2 model states that the test smell is present in the highlighted sentence, but its explanation is inconsistent. It claims that “*there are actions in the verification list ('Open some windows') that are not part of a typical user command or instruction*”, yet it does not specify which verifications are being referenced in its response.

In the responses from Phi-4, we also observed a limitation in identifying the *Misplaced Precondition* type from a single test sentence. For the sentence “*Make sure auto-hide is enabled*”, the model indicated that without additional context to show that this statement is indeed the first action, it is not possible to definitively classify it as an occurrence of the *Misplaced Precondition* smell. In future evaluations, we will assess the deeper detection of these types of smells using a set of sentences or even a complete test case.

#### 4.4 Can We Do It Better? The Role of Meta Prompting

Given the limitations of the models in detecting test smells, we conducted a deeper analysis using a more refined and targeted prompt. We considered the relevance of applying the Meta Prompting strategy [8] to a subset of sentences. This strategy refers to a prompting technique used to guide a language model in generating or refining other prompts.

Based on the configuration defined for the Phi-4 model, we asked it to generate an ideal prompting to detect the test smells *Ambiguous Test*, *Conditional Test*, and *Eager Action*. Only these three types of test smells were evaluated with Meta Prompting, as the Phi-4 model required more time to process the instructions and evaluate the sentences from the *False Negatives* subset. We chose the Phi-4 model due to its superior performance in the test smell detection process, presented in Section 3.

The Meta Prompting design was generated by Phi-4 based on the following query: *I need to create a prompt for scientific research purposes. Specifically, I need a prompt to detect 'Ambiguous Test' smells in manual testing. How would you design a prompt for an SLM to perform this task?* In response, the model provided a specific prompt design aimed at detecting the *Ambiguous Test* smell. This design included an explanation of the test smell and instructions to guide the detection of issues in test sentences. The query was later adapted to address the *Conditional Test* and *Eager Action*.

We evaluated the first output generated by the model for each type of test smell, organized the textual structures, and preserved the order and composition of the instructions created by the Phi-4 model to define the final version of the prompt (see Figure 4).




Figure 4: Improving Prompt Quality for Test Smell Detection Using Meta-Prompting.

The instructions generated by the Phi-4 model itself enabled an improvement in the detection of the test smells. For the *Ambiguous Test* type, the prompt was structured to analyze *unclear purpose*, *overly complex logic*, *hardcoded values*, *lack of descriptive naming*, and *redundant assertions*. For the *Conditional Test* type, the instructions were aimed at *verifying tests with deeply nested if-else statements*, *tests relying on specific input values leading to different execution paths*, and *tests using numerous logical operators (AND, OR) in a single condition*. Meanwhile, for the *Eager Action* type, the model was instructed to evaluate *signs of premature execution*, *lack of proper setup*, or *execution of tests without initializing the necessary components*.

We employed Meta Prompting to re-evaluate 19 sentences using the Phi-4 model. The model detected *Ambiguous Test* in six sentences (6/6) and *Eager Action* in 11 sentences (11/11). The re-evaluated sentences were *False Negatives* that the model had previously missed the smells with DSP. We highlight one of the model’s responses in the *Example of Phi-4 explanations with Meta Prompting* box.

**Example of Phi-4 explanations with Meta Prompting**

**Output:** The provided test sentence “The slideshow is entirely in your language” exhibits several ambiguous test smells based on the definitions provided:

- Unclear Purpose: The test sentence does not specify what aspect of the slideshow it is testing or under what conditions [...].

Regarding *Conditional Test*, the model generated negative responses for the presence of a smell in the sentence “*When ubiquity starts select your language in the left column.*” This sentence is registered in two different test cases: *image/1303\_Live Session* and *image/1497\_Install (live session)*. According to the model: “*The test sentence is simple and clear, lacking any elements that would typically contribute to Conditional Test Logic smells*”. However, the structure “*When [event] [action]*” represents a temporal condition, where selecting the language only occurs after the condition ‘ubiquity starts’ is met. The sentence does not include the common term ‘if’, but there is still a time dependency, as the second event only happens if the first one occurs.

