An MDSS Architecture for Customer Experience Decision Support in E-Commerce Using Data-Driven Decision-Making

Adrián Sánchez Ortega[1]*, Tania Ofelia López Zeron[2], Luis Alejandro Santana Valadez[3], Efraín Mohedano Torres[4], Marcelo Quezada Quiterio[5], Arizbeth Escudero Hernández [6]

Adrián Sánchez Ortega

***Corresponding Author**

DOI: https://doi.org/10.47772/IJRISS.2026.102500012

Received: 28 October 2026; Accepted: 03 November 2026; Published: 14 November 2026

ABSTRACT

Direct product reviews on e-commerce platforms provide information about consumer experiences, but their use in marketing decisions requires connecting heterogeneous inputs, aspect-level analysis, and interpretable outputs. This study proposes a conceptual marketing decision support system (MDSS) architecture grounded in a systematic literature review reported with support from the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 guidelines. The research adopts a descriptive and analytical documentary approach that connects published evidence with architectural decisions. The synthesis includes 28 studies and examines data management, opinion analysis, indicator construction, and assisted interpretation in relation to decision support. The resulting architecture organizes automated daily acquisition, input normalization and validation, aspect-based sentiment analysis, polarity aggregation in Python, and interactive querying through a large language model (LLM). A common JSON model accommodates platform connectors and exported files, while validation states determine which records proceed to analysis. Valid aspect mentions retain their product and review identifiers and supporting textual evidence. Mean polarity summarizes classifications encoded as negative, neutral, and positive, together with the number of valid mentions. The LLM layer uses predefined Python tools to query aspects, compare products or publication periods, and retrieve evidence. Its principal output is a decision-oriented report connecting findings, supporting evidence, marketing implications, possible actions, and information still needed for assessment. An application programming interface exposes queries, reports, and processing status. The contribution is an architectural specification that connects consumer opinions with interpretable decision support while preserving human judgment. Data-driven decision-making guides the link between findings and evidence, and Marketing 7.0 provides a conceptual context for consumer understanding and human interpretation. Subsequent development would examine processing quality, reporting fidelity, integration, and resource consumption.

Keywords: Marketing Decision Support Systems, Aspect-Based Sentiment Analysis, Large Language Models, Data-Driven Decision-Making, E-Commerce

INTRODUCTION

The growth of e-commerce has expanded opportunities to understand consumer experiences through the opinions posted on digital platforms. Product reviews contain assessments and descriptions of attributes, conditions of use, and purchasing experiences that may inform marketing management. Baier et al. (2025) identify applications of user-generated content in brand management, quality improvement, and product development. Using this content in a decision requires organizing and analyzing it and connecting it with specific management problems.

Marketing decision support systems (MDSS) bring together data and analytical models to support decision-makers. Their development predates the incorporation of generative models. In their review of applications for new product design, Figueroa-Perez et al. (2019) describe different methods for incorporating consumer information into decisions. Nave et al. (2018) propose a system that combines opinion mining, dimensional

models, and visualization to examine user comments in tourism. These studies position review analysis within analytical support for management, whose usefulness depends on the relationship between results and users' information needs.

Natural language processing provides methods for identifying the attributes mentioned in opinions and the assessments associated with them. These include aspect-based sentiment analysis (ABSA), whose applications are addressed in the reference model presented by D'Aniello et al. (2022). In marketing, Gooljar et al. (2024) examine the integration of sentiment analysis and predictive modeling into the purchasing journey and identify limitations associated with data and the diversity of opinions. The proposal incorporates ABSA to organize consumer opinions by the attributes evaluated and their associated assessments. This information provides the analytical basis for generating indicators that support the identification of problems and opportunities for improvement from a marketing perspective.

Large language models (LLMs) also appear in recent proposals for review analysis. Özbay et al. (2026), for example, integrate sentiment classification, LLM-generated summaries, and fuzzy evaluation into an e-commerce system. The combination of analytical and generative functions therefore has precedents. A new proposal needs to specify each component's responsibilities, the information exchanged, and the link between results and the reviews that support them.

The information used also requires a temporal perspective. Hägele et al. (2025) present a measure of aspect-level review currency that accounts for changes in the characteristics being evaluated. This precedent supports considering data updates and provenance within the proposed information flow. Daily acquisition would allow new reviews to be incorporated and analysis periods to be defined; querying a result would require identifying the dataset from which its indicator was derived.

The discussion of Marketing 7.0 provides a recent context for examining automation and consumer understanding. In an essay on this approach, Carneiro (2026) highlights the risk of focusing management on immediate performance metrics and displacing human interpretation of context. Negricea and Purcarea (2026) emphasize governance and traceability when examining the incorporation of AI into marketing management. These perspectives contextualize the system's role in this article: organizing evidence and supporting its interpretation while leaving the assessment of alternatives to the marketing manager.

The design problem is to connect the stages that turn heterogeneous reviews into queryable information to support marketing decisions. The proposal includes automated daily acquisition, transformation into a common data model, aspect-level analysis, and an explicit procedure for calculating mean polarity by aspect. An LLM-assisted query layer would allow users to request explanations, compare results, and explore possible actions. Its responses would need to distinguish analytical results from suggestions and retain their connection with the available evidence. Users would remain responsible for assessing alternatives and making decisions.

This article aims to propose a conceptual MDSS architecture for analyzing direct product reviews on e-commerce platforms, grounded in a systematic literature review reported according to PRISMA 2020. The contribution specifies how data acquisition and normalization, textual analysis, polarity aggregation, and decision-oriented reporting are connected, together with their interfaces and traceability requirements. The design characteristics will be compared with the reviewed studies through documentary analysis. The scope covers the architecture's rationale and design; implementation and empirical evaluation are proposed for subsequent work.

METHODOLOGY

The study adopts a descriptive and analytical documentary research approach to designing a conceptual architecture for marketing decision support. It examines contributions on decision support systems, review analysis, and language models to establish a rationale for the MDSS components and their relationships. This approach draws on Bernal Torres (2010), who describes documentary research as the analysis of written information to establish relationships, differences, and perspectives on a topic.

The literature review connects prior studies with design decisions. Hernández Sampieri et al. (2014) indicate that developing processes, algorithms, and technologies involves examining previous solutions through such a review. This study analyzes the functions performed by the reviewed systems, how they transform consumer data into information for decision-makers, and the limitations reported by their authors. This examination supports the specification of the information flow and component responsibilities.

Study identification and selection are reported with support from PRISMA 2020 and its explanation and elaboration document, which promote transparent reporting of systematic review methods and results (Page et al., 2021). The technical corