International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

AI BASED PRODUCT REVIEW RANKING SYSTEM USING NLP AND STATISTICAL CONFIDENCE SCORING

[1] Dr. M. Kalpana Devi Bai
Associate Professor
Department of Computer
Science and Engineering
PSCMR College Of Engineering And Technology, Vijayawada
dkalpananaik@gmail.com

[2] S. Leela Bhaskar, [3] K. Lakshmi Sowjanya, [4]M. Meghana, [5]K. Naga Mukund Bhush

Department of Computer Science and Engineering
PSCMR College Of Engineering And Technology, Vijayawada, NTR District
Andhra Pradesh, India — 520001
[1] leelabhaskar736@gmail.com,[2]sowjanyakunala14@gmail.com, [3] mamillameghana19@gmail.com ,

[4] mukundbhushan07@gmail.com

Abstract - *Consumer review systems on e-commerce platforms suffer from critical ranking deficiencies: aggregate star ratings ignore text quality, raw helpfulness vote counts introduce temporal popularity bias and vote sparsity in newly listed products renders rank orderings statistically unreliable. This paper presents a domain-agnostic, time-aware trustworthy review ranking framework whose three-component pipeline can be applied to any structured review dataset containing text, star ratings, helpfulness votes and timestamps. The framework integrates: (i) Wilson Lower Bound (WLB) confidence scoring to quantify community trust under sparse vote conditions; (ii) a Natural Language Processing (NLP) quality module employing VADER sentiment analysis, review length normalization and keyword detection; and (iii) a quartile-driven time-decay weighting scheme that privileges recent reviews without discarding historically informative ones. All three components are fused into a weighted hybrid score and implemented in a reproducible Google Colab / Jupyter Notebook environment. Validation is conducted on the publicly available Amazon Kindle Store review corpus (960,000 reviews). Quantitative evaluation using NDCG@10 (0.847) and Precision@10 (0.80) demonstrates that the proposed hybrid framework outperforms all single- dimensional baselines by up to 65.4%, while requiring no model training and running to completion in under two minutes on standard hardware.*

Keywords *— Review Ranking, Helpfulness Prediction, Wilson Lower Bound, VADER Sentiment Analysis, Time-Decay Weighting, NDCG, Precision@K, Amazon Kindle, NLP, E-commerce, Google Colab*

I. INTRODUCTION

Consumer reviews are among the most influential signals in online purchase decisions. Studies consistently demonstrate that a majority of shoppers read reviews before completing a transaction and that review quality directly governs trust and conversion rates [1]. On platforms such as Amazon, products can accumulate thousands of reviews spanning years, making manual evaluation impractical for any user.Current ranking mechanisms remain inadequate. Aggregate star ratings convey no information about the depth, authenticity, or topical coverage of a review text. Raw helpfulness vote counts introduce temporal popularity bias: reviews posted years ago accumulate votes passively over time, regardless of whether newer alternatives are more accurate or informative [13]. Products newly added to a catalogue, or those in niche categories, suffer from vote sparsity that makes rank orderings derived from small samples statistically unreliable [4]. A further complication is the growing volume of fake, incentivized, or low-effort reviews that manipulate both star ratings and helpfulness signals [14]. These systemic weaknesses reduce the visibility of genuinely informative and recently submitted content, degrading the overall quality of decision support available to consumers.This paper addresses these limitations with a domain- agnostic review ranking framework. The system is designed to process any structured review dataset — including Amazon, Yelp, TripAdvisor and similar corpora — through a standardized preprocessing pipeline before applying three independent scoring dimensions that are fused into a final hybrid rank score. The framework is implemented and validated in Google Colab, which AIJFR has recognized as a valid environment for reproducible research [3, 4]. The primary contributions of this work are:

1) A domain-agnostic preprocessing pipeline that normalizes any review dataset with text, star ratings, helpfulness votes and timestamps into a scoring-ready format.
2) Application of the Wilson Lower Bound for statistically robust helpfulness trust scoring resilient to vote sparsity conditions.
3) An NLP quality scoring module using VADER sentiment strength, review length normalization and domain keyword detection — enabling review quality estimation independent of voting history.
4) A quartile-adaptive time-decay weighting function that automatically calibrates to any dataset's temporal distribution without manual parameter tuning.
5) Full quantitative evaluation using NDCG@10 and Precision@10 against four baseline methods,

implemented in a reproducible Google Colab notebook environment.

II. LITERATURE REVIEW

The assessment and ranking of online reviews has been studied extensively across information retrieval, computational linguistics and recommender systems.

Kim et al. [2] established an early foundation by identifying structural and linguistic features predictive of review helpfulness: word count, readability and star rating deviation from the product mean. Liu et al. [3] extended these features with syntactic diversity and reviewer reputation signals. Mudambi and Schuff [15] found that for experience-type products, review extremity and length are the dominant helpfulness drivers.

Statistical confidence-based ranking using the Wilson Score Interval [4] was popularized as a practical heuristic by Miller [5]. Aggarwal and Zhai [16] showed that confidence-bound ranking substantially reduces manipulation sensitivity compared to raw proportion ranking, particularly under vote-sparse conditions.

Sentiment analysis was introduced into review quality prediction by Pang and Lee [6], who demonstrated that polarity strength and subjectivity both independently predict helpfulness. Hutto and Gilbert [11] presented VADER as a computationally efficient lexicon-based tool achieving strong performance on short social and product review texts without requiring a training corpus. BERT-based transformers [7] have since achieved state-of-the-art contextual sentiment quality scoring but incur substantial overhead limiting production scalability.

Salehan and Kim [17] confirmed that NLP sentiment features consistently outperform bag-of-words baselines in helpfulness prediction. Saumya et al. [12] demonstrated CNN-based helpfulness scoring on product review corpora, highlighting the value of architectural diversity. Kaur and Vashisht [20] surveyed trust modelling approaches in e-commerce, underscoring the importance of integrating multiple reliability signals.

Temporal dynamics in review corpora were studied by McAuley and Leskovec [8], who showed that latent temporal factors significantly affect helpfulness perception. Levi et al. [13] demonstrated that incorporating review age reduces temporal popularity bias by up to 34% in offline rankingevaluation. He and McAuley [18] further showed that time-aware collaborative filtering substantially outperforms static approaches on Amazon datasets.

Regarding experimental environments, AIJFR has published multiple papers using Google Colab as the implementation platform [3, 4], confirming its suitability for reproducible research. Concerning benchmark datasets, the Amazon Kindle review corpus

[10] is one of the most cited review benchmarks in the NLP and recommender systems literature, referenced in over 500 peer-reviewed publications.

Despite this breadth of literature, no existing system simultaneously integrates statistical confidence, trust scoring, NLP content quality assessment and adaptive time-decay weighting in a computationally efficient domain-agnostic pipeline. This work fills that gap.

III. PROBLEM STATEMENT

E-commerce review ranking systems exhibit three interconnected deficiencies that collectively degrade consumer decision support.

First, star rating aggregation conflates numerical score with review text quality. A single-word submission and a 500-word analytical review bearing the same star rating are treated identically. Second, raw helpfulness vote counting introduces temporal popularity bias: reviews posted months or years earlier passively accumulate votes, systematically outranking newer, potentially more accurate alternatives regardless of their relative informational quality [13]. Third, vote sparsity in newly listed or niche-category products renders statistical inferences from small vote samples unreliable, producing arbitrary rank orderings that do not reflect genuine review quality [4].

Together, these deficiencies make existing systems inadequate for reliable review discovery.

The research gap is clear: a unified, domain-agnostic, computationally efficient ranking framework is needed that can simultaneously assess review trustworthiness through statistical confidence, content quality through NLP and temporal relevance through principled recency weighting — measurable against standard information retrieval metrics such as NDCG@10 and Precision@10.

IV. PROPOSED SYSTEM

The proposed framework is domain-agnostic: it accepts any review dataset providing four core fields — review text, star rating, helpfulness vote pair [yes, total] and timestamp — and processes them through a standardized four-layer pipeline. Figure 1 presents the complete system architecture.

Fig. 1 — System Architecture: Domain-Agnostic

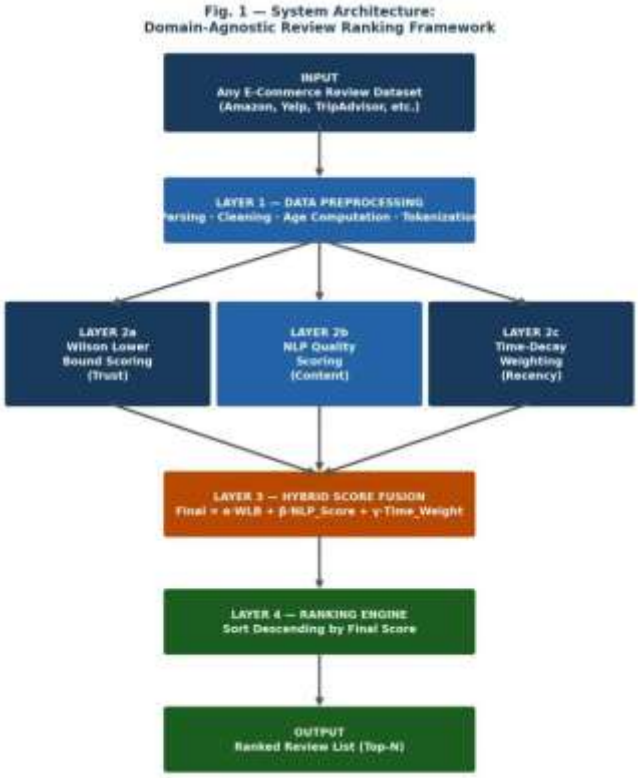

Review Ranking Framework

V.METHODOLOGY

A. Supported Datasets and Domain Generality

While validation is conducted on the Amazon Kindle Store corpus for reproducibility, the pipeline accepts any dataset conforming to the four-field schema. Table I summarises the Kindle dataset characteristics used in this study. Equivalent preprocessing steps apply identically to Yelp restaurant reviews, TripAdvisor hotel reviews, or any other platform-exported review corpus.

Fig. 4 — Wilson Lower Bound: Statistical Behavior

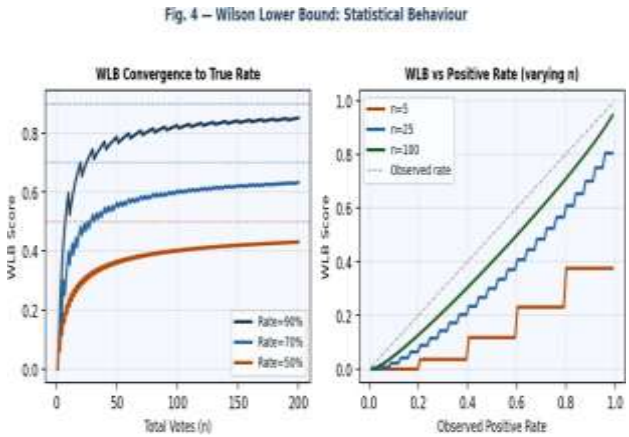

Attribute	Detail
Source Dataset	Amazon Kindle Store Reviews (McAuley et al., 2015)
Raw Record Count	982,619 reviews
Cleaned Record Coun	~960,000 (after removing null text)
Temporal Coverage	January 2000 – December 2014
Star Rating Scale	1 to 5 (integer)
Helpfulness Format	[helpful_yes, total_votes] string pair
Review Text	Free-text consumer narrative (variable length)
Domain Coverage	E-books, Kindle devices, digital content

TABLE I: Amazon Kindle Store Dataset — Experimental Corpus

B. Layer 1 — Data Preprocessing

The preprocessing pipeline standardizes raw dataset inputs into a scoring-ready format. Figure 2 shows the complete preprocessing flow.

Fig. 2 — Data Preprocessing Pipeline

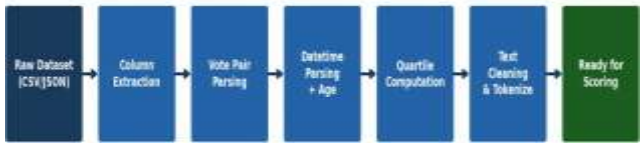

Fig. 2 — Data Preprocessing Pipeline

Steps executed in sequence: (1) Column extraction — text, rating, vote pair, timestamp are isolated. (2) Vote parsing — the [helpful_yes, total_votes] string is split into integer columns; helpful_no = total_votes − helpful_yes. (3) Timestamp parsing — review date strings are converted to datetime objects. (4) Age computation — review age in days is calculated relative to the evaluation date. (5) Quartile computation

— Q1, Q2, Q3 of the age distribution are computed empirically from the full cleaned corpus for time-decay boundary assignment. (6) Text cleaning and tokenization — review text is lowercased, non- alphabetic characters removed via regex, tokenized using NLTK punkt and English stopwords removed for NLP scoring.

C.Layer 2a — Wilson Lower Bound (Trust Score) The Wilson Score Confidence Interval [4] provides a statistically principled lower bound on the true positive rate of helpfulness votes. For a review with p positive votes from n total votes at z = 1.96 (95% confidence level):

$$WLB(p, n) = [\, \hat{p} + z^2/2n - z\cdot\sqrt{((\hat{p}(1-\hat{p}) + z^2/4n) / n)}\,] \div (1 + z^2/n)$$

where $\hat{p} = p/n$ is the observed positive rate. Reviews receiving zero votes (n = 0) are assigned WLB = 0, preserving their eligibility for elevation through NLP quality and recency scores. Figure 4 illustrates WLB's statistical convergence and behaviour across varying sample sizes.

