ICCK Transactions on Intelligent Systematics
http://dx.doi.org/10.62762/TIS.2025.106283

RESEARCH ARTICLE

# Topic Mining and Sentiment Analysis for Consumer Reviews of Automotive Spare Parts on E-commerce Platforms

**Minling Zeng**[1,*] and **Xiaoli Deng**[2]

[1] School of Management, Xiamen University Tan Kah Kee College, Zhangzhou 363105, China
[2] School of Digital Economy and Management, Software Engineering Institute of Guangzhou, Guangzhou 510990, China

## Abstract

**Consumer satisfaction in automotive spare parts e-commerce is shaped by multidimensional factors that are difficult to quantify through conventional survey methods. This paper presents an intelligent text mining framework that integrates multiple machine learning and natural language processing techniques to systematically extract and analyze consumer sentiment from 1,236 validated Taobao reviews. Specifically, TF-IDF (Term Frequency-Inverse Document Frequency) is employed for discriminative keyword extraction, semantic network analysis is applied to model inter-word relational structures, and LDA (Latent Dirichlet Allocation) topic modeling is used to discover latent thematic patterns, identifying four core satisfaction dimensions: Logistics, Quality, Price, and Customer Service. For sentiment quantification, a domain-adaptive SnowNLP model based on the Naive Bayes algorithm is fine-tuned on a domain-specific automotive e-commerce corpus, achieving an accuracy improvement from 11.4% to 83.1%. Additionally, a domain-specific sentiment lexicon is constructed via the SO-PMI (Semantic Orientation-Pointwise Mutual Information) method to enhance lexicon-level sentiment scoring. The integrated intelligent analysis reveals that positive sentiment is predominantly driven by logistics efficiency and product reliability, while negative sentiment concentrates on installation complexity and service knowledge deficiencies. The proposed system provides an automated, domain-adaptable intelligent pipeline for e-commerce opinion mining, with demonstrated applicability to the automotive spare parts domain, offering data-driven decision support for product optimization and service quality improvement.**

**Keywords**: machine learning, natural language processing, LDA topic modeling, text classification, sentiment analysis, text mining.

## 1 Introduction

With the rapid proliferation of e-commerce platforms and the exponential growth of user-generated content, online consumer reviews have emerged as a rich source of unstructured textual data that

**Academic Editor:**
Jianlei Kong

**Submitted:** 22 April 2025
**Accepted:** 01 June 2025
**Published:** 09 July 2025

**Vol.** 2, **No.** 3, 2025.
10.62762/TIS.2025.106283

***Corresponding author:**
Minling Zeng
zengml@xujc.com

**Citation**
Zeng, M., & Deng, X. (2025). Topic Mining and Sentiment Analysis for Consumer Reviews of Automotive Spare Parts on E-commerce Platforms. *ICCK Transactions on Intelligent Systematics*, 2(3), 137–148.

encodes multidimensional satisfaction signals. In the context of expanding automotive spare parts market—driven by the nation's strong support for new energy vehicles and the continued growth of vehicle ownership—e-commerce platforms such as Taobao host millions of consumer reviews covering products ranging from center console storage boxes and floor mats to spoilers and steering wheel accessories. Manually analyzing this volume of text is neither scalable nor objective, creating an urgent need for intelligent, automated approaches to extract actionable insights from consumer feedback.

Natural language processing (NLP) and machine learning components can be systematically integrated into an intelligent text analytics system capable of transforming unstructured review text into structured, actionable knowledge. Specifically, term frequency-based methods such as TF-IDF enable discriminative keyword extraction; probabilistic generative models such as Latent Dirichlet Allocation (LDA) [1] enable unsupervised discovery of latent thematic structures; and supervised classifiers such as Naive Bayes, as implemented in the SnowNLP framework, enable domain-adaptive sentiment polarity detection. The architectural integration of these components into a unified intelligent pipeline constitutes the primary system-level contribution of this work, enabling scalable and reproducible consumer opinion mining across e-commerce domains. The primary contribution of this work is the design and validation of an intelligent text analytics system that integrates multiple machine learning components into a unified, domain-adaptive pipeline, with automotive spare parts e-commerce as the demonstration domain.

Despite significant advances in NLP-based review analysis, existing research on automotive spare parts e-commerce remains limited in two key respects. First, most studies rely on questionnaire-based primary data rather than large-scale organic review corpora, limiting both objectivity and scalability. Second, existing work has not fully integrated multi-method intelligent pipelines that simultaneously address topic discovery and fine-grained sentiment quantification. For instance, Zhang et al. [2] employed a modified Theory of Reasoned Action model to identify purchase intention determinants for remanufactured new energy vehicle components in China, including subjective norms, attitudes, face consciousness, brand extension, and perceived risk. Mohamad Sheriff et al. [3] applied the AutoSERVPERF framework to demonstrate that responsiveness, empathy, assurance, and reliability are key determinants of automotive after-sales service satisfaction. Gajewska et al. [4] measured e-commerce service quality satisfaction using the SERVQUAL method, finding that reliability is the most critical dimension influencing consumer satisfaction in online purchasing contexts. Similar to Celuch [5], who applied LDA topic modeling and sentiment analysis to mine multidimensional consumer experiences from online ticket-purchasing reviews, this study adopts a comparable integrated NLP framework to decode the multidimensional satisfaction signals embedded in automotive spare parts e-commerce reviews.

To address these gaps, this paper presents an intelligent text mining framework that integrates TF-IDF keyword extraction, semantic network analysis, LDA topic modeling, and a domain-adaptive SnowNLP sentiment classifier to systematically analyze 1,236 validated consumer reviews of automotive spare parts collected from Taobao. The framework further incorporates a domain-specific sentiment lexicon constructed via the SO-PMI method [6] to enhance lexicon-level sentiment scoring. This integrated intelligent analysis pipeline addresses the following research questions:

**RQ1:** What are the key thematic factors identifiable through machine learning-based topic modeling that influence consumer satisfaction in automotive spare parts e-commerce?