The Meta Prompting strategy enhances the detection of test smells by generating more accurate and detailed explanations for the identified issues. These model-generated explanations play a key role in helping developers understand problems more deeply, facilitating the correction of flaws and the improvement of manual test cases. From the perspective of applying Meta Prompting to the entire dataset, a significant improvement in the performance of the Phi-4 model can be anticipated for the detection of all test smell occurrences.

## 4.5 Threats to Validity

There are some threats to validity that could impact the results and interpretations [23]. We conducted the tests by applying DSP referenced in the literature, manipulating a dataset of natural language sentences that was carefully reviewed by researchers in a previous study. To ensure the validity of this study, the model's responses were reviewed by two authors of the paper to avoid incorrect interpretations in the explanations of the results.

The reproducibility of the results depends on the consistent behavior of SLMs, which may vary with updates or changes in the models. Therefore, we emphasize both the setup of the models and the execution period of the study. Additionally, two authors manually reviewed the responses generated by the SLMs to assess potential hallucinations regarding the test smells. This task is challenging, as the models tend to produce lengthy explanations for each evaluated sentence.

The study demonstrates the potential of SLMs in detecting test smells, the conclusions drawn are based on a specific set of natural language sentences. Further testing with a larger dataset of natural language examples is needed to comprehensively validate the findings. The conclusions of the study should be interpreted considering the limitations mentioned earlier and the need for ongoing evaluation and refinement of the capabilities of SLMs.

## 5 Related Work

Aranda *et al.* [3] proposed a catalog of transformations to automatically remove seven types of test smells from tests written in natural language, such as Ambiguous Test and Eager Action. The authors also developed the Manual Test Alchemist tool, which applies these transformations using NLP techniques. Their evaluation, based on real-world test cases from the Ubuntu OS and a survey with professionals, demonstrated high acceptance (91.4%) and a promising F-Measure rate of 83.7%. In the study by Soares *et al.* [24], three systems from different domains were explored, and the authors also highlight promising results from using NLP strategies to reduce the effort required in analyzing tests written in different languages.

Viezaga *et al.* [29] proposed an automated system to detect ambiguous or problematic linguistic patterns in requirements written in natural language. The authors highlight that contextualized rephrasing suggestions can assist analysts in improving requirements quality.

In the context of test smells in natural language and following the study by Hauptmann *et al.* [7], Rajkovic and Enouï [22] developed the NALABS tool to detect smells in requirements and test specifications written in natural language, using keyword lists to assess metrics such as vagueness, referencability, optionality, subjectivity, and fragility.

Concerning programming languages, Yang *et al.* [31] addressed the challenges and discussed systematic methods for exploring new types of test smells. In the system proposed by the authors, six detection methods were applied to identify different types of test smells. Similarly, from the

perspective of analyzing various types of test smells in programming languages, Lucas *et al.* [16] applied the use of LLMs to automatically detect 30 types of test smells. The authors emphasize the potential of language models in identifying test smells, highlighting the models' capabilities to detect and offer improvement suggestions for software tests.

In the context of smell analysis using language models, Wu *et al.* [30] developed the *iSMELL* system, which integrates LLMs with a set of specialized tools to detect and refactor code smells. The proposed system dynamically selects the most suitable tools to identify three types of complex smells, and the authors emphasize that the tool stands out for its scalability, flexibility, and optimized performance, representing a promising advancement in AI-assisted software engineering.

In our work, we investigate the potential of SLM for automatically detecting test smells, exploring their applicability and effectiveness in identifying issues in natural language test cases. Our goal is to provide a lightweight, scalable, and accessible solution that complements existing methodologies. Phi-4 successfully detects 97% of the test smells, demonstrating promising results. The Gemma3 and Llama3.2 models performed less effectively, detecting 91% and 90% of the smells, respectively.

## 6 Conclusion

We evaluated the ability of SLMs to detect test smells in natural language test sentences from real-world test cases. Phi-4 outperformed the others, detecting 97% of the tests with smells, with a *pass@2* of 97%, while Gemma3 and Llama3.2 detected 91%. Gemma 4B and Llama 3B detected around 91% of the cases, which is promising considering they have three to four times fewer parameters than Phi-4 14B. This suggests they may be useful when a lighter and faster model is needed. All data generated by the models is available online [15].

SLMs demonstrate strong potential for detecting test smells by providing explanations and suggestions that support refactoring, based only on conceptual definitions of smells and reducing the need for extensive syntactic specifications. This potential is further enhanced by the use of Meta Prompting, which can boost model performance in identifying an even greater number of test smell occurrences in test cases, at a low computational cost.

