

Enhancing Recommendation Systems via Aspect-Level Sentiment Analysis for Identifying Unreliable Reviews

El Mehdi Lghaouch¹, Fatima Zahra Abbour¹, Soumaya Ounacer¹, Soufiane Ardchir², Mohamed Azzouazi¹

¹Laboratory of Information Technology and Modeling, Faculty of Sciences Ben M'Sik, Hassan II University, Casablanca, Morocco

²National School of Management and Marketing, Hassan II University, Casablanca, Morocco

Article Info

Article history:

Received Feb 2, 2026

Revised Mar 27, 2026

Accepted Apr 16, 2026

Keywords:

Artificial Intelligence

Sentiment Analysis

Review Filtering

Recommendation Systems

Aspect Analysis

E-Commerce Reviews

ABSTRACT

User-generated reviews are an important source of information for e-commerce recommendation systems. However, reviews may contain conflicting opinions across product aspects such as quality, price, and delivery, which can weaken both review-level trust and aspect-level signals. To address this issue, we propose **AspectGuard**, a filtering framework based on Aspect-Based Sentiment Analysis (ABSA) and a Sentiment Divergence Score (SDS) that quantifies disagreement across aspect-level sentiments within a review. Reviews with high divergence are identified as potentially unreliable and excluded before constructing recommendation profiles. To isolate the contribution of the proposed filtering mechanism, we compare the same content-based recommendation model with and without AspectGuard filtering. Experiments conducted on a subset of the Amazon Reviews 2023 Electronics dataset show that the proposed filtering strategy improves both rating prediction and ranking quality. These results indicate that aspect-level sentiment divergence can serve as an effective signal for identifying unreliable reviews and improving the reliability of recommendation systems.

Copyright © 2026 Institute of Advanced Engineering and Science.

All rights reserved.

Corresponding Author:

El Mehdi Lghaouch

Laboratory of Information Technology and Modeling, Faculty of Sciences Ben M'Sik, Hassan II University
Casablanca, Morocco

Email: elmehdi.lghaouch-etu@etu.univh2c.ma

1. INTRODUCTION

As e-commerce continues to expand, Artificial Intelligence (AI) and machine learning are increasingly used in retail environments to better understand customer behavior, improve personalization, and support predictive decision-making [1]. In this context, user-generated reviews have become an important form of social proof that influences purchasing decisions and supports recommendation systems [2, 3]. These systems often exploit review text to go beyond traditional collaborative filtering and build richer user and item representations, thereby improving personalization. However, the trustworthiness of this information remains a major challenge. A review may express conflicting opinions across different product attributes. For example, a user may praise the design of a product while criticizing its delivery, producing a mixed signal that can reduce confidence in both the review itself and the recommendations derived from it [4]. Traditional sentiment analysis, which typically assigns a single overall polarity to the full review, is not well suited to this level of complexity. By collapsing multiple opinions into one score, it may overlook important contradictions expressed toward individual aspects [5]. As a result, a review with strong internal disagreement may be interpreted as neutral or weakly polarized, thereby introducing noise into downstream recommendation tasks [6].

Aspect-Based Sentiment Analysis (ABSA) offers a more fine-grained alternative by identifying sentiments associated with specific product aspects such as price, quality, or delivery [7]. This level of analysis provides richer and more interpretable information than document-level sentiment because it captures the diversity of opinions expressed within the same review [8]. Previous work has incorporated review sentiment into recommendation models, for example by weighting ratings with sentiment scores or modeling topics from review text [9]. The Review Polarity-wise Recommender, for instance, constructs separate positive and negative user profiles from review content [10]. However, such approaches do not explicitly quantify contradictions among aspect-level sentiments within a single review, and may therefore miss cases of internally inconsistent feedback.

To address this gap, we propose **AspectGuard**, a framework that integrates aspect-level sentiment divergence into the review filtering process. AspectGuard identifies product aspects mentioned in a review, estimates the sentiment associated with each aspect, and computes a *Sentiment Divergence Score (SDS)* that measures disagreement across those aspect-level sentiments. Reviews with high divergence are considered potentially unreliable and are excluded before constructing recommendation profiles. In this way, the recommendation process relies on reviews that are more internally consistent, which can improve both recommendation quality and interpretability.