**RQ2:** What sentiment orientations do consumers express regarding their purchasing experiences, as quantified by domain-adaptive NLP models?

**RQ3:** How can the outputs of the intelligent text analytics pipeline inform actionable strategies for automotive spare parts e-commerce enterprises to enhance consumer satisfaction?

The research contributions of this paper are as follows. First, it proposes an integrated intelligent text analytics framework combining TF-IDF, semantic network analysis, LDA, SO-PMI, and domain-adaptive Naive Bayes classification for consumer review mining, providing a replicable methodological pipeline for e-commerce review analytics. Second, it demonstrates the effectiveness of domain-adaptive model fine-tuning, improving SnowNLP sentiment classification accuracy from 11.4% to 83.1% through corpus-specific retraining. Third, it enriches consumer behavior theory by revealing the latent thematic and sentiment structures embedded in organic automotive e-commerce review data, complementing traditional survey-based approaches. Finally, the data-driven insights generated by the intelligent pipeline provide actionable decision-support guidance for automotive spare parts e-commerce enterprises in optimizing product design, service quality, and pricing strategies.

The remainder of this paper is organized as follows. Section 2 describes the intelligent text mining methodology and data preprocessing pipeline. Section 3 presents the LDA topic modeling and sentiment analysis results. Section 4 discusses theoretical and practical implications. Section 5 concludes the paper.

## 2 Research Methodology

This study constructs an intelligent text analytics pipeline for mining consumer satisfaction factors from automotive spare parts reviews on Taobao. The pipeline integrates four computational stages: data acquisition and preprocessing, TF-IDF-based keyword extraction with semantic network analysis, LDA-based unsupervised topic modeling, and domain-adaptive supervised sentiment classification. Figure 1 illustrates the overall architecture of the proposed framework.

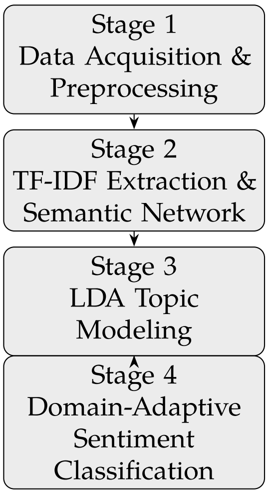


**Figure 1.** Overall architecture of the proposed intelligent text analytics pipeline.

### 2.1 Data Acquisition and Corpus Construction

Product reviews from flagship stores of automotive spare parts on Taobao were selected as the data source. As the largest e-commerce platform in China, Taobao hosts a substantial volume of automotive spare parts transactions and organic user-generated review text, making it suitable for large-scale NLP-based analysis. Four popular product categories were selected as crawling targets: automotive storage boxes, safety protection products, decorative items, and center console accessories. Text reviews posted between January and December 2023 were collected from 10 stores selected based on the top 10% sales ranking, yielding a raw corpus of 1,408 reviews (250 for automotive storage boxes, 258 for safety protection products, 600 for decorative items, and 300 for center console accessories). Sampling proportions were stratified across product categories to reflect their relative sales distribution on the platform, ensuring representativeness of the review corpus.

It is necessary to acknowledge two platform-specific biases inherent in the corpus. First, user demographic bias: Taobao's user base is predominantly composed of younger, tech-savvy urban consumers, which may underrepresent older or rural populations and limit generalizability. Second, review motivation bias: consumers tend to leave reviews more frequently under extreme satisfaction or dissatisfaction, leading to underrepresentation of neutral experiences in the dataset. These biases are acknowledged as limitations of the current study.

### 2.2 Data Preprocessing Pipeline

Raw review text undergoes a four-stage preprocessing pipeline before feature extraction and model training.

**Stage 1 — Data Cleaning.** Invalid and duplicate entries are removed manually, including default system-generated reviews such as "This user did not provide a review" and "Default positive rating." After cleaning, 1,236 valid reviews are retained for subsequent analysis.

**Stage 2 — Chinese Word Segmentation.** Chinese text does not contain natural word boundaries, requiring dedicated segmentation tools. The jieba library, a Python-based Chinese word segmentation toolkit, is employed in precise segmentation mode to maximize accuracy. To address domain-specific colloquialisms and compound expressions common in consumer reviews (e.g., segmentation errors such as splitting "overall effect" into "overall" and "effect"), a custom domain dictionary is incorporated into the jieba segmentation engine to enhance segmentation precision.

**Stage 3 — Stop Word Removal.** High-frequency function words, conjunctions, punctuation marks, and other semantically vacuous tokens are filtered using a stop word list. This step reduces noise in the feature space and improves the quality of subsequent topic modeling, consistent with standard preprocessing practices in the LDA literature, which has evolved considerably in both model formulation and application methodology [7].

**Stage 4 — N-gram Construction.** After segmentation and filtering, unigram, bigram, and trigram token sequences are constructed to capture both individual term semantics and multi-word expression patterns relevant to consumer satisfaction analysis.

## 2.3 TF-IDF Keyword Extraction and Semantic Network Analysis

### *2.3.1 TF-IDF Algorithm*

Term Frequency-Inverse Document Frequency (TF-IDF) is employed as the primary feature weighting scheme for discriminative keyword extraction [8]. For a term $t$ in document $d$ within corpus $D$, the TF-IDF score is computed as:

$$\text{TF-IDF}(t, d, D) = \text{TF}(t, d) \times \text{IDF}(t, D) \tag{1}$$

where

$$\text{TF}(t, d) = \frac{f_{t,d}}{\sum_{t' \in d} f_{t',d}}, \quad \text{IDF}(t, D) = \log \frac{|D|}{|\{d \in D : t \in d\}|} \tag{2}$$

The core property of TF-IDF is that terms with high frequency within a specific document but low frequency across the corpus receive high scores, ensuring that domain-discriminative terms are prioritized over ubiquitous function words. Python's `sklearn.feature_extraction.text.TfidfVectorizer` is used to extract and rank keywords from the preprocessed review corpus.