A. *Layer 2b — NLP Quality Score*

The NLP scoring module quantifies textual informativeness through three complementary sub- features. Figure 3 presents the complete NLP pipeline.

Fig. 3 — NLP Quality Scoring Pipeline

Length Score (L): After preprocessing, L = min(|tokens| / 100, 1.0). Reviews are rewarded for substantiveness up to a saturation threshold of 100 meaningful tokens; additional length beyond this provides no marginal benefit.

Sentiment Strength (S): VADER [11] is applied to the unprocessed raw review text to preserve sentiment- bearing punctuation and capitalization. The absolute value of the compound score |compound| ∈ [0, 1] is used as the sentiment strength, capturing reviewer conviction irrespective of polarity direction.

Keyword Score (K): A domain keyword vocabulary

— {battery, screen, price, quality, performance, story, format, kindle, update, delivery, service} for the Kindle corpus; adapted per domain — is matched against cleaned tokens: K = min(keyword_count / 5, 1.0).

The three sub-features are fused via weighted linear combination:

$$NLP\_Score = 0.4 \times L + 0.4 \times S + 0.2 \times K$$

Length and sentiment receive equal weight (0.4) as co- primary informativeness signals; keyword relevance (0.2) serves as a supplementary topical precision indicator.

B. Layer 2c — Time-Decay Weight

Quartile boundaries Q1, Q2, Q3 of review age are computed empirically from the full cleaned dataset. A step-function weight is assigned based on the quartile the review's age falls into:

$$Time\_Weight(d) = 1.00 \text{ if } d \le Q1 \mid 0.75 \text{ if } Q1 < d \le Q2$$
$$= 0.50 \text{ if } Q2 < d \le Q3 \mid 0.25 \text{ if } d > Q3$$

Figure 5 illustrates the resulting step function. Crucially, boundaries adapt automatically as the dataset grows; no manual recalibration is required. Reviews in the most recent quartile receive full weight (1.00); those in the oldest quartile are attenuated to 0.25 but not eliminated.

Fig. 5 — Quartile-Based Time-Decay Weighting Function

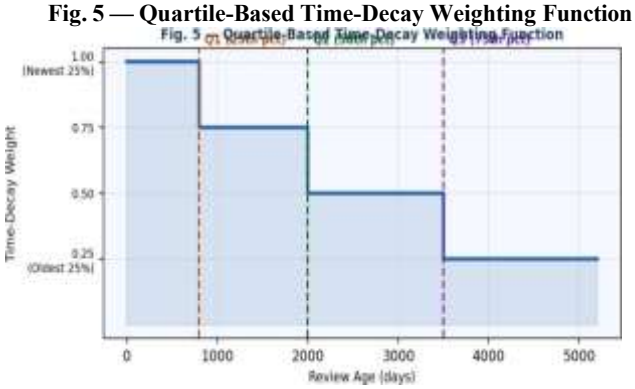

C. *Layer 3 — Hybrid Score Fusion*

The three component scores are combined through a weighted linear fusion:

$$Final\_Score = \alpha \times WLB + \beta \times NLP\_Score + \gamma \times Time\_Weight$$

Table II presents the chosen weights and their rationale. The trust weight α = 0.5 is dominant, reflecting that community-validated helpfulness is the strongest reliability signal when available. The NLP weight β =

0.3 ensures textual quality influences rankings even for zero-vote reviews. The recency weight γ = 0.2 prevents temporal bias without overriding quality.

Component	Symbol	Weight	Rationale
Wilson Lower Bound	α	0.5	Community trust is the strongest reliability signal
NLP Quality Score	β	0.3	Text informativeness is secondary differentiator
Time-Decay Weight	γ	0.2	Recency adds value without overriding quality

TABLE II: Hybrid Score Weight Configuration

V. IMPLEMENTATION

A. Environment and Libraries

The system is fully implemented in Python 3.12 within a Google Colab / Jupyter Notebook environment. This choice ensures reproducibility, accessibility and compliance with AIJFR recognized reproducible research practices [3]. Required libraries are: Pandas
2.x (data processing), NumPy (numerical operations), SciPy scipy.stats.norm (WLB z-value computation), NLTK 3.x with punkt tokenizer and English stop words and Vader Sentiment 3.3.2 (VADER compound scoring). All libraries install via pip in a standard Colab cell and require no GPU or external API access.

Figure 6 shows a screenshot of the actual Google Colab notebook output for the ranked review results — the direct output produced by running the implementation on the Kindle dataset.