Future work will extend our approach to detect additional types of test smells and will incorporate agentic methods [17] to assess model performance on larger and more diverse datasets. We are also interested in investigating how foundation models can enable automated refactoring of test cases.

## ARTIFACT AVAILABILITY

All study artifacts are available online [15].

## ACKNOWLEDGMENTS

We thank the reviewers for their feedback. This work was partially supported by CNPq (403719/2024-0, 310313/2022-8, 404825/2023-0, 443393/2023-0, 312195/2021-4), FAPESQ-PB (268/2025).

## REFERENCES

- [1] Marah Abidin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, et al. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905* (2024).
- [2] David H Ackley, Geoffrey E Hinton, and Terrence J Sejnowski. 1985. A learning algorithm for Boltzmann machines. *Cognitive science* 9, 1 (1985), 147–169.
- [3] Manoel Aranda, Naelson Oliveira, Elvys Soares, Márcio Ribeiro, Davi Romão, Ulyanne Patriota, Rohit Gheyi, Emerson Souza, and Ivan Machado. 2024. A Catalog of Transformations to Remove Smells From Natural Language Tests. In *International Conference on Evaluation and Assessment in Software Engineering*. ACM, 7–16.
- [4] Manoel Aranda, Naelson Oliveira, Elvys Soares, Márcio Ribeiro, Davi Romão, Ulyanne Patriota, Rohit Gheyi, Emerson Souza, and Ivan Machado. 2024. Manual Test Alchemist. <https://github.com/easy-software-ufal/manual-test-alchemist/tree/main/testcases> Retrieved April 19, 2024.
- [5] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374* (2021).
- [6] Benedikt Hauptmann. 2016. *Reducing system testing effort by focusing on commonalities in test procedures*. Ph. D. Dissertation. Technische Universität München.
- [7] Benedikt Hauptmann, Maximilian Junker, Sebastian Eder, Lars Heinemann, Rudolf Vaas, and Peter Braun. 2013. Hunting for smells in natural language tests. In *International Conference on Software Engineering*. IEEE Computer Society, 1217–1220.
- [8] Yutai Hou, Hongyuan Dong, Xinghao Wang, Bohan Li, and Wanxiang Che. 2022. MetaPrompting: Learning to learn better prompts. *arXiv preprint arXiv:2209.11486* (2022).
- [9] Ollama Inc. 2025. *Ollama: Run large language models locally*. <https://ollama.com>
- [10] Katharina Juhnke, Alexander Nikic, and Matthias Tichy. 2021. Clustering Natural Language Test Case Instructions as Input for Deriving Automotive Testing DSLs. *J. Object Technol.* 20, 3 (2021), 5–1.
- [11] Nildo Silva Junior, Luana Martins, Larissa Rocha, Heitor Costa, and Ivan Machado. 2021. How are test smells treated in the wild? A tale of two empirical studies. *Journal of Software Engineering Research and Development* 9 (2021), 9–1.
- [12] Nildo Silva Junior, Larissa Rocha, Luana Almeida Martins, and Ivan Machado. 2020. A survey on test practitioners’ awareness of test smells. In *Iberoamerican Conference on Software Engineering*. Curran Associates, 462–475.
- [13] Zekun Li, Baolin Peng, Pengcheng He, Michel Galley, Jianfeng Gao, and Xifeng Yan. 2023. Guiding Large Language Models via Directional Stimulus Prompting. *arXiv:2302.11520 [cs.CL]* <https://arxiv.org/abs/2302.11520>
- [14] Zhenyan Lu, Xiang Li, Dongqi Cai, Rongjie Yi, Fangming Liu, Xiwen Zhang, Nicholas D Lane, and Mengwei Xu. 2024. Small language models: Survey, measurements, and insights. *arXiv preprint arXiv:2409.15790* (2024).
- [15] Keila Lucas, Rohit Gheyi, Márcio Ribeiro, Fabio Palomba, Luana Martins, and Elvys Soares. 2025. Investigating the Performance of Small Language Models in Detecting Test Smells in Manual Test Cases (artifacts). <https://doi.org/10.5281/zenodo.15484997>.