The contribution of this work lies in introducing a review filtering framework that exploits aspect-level sentiment divergence to improve recommendation reliability. Rather than proposing a new ABSA backbone, the study focuses on how fine-grained sentiment signals extracted from review text can be transformed into a practical filtering mechanism for downstream recommendation. In particular, we formalize the notion of intra-review inconsistency through the Sentiment Divergence Score (SDS), which quantifies disagreement among aspect-level sentiments within a review, and we integrate this signal into a controlled recommendation setting. By comparing the same content-based recommendation model with and without AspectGuard filtering, we show that the proposed mechanism contributes to cleaner review-based representations and leads to improved recommendation performance.

The remainder of this paper is organized as follows. Section 2 reviews related work on sentiment-aware recommendation systems and review reliability. Section 3 presents the proposed methodology, including aspect extraction, sentiment classification, sentiment divergence analysis, review filtering, recommendation integration, and evaluation metrics. Section 4 describes the dataset, experimental setup, and results. Section 5 discusses the findings and limitations. Section 6 concludes the paper.

2. RELATED WORK

The authenticity of e-commerce review systems is often threatened by fake or emotionally inconsistent content, and this can negatively impact the performance of recommendation systems. Typically, fake review detection (opinion spam) depends on language features and reviewer profile features, such as posting frequency and rating distributions, to separate fake content from genuine reviews [11–13]. Beyond outright fraud, another dimension of review quality is ‘helpfulness,’ where studies have shown that features like review length and sentiment extremity can predict how useful a review is to other users. While effective, these approaches often overlook the more subtle issue of fine-grained sentiment inconsistencies within a single, genuine review.

Progress in artificial intelligence, specifically deep learning, has brought forth more complex methods for detecting problematic content. For example, Phukon (2024) devised FakeDetectionGCN, which incorporates aspect-based sentiment analysis and Graph Convolutional Networks to categorize authentic reviews from fake content [14]. Meanwhile, Periasamy et al. (2024) suggested a hybrid ensemble model combining multiple classifiers for improved fake-review detection [15]. In addition, Karmakar and Hawkins (2024) presented a platform-agnostic framework that introduces sentiment inconsistency analysis, identifying unreliable reviews based on misalignment between textual sentiment and star ratings [16].

At the same time, foundational research on Aspect-Based Sentiment Analysis (ABSA) provides an essential basis for our work. One of the earliest studies in ABSA, Hu and Liu (2004), developed methodologies for extracting product aspects (primarily frequent nouns) and sentiment from review text [7]. More recent surveys provide taxonomies of ABSA tasks and highlight pretrained language models as a method to improve performance [8]. Tun Thura Thet et al. (2021) assess aspect-level sentiment and identify contrast cases (positive overall tone with negative sentiment toward specific features), also highlighting ABSA as an interpretability method [17]. Kumar et al. (2025) further emphasize the necessity of aspect-level sentiment analysis in e-

commerce reviews across multiple granularities [18].

The integration of review sentiment has also been considered in improving recommendation models. McAuley and Leskovec (2013) show that utilizing ratings latent factors and latent topics from review text through joint modeling can improve recommendation accuracy [19]. The development of Neural Collaborative Filtering (NCF) provides a basis for modern recommender models and has led to many content-aware and sentiment-aware extensions [20]. Based on NCF, Chen et al. (2018) proposed a multi-aspect user-interest model incorporating sentiment analysis of features [21]. Ren et al. (2019) introduced a knowledge-enhanced recommender model using aspect-level information formally in recommendations [22]. In a different context, Chan et al. (2020) proposed inconsistency loss in review summarization, showing the benefit of analyzing divergence as a data-quality metric [23]. A comprehensive review by Gupta et al. (2024) provides further context for ensuring review integrity [24].

Despite these advances, limited work has directly applied ABSA to review moderation as a filtering step for improving recommendation quality. While fake-review detection can incorporate ABSA (e.g., Phukon [14]) and sentiment inconsistency tools continue to evolve (e.g., Karmakar & Hawkins [16]), fewer studies investigate internal aspect-level disagreement within a single review as a direct signal for improving recommender data quality. This study aims to address that gap.

3. METHODOLOGY

This section presents the AspectGuard framework for identifying aspect-specific sentiment inconsistencies in user reviews and incorporating this signal into a recommendation system. The objective is to ensure that only sufficiently consistent reviews influence the recommendation process, thereby improving the reliability and interpretability of the generated recommendations. AspectGuard consists of three main components: aspect extraction and sentiment classification, sentiment divergence analysis and review filtering, and recommendation system integration. A high-level representation of the framework is shown in Figure 1.