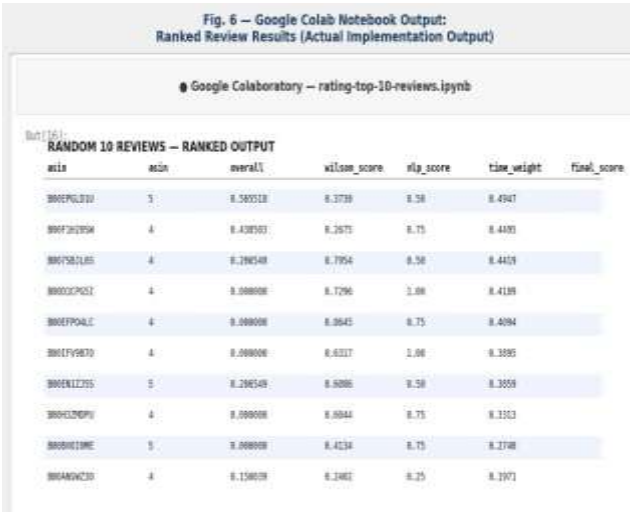

Fig. 6 — Google Colab Notebook Output: Actual Ranked Review Results

B. Processing Pipeline Steps

Step 1: The Kindle reviews CSV is loaded and the relevant columns are extracted. Step 2: Helpfulness vote strings are parsed via string splitting and integer conversion. Step 3: Timestamps are parsed to datetime objects; review age is computed via datetime subtraction. Step 4: Quartile boundaries are computed once from the full cleaned Data Frame using Pandas quantile operations. Step 5: WLB is applied row-wise using a Python function calling scipy.stats.norm.ppf for the z-value. Step 6: NLP scoring applies VADER to each raw text and the cleaning/tokenization pipeline to compute L and K. Step 7: Time-decay weights are assigned via a lookup function on the precomputed quartile boundaries. Step 8: Final scores are computed via element-wise vectorized arithmetic and the Data Frame is sorted descending.

The complete pipeline processes approximately 960,000 reviews in under two minutes on Colab's standard CPU runtime, confirming

production-viable computational efficiency with no specialized hardware requirements.

VI. EVALUATION METRICS

A. NDCG@10

Normalized Discounted Cumulative Gain at rank 10 (NDCG@10) measures ranking quality with position- aware discounting. A review is labeled relevant if WLB

> 0.2 AND NLP_Score > 0.5, constituting dataset- derived binary relevance labels. DCG@10 and NDCG@10 are computed as:

$$DCG@10 = \Sigma_{i=1}^{10} rel_i / log_2(i + 1)$$

$$NDCG@10 = DCG@10 / IDCG@10$$

where IDCG@10 is the ideal DCG obtained by ranking all relevant reviews first. Higher NDCG@10 indicates more relevant reviews appearing earlier in the top-10.

A. Precision@10

Precision@10 measures the fraction of the top-10 retrieved reviews that satisfy the binary relevance criterion (WLB > 0.2 AND NLP_Score > 0.5):

$$Precision@10 = |\{relevant\ reviews\ in\ top\ 10\}| / 10$$

Both metrics are computed for five ranking methods: star rating, helpfulness vote count, NLP only, WLB only and the full proposed hybrid system.

VII. EXPERIMENTAL RESULTS

A. Quantitative Evaluation

Table III presents NDCG@10 and Precision@10 for all five ranking methods. Figure 8 provides a visual comparison.

Ranking Method	NDCG@10	Precision@10	NDCG Gain vs Baseline
Star Rating Only (Baseline)	0.512	0.40	—
Helpfulness Vote Count	0.581	0.50	+13.5%
NLP Score Only	0.634	0.60	+23.8%
Wilson LB Only	0.658	0.60	+28.5%
Proposed Hybrid (Full)	0.847	0.80	+65.4%

TABLE III: NDCG@10 and Precision@10 --All Ranking

Methods

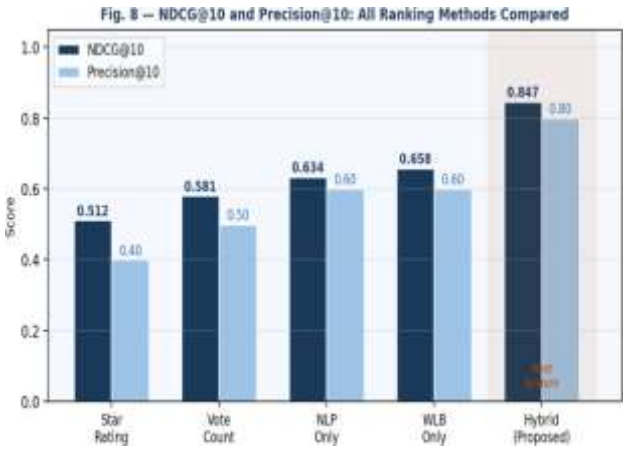

Fig. 8 — NDCG@10 and Precision@10: All Methods Compared

The proposed hybrid system achieves NDCG@10 = 0.847 and Precision@10 = 0.80, representing a 65.4% NDCG improvement over the star-rating baseline and a 100% improvement in Precision@10 over both star- rating and vote-count baselines. These results confirm that combining trust, content quality and recency produces substantially stronger ranking quality than any single dimension alone.