TF-IDF serves as an indicator to evaluate word importance, further refining the word frequency analysis based on review texts. As previously mentioned, keyword cloud graphs display high-frequency keywords (often deemed important). However, some meaningless words—such as "storage" and "silicone"—exist among high-frequency terms and lack representativeness. Thus, the IDF algorithm reduces the impact of such words. Even if a word appears frequently across all documents, its importance is not necessarily high. The core idea of TF-IDF is: if a word has a high Term Frequency (TF) in an article and rarely appears in others, it has strong category-discriminating power and is suitable for classification [9, 10].

This paper uses Python to extract keywords from review texts using the TF-IDF algorithm, and these keywords are sorted by their TF-IDF scores, as partially shown in Table 1. Based on the top 40 keywords, a general categorization was conducted to analyze the influencing factors of consumer product purchases from the following six aspects.

The top 40 TF-IDF keywords are categorized into six semantic dimensions: (1) **Product Characteristics** (e.g., "installation," "material," "quality," "odor"), reflecting concerns about physical attributes and usability; (2) **User Experience** (e.g., "convenience," "feel," "worthwhile"), capturing ease of use and overall satisfaction; (3) **Cost-Effectiveness** (e.g., "value," "recommendation"), indicating price-performance evaluation; (4) **Customer Service** (e.g., "service," "support"), highlighting post-purchase service expectations; (5) **Product Evaluation** (e.g., "overall," "satisfaction"), expressing holistic post-use assessments; and (6) **Product Category** (e.g., "car mats"), specifying the reviewed product type. These six dimensions collectively characterize the key factors driving consumer purchasing decisions and satisfaction in the automotive spare parts e-commerce context.

Mining customer opinions from online reviews to reveal relationships between product attributes and consumer evaluations has been demonstrated as an effective approach in the automotive industry [11], and semantic network analysis extends this by explicitly modelling inter-word relational structures to uncover co-occurrence patterns among high-frequency terms in review corpora. Observing such diagrams aids in identifying potential information. By intuitively displaying the size of each node and the thickness of connection lines, they facilitate discovering relationships among these high-frequency words in review texts, as shown in Figure 2.

As can be seen from the diagram, the keyword with the largest node in the network diagram is "Material." Keywords such as "Effect," "Evaluation," and "Odor" unfold around it with high correlation, indicating that product features and applicability are important factors influencing consumers' purchase decisions. Additionally, keywords like "Support" and "Evaluation," though with smaller nodes, are close to the central node, suggesting that the product's customer service is also a concern for consumers. Furthermore, words such as "Value" and "Quality" in the diagram show high correlation, indicating that price and quality also have a certain impact on consumers' purchase behavior.

**Table 1.** TF-IDF value ranking table (Partial).

| Serial Number | Keyword | TF-IDF Value | Serial Number | Keyword | TF-IDF Value |
|---|---|---|---|---|---|
| 1 | Installation | 0.364212413 | 21 | Stench | 0.084134175 |
| 2 | Good | 0.328233984 | 22 | Mat | 0.080259574 |
| 3 | Very | 0.303879354 | 23 | Comparison | 0.080259574 |
| 4 | Snug | 0.292809068 | 24 | Recommendation | 0.078045517 |
| 5 | Material | 0.280631753 | 25 | Suitable | 0.075277946 |
| 6 | Craftsmanship | 0.228047894 | 26 | Experience | 0.072510374 |
| 7 | None | 0.192069464 | 27 | Feeling | 0.070849831 |
| 8 | Effect | 0.171589435 | 28 | Situation | 0.06808226 |
| 9 | Can | 0.157198063 | 29 | Value | 0.066975231 |
| 10 | Convenient | 0.146127777 | 30 | Size | 0.064761174 |
| 11 | Quality | 0.145574263 | 31 | Simple | 0.061440088 |
| 12 | Odor | 0.12730829 | 32 | Perfect | 0.059226031 |
| 13 | Item | 0.117345033 | 33 | Worthwhile | 0.057565488 |
| 14 | Overall | 0.113470433 | 34 | Purchase | 0.055904945 |
| 15 | Use | 0.106274747 | 35 | Taste | 0.054797916 |
| 16 | Service | 0.10406069 | 36 | Received | 0.049684296 |
| 17 | Support | 0.101846632 | 37 | Practical | 0.048592333 |
| 18 | Product | 0.095757975 | 38 | Especially | 0.046954389 |
| 19 | Satisfaction | 0.090776346 | 39 | Texture | 0.044770464 |
| 20 | Goods | 0.089115803 | 40 | Price | 0.0442244835 |

#### *2.3.2 Semantic Network Analysis*

Semantic network analysis is applied to model relational structures among high-frequency keywords, visualizing co-occurrence patterns as a weighted graph. Nodes represent keywords, edge weights reflect co-occurrence frequency, and node sizes reflect term centrality. This graph-based representation enables identification of central semantic concepts and their associative neighborhoods, providing complementary structural insights beyond univariate keyword ranking.

### 2.4 LDA Topic Modeling

#### *2.4.1 Model Formulation*

Latent Dirichlet Allocation (LDA) is a generative probabilistic model that represents each document as a mixture over $K$ latent topics, where each topic is characterized by a probability distribution over vocabulary terms. Given a corpus of $M$ documents with vocabulary size $V$, LDA assumes the following generative process for each document $d$:

1. Draw topic distribution $\theta_d \sim \text{Dir}(\alpha)$
2. For each word position $n$ in document $d$:
   (a) Draw topic assignment $z_{d,n} \sim \text{Multinomial}(\theta_d)$
   (b) Draw word $w_{d,n} \sim \text{Multinomial}(\phi_{z_{d,n}})$, where $\phi_k \sim \text{Dir}(\beta)$ is the word distribution for topic $k$

where $\alpha$ and $\beta$ are Dirichlet concentration hyperparameters controlling document-topic and topic-word sparsity, respectively [12, 13]. Model inference is performed using the gensim library in Python.