Figure 1. Conceptual diagram of the AspectGuard pipeline.

3.1. Aspect Extraction and Sentiment Classification

The first stage of the framework captures fine-grained opinions expressed in reviews. A single review may contain multiple product facets, such as *quality*, *price*, or *delivery*, and each facet may carry a different sentiment. To systematically identify these facets, we use spaCy to extract candidate aspects through noun and noun-phrase detection. Generic expressions (e.g., ‘this device’) are removed using a stoplist. The remaining phrases are normalized using a curated synonym dictionary in order to obtain domain-relevant aspect categories.

More specifically, aspect extraction is performed using spaCy-based noun and noun-phrase detection, followed by normalization with a curated synonym dictionary and clustering based on TF-IDF-weighted Sentence-BERT embeddings. Aspect-level sentiment classification is then performed using the pretrained model yangheng/deberta-v3-base-absa-v1.1. In this work, we use this ABSA backbone as an

off-the-shelf sentiment component; the contribution of the paper lies not in introducing a new ABSA architecture, but in exploiting aspect-level sentiment divergence as a filtering signal for recommendation reliability.

For each tracked aspect a in review r , we associate a continuous sentiment score:

$$s_a \in [-1, 1], \quad (1)$$

where $s_a > 0$ indicates positive sentiment, $s_a < 0$ indicates negative sentiment, and $s_a = 0$ indicates *no detected aspect-level sentiment signal* (e.g., the aspect is absent from the relevant review sentences or the evidence is below a confidence threshold). The output is a vector of aspect sentiment scores for each review:

$$\mathbf{S}_r = [s_{a_1}, s_{a_2}, \dots, s_{a_n}], \quad (2)$$

where n is the number of tracked aspects.

3.2. Sentiment Divergence Analysis and Review Filtering

Once aspect-level sentiment vectors are generated, we assess the internal consistency of each review. Reviews that express conflicting sentiments across aspects—for example, praising *sound quality* while criticizing *battery life*—may provide misleading signals if used directly in recommendation models. We therefore define the Sentiment Divergence Score (SDS) as a **weighted variance** of aspect sentiment scores:

$$\text{SDS}(r) = \sum_{i=1}^n w_i (s_{a_i} - \bar{s}_w)^2, \quad \bar{s}_w = \sum_{i=1}^n w_i s_{a_i}, \quad (3)$$

where $w_i \geq 0$ are normalized aspect weights such that $\sum_{i=1}^n w_i = 1$. In our implementation, we set

$$w_i = \frac{m_i}{\sum_{k=1}^n m_k}, \quad (4)$$

where m_i is the empirical mention rate of aspect a_i in the dataset, defined as the percentage of reviews for which $s_{a_i} \neq 0$. A higher SDS indicates stronger disagreement across aspect-level sentiments and therefore a higher likelihood that the review contains internally inconsistent feedback. A threshold τ is then applied to classify reviews as reliable or unreliable:

$$\text{Reliable}(r) = \begin{cases} 1, & \text{if } \text{SDS}(r) \leq \tau, \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

3.3. Recommendation System Integration

In this work, the downstream recommendation module is implemented as a content-based recommender, where user and product profiles are constructed from aspect-level sentiment representations extracted from review text. After unreliable reviews are filtered out, we incorporate aspect-level sentiment into user and product representations. AspectGuard aggregates only *reliable* aspect vectors so that user preferences and item characteristics are not distorted by highly divergent opinions. For each product j , the aggregated aspect sentiment vector $\mathbf{P}_j$ is computed as the mean of all reliable review vectors associated with that product:

$$\mathbf{P}_j = \frac{1}{|R_j|} \sum_{r \in R_j} \mathbf{S}_r, \quad (6)$$

where R_j denotes the set of filtered, reliable reviews for product j . Similarly, each user u is represented by a vector $\mathbf{U}_u$ derived from the aspect vectors of the reliable reviews authored by that user:

$$\mathbf{U}_u = \frac{1}{|R_u|} \sum_{r \in R_u} \mathbf{S}_r, \quad (7)$$

where R_u denotes the set of reliable reviews written by user u .