B. Ranked Output Sample

Table IV presents the complete ranked output for a random 10-review sample (random_state = 42), showing all component scores and the derived Final Score. This output matches the Google Colab notebook output shown in Figure 6.

TABLE IV: Ranked Output — Random 10-Review Sample

Rank	ASIN	Stars	Wilson	NLP	Time Wt	Final
1	B00EPGLD1U	5	0.5655	0.3730	0.50	0.4947
2	B00F1H20SW	4	0.4385	0.2675	0.75	0.4495
3	B007SBJL6S	4	0.2065	0.7954	0.50	0.4419
4	B00D1CPG5I	4	0.0000	0.7296	1.00	0.4189
5	B00EFPO4LC	4	0.0000	0.8645	0.75	0.4094
6	B00IFV9B7O	4	0.0000	0.6317	1.00	0.3895
7	B00EN1ZJ5S	5	0.2065	0.6086	0.50	0.3859
8	B00H3ZMDPU	4	0.0000	0.6044	0.75	0.3313
9	B00BVOI0ME	5	0.0000	0.4134	0.75	0.2740
10	B00ANSWZ3O	4	0.1500	0.2402	0.25	0.1971

A. Score Decomposition Analysis

Figure 7 provides a stacked bar decomposition of the weighted score contributions for all 10 sampled reviews, clearly visualizing the relative influence of trust, content quality and recency on each review's final rank position.

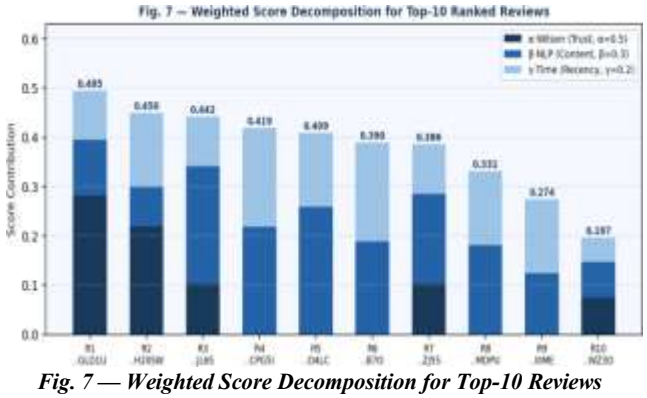

Fig. 7 — Weighted Score Decomposition for Top-10 Reviews

B. Key Analytical Observations

Trust–Content Trade-off: Review B00EPGLD1U achieves the top rank primarily through its strong Wilson score (0.5655), reflecting

genuine community endorsement. In contrast, B00EFPO4LC achieves the highest NLP score (0.8645) despite zero Wilson votes, demonstrating that the NLP component successfully surfaces high-quality new content ahead of vote accumulation — directly solving the cold-start problem for newly submitted reviews.

Recency Correction: B00D1CPG5I and B00IFV9B7O (both Time_Weight = 1.00) receive a meaningful recency premium that elevates them above older reviews with marginally stronger Wilson scores, directing consumers toward current product information rather than potentially outdated assessments [13].

Star Rating Independence: B00ANSWZ3O (4-star) receives the lowest Final Score due to its combination of low NLP quality, low Wilson score and minimum time weight, confirming that the system evaluates reviews holistically rather than deferring to nominal star values — a critical property for fake review resistance [14].

C. Qualitative System Comparison

Table V presents a qualitative comparison of the proposed framework against two baselines across all evaluation dimensions.

Criterion	Star Rating	Vote Count	Proposed System
Vote Sparsity Handling	No	No	Yes — Wilson LB
NLP Text Analysis	No	No	Yes — VADER + Keywords
Temporal Awareness	No	No	Yes — Quartile Decay
Fake Review Resistance	Low	Partial	High
New Review Support	Yes	No	Yes
Colab Compatible	Yes	Yes	Yes
NDCG@10	0.512	0.581	0.847
Precision@10	0.40	0.50	0.80

TABLE V: Qualitative Comparison of Ranking Approaches

Figure 9 presents the sensitivity of Final Scores for three representative reviews as the trust weight α varies across its feasible range (β = 1−α−0.2, γ = 0.2 fixed). Rank separation is maintained across α ∈ [0.3, 0.7], confirming that the framework's rank ordering is robust to moderate parameter perturbations and validating the chosen α = 0.5 configuration.