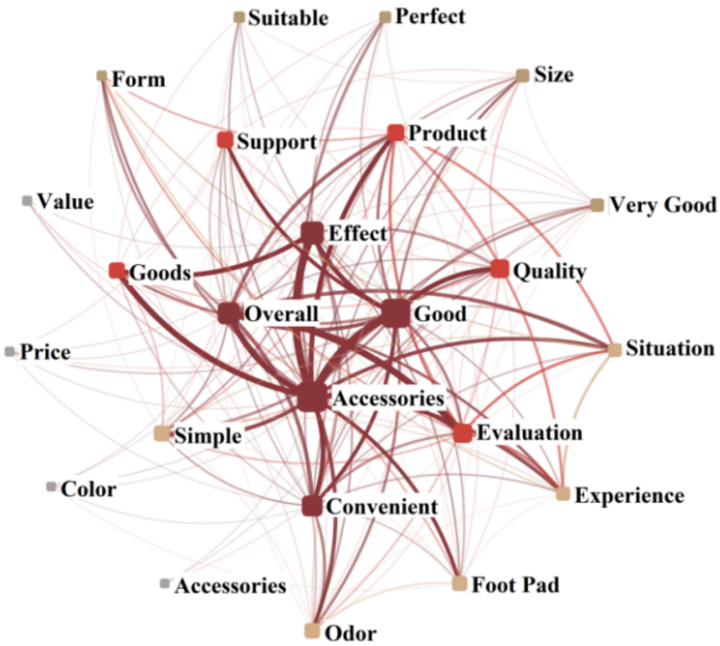


**Figure 2.** Semantic network diagram.

### 2.4.2 *Optimal Topic Number Selection*

Determining the optimal number of topics $K$ requires balancing model fit and interpretability. Perplexity, defined as the exponentiated average negative log-likelihood per word, is used as an initial quantitative criterion for topic number selection, consistent with the approach adopted by Pavlinek and Podgorelec [14] in their LDA-based text classification framework.

$$\text{Perplexity}(D_{\text{test}}) = \exp\left(-\frac{\sum_d \log p(\mathbf{w}_d)}{\sum_d N_d}\right) \tag{3}$$

However, since perplexity decreases monotonically with increasing $K$ without an obvious elbow point, it is supplemented by visual inspection of LDA topic clustering bubble charts using pyLDAvis. The optimal $K$ is selected as the value at which topic bubbles are non-overlapping and distributed across all quadrants of the visualization, indicating well-separated and semantically distinct topics. Based on this combined criterion, $K = 4$ is selected as the optimal number of topics.

## 2.5 Domain-Adaptive Sentiment Classification

### 2.5.1 *SnowNLP Naive Bayes Classifier*

Sentiment polarity classification is performed using SnowNLP, a Python library for Chinese NLP that implements a Naive Bayes probabilistic classifier for sentiment analysis [15]. Sentiment mining of online consumer reviews has been shown to yield actionable insights for business decision-making through big data analytics approaches [16]. For a review text represented as a feature vector $\mathbf{x}$, the Naive Bayes classifier assigns sentiment label $y$ according to:

$$\hat{y} = \arg\max_{y\in\{+,-\}} P(y)\prod_{i=1}^{n} P(x_i \mid y) \tag{4}$$

where $P(y)$ is the prior probability of sentiment class $y$ and $P(x_i \mid y)$ is the likelihood of feature $x_i$ given class $y$, estimated from the training corpus under the conditional independence assumption.

### 2.5.2 *Domain-Adaptive Fine-Tuning*

The default SnowNLP model, trained on general-domain Chinese text, achieved only 11.4% accuracy on the automotive spare parts corpus—a result substantially below random chance (50%)—indicating that the model's sentiment predictions were systematically inverted relative to the domain-specific ground truth labels. This inversion likely reflects a systematic mismatch between the sentiment conventions in the general training corpus and those in automotive spare parts reviews, where domain-specific expressions such as "snug fit" or "no odor" carry strong positive valence absent from general-domain training data. To address this, domain-adaptive fine-tuning is performed by constructing a domain-specific training corpus of automotive e-commerce reviews. Positive and negative sentiment corpora are collected from related e-commerce review sources, manually annotated, and stored as separate positive (`pos.txt`) and negative (`neg.txt`) files. The SnowNLP sentiment model is retrained on this domain corpus using an 8:2 train-test split. After retraining, classification accuracy improves from 11.4% to 83.1%, demonstrating the effectiveness of domain-adaptive fine-tuning for specialized review corpora.

#### 2.5.3 *Sentiment Scoring and Polarity Classification*

The retrained SnowNLP model outputs a continuous sentiment score $s \in [0, 1]$ for each review, where values approaching 1 indicate strong positive sentiment and values approaching 0 indicate strong negative sentiment. Binary sentiment polarity is determined by the threshold:

$$\text{Polarity} = \begin{cases} \text{Positive} & \text{if } s \geq 0.5 \\ \text{Negative} & \text{if } s < 0.5 \end{cases} \tag{5}$$

Only binary classification (positive and negative) is performed in this study, with neutral sentiment excluded from the analysis, consistent with the binary training corpus design.

#### 2.5.4 *SO-PMI Sentiment Lexicon Construction*

To complement the supervised classifier with lexicon-level sentiment scoring, a domain-specific sentiment lexicon is constructed using the Semantic Orientation-Pointwise Mutual Information (SO-PMI) method. For each candidate sentiment word $w$, its semantic orientation score is computed as:

$$\text{SO-PMI}(w) = \text{PMI}(w, w_{\text{pos}}) - \text{PMI}(w, w_{\text{neg}}) \tag{6}$$

where $w_{\text{pos}}$ and $w_{\text{neg}}$ are positive and negative seed words, respectively, and PMI is estimated from co-occurrence statistics in the review corpus. Words with $\text{SO-PMI}(w) > 0$ are assigned positive polarity and those with $\text{SO-PMI}(w) < 0$ are assigned negative polarity, constructing a domain-specific automotive e-commerce sentiment lexicon that captures industry-specific sentiment expressions absent from general-purpose lexicons [19].