Ranking score The similarity between user preferences and product characteristics is computed using cosine similarity:

$$s_{uj} = \cos(\mathbf{U}_u, \mathbf{P}_j) = \frac{\mathbf{U}_u \cdot \mathbf{P}_j}{\|\mathbf{U}_u\| \|\mathbf{P}_j\|}. \quad (8)$$

Predicted rating from cosine similarity Cosine similarity is a matching score and must be converted into a predicted star rating for RMSE/MAE evaluation. We use a user-calibrated linear mapping:

$$\hat{r}_{uj} = \text{clip}_{[1,5]}(\mu_u + \alpha s_{uj}), \quad (9)$$

where μ_u is the mean rating given by user u on the training set, computed on reliable reviews, and α is a scaling parameter tuned on a validation split through grid search to minimize MAE. Clipping ensures $\hat{r}_{uj} \in [1, 5]$. For users with insufficient history, μ_u is replaced by the global mean rating.

3.4. Evaluation Metrics

To assess the effectiveness of the proposed framework, we employ standard recommender-system metrics that capture both ranking quality and rating prediction accuracy.

Precision@K measures the proportion of relevant items among the top- K recommendations:

$$\text{Precision@K} = \frac{|\text{Relevant} \cap \text{Recommended@K}|}{K}. \quad (10)$$

Recall@K measures the proportion of relevant items that are successfully retrieved within the top- K recommendations:

$$\text{Recall@K} = \frac{|\text{Relevant} \cap \text{Recommended@K}|}{|\text{Relevant}|}. \quad (11)$$

Normalized Discounted Cumulative Gain (NDCG@K) accounts for the ranking position of relevant items:

$$\text{NDCG@K} = \frac{\text{DCG@K}}{\text{IDCG@K}}, \quad \text{DCG@K} = \sum_{i=1}^K \frac{2^{\text{rel}_i} - 1}{\log_2(i + 1)}. \quad (12)$$

where rel_i denotes the relevance score of the item at position i , and IDCG@K is the ideal DCG for the top- K relevant items.

For rating prediction, we report **Root Mean Squared Error (RMSE)** and **Mean Absolute Error (MAE)**:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{r}_i - r_i)^2}, \quad (13)$$

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |\hat{r}_i - r_i|, \quad (14)$$

where r_i is the ground-truth star rating, $\hat{r}_i$ is the predicted rating computed from the cosine-based predictor, and N is the total number of evaluated interactions.

4. EXPERIMENTS AND RESULTS

4.1. Dataset and Preprocessing

The evaluation used a subset of the Amazon Reviews 2023 dataset, specifically from the Electronics category [25]. We first removed entries with missing review text and then sampled **50,000 clean reviews** from the remaining data to ensure a fixed dataset size. Each record includes a user ID, product ID, rating (1–5), and review text. Ratings were converted to floating-point values. Reviews were truncated to 1,000 characters for computational efficiency during aspect extraction. The dataset was divided into an 80% training set (40,000 reviews) and a 20% test set (10,000 reviews).

4.2. Experimental Setup

The implementation used Python 3.11 with `transformers`, `spacy`, `scikit-learn`, `pandas`, and `sentence-transformers`.

Aspect-level sentiment analysis was performed with `yangheng/deberta-v3-base-absa-v1.1` [26] on a GPU environment using batched processing (batch size = 32). Aspect extraction filtered generic terms using a stoplist and used TF-IDF-weighted Sentence-BERT embeddings for clustering, followed by mapping to a curated synonym dictionary. Reviews were filtered using SDS with threshold $\tau = 0.15$.

For the controlled comparison, we used the same content-based recommendation architecture in both settings. The only difference is whether AspectGuard filtering is applied before constructing user and product aspect profiles. In the **without filtering** setting, all reviews are used. In the **with AspectGuard filtering** setting, reviews with high sentiment divergence ($\text{SDS} > 0.15$) are excluded. This design allows us to isolate the contribution of the proposed filtering mechanism.

Statistical significance. To assess whether the observed improvements are statistically significant, we applied the **Wilcoxon signed-rank test** to paired outcomes between the two settings. For ranking metrics (Precision@10, Recall@10, NDCG@10), we computed each metric **per user** and tested the paired distributions. For rating prediction (RMSE/MAE), we tested paired absolute errors $|\hat{r}_i - r_i|$ across test interactions. Reported improvements are considered significant when $p < 0.01$.

