

# PRAISE: Enhancing Product Descriptions with LLM-Driven Structured Insights

Adnan Qidwai<sup>1\*</sup>, Srija Mukhopadhyay<sup>1\*</sup>, Prerana Khatiwada<sup>2\*</sup>  
Dan Roth<sup>3</sup>, Vivek Gupta<sup>4†</sup>

<sup>1</sup>IIIT Hyderabad, <sup>2</sup>University of Delaware,

<sup>3</sup>University of Pennsylvania, <sup>4</sup>Arizona State University

{adnan.qidwai@student, srija.mukhopadhyay@research}.iiit.ac.in; preranak@udel.edu;

danroth@seas.upenn.edu; vgupt140@asu.edu

## Abstract

Accurate and complete product descriptions are crucial for e-commerce, yet seller-provided information often falls short. Customer reviews offer valuable details but are laborious to sift through manually. We present PRAISE: Product Review Attribute Insight Structuring Engine, a novel system that uses Large Language Models (LLMs) to automatically extract, compare, and structure insights from customer reviews and seller descriptions. PRAISE provides users with an intuitive interface to identify missing, contradictory, or partially matching details between these two sources, presenting the discrepancies in a clear, structured format alongside supporting evidence from reviews. This allows sellers to easily enhance their product listings for clarity and persuasiveness, and buyers to better assess product reliability. Our demonstration showcases PRAISE’s workflow, its effectiveness in generating actionable structured insights from unstructured reviews, and its potential to significantly improve the quality and trustworthiness of e-commerce product catalogs.

## 1 Introduction

In the rapidly expanding e-commerce landscape, platforms like Amazon heavily rely on detailed product descriptions to drive purchasing decisions (Vandic et al., 2018) and build customer trust (Reibstein, 2002). However, seller-provided descriptions frequently suffer from incompleteness or inaccuracies. While customer reviews often contain rich, factual information about product attributes and performance (Askalidis and Malthouse, 2016), manually extracting these details and reconciling them with seller descriptions is tedious and error-prone (Hu and Liu, 2004b). This gap highlights the need for automated tools to bridge information sources.

\*equal contribution †corresponding author

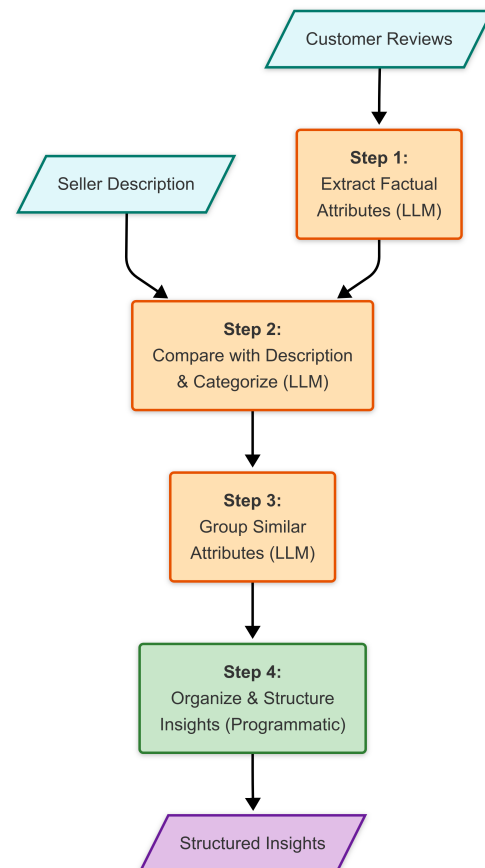


Figure 1: End-to-End Pipeline of PRAISE for Attribute Extraction and Structuring

Recent advances in LLMs, with their proficiency in natural language understanding and generation (Roumeliotis et al., 2024; Zhou et al., 2023; Soni, 2023), offer a powerful means to address this challenge. Building on this potential, we developed PRAISE, an interactive system designed to automatically enhance product descriptions using insights obtained from customer reviews.

PRAISE’s LLM-driven pipeline: (1) **Extracts** factual attributes from reviews (filtering opinions); (2) **Compares** attributes to seller descriptions; (3) **Categorizes** discrepancies (Missing, Contradictory, Partially-matching) with justifications; and (4) **Structures** findings for action. This allows users to

quickly identify areas where product descriptions can be improved for accuracy and completeness. Our main contributions are:

1. **The PRAISE System:** A novel, publicly accessible system demonstrating the use of LLMs for structured comparison of product descriptions and reviews. Refer to Figure 1 for the complete system pipeline.
2. **Evaluation and Insights:** An analysis of the system’s performance, highlighting its strengths and current limitations in processing real-world e-commerce data.

We invite readers to explore PRAISE’s capabilities through the following resources:

- **Project Page:** [project-praise.github.io](https://project-praise.github.io)
- **Demo Video:** [project-praise.github.io/demo/](https://project-praise.github.io/demo/)
- **Try It Out:** [project-praise.github.io/tryout/](https://project-praise.github.io/tryout/)

The system showcases a practical application of LLMs for tangible improvements in e-commerce information quality.

## 2 Related Work

Analyzing customer reviews for insights like sentiment and feature extraction has a rich history (Hu and Liu, 2004a; Popescu and Etzioni, 2005; Mabrouk et al., 2021). The advent of LLMs has significantly advanced capabilities in processing review data for tasks such as description generation, product categorization, and search refinement (Liu et al., 2023; Roumeliotis et al., 2024; Wang et al., 2024b; Choudhary et al., 2024). Techniques like prompt engineering are vital for optimizing LLM outputs for specific tasks like information extraction and bias mitigation (Marvin et al., 2023; Russe et al., 2024; Wang et al., 2024a). While these works provide foundational techniques, PRAISE distinguishes itself by implementing a *structured comparison* pipeline specifically designed to identify and categorize discrepancies (*missing, contradictory, partial*) between review facts and seller descriptions, presenting them in an actionable format through an interactive system. Our focus is on demonstrating this end-to-end system for refining product catalog quality.

## 3 The PRAISE System: Architecture and Workflow

PRAISE employs a multi-step pipeline, primarily leveraging LLMs, combined with programmatic

orchestration to enrich product descriptions. The system processes customer reviews to extract pertinent descriptive information and systematically compares it against the seller-provided description. Our approach utilizes the advanced language understanding capabilities of LLMs for analysis, leveraging their ability to generate responses adhering to predefined structured formats (like JSON) where applicable, while integrating programmatic steps for structuring and organizing the results effectively. The following steps detail this workflow:

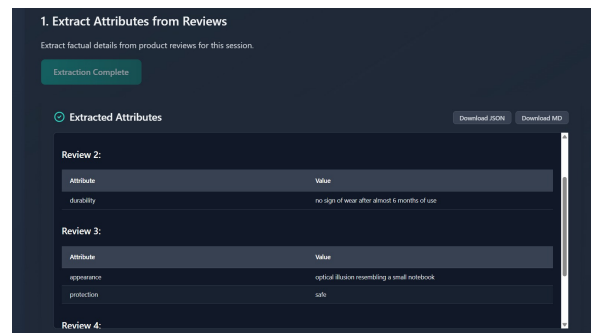


Figure 2: Attribute extraction from product reviews (Step 1 of PRAISE)

**Step 1: Extracting Descriptive Details from Reviews.** The initial step focuses on analyzing each customer review individually to identify and isolate objective, factual information about the product (Figure 2). An LLM is guided to distinguish these descriptive details (like materials, dimensions, or specific features) from subjective opinions, personal anecdotes, or irrelevant commentary. The core purpose is to filter out noise and retain only the verifiable, product-specific facts mentioned by reviewers. The output of this step is a collection of factual attribute-value pairs derived from each processed review.

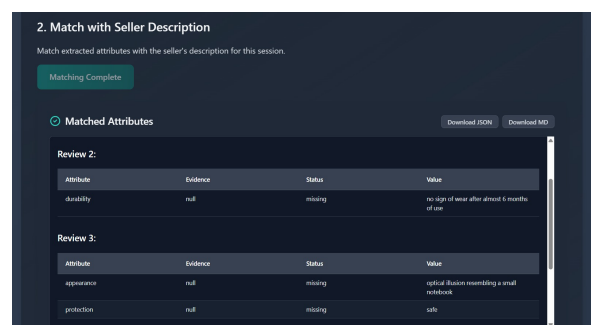


Figure 3: Matching Extracted Attributes with Seller Descriptions (Step 2 of PRAISE)

**Step 2: Comparison with Seller Description and Categorization.** Next, the system takes the fac-

tual attributes corresponding to each review extracted in the previous step and performs a comparative analysis against the official seller-provided product description (Figure 3). An LLM examines each attribute derived from the reviews, determines its presence and consistency within the seller’s text, and assigns a category based on the comparison. The key categories identify whether the information from the review is *Missing* from the seller description, *Contradictory* to it, *Matching*, or only *Partially-matching*. The process also includes providing reasoning or evidence from the seller’s description where applicable. The outcome is a structured set of comparison results for each review, detailing how the factual points align or conflict with the seller’s claims.

**Step 3: Grouping of Similar Attributes.** To improve the organization of the findings, this step focuses on categorizing the diverse attributes identified across all reviews (Figure 4). Based on the attribute names (like ‘weight’, ‘color’, ‘battery life’), an LLM assigns each unique attribute to a broader, intuitive category (such as “Physical Attributes”, “Appearance”, “Performance”). This grouping is based on the semantic similarity of the attribute concepts themselves, rather than their specific values, aiming for a generalized and user-friendly classification. The result of this step is a mapping that assigns a logical category to each type of attribute encountered.

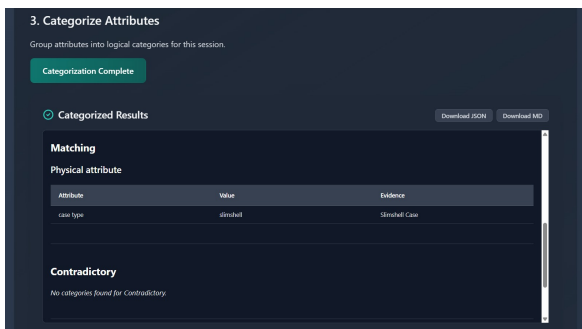


Figure 4: Grouping and Structuring Attributes into Logical Categories (Steps 3 & 4 of PRAISE)

**Step 4: Organizing and Presenting Structured Insights.** In the final step, the system consolidates the comparison results from Step 2 and applies the category labels generated in Step 3. It programmatically structures this information through a *rule-based* method, primarily highlighting the actionable insights – instances where review information was *Missing*, *Contradictory*, or *Partially-*

*matching* compared to the seller description. The findings are organized logically, grouped first by the comparison status and then by the attribute categories. This produces a final, structured output that presents the key discrepancies and alignments in an easy-to-navigate format, allowing users to quickly understand which aspects of the product description may require revision or verification based on customer feedback.

**Pipeline Efficiency, Cost, and Optimization.**

The system’s architecture is designed to balance performance with output quality. For a product with  $R$  reviews, the pipeline makes approximately  $2R + 1$  LLM API calls, with the review-level extraction (Step 1) and comparison (Step 2) tasks being highly parallelizable. We used Python’s `ThreadPoolExecutor` to execute these steps concurrently, significantly reducing wall-clock time. Robust error handling and retry logic are implemented for each LLM interaction to ensure resilience.

This modular, multi-step design has direct implications for operational cost. While the linear scaling with the number of reviews is predictable, the cost can become a factor for products with extensive feedback. This design represents a deliberate trade-off: as demonstrated in our ablation analysis (Section 6), aggregating steps into a single, cheaper API call (our end-to-end baseline) led to a significant degradation in quality, with higher rates of hallucination and incorrect categorization. The higher call volume of our modular pipeline is therefore a necessary investment to achieve reliable and actionable insights.