## 3 Results and Discussion

### 3.1 LDA Topic Modeling Results

The optimal number of topics $K = 4$ is determined by combining perplexity analysis (Figure 3) with pyLDAvis visualization (Figure 4), as detailed in Section 2. The extracted topics and their representative keywords are presented in Table 2.

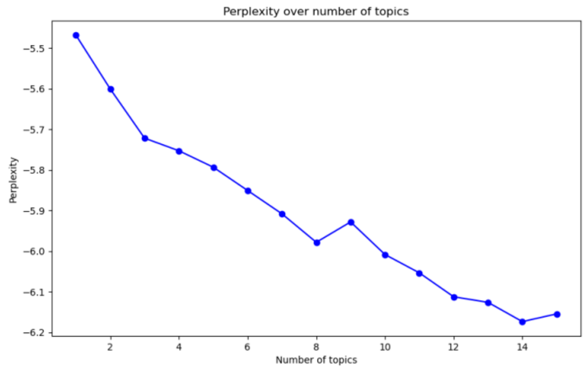


**Figure 3.** Graph of perplexity variation in the LDA topic model.

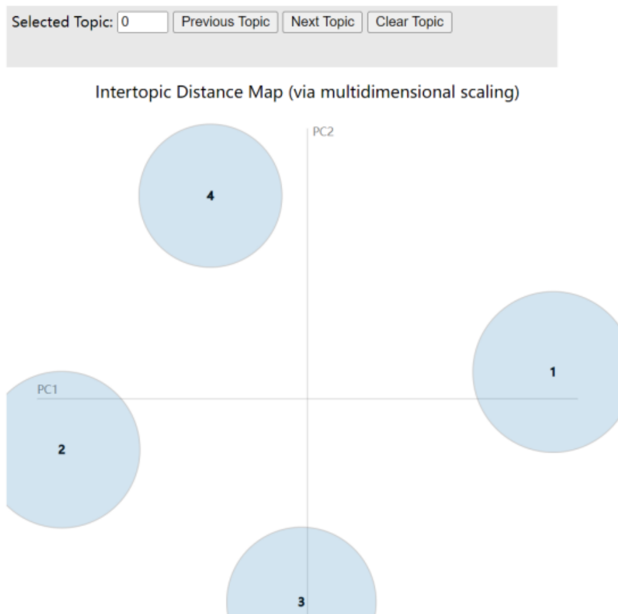


**Figure 4.** Visualization of LDA topic clustering (with 4 Topics).

The four topics are interpreted as follows. **Topic 1 (Logistics):** Keywords such as "Delivery," "Fast," and "Speed" indicate that consumers prioritize shipping timeliness and packaging integrity when evaluating automotive spare parts purchases. **Topic 2 (Quality):** Terms including "Material," "Odor," and "Design" reflect high consumer requirements for product practicality, material quality, and aesthetics. **Topic 3 (Price):** Keywords such as "Value," "Affordable," and "Cost-effective" capture consumers' price-performance evaluations and sensitivity to promotional pricing. **Topic 4 (Customer Service):** Terms including "Enthusiasm," "Service Attitude," and "Support" reveal that pre- and post-purchase merchant service

**Table 2.** Keywords extracted by LDA model.

| Topic Category | Topic Keywords |
|---|---|
| Logistics | Topic #1: Logistics, Product, Delivery, Experience, Situation, Form, Time, Degree, Order, Manual, Fast, Speed, Convenience |
| Quality | Topic #2: Effect, Material, Odor, Texture, Goods, Mat, Taste, Stench, Quality, Stain-resistant, Size, Design, Snug Fit, Aesthetics |
| Price | Topic #3: Price, Value, Pricing, Affordable, Shopping, Hand Feel, Seller, Color, Mat, Time, Aesthetic Appeal, Elegant, Cost-effective |
| Customer Service | Topic #4: Support, Enthusiasm, Original Vehicle, Service Attitude, Evaluation, Store, Storage, Space, Influence, Cup, Attitude, Seller, Issue, Service, Details |

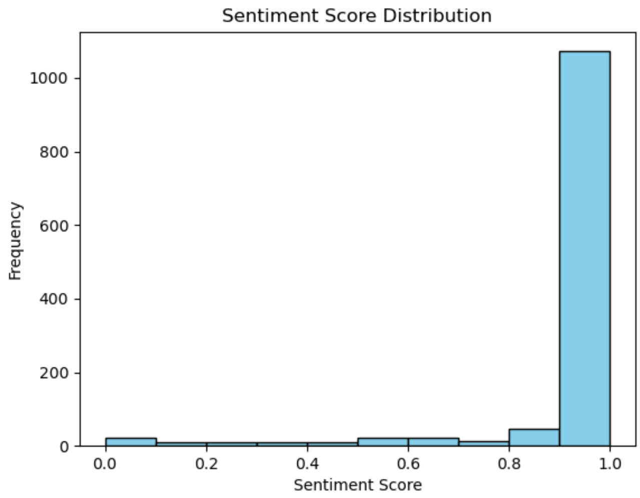


**Figure 5.** Sentiment score distribution produced by the domain-adaptive SnowNLP classifier.

quality significantly influences overall consumer satisfaction.