4.3. Results

4.3.1. Aspect-Level Sentiment Analysis

Aspect extraction and mapping yielded five domain-relevant aspects: *sound quality*, *battery life*, *price*, *ease of use*, and *durability*. Aspect-level sentiment was assigned only when an aspect (or one of its synonyms) was detected in relevant sentences; otherwise, the aspect sentiment score was set to 0. Overall, 17,483 out of 50,000 reviews (34.97%) contained at least one non-zero aspect sentiment score. Table 1 presents the aspect-level sentiment statistics across all reviews. ‘Mentions (%)’ denotes the percentage of reviews in which an aspect received a non-zero sentiment score. Therefore, ‘Neutral/Not detected (%)’ includes both explicitly neutral predictions and cases where the aspect was not detected (score 0).

Table 1. Aspect-level sentiment distribution over 50,000 reviews. ‘Mentions (%)’ denotes the percentage of reviews in which the aspect receives a non-zero sentiment score. ‘Neutral/Not detected (%)’ includes cases where the aspect is not detected (score 0).

Aspect	Positive (%)	Negative (%)	Neutral/Not detected (%)	Mentions (%)
Battery Life	11.652	2.206	86.142	13.858
Sound Quality	8.582	1.504	89.914	10.086
Price	12.026	2.460	85.514	14.486
Ease of Use	5.042	0.274	94.684	5.316
Durability	1.696	0.552	97.752	2.248

4.3.2. Sentiment Divergence Detection

The Sentiment Divergence Score (SDS) was computed for each review from its aspect-level sentiment vector. Using threshold $\tau = 0.15$, 27.59% of reviews were flagged as high-divergence ($\text{SDS} > 0.15$). Consequently, 36,203 reviews remained after filtering ($\text{SDS} \leq 0.15$). Table 2 summarizes this outcome.

Table 2. Filtering outcome using SDS threshold $\tau = 0.15$

Category	Number of Reviews	Percentage (%)
Reliable reviews ($\text{SDS} \leq 0.15$)	36,203	72.41
High-divergence reviews ($\text{SDS} > 0.15$)	13,797	27.59

4.3.3. Recommendation System Performance

The recommendation performance was evaluated on the dataset split (80% training, 20% test) under the two settings described above. The **without filtering** setting uses all reviews to build user and product aspect profiles, whereas the **with AspectGuard filtering** setting removes reviews with high sentiment divergence before profile construction. This comparison directly measures the contribution of the proposed filtering mechanism.

On the error metrics (Fig. 2), the recommender with AspectGuard filtering achieves $\text{RMSE} = 0.92$ and $\text{MAE} = 0.68$, compared with $\text{RMSE} = 1.20$ and $\text{MAE} = 0.92$ without filtering, corresponding to reductions of 23.3% and 26.1%, respectively ($p < 0.01$). On the ranking metrics (Fig. 3), the filtered setting attains




Figure 2. Comparison of content-based recommender without filtering and with AspectGuard filtering on error metrics.

Precision@10 = 0.82, Recall@10 = 0.85, and NDCG@10 = 0.60, outperforming the unfiltered setting (0.68, 0.75, and 0.45) with relative gains of 20.6%, 13.3%, and 33.3%, respectively ($p < 0.01$).




Figure 3. Comparison of content-based recommender without filtering and with AspectGuard filtering on ranking metrics.

These results show that filtering reviews with high aspect-level sentiment divergence improves both rating prediction and ranking quality.

5. DISCUSSION

The observed RMSE reduction (0.92 versus 1.20) and the gain in Precision@10 (0.82 versus 0.68) indicate that AspectGuard improves the performance of the recommendation system in the studied setting. Filtering 27.59% of reviews (13,797 out of 50,000) with high SDS (> 0.15) reduces noise and supports cleaner user and product profiles. In Table 1, *price* (Mentions 14.486%) and *battery life* (Mentions 13.858%) appear to be among the most salient aspects in Electronics reviews. In addition, aspect-level explanations, such as alignment on *sound quality*, may enhance user trust by improving interpretability.