For production-scale deployment, several optimizations could further mitigate these costs without compromising quality:

- **Model Tiering:** A tiered strategy could employ a powerful model (e.g., Gemini Pro) for the nuanced extraction and comparison tasks, while using a smaller, faster, and more cost-effective model (like the Gemini Flash model used in our experiments) for the simpler attribute grouping task (Step 3).
- **Caching:** Implementing a caching layer to store results for previously processed reviews and products would eliminate redundant API calls and computations entirely.

Thus, while our current implementation prioritizes demonstrative clarity and accuracy, a clear path

exists for optimizing its cost-efficiency for large-scale use.

**Implementation and Accessibility:** Our public demonstration of PRAISE is powered by the Gemini API. To ensure the system remains freely accessible for experimentation, users are required to provide their own Gemini API key, which is available at no cost and includes a generous free usage tier. To ensure user security, the provided key is handled exclusively on the server-side for the duration of the API calls and is never permanently stored or exposed to the client. This approach balances the practicalities of hosting a public LLM-powered demo with user accessibility and security.

## 4 Methodology and Evaluation

In this section, we describe the methodology behind the PRAISE system, detailing its multi-step pipeline for extracting, comparing, and structuring product insights using LLMs.

**Experimental Setup.** To evaluate the performance of the PRAISE pipeline, we generated outputs using the implementation described in Section 3, primarily using the Gemini 2.0 Flash model. We selected Gemini 2.0 Flash for its strong balance of performance, inference speed, and cost-effectiveness, making it suitable for a scalable, interactive system. The evaluation dataset was derived from Amazon Reviews (Hou et al., 2024), encompassing 9 diverse product categories which included appliances, arts and crafts, beauty, books, CDs and vinyl records, cell phone accessories, clothing, shoes and jewelry, digital music and electronics. We selected around 10 products per category, with each product containing 7-9 reviews. We manually selected the reviews to make sure they had both opinions and descriptive details. This helped us accurately test how well our proposed system works.

**Evaluation Protocol.** A panel of three research team members manually verified the outputs of the pipeline’s core LLM-driven stages: Step 1 (Attribute Extraction), Step 2 (Review-Seller Comparison), and Step 3 (Attribute Grouping). Evaluators used the original reviews and seller descriptions as ground truth and followed detailed, pre-defined rubrics to ensure consistent assessment.

Each identified error was counted as one point, enabling a quantitative analysis of error frequencies and types. Annotators followed a detailed

evaluation rubric to identify specific error categories across all stages of the pipeline. In Step 1 (Extraction), common issues included incorrect attribute-value extraction, failure to filter out opinions, inclusion of irrelevant (non-product) information, and omission of valid attributes. In Step 2 (Comparison), evaluation focused on the correctness of the assigned status—Missing, Matching, Contradictory, or Partially Matching—as well as the validity of the accompanying justifications, with particular attention to misclassifications between these categories. In Step 3 (Grouping), errors involved inappropriate category naming, incorrect assignment of attributes to categories, and issues with grouping granularity, including both over-splitting and under-splitting of semantically related attributes.

## 5 Results and Analysis

### 5.1 Key Observations from Error Analysis

The quantitative evaluation provided specific insights into the performance and challenges of each pipeline stage using the Gemini 2.0 Flash model.




Figure 5: Distribution of Error Types in Step 1

**Step 1 (Extraction):** This initial step, focused on extracting factual attributes from raw reviews, exhibited the highest error frequency.

- *Strengths:* Despite these challenges, the system successfully extracted many basic factual attribute-value pairs, forming the necessary input for subsequent steps. The low count for ‘Other errors’ (2) suggests the defined rubric categories comprehensively captured the types of issues encountered.
- *Weaknesses:* The most prominent issue was the inclusion of Irrelevant Information (403 instances), where the model struggled to differentiate product-specific details from

user context or opinions disguised as facts. Assigning accurate attribute names also proved difficult, leading to numerous Incorrect Normalization errors (224).

Figure 5 demonstrates the major categories showcasing the errors made by the system, as found out by our robust evaluation methodology.




Figure 6: Error Distribution in Step 2—Matching Extracted Attributes with Seller Descriptions

**Step 2 (Comparison):** Figure 6 shows the errors encountered during Step 2. This step compared extracted attributes to the seller description, showing a pretty low number of errors (approx. 237 total) compared to extraction. This also shows the ability of the system to stop errors from cascading onto further steps, enabled by our robust prompt engineering.