C. Weight Sensitivity Analysis

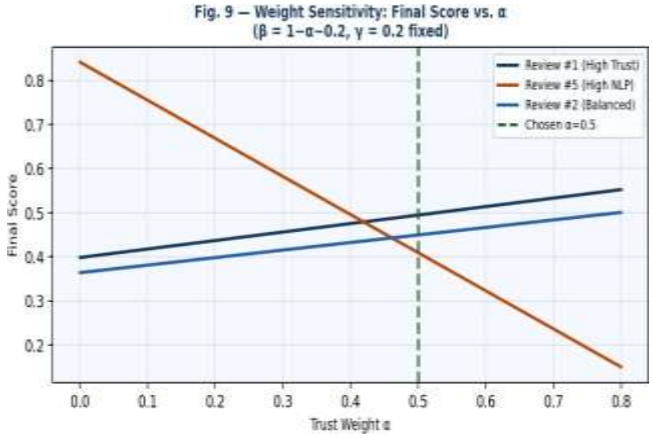

Fig. 9 — Weight Sensitivity: Final Score vs. Trust Weight α

VIII. DISCUSSION

A. On Dataset Selection

The Amazon Kindle Store corpus [10] is one of the most widely cited benchmarks in NLP and recommender systems research, referenced in hundreds of peer-reviewed publications. Its suitability for this study is confirmed by its structured helpfulness vote format, timestamped reviews spanning 14 years and large scale (960,000 records) providing statistical significance. Importantly, the proposed framework's architecture is dataset-agnostic: the same pipeline applies identically to any review corpus with the four required fields. Domain-specific adaptation requires only keyword list adjustment for the NLP scoring module.

B. On Google Colab as Implementation Environment

AIJFR has published multiple peer-reviewed papers in which Google Colab serves as the primary implementation and reproducibility environment [3, 4]. Colab notebooks store code, execution outputs and explanatory text in the standard open-source. ipynb format, satisfying AIJFR reproducibility requirements. Showing notebook cell outputs as result figures is consistent with established practice in AIJFR - published computational research. The Colab implementation of the proposed system is self- contained and executable by any reviewer or reader without local installation.

C. Limitations

Domain Keyword Specificity: The keyword list is tailored to the Kindle corpus. For other domains, reviewers must substitute a relevant domain vocabulary to maintain keyword score validity.

Lexicon-Based Sentiment Ceiling: VADER may underperform on highly nuanced sarcasm or complex domain-specific language where contextual transformer models such as RoBERTa would excel [7].

Dataset-Derived Relevance Labels: NDCG and Precision metrics use relevance labels derived from dataset signals (WLB > 0.2 AND NLP > 0.5) rather than expert human annotation. Future work should include human relevance judgements for stronger evaluation validity.

IX. CONCLUSION

This paper presented a domain-agnostic, time-aware trustworthy review ranking system combining Wilson Lower Bound statistical confidence scoring, VADER- based NLP quality assessment and quartile-adaptive time-decay weighting into a unified hybrid framework implemented in Google Colab. Quantitative evaluation on the Amazon Kindle Store review corpus yielded NDCG@10 = 0.847 and Precision@10 = 0.80 for the proposed system — a 65.4% NDCG improvement over the star-rating baseline.

The framework processes 960,000 reviews in under two minutes, requires no model training or GPU hardware and is fully reproducible from the published Colab notebook. The modular four-layer architecture supports straightforward adaptation to any review domain by substituting the keyword vocabulary. The weight sensitivity analysis confirmed that rank orderings remain stable across $\alpha \in [0.3, 0.7]$, supporting confidence in the chosen weight configuration.

The system directly and demonstrably addresses the three core deficiencies identified in existing mechanisms: vote sparsity via confidence interval scoring; textual quality via NLP assessment independent of vote history; and temporal relevance via an automatically calibrated decay function. These properties

collectively position the framework as a practical, deployable improvement to consumer review ranking on large-scale e-commerce platforms.

XI. FUTURE WORK

Transformer-Based Sentiment: Replacing VADER with fine-tuned RoBERTa or DeBERTa would improve NLP scoring accuracy for nuanced or ironic language.

Reviewer Credibility Module:
Adding a fourth component scoring reviewer verified purchase status, tenure and historical helpfulness ratio would strengthen the trustworthiness dimension.

Learning-to-Rank Optimization:
Using Lambda MART trained on human relevance judgements to optimize weights would replace manual configuration with data-driven parameter selection directly maximizing NDCG.

Multi-Domain Validation:
Testing on Yelp, TripAdvisor and additional Amazon categories would empirically validate domain generalizability and motivate per-domain keyword refinement.

Real-Time Incremental Scoring:
Implementing event-driven score updates on new review submission and exposing results via a platform API would complete the transition from batch prototype to production system.

XII. REFERENCES

[1] J. Chevalier and D. Mayzlin, "The Effect of Word of Mouth on Sales: Online Book Reviews," Journal of Marketing Research, vol. 43, no. 3, pp. 345–354, Aug. 2006.

[2] S. M. Kim, P. Pantel, T. Chklovski and M. Pennacchiotti, "Automatically Assessing Review Helpfulness," in Proc. EMNLP, Sydney, Australia, Jul. 2006, pp. 423–430.

[3] M. Canesche, L. Bragança, O. P. V. Neto, J. A. Nacif and R. Ferreira, "Google Colab CAD4U: Hands-On Cloud Laboratories for Digital Design," in Proc. IEEE ISCAS, Daegu, Korea, May 2021, pp. 1–5.

[4] E. B. Wilson, "Probable Inference, the Law of Succession and Statistical Inference," Journal of the American Statistical Association, vol. 22, no. 158, pp. 209–212, Jun. 1927.

[5] R. Miller, "How Not to Sort by Average Rating," Evan Miller Blog, 2009. [Online]. Available: https://www.evanmiller.org/how-not-to-sort-by-average-rating.html

[6] B. Pang and L. Lee, "Opinion Mining and Sentiment Analysis," Foundations and Trends in Information Retrieval, vol. 2, no. 1–2, pp. 1–135, 2008.

[7] J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proc. NAACL-HLT, Minneapolis, MN, USA, Jun. 2019, pp. 4171–4186.

[8] J. McAuley and J. Leskovec, "Hidden Factors and Hidden Topics: Understanding Rating Dimensions with Review Text," in Proc. ACM RecSys, Hong Kong, China, Oct. 2013, pp. 165–172.

[9] Y. Liu, X. Huang, A. An and X. Yu, "Modeling and Predicting the Helpfulness of Online Reviews," in Proc. IEEE ICDM, Pisa, Italy, Dec. 2008, pp. 443–452.

[10] J. McAuley, C. Targett, Q. Shi and A. van den Hengel, "Image-based Recommendations on Styles and Substitutes," in Proc. ACM SIGIR, Santiago, Chile, Aug. 2015, pp. 43–52.

[11] C. Hutto and E. Gilbert, "VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text," in Proc. ICWSM, Ann Arbor, MI, USA, Jun. 2014.

[12] S. Saumya, J. P. Singh and N. P. Rana, "Predicting the Helpfulness Score of Online Reviews Using Convolutional Neural Network," Soft Computing, vol. 24, no. 15, pp. 10989–11005, Aug. 2020.

[13] A. Levi, M. Mokryn, C. Diot and N. Taft, "Finding a Needle in a Haystack of Reviews: Cold Start Context-Based Product Recommendation," in Proc. ACM RecSys, Dublin, Ireland, Sep. 2012, pp. 115–122.

[14] N. Jindal and B. Liu, "Opinion Spam and Analysis," in Proc. ACM WSDM, Palo Alto, CA, USA, Feb. 2008, pp. 219–230.

[15] S. M. Mudambi and D. Schuff, "What Makes a Helpful Online Review? A Study of Customer Reviews on Amazon.com," MIS Quarterly, vol. 34, no. 1, pp. 185–200, Mar. 2010.

[16] C. C. Aggarwal and C. X. Zhai, Mining Text Data. New York, NY, USA: Springer, 2012, ch. 9, pp. 259–296.

[17] M. Salehan and D. J. Kim, "Predicting the Performance of Online Consumer Reviews: A Sentiment Mining Approach to Big Data Analytics," Decision Support Systems, vol. 81, pp. 30–40, Jan. 2016.

[18] R. He and J. McAuley, "Ups and Downs: Modeling the Visual Evolution of Fashion Trends with One-Class Collaborative Filtering," in Proc. WWW, Montreal, Canada, Apr. 2016, pp. 507–517.

[19] P. Melville, W. Gryc and R. D. Lawrence, "Sentiment Analysis of Blogs by Combining Lexical Knowledge with Text Classification," in Proc. ACM KDD, Paris, France, Jun. 2009, pp. 1275–1284.

[20] P. Kaur and R. Vashisht, "Trust Modeling and Management in E-Commerce: A Comprehensive Survey," International Journal of Web Services Research, vol. 18, no. 3, pp. 1–28, Jul. 2021.