AspectGuard is related to existing approaches such as the Review Polarity-wise Recommender [10] and multi-aspect recommender models [21], but differs in that it explicitly targets intra-review aspect sentiment divergence. Unlike approaches that focus on text-rating misalignment [16], AspectGuard examines contradictions inside the review itself, which is directly relevant to review-based representation learning and moderation. The empirical evaluation was intentionally conducted in a text-driven content-based recommendation setting, since the proposed framework is built around review-text analysis and aspect-level sentiment representations. This choice provides the most direct environment for assessing the contribution of the proposed filtering mechanism. Nevertheless, collaborative filtering remains a standard recommendation paradigm, and extending the

evaluation to collaborative filtering baselines would further strengthen the empirical positioning of the work.

Another point concerns the interpretation of unreliable review detection. In the current framework, SDS is used as an unsupervised filtering signal for identifying reviews that are potentially problematic for downstream recommendation. However, the dataset used in this study does not provide gold-standard labels indicating whether a review is truly reliable or unreliable. Therefore, a direct Precision–Recall evaluation is outside the scope of the current dataset and experimental design. Accordingly, the proposed method was assessed primarily through its downstream effect on recommendation performance, which is the main target application of the framework. Constructing a manually validated benchmark for unreliable review detection would provide a stronger basis for direct classification-oriented evaluation in future work.

Several limitations should also be noted. First, the high neutral/not-detected percentages for some aspects (e.g., 97.752% for *durability*) may indicate conservative sentiment assignment or limited aspect coverage. Second, the SDS threshold may also flag reviews that contain useful pros-and-cons information rather than genuinely unreliable content. Third, ABSA inference introduces an additional computational cost. Future work may investigate stronger ABSA models, adaptive thresholds, broader cross-domain validation, and more comprehensive evaluation settings across alternative recommendation architectures.

6. CONCLUSION

This study introduced AspectGuard, a framework that integrates Aspect-Based Sentiment Analysis (ABSA) into the review filtering process to improve the quality and trustworthiness of e-commerce recommendation systems. AspectGuard extracts fine-grained sentiment signals for key product aspects and filters reviews using a Sentiment Divergence Score (SDS). On 50,000 Amazon Reviews 2023 Electronics reviews, 27.59% of reviews were flagged as high-divergence using $\tau = 0.15$, leaving 36,203 reviews for building more reliable recommendation profiles. A controlled comparison using the same content-based recommendation model with and without filtering showed that the proposed filtering mechanism improves both rating prediction and ranking performance. AspectGuard also supports explainability by providing aspect-level justifications for recommendations. Future work will focus on improving aspect coverage, reducing neutral bias, optimizing computation, and validating the approach across additional domains and user behavior signals.

ACKNOWLEDGEMENTS

The authors express their gratitude for the support received from the Laboratory of Information Technologies and Modeling at the Faculty of Sciences Ben M'Sik, Hassan II University, Casablanca, Morocco, which provided the necessary facilities for conducting this research. The authors also acknowledge the support of the CNRST under the ‘PhD-AS-sociate Scholarship - PASS’ Program.

DECLARATIONS

Competing interests: The authors declare that they have no competing interests.

Data availability: The dataset used in this study is publicly available: Amazon Reviews 2023 [25]. The specific experimental subset is described in Section 4.

Code availability: Code used to reproduce the experiments is available from the corresponding author upon reasonable request.

Funding: This work was supported in part by the CNRST under the ‘PhD-AS-sociate Scholarship - PASS’ Program.

Author contributions: *El Mehdi Lghaouch*: Conceptualization, Methodology, Software, Writing—original draft. *Fatima Zahra Abbour*: Software, Validation, Writing—review & editing. *Soumaya Ounacer*: Supervision, Writing—review & editing. *Soufiane Ardchir*: Validation, Writing—review & editing. *Mohamed Azzouazi*: Supervision, Project administration.

Ethics approval: Not applicable (the study uses publicly available, anonymized user review data).

REFERENCES

- [1] M. Meftah, S. Ounacer, and M. Azzouazi, ‘Enhancing customer engagement in loyalty programs through AI-powered market basket prediction using machine learning algorithms,’ in *Synthesis Lectures on Engi-*