- *Strengths:* The system performed considerably better here than in Step 1, successfully categorizing many attributes. Notably, the misclassification rate for Contradictory status was very low (7 instances), suggesting the model is relatively conservative or accurate when identifying direct conflicts, which is valuable for highlighting critical discrepancies.
- *Weaknesses:* The primary difficulty lay in accurately classifying the relationship between review and description attributes. Misclassifications where the model incorrectly identified attributes as Partially Matching (73 instances) or Missing (55 instances) were most common, indicating struggles with nuanced semantic differences versus true absences.

**Step 3 (Grouping):** Tasked with grouping related attributes based on their keys, this step was the most robust, exhibiting the fewest errors (approx.




Figure 7: Breakdown of Error Types in Step 3—Attribute Grouping and Categorization

187 total). This further shows how our system produces helpful end results.

- *Strengths:* The relatively low overall error count suggests that grouping based primarily on attribute keys is an effective strategy for this LLM. It successfully organized the majority of attributes into reasonable clusters. The low count for 'Other errors' (2) again implies the rubric covered the main issues.
- *Weaknesses:* The most frequent error was Incorrect Category Naming (70 instances), where the LLM generated labels that were not optimally descriptive or semantically appropriate for the grouped attributes. Issues with grouping granularity were also present, including Missing Category (under-splitting, 43 instances) where distinct concepts were wrongly merged, and Incorrect Splitting (over-splitting, 34 instances) where attributes were unnecessarily separated.

We show a more in-depth analysis of errors encountered during Step 3 in Figure 7.

## 5.2 Product-Wise Error Analysis

The evaluation revealed significant variations in pipeline performance across different product categories, as measured by precision, recall, and F1 score. This indicates that the typical language used in the reviews of products heavily influence the system's ability to accurately identify relevant attributes. Table 1 presents these detailed performance metrics for each category.

**High-Performance Categories:** Models performed the best in *Arts and Crafts* with the highest F1 score (0.82), driven by excellent recall (0.96) and good precision (0.72). This suggests the system is highly effective at identifying the vast majority

Category	Precision	Recall	F1 Score
Appliances	0.41	0.84	0.55
Arts and Crafts	0.72	0.96	0.82
Beauty	0.6	0.99	0.75
Books	0.23	0.79	0.36
CD Vinyl	0.52	0.6	0.56
Cell Phone and Accessories	0.56	0.82	0.66
Clothes, Shoes, Jewellery	0.45	0.79	0.57
Digital music	0.46	0.71	0.56
Electronics	0.47	0.86	0.61

Table 1: Performance Metrics by Product Category for Attribute Selection Using PRAISE

of true attributes for this category, while also maintaining a relatively low rate of incorrectly selecting non-attributes.

A similar performance was noticed for *Beauty* products with an F1 score of 0.75. Notably, this category achieved near-perfect recall (0.99), indicating the system rarely misses a relevant beauty attribute. However, its precision (0.60) is moderate, suggesting that while comprehensive, the system may also select a fair number of terms that are not true attributes, possibly due to the highly descriptive and often subjective language in beauty reviews (e.g., "silky," "glow," "subtle scent").

**Challenging Categories:** This group, consisting *Clothes, Shoes, and Jewellery* (F1 0.57), *CD & Vinyl* (F1 0.56), *Digital Music* (F1 0.56), *Appliances* (F1 0.55), and most notably *Books* (F1 0.36), caused significant hurdles in attribute selection, primarily due to low precision.

For categories like *Appliances* (P 0.41, R 0.84) and *Clothes, Shoes, and Jewellery* (P 0.45, R 0.79), good recall was undermined by very low precision where the system identified most true attributes but also incorrectly selected many non-attribute terms due to complex technical/usage details (*Appliances*) or subjective and overly descriptive language.

The music categories (*Digital Music*: P 0.46, R 0.71; *CD & Vinyl*: P 0.52, R 0.60) showed modest overall scores, struggling with precise extraction despite standardized metadata, likely because reviews often prioritize opinions over discrete factual features. *Books* was the most problematic, with acceptable recall (0.79) but exceptionally low precision (0.23). This starkly indicates that the high volume of subjective commentary, thematic discussions, and fewer distinct factual attributes in book reviews makes it exceedingly difficult for the system to distinguish true attributes from textual noise, leading to a very high rate of false positives.

### Overall Implications for Attribute Selection:

The category-specific analysis reveals distinct performance profiles. While some categories (*Arts and Crafts, Beauty*) achieve high recall, effectively identifying most true attributes, precision is a primary challenge across many others. This is particularly true for categories characterized by complex technical language (e.g., *Appliances, Electronics*), high subjectivity (e.g., *Beauty, Books*), or a lot of descriptive text (e.g., *Clothes*). The starkly low precision for *Books*, despite its standardized metadata, illustrates how a high volume of subjective or descriptive text can severely affect accurate attribute selection. Future efforts must prioritize enhancing the system’s discrimination between true attributes and textual noise, particularly for these low-precision categories.

## 6 Baseline and Ablation Analysis

To evaluate the contribution of our multi-step pipeline design, we conducted comparative analyses against two simpler approaches. We assessed performance based on several criteria reflecting the quality of the final structured output: the ability to retain important product details, exclude subjective opinions, maintain focus on product-specific information, and accurately categorize information (e.g., as missing or contradictory). Performance differences are illustrated by comparing counts across these criteria, as shown in Figure 8 and Figure 9.

### Baseline Comparison: End-to-End Prompting.

We compared PRAISE against a *Baseline* system. This baseline used a single, direct prompt instructing the LLM to perform the entire task end-to-end – taking raw reviews and seller description as input and generating the final structured output format without the intermediate processing steps defined in our pipeline. This comparison establishes a benchmark against a less structured, single-shot approach.

As shown in Figure 8, the full PRAISE pipeline consistently outperformed the baseline across nearly all evaluated metrics, particularly in areas like correct categorization and opinion exclusion. This demonstrates the significant value added by the structured intermediate steps in maintaining accuracy and information fidelity compared to direct end-to-end generation.

### Ablated System Comparison: Isolating Later Stages.

Second, we evaluated an *Ablated System*.

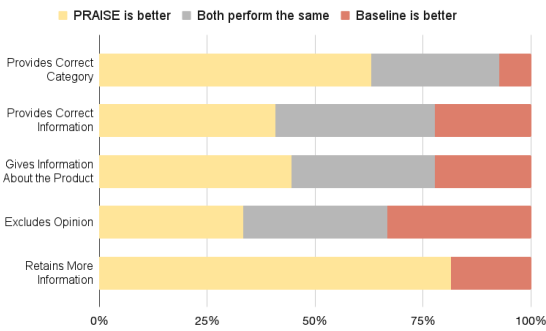


Figure 8: Comparison of the full PRAISE pipeline against the single-prompt Baseline.

This approach performed only Step 1 of our pipeline (structured attribute extraction) and then used a direct prompt to generate the final categorized output directly from these extracted attributes, bypassing the explicit comparison, grouping, and organization stages (Steps 2-4). This comparison isolates the contribution of these later structured processing steps, given the initial




Figure 9: Comparison of the full PRAISE pipeline against the Ablated System (Step 1 + Direct Prompt).

Figure 9 clearly indicates that the full PRAISE pipeline significantly outperformed this ablated system across all evaluation criteria. While structured extraction (Step 1) provides a better starting point than raw text (as suggested by the Baseline comparison), the results here highlight that the subsequent dedicated steps for comparison (Step 2), grouping (Step 3), and organization (Step 4) are crucial for achieving the highest level of accuracy and generating the most reliable structured insights. Together, these comparisons validate the effectiveness of our complete multi-step pipeline design.