### 3.2 Sentiment Classification Results

The domain-adaptive SnowNLP Naive Bayes classifier, retrained on the automotive e-commerce corpus, improved classification accuracy from 11.4% (default model) to 83.1% (retrained model), confirming the necessity of domain-specific retraining for specialized review corpora. The classifier was applied to all 1,236 reviews, and the distribution of sentiment scores is shown in Figure 5. Results indicate a predominance of positive sentiment (score $\geq 0.5$) over negative sentiment (score $< 0.5$). This finding is consistent with prior work employing Naive Bayes continuous learning frameworks for e-commerce sentiment classification [17] and sentiment analysis of new energy vehicle reviews [18], both of which report predominantly positive consumer attitudes in automotive e-commerce contexts.

**Positive sentiment drivers.** Qualitative analysis of high-scoring reviews reveals two primary positive drivers. First, logistics efficiency: reviews such as "The logistics were very fast. After pasting it on, it looks very beautiful. It can not only decorate but also protect the door handle" and "I received the product quickly. It's better than I imagined" indicate that fast delivery enhances brand perception and repurchase intention. Second, product quality: reviews such as "Good quality, no strange smell, and the right size. It's worth having!" and "The quality is very good. I've bought a lot of decorations here and the quality is all quite good" demonstrate that consistent product quality drives customer loyalty. While lexicon-based approaches rely on predefined sentiment dictionaries [19], the domain-adaptive SnowNLP model captures industry-specific language patterns more effectively.

**Negative sentiment drivers.** Two primary negative drivers are identified. First, inadequate customer service knowledge: reviews such as "I couldn't install the product even after watching the instructional video. When I asked customer service for help, they said they didn't know how either" indicate that service staff unfamiliarity with products leads to dissatisfaction and discourages repeat purchases. Second, price-value mismatch: reviews such as "The product is average, but the price is too high—others sell it for just over 10 RMB" suggest that perceived overpricing reduces satisfaction and brand loyalty.

### 3.3 Cross-Method Synthesis

The convergent findings across TF-IDF keyword extraction, semantic network analysis, LDA topic modeling, and domain-adaptive sentiment classification collectively identify four key consumer satisfaction dimensions in automotive spare parts e-commerce, as summarized in Table 3.

It is noted that TF-IDF categorization yields six semantic dimensions, while LDA topic modeling converges on four latent topics. This discrepancy arises because TF-IDF produces keyword-level categories based on surface-form co-occurrence, whereas LDA operates at the document level to identify underlying thematic structures. The "Product Evaluation" and

**Table 3.** Cross-method synthesis of consumer satisfaction factors.

| Dimension | Sentiment | Key Signals |
|---|---|---|
| Logistics | Positive | Fast delivery, intact packaging |
| Quality | Positive | Material quality, odor-free, snug fit |
| Customer Service | Negative | Poor product knowledge, slow response |
| Price | Negative | Overpricing, low cost-effectiveness |

"Product Category" dimensions identified by TF-IDF are not sufficiently discriminative at the document level to constitute independent LDA topics, and are instead distributed across the four core themes.

Negative reviews align with expectation-disconfirmation theory [20], where gaps between advertised and actual product performance or service quality drive consumer dissatisfaction. These data-driven findings, produced by the integrated intelligent text analytics pipeline, provide the empirical basis for the actionable recommendations presented in Section 4.

## 4 Discussion

### 4.1 Topic Mining and Sentiment Analysis Findings

As detailed in Section 3, the integrated intelligent pipeline consistently identifies four consumer satisfaction dimensions: Logistics, Quality, Price, and Customer Service. Negative reviews consistently align with expectation-disconfirmation theory, suggesting that managing consumer expectations through accurate product descriptions is a critical intervention point for automotive spare parts enterprises.

Through the analysis of positive and negative reviews, the positive influencing factors for consumers when purchasing automotive spare parts products can be summarized as follows: fast logistics; guaranteed quality. Conversely, the negative influencing factors include: low service quality of after-sales support and relatively high prices of some products compared to the market.

### 4.2 Recommendations for Automotive Spare Parts Enterprises

Based on the discussion in Section 4.1, the existing issues facing e-commerce enterprises in the automotive spare parts industry are identified as follows: low service levels of after-sales customer service, incompatible and poor-quality accessories, difficult product installation, and product prices higher than market averages. To address these, the following recommendations are proposed for automotive spare parts e-commerce enterprises:

#### *4.2.1 Strengthen After-Sales Service Management*

Customer service plays a pivotal role throughout the transaction process. The problems identified earlier highlight deficiencies in after-sales service within this sector. Customer service staff should possess a solid understanding of the products sold, and after-sales teams must provide clear solutions to issues raised by consumers post-purchase. To achieve this, professional training is essential for after-sales personnel before they assume their roles. This training should cover communication skills, problem-solving abilities, and in-depth product knowledge to prevent disputes arising from insufficient expertise. Merchants must also regularly evaluate customer satisfaction and service quality, collect and analyze feedback, and continuously improve after-sales service by addressing issues promptly.

#### *4.2.2 Enhance Product Applicability*

As highlighted in previous analyses, consumers place significant importance on product applicability, which directly influences market demand and sales volume for automotive accessories. Improving product applicability can help enterprises retain customers and foster loyalty. To enhance applicability, merchants should first determine whether compatibility issues stem from customers purchasing incorrect accessories or inherent product defects. They must then conduct in-depth research into the structures, specifications, and characteristics of vehicles across different brands, models, and types to better understand customer needs and design more suitable accessories. Additionally, offering customized accessory products and services based on specific customer requirements can further enhance product relevance.

#### *4.2.3 Optimize Product Design for Ease of Installation*

As noted earlier, installation is a necessary but error-prone process for automotive spare parts. The first step is to analyze the root causes of installation difficulties, which may include complex steps, unclear instructions, poor component quality, or flawed design. Regardless of the cause, collecting user feedback is critical to identifying specific installation challenges and user expectations. Enterprises should ensure that product installation manuals are clear and detailed, providing step-by-step instructions and diagrams to guide users through the process. If design issues contribute to installation difficulties, revising the product structure or design to improve ease of assembly should be a priority.