- neering, Science, and Technology.. Cham: Springer, 2024, pp. 319–337. doi: 10.1007/978-3-031-50300-9_18.
- [2] C. C. Aggarwal, *Recommender Systems: The Textbook*. Springer, 2016.
- [3] H. Aitelbour, S. Ounacer, Y. Elghomari, H. Jihal, and M. Azzouazi, ‘A crime prediction model based on spatial and temporal data,’ *Periodicals of Engineering and Natural Sciences*, vol. 6, no. 2, pp. 360–364, 2018.
- [4] A. M. S. Osman, Z. A. El-Gohary, and A. E. D. M. El-Nakeeb, ‘The role of trust in e-commerce recommendation systems: a systematic literature review,’ *International Journal of Computer Applications*, vol. 174, no. 43, pp. 24–32, 2021.
- [5] S. M. AL-Ghuribi and S. A. Mohd Noah, ‘A review of aspect-based sentiment analysis: taxonomy, methods, datasets, and challenges,’ *arXiv preprint arXiv:2107.03433*, 2021.
- [6] E. M. Lghaouch, S. Ounacer, S. Ardchir, and M. Azzouazi, ‘Enhancing sentiment analysis through topic modeling: Comprehensive overview,’ in *Studies in Systems, Decision and Control*, vol. 565. Cham: Springer, 2024, pp. 161–179. doi: 10.1007/978-3-031-70996-8_8.
- [7] M. Hu and B. Liu, ‘Mining and summarizing customer reviews,’ in *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2004, pp. 168–177.
- [8] W. Zhang, S. A. G. M. Stilo, and T. L. A. van der Weide, ‘A survey of aspect-level sentiment analysis: methods, performance, and applications,’ *Knowledge and Information Systems*, vol. 64, pp. 2439–2486, 2022.
- [9] L. Qu, G. Ifrim, and G. Weikum, ‘The bag-of-opinions method for review rating prediction from sparse text,’ in *Proceedings of the 23rd International Conference on Computational Linguistics (COLING)*, 2010, pp. 913–921.
- [10] J. Liu, Y. Zhang, Y. Wang, L. Guo, and B. Li, ‘A review polarity-wise recommender,’ in *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR)*, 2021, pp. 2165–2169.
- [11] P. Resnick and H. R. Varian, ‘Recommender systems,’ *Communications of the ACM*, vol. 40, no. 3, pp. 56–58, 1997.
- [12] N. Jindal and B. Liu, ‘Opinion spam and analysis,’ in *Proceedings of the 2008 International Conference on Web Search and Data Mining (WSDM)*, 2008, pp. 219–230.
- [13] S. M. Kim, P. Pantel, T. Chklovski, and M. Pennacchiotti, ‘Automatically assessing review helpfulness,’ in *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2006, pp. 423–430.
- [14] P. Phukon, ‘Detecting fake reviews using aspect based sentiment analysis and graph,’ Master’s Project, San Jose State University, 2024.
- [15] A. Periasamy, R. Karthik, and V. Subramanian, ‘Finding fake reviews in e-commerce platforms by using hybrid algorithms,’ *arXiv preprint arXiv:2404.06339*, 2024.
- [16] A. Karmakar and J. Hawkins, ‘Enhanced review detection and recognition: A platform-agnostic approach with application to online commerce,’ *arXiv preprint arXiv:2405.06704*, 2024.
- [17] T. T. Thet, C. W. Oo, and S. Phyoo, ‘Systematic reviews in sentiment analysis: A tertiary study,’ *Artificial Intelligence Review*, vol. 54, no. 8, pp. 5999–6033, 2021.
- [18] P. Kumar, A. Sharma, and D. K. Sharma, ‘Aspect-based sentiment classification of user reviews to understand customer satisfaction of e-commerce platforms,’ *Electronic Commerce Research*, published online Feb. 20, 2025. doi: 10.1007/s10660-025-09721-3.