## 7 Model Effectiveness Comparison

**Multi-Step vs Automatic Prompt.** To further validate our method, we conducted supplementary analyses for establishing its credibility and demonstrating its advantages over alternative approaches. We first tested a single, all-in-one

prompt to complete the entire task. This helped us show that our multi-step approach is more accurate, consistent, and easier to understand. We then bench marked against automatic prompting. We used the intent-based prompt calibration technique outlined in (Levi et al., 2024) to automatically calibrate prompts for the task of catalog expansion. This comparison highlights how our well-designed prompts and step-by-step approach lead to better performance.

When attempting to execute the entire process in a single step, we observed a significant increase in hallucination from the model. This suggests that the model struggled to differentiate between the various stages of the task and consequently lost crucial information. Similarly, the automatic prompt generation approach yielded suboptimal results, likely due to the inherent complexity of the task and the lack of a well-defined evaluation metric for this specific application.

**System License Details.** LLMs used in this study are licensed by their creators, while our platform uses the Apache 2.0 License. This license permits any use, distribution, modification, and commercial use of the software, including sublicensing and adding warranties.

## 8 Conclusion

Our work shows that LLMs can improve e-commerce product listings by integrating insights from customer reviews. We used a structured approach to filter out opinions and keep factual information, allowing us to accurately compare reviews with seller descriptions. This helped identify matches, gaps, and contradictions, and organize information into clear tables for easier analysis. Our results indicate that while LLMs are effective at summarizing and correcting spellings, they can struggle with filtering subjective opinions. Overall, LLMs are useful for improving product descriptions but have varying performance depending on the task, the model and the quality of reviews and descriptions provided.

Future work will focus on improving attribute relevance filtering and refining key-value alignment to enhance precision. Expanding attribute categorization and incorporating Q&A data will boost insight quality and coverage. Finally, we plan to test the system across more product categories and extend support to multilingual inputs.

## Limitations

While PRAISE demonstrates strong performance in structuring review-based insights, it is limited by its exclusive reliance on customer reviews, which may vary in quality and clarity, leading to inconsistent extractions. The current system does not incorporate other user-generated content such as Q&A pairs, which could enhance coverage. It is restricted to English-language inputs, reducing its applicability in multilingual settings. Evaluation depends on human judgment, which introduces subjectivity and limits scalability. The model may struggle with nuanced semantic distinctions, especially in subjective or mixed-content reviews, and can extract irrelevant or noisy information. Moreover, the use of a fixed prompt-based pipeline with a single LLM may constrain adaptability across diverse product categories. Finally, the absence of end-user feedback mechanisms limits our ability to assess real-world utility and usability.

## Ethics Statement

PRAISE employs large language models (LLMs) to extract factual information from customer reviews, introducing several ethical considerations. While the system incorporates step-wise prompting and structured evaluation to mitigate common failure modes, LLMs remain susceptible to producing biased or inaccurate outputs due to limitations in their training data. The reviews processed are publicly available and free of personally identifiable information; however, any future extensions involving private or sensitive data must ensure robust privacy protections. The reliance on proprietary models constrains transparency and interpretability, which we partially address through systematic documentation and error analysis. The system may also be vulnerable to misuse, such as selectively emphasizing favorable attributes or suppressing critical ones. To reduce this risk, PRAISE is explicitly designed to extract verifiable content and highlight missing or contradictory details. The system is released under the Apache 2.0 license, and users must supply their own Gemini API keys, which are not stored or logged. Responsible deployment of PRAISE requires human oversight to safeguard fairness, ensure accountability, and prevent potential misuse.

## Acknowledgement

This work was primary sponsored by the Army Research Office under Grant Number W911NF-

20-1-0080. The views and conclusions expressed herein are those of the authors and do not necessarily reflect the official policies, either expressed or implied, of the Army Research Office or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notice herein. This research was also partially supported by ONR Contract N00014-19-1-2620.

We thank the CogComp group (UPenn) for their insights; Mahika Vajpeyi (UPenn) for initial discussions; and Preethi Kallimuddanahalli Suresh (ASU) for help with dataset wrangling. We are also grateful to our anonymous reviewers for their thoughtful feedback and to the CoRAL Lab (ASU) for computational support. Special thanks to our lab cat, Coco, for keeping Prof. V.G. creatively unhinged during deadlines.

## References

- Georgios Askalidis and Edward C Malthouse. 2016. The value of online customer reviews. In *Proceedings of the 10th ACM Conference on Recommender Systems*, pages 155–158.
- Nurendra Choudhary, Edward W Huang, Karthik Subbian, and Chandan K Reddy. 2024. An interpretable ensemble of graph and language models for improving search relevance in e-commerce. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 206–215.
- Yupeng Hou, Jiacheng Li, Zhankui He, An Yan, Xiusi Chen, and Julian J McAuley. 2024. Bridging language and items for retrieval and recommendation. *CoRR*.
- Minqing Hu and Bing Liu. 2004a. [Mining and summarizing customer reviews](#). In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '04, page 168–177, New York, NY, USA. Association for Computing Machinery.
- Minqing Hu and Bing Liu. 2004b. Mining opinion features in customer reviews. In *AAAI*, volume 4, pages 755–760.
- Elad Levi, Eli Brosh, and Matan Friedmann. 2024. [Intent-based prompt calibration: Enhancing prompt optimization with synthetic boundary cases](#). In *ICLR 2024 Workshop on Navigating and Addressing Data Problems for Foundation Models*.
- Yiheng Liu, Tianle Han, Siyuan Ma, Jiayue Zhang, Yuanyuan Yang, Jiaming Tian, Hao He, Antong Li, Mengshen He, Zhengliang Liu, et al. 2023. Summary of chatgpt-related research and perspective towards the future of large language models. *Meta-Radiology*, page 100017.



- Alhassan Mabrouk, Rebeca P Díaz Redondo, and Mohammed Kayed. 2021. Seopinion: summarization and exploration of opinion from e-commerce websites. *Sensors*, 21(2):636.
- Ggaliwango Marvin, Nakayiza Hellen, Daudi Jjingo, and Joyce Nakatumba-Nabende. 2023. Prompt engineering in large language models. In *International conference on data intelligence and cognitive informatics*, pages 387–402. Springer.
- Ana-Maria Popescu and Oren Etzioni. 2005. [Extracting product features and opinions from reviews](#). In *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing*, pages 339–346, Vancouver, British Columbia, Canada. Association for Computational Linguistics.
- David J Reibstein. 2002. What attracts customers to online stores, and what keeps them coming back? *Journal of the academy of Marketing Science*, 30:465–473.
- Konstantinos I Roumeliotis, Nikolaos D Tselikas, and Dimitrios K Nasiopoulos. 2024. Llms in e-commerce: a comparative analysis of gpt and llama models in product review evaluation. *Natural Language Processing Journal*, 6:100056.
- M. F. Russe, M. Reisert, F. Bamberg, and A. Rau. 2024. [Improving the use of llms in radiology through prompt engineering: from precision prompts to zero-shot learning](#). *RöFo - Fortschritte Auf Dem Gebiet Der Röntgenstrahlen Und Der Bildgebenden Verfahren*.
- Vishvesh Soni. 2023. Large language models for enhancing customer lifecycle management. *Journal of Empirical Social Science Studies*, 7(1):67–89.
- Damir Vandic, Flavius Frasinca, and Uzay Kaymak. 2018. A framework for product description classification in e-commerce. *Journal of Web Engineering*, pages 001–027.
- L. Wang, X. Chen, X. Deng, H. Wen, M. You, W. Liu, Q. Li, and J. Li. 2024a. [Prompt engineering in consistency and reliability with the evidence-based guideline for llms](#). *NPJ Digital Medicine*, 7.
- Ningfei Wang, Yupin Huang, Han Cheng, Jiri Gesi, Xiaojie Wang, and Vivek Mittal. 2024b. Towards robustness analysis of e-commerce ranking system. In *Companion Proceedings of the ACM on Web Conference 2024*, pages 364–373.
- Jianghong Zhou, Bo Liu, Jhalak Acharya, Yao Hong, Kuang-Chih Lee, and Musen Wen. 2023. [Leveraging large language models for enhanced product descriptions in eCommerce](#). In *Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 88–96, Singapore. Association for Computational Linguistics.