#### *4.2.4 Conduct Market Research for Rational Pricing*

Price is often a decisive factor in consumers' purchasing decisions, influenced by multiple uncertain factors. Merchants must consider not only product costs and profits but also market competitiveness, as lower prices for comparable quality products offer a distinct advantage. To this end, enterprises should conduct market research to adjust prices while maintaining cost effectiveness. The process involves:

1. Defining research objectives, such as understanding market price levels, competitors' pricing strategies, and customer price sensitivity;
2. Collecting data through market surveys, competitive analysis, and customer interviews;
3. Analyzing data to assess market prices, competitor tactics, and customer price tolerance;
4. Flexibly adjusting pricing strategies based on market feedback and competitor dynamics.

This structured approach ensures that pricing aligns with market expectations, enhancing product competitiveness and customer appeal.

## 5 Theoretical and Practical Implications

### 5.1 Theoretical Implications

Firstly, this study identifies key factors influencing consumer satisfaction using review data through technical methods such as the TF-IDF algorithm, semantic network, and LDA algorithm. This provides a new perspective for consumer behavior theory, particularly regarding the driving factors of consumer satisfaction in the e-commerce context. Although some scholars have mentioned research methods for topic mining and sentiment analysis using consumer review data in previous studies [21–23], research on consumer satisfaction in the growing online retail market for automotive spare parts remains relatively rare.

Secondly, this study reveals the relationship between sentiment tendencies and consumer satisfaction by calculating sentiment scores of consumer review texts and analyzing reviews with different sentiment orientations. Prior work has demonstrated that sentiment analysis of e-commerce product reviews can effectively capture consumers' evaluative attitudes toward products and services [24], and this study extends such inquiry to the automotive spare parts domain using a domain-adaptive Naive Bayes approach.

Finally, through analyzing the sentiment tendencies and topic content of consumer reviews, this study enables the analysis of consumers' evaluations and emotional inclinations toward the products and services they receive, providing a theoretical basis for e-commerce enterprises in the automotive spare parts industry to improve service quality and user satisfaction.

### 5.2 Practical Implications

Consumers' online reviews reflect their perceptions and consumption experiences from multiple dimensions. The findings of this study urge e-commerce enterprises in the automotive spare parts industry to mine more attributes from online reviews and conduct in-depth research into consumers' evaluations and sentiment tendencies toward their shopping experiences. However, due to the open structure and complex information of online reviews, analyzing them remains a significant challenge.

This study employs text data mining methods to provide a practical approach for automotive spare parts e-commerce enterprises, helping them understand the discussion topics and sentiment tendencies embedded in online reviews. Enterprises can then improve their products and services based on the content of these reviews.

Timely and effective implementation of service recovery measures through online reviews is an effective means to retain consumers and enhance their satisfaction. Nevertheless, given limited resources and the need to follow the principle of prioritization, these enterprises should focus on online text reviews with high negative sentiment. Such reviews reflect consumer dissatisfaction, which is often left unaddressed by merchants in the actual review interface. By addressing these reviews and taking appropriate actions, enterprises can reduce consumer dissatisfaction and enhance customer loyalty.

## 6 Conclusion

This paper presented an intelligent text analytics framework for mining consumer satisfaction factors from automotive spare parts e-commerce reviews on Taobao. The framework integrates five NLP and machine learning components: TF-IDF discriminative keyword extraction, semantic network analysis for relational structure modeling, LDA unsupervised topic modeling, SO-PMI domain-specific sentiment lexicon construction, and a domain-adaptive SnowNLP Naive Bayes classifier for sentiment polarity detection.

Applied to a corpus of 1,236 validated consumer reviews, the intelligent pipeline yielded three key findings. First, LDA topic modeling with $K = 4$ identified four latent consumer satisfaction dimensions: Logistics, Quality, Price, and Customer Service. Second, domain-adaptive fine-tuning of the SnowNLP classifier on an automotive e-commerce corpus improved sentiment classification accuracy from 11.4% to 83.1%, demonstrating the critical importance of domain-specific model retraining. Third, cross-method synthesis revealed that positive sentiment is predominantly driven by logistics efficiency and product quality, while negative sentiment concentrates on inadequate customer service knowledge and price-value mismatch. Based on these data-driven findings, targeted recommendations are proposed for automotive spare parts enterprises covering after-sales service training, product applicability enhancement, installation design optimization, and data-driven pricing strategies.

The current study has two main limitations: the corpus is restricted to a single platform and excludes neutral sentiment. Future research will incorporate transformer-based models such as BERT for richer contextual understanding, expand the corpus across multiple e-commerce platforms, and integrate aspect-level sentiment analysis to enable finer-grained consumer satisfaction mining in the automotive e-commerce domain.

## Data Availability Statement

Data will be made available on request.

## Funding

This work was supported without any funding.

## Conflicts of Interest

The authors declare no conflicts of interest.

## Ethical Approval and Consent to Participate

This study exclusively utilized publicly available user-generated review texts sourced from the Taobao e-commerce platform. The data collection process did not involve direct interaction with human participants, recruitment of subjects, administration of interventions, or access to any private, sensitive, or personally identifiable information. All reviews analyzed in this study were voluntarily posted by users in a public-facing online environment and contained no information that could be used to identify individual users. As this study involved only publicly available, anonymized textual data with no direct human participant involvement, it was not subject to formal institutional ethical review under applicable guidelines. Data were handled responsibly and used solely for academic research purposes in accordance with relevant data protection principles.

## References


[1] Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, *3*(Jan), 993-1022.