- [16] Keila Lucas, Rohit Gheyi, Elvys Soares, Márcio Ribeiro, and Ivan Machado. 2024. Evaluating Large Language Models in Detecting Test Smells. In *Brazilian Symposium on Software Engineering*. 672–678.
- [17] Rian Melo, Pedro Simões, Rohit Gheyi, Marcelo d’Amorim, Márcio Ribeiro, Gustavo Soares, Eduardo Almeida, and Elvys Soares. 2025. Agentic SLMs: Hunting Down Test Smells. *arXiv:2504.07277 [cs.SE]* <https://arxiv.org/abs/2504.07277>
- [18] Meta AI. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/> Accessed: 2025-05-20.
- [19] Shuyin Ouyang, Jie M Zhang, Mark Harman, and Meng Wang. 2025. An empirical study of the non-determinism of chatgpt in code generation. *ACM Transactions on Software Engineering and Methodology* 34, 2 (2025), 1–28.
- [20] Max Peepkorn, Tom Kouwenhoven, Dan Brown, and Anna Jordanous. 2024. Is temperature the creativity parameter of large language models? *arXiv preprint arXiv:2405.00492* (2024).
- [21] Myron David Lucena Campos Peixoto, Davy de Medeiros Baia, Nathalia Nascimento, Paulo Alencar, Balduino Fonseca, and Márcio Ribeiro. 2024. On the Effectiveness of LLMs for Manual Test Verifications. *arXiv preprint arXiv:2409.12405* (2024).
- [22] Kostadin Rajkovic and Eduard Enoiu. 2022. Nalabs: Detecting bad smells in natural language requirements and test specifications. *arXiv preprint arXiv:2202.05641* (2022).
- [23] June Sallou, Thomas Durieux, and Annibale Panichella. 2024. Breaking the Silence: the Threats of Using LLMs in Software Engineering. In *International Conference on Software Engineering - New Ideas and Emerging Results*. ACM/IEEE.
- [24] Elvys Soares, Manoel Aranda, Naelson Oliveira, Márcio Ribeiro, Rohit Gheyi, Emerson Souza, Ivan Machado, André L. M. Santos, Balduino Fonseca, and Rodrigo Bonifácio. 2023. Manual Tests Do Smell! Cataloging and Identifying Natural Language Test Smells. In *International Symposium on Empirical Software Engineering and Measurement*. IEEE, 1–11.
- [25] Gabriela Soares, Vanessa Santos, Márcio Ribeiro, Luana Martins, Valeria Pontillo, Manoel Aranda III, Rohit Gheyi, Ivan Machado, and Fabio Palomba. 2025. On the Harmfulness of Test Smells in Manual System Testing: A Controlled Experiment. In *International Symposium on Empirical Software Engineering and Measurement*. IEEE.
- [26] Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. 2025. Gemma 3 Technical Report. *arXiv preprint arXiv:2503.19786* (2025).
- [27] Ubuntu. 2024. *Ubuntu Manual Tests in Launchpad*. <https://launchpad.net/ubuntu-manual-tests>
- [28] Chien van Nguyen, Xuan Shen, Ryan Aponte, Yu Xia, Samyadeep Basu, Zhengmian Hu, Jian Chen, Mihir Parmar, Sasidhar Kunapuli, Joe Barrow, et al. 2024. A survey of small language models. *arXiv preprint arXiv:2410.20011* (2024).
- [29] Alvaro Veizaga, Seung Yeob Shin, and Lionel C Briand. 2024. Automated smell detection and recommendation in natural language requirements. *IEEE Transactions on Software Engineering* 50, 4 (2024), 695–720.
- [30] Di Wu, Fangwen Mu, Lin Shi, Zhaoqiang Guo, Kui Liu, Weiguang Zhuang, Yuqi Zhong, and Li Zhang. 2024. iSMELL: Assembling LLMs with Expert Toolsets for Code Smell Detection and Refactoring. In *Proceedings of the 39th IEEE/ACM International Conference on Automated Software Engineering*. 1345–1357.
- [31] Yanming Yang, Xing Hu, Xin Xia, and Xiaohu Yang. 2024. The lost world: Characterizing and detecting undiscovered test smells. *ACM Transactions on Software Engineering and Methodology* 33, 3 (2024), 1–32.