- [19] J. McAuley and J. Leskovec, ‘Hidden factors and hidden topics: Understanding rating dimensions with review text,’ in *Proceedings of the 7th ACM Conference on Recommender Systems (RecSys)*, 2013, pp. 165–172.
- [20] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, ‘Neural collaborative filtering,’ in *Proceedings of the 26th International Conference on World Wide Web (WWW)*, 2017, pp. 173–182. doi: 10.1145/3038912.3052569.
- [21] J. Chen, H. Huang, Y. Tian, C. Qu, and L. Wu, ‘Multi-aspect user-interest model based on sentiment analysis and uncertainty theory for recommender systems,’ *Electronic Commerce Research and Applications*, vol. 28, pp. 126–140, 2018.
- [22] J. Ren, S. Ma, H. Chen, and L. Liu, ‘A knowledge-enhanced recommender system with aspect-level information,’ *Expert Systems with Applications*, vol. 138, p. 112810, 2019.
- [23] H. P. Chan, W. Chen, and I. King, ‘A unified dual-view model for review summarization and sentiment classification with inconsistency loss,’ *arXiv preprint arXiv:2006.01592*, 2020.
- [24] R. Gupta, A. Kumar, P. K. Sharma, and H. K. Tripathi, ‘Recent state-of-the-art of fake review detection: a comprehensive review,’ *The Knowledge Engineering Review*, vol. 39, p. e8, 2024. doi: 10.1017/S0269888924000054.
- [25] J. McAuley *et al.*, ‘Amazon Reviews 2023 Dataset,’ 2023. Available: <https://huggingface.co/datasets/McAuley-Lab/Amazon-Reviews-2023>.
- [26] P. He, X. Liu, J. Gao, and W. Chen, ‘DeBERTa: Decoding-enhanced BERT with Disentangled Attention,’ *arXiv preprint arXiv:2006.03654*, 2021.
- [27] M. Honnibal, I. Montani, S. Van Landeghem, and A. Boyd, ‘spaCy: Industrial-strength Natural Language Processing in Python,’ 2020. Available: <https://spacy.io>.


BIOGRAPHIES OF AUTHORS



El Mehdi Lghaouch    received his Bachelor’s degree in Mathematics and Computer Science from the Faculty of Sciences Ben M’sik (FSBM), Hassan II University, Casablanca, Morocco, in 2021. He obtained his Master’s degree in Data Science from the same faculty in 2023. He is currently pursuing a Ph.D. in Data Science at the Faculty of Sciences Ben M’sik, Hassan II University. His research interests include data science, machine learning, and recommender systems. He can be contacted at email: elmehdi.lghaouch-etu@etu.univh2c.ma



Fatima Zahra Abbour    received her Bachelor’s degree in Mathematics and Computer Science from the Faculty of Sciences Ben M’sik (FSBM), Hassan II University of Casablanca, Morocco, in 2022. She obtained her Master’s degree in Data Science and Big Data from the same university in 2024. She is currently pursuing a Ph.D. at Hassan II University of Casablanca. Her research interests lie at the intersection of machine learning, deep learning, and natural language processing, with a particular focus on applying these technologies to recruitment-related challenges. She is an active member of the LTIM Laboratory, contributing to research and collaborative projects. She can be contacted at email: fatima.abbour-etu@etu.univh2c.ma



Soumaya Ounacer    received her Bachelor’s degree in Mathematics and Computer Science from the Faculty of Sciences Ben M’sik (FSBM), Hassan II University, Casablanca, Morocco, in 2008. She obtained her Master’s degree in Computer Engineering from ENSIAS, Mohammed V University, Rabat, in 2010, and her Doctorate in Big Data and Machine Learning from FSBM in 2020. She has been an Associate Professor at FSBM, Hassan II University, since 2018, and is a member of the LTIM Laboratory. She has authored over 20 scientific papers and supervised 4 theses. Her research interests include machine learning, data science, big data, education, and mathematics. She can be contacted at email: soumayaounacer@gmail.com



Soufiane Ardchir    received his Bachelor's degree in Computer Science from the Faculty of Sciences Ben M'sik, Hassan II University, Casablanca, in 2003. He obtained his Master's degree in Networks, Telecommunications, and Multimedia from ENSIAS, Mohammed V University, Rabat, in 2007, and his Doctorate in Computer Engineering in 2020. He has been an Associate Professor at ENCG, Hassan II University, Casablanca, since 2020. He is affiliated with the LTIM Laboratory. He has published 17 scientific papers and supervised 2 theses. His research interests include artificial intelligence, machine learning, and deep learning. He can be contacted at email: s.ardchir@encgcasa.ma



Mohamed Azzouazi    received his Bachelor's degree in Mathematics and Applications from the Faculty of Sciences Ben M'sik (FSBM), Hassan II University, Casablanca, Morocco, in 1990. He obtained a DESA in Information Systems from ENSIAS, Mohammed V University, Rabat, in 1994, and his Doctorate in Big Data Analytics from FSBM in 2017. He has been an Associate Professor at FSBM, Hassan II University, since 2000, and is a member of the LTIM Laboratory. His research interests include artificial intelligence, machine learning, deep learning, and big data solutions. He can be contacted at email: azouazii@gmail.com