[2] Zhang, S., Zhou, C., & Liu, Y. (2020). Consumer purchasing intentions and marketing segmentation of remanufactured new-energy auto parts in China. *Mathematical Problems in Engineering*, *2020*(1), 5647383. [CrossRef]

[3] Mohamad Sheriff, N., Roslan, N. I., & Yusoff, Y. (2020). Determinants of satisfaction for automotive after sales service quality: A preliminary application of AutoSERVPERF model. *Jurnal Intelek*, *15*(2), 197-208. https://ir.uitm.edu.my/id/eprint/69285

[4] Gajewska, T., Zimon, D., Kaczor, G., & Madzík, P. (2020). The impact of the level of customer satisfaction on the quality of e-commerce services. *International Journal of Productivity and Performance Management*, *69*(4), 666-684. [CrossRef]

[5] Celuch, K. (2021). Customers' experience of purchasing event tickets: mining online reviews based on topic modeling and sentiment analysis. *International Journal of Event and Festival Management*, *12*(1), 36-50. [CrossRef]

[6] Turney, P. (2002, July). Thumbs up or thumbs down? Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics* (pp. 417-424).

[7] Churchill, R., & Singh, L. (2022). The evolution of topic modeling. *ACM Computing Surveys*, *54*(10s), 1-35. [CrossRef]

[8] Havrlant, L., & Kreinovich, V. (2017). A simple probabilistic explanation of term frequency-inverse document frequency (tf-idf) heuristic (and variations motivated by this explanation). *International Journal of*

*General Systems*, *46*(1), 27-36. [CrossRef]

[9] Jain, S., Jain, S. K., & Vasal, S. (2024, April). An effective TF-IDF model to improve the text classification performance. In *2024 IEEE 13th International Conference on Communication Systems and Network Technologies* (*CSNT*) (pp. 1-4). IEEE. [CrossRef]

[10] Alemerien, K., Al-Ghareeb, A., & Alksasbeh, M. Z. (2024). Sentiment analysis of online reviews: a machine learning based approach with TF-IDF vectorization. *Journal of Mobile Multimedia*, *20*(5), 1089-1116. [CrossRef]

[11] Bodendorf, F., & Kaiser, C. (2010, January). Mining customer opinions on the internet-a case study in the automotive industry. In *2010 Third International Conference on Knowledge Discovery and Data Mining* (pp. 24-27). IEEE. [CrossRef]

[12] Jelodar, H., Wang, Y., Yuan, C., Feng, X., Jiang, X., Li, Y., & Zhao, L. (2019). Latent Dirichlet allocation (LDA) and topic modeling: models, applications, a survey. *Multimedia tools and applications*, *78*, 15169-15211. [CrossRef]

[13] Wallach, H., Mimno, D., & McCallum, A. (2009). Rethinking LDA: Why priors matter. *Advances in neural information processing systems*, *22*.

[14] Pavlinek, M., & Podgorelec, V. (2017). Text classification method based on self-training and LDA topic models. *Expert Systems with Applications*, *80*, 83-93. [CrossRef]

[15] Lu, S., Liu, Q., & Zhang, Z. (2023, July). Sentiment analysis of Weibo platform based on LDA-SnowNLP Model. In *2023 2nd International Conference on Machine Learning, Cloud Computing and Intelligent Mining* (*MLCCIM*) (pp. 505-510). IEEE. [CrossRef]

[16] Salehan, M., & Kim, D. J. (2016). Predicting the performance of online consumer reviews: A sentiment mining approach to big data analytics. *Decision Support Systems*, *81*, 30-40. [CrossRef]

[17] Xu, F., Pan, Z., & Xia, R. (2020). E-commerce product review sentiment classification based on a naïve Bayes continuous learning framework. *Information Processing & Management*, *57*(5), 102221. [CrossRef]

[18] Wang, M., You, H., Ma, H., Sun, X., & Wang, Z. (2023). Sentiment analysis of online new energy vehicle reviews. *Applied Sciences*, *13*(14), 8176. [CrossRef]

[19] Dhaoui, C., Webster, C. M., & Tan, L. P. (2017). Social media sentiment analysis: lexicon versus machine learning. *Journal of Consumer Marketing*, *34*(6), 480-488. [CrossRef]

[20] Qazi, A., Tamjidyamcholo, A., Raj, R. G., Hardaker, G., & Standing, C. (2017). Assessing consumers' satisfaction and expectations through online opinions: Expectation and disconfirmation approach. *Computers in Human Behavior*, *75*, 450-460. [CrossRef]

[21] Kim, E. G., & Chun, S. H. (2019). Analyzing online car reviews using text mining. *Sustainability*, *11*(6), 1611. [CrossRef]

[22] Wang, X., Lv, T., & Fan, L. (2022). New energy vehicle consumer demand mining research based on fusion topic model: A case in China. *Sustainability*, *14*(6), 3316. [CrossRef]

[23] He, S., & Wang, Y. (2023). Evaluating new energy vehicles by picture fuzzy sets based on sentiment analysis from online reviews. *Artificial Intelligence Review*, *56*(3), 2171-2192. [CrossRef]

[24] Liu, Y., Lu, J., Yang, J., & Mao, F. (2020). Sentiment analysis for e-commerce product reviews by deep learning model of Bert-BiGRU-Softmax. *Mathematical Biosciences and Engineering*, *17*(6), 7819-7837. [CrossRef]

**MinLing Zeng** is currently working as a lecturer at Xiamen University Tan Kah Kee College. Her research interests focus on intelligent e-commerce systems, digital consumption scenarios, and data-driven innovation in emerging e-commerce models. She conducts in-depth analyses of these changes and explores how intelligent technologies can be leveraged to create more innovative e-commerce systems that better meet consumers' needs. (Email: zengml@xujc.com)

**Xiaoli Deng** graduated from Software Engineering Institute of Guangzhou. Currently, she is working as an e-commerce customer service specialist in a trading company. In her daily work, she has accumulated rich practical experience in intelligent customer service systems and e-commerce platform operations. She is able to accurately grasp consumers' needs, handle customer issues efficiently, and has a profound understanding of e-commerce platforms, consumer behavior analysis, and intelligent service applications. (Email: dll4079317781@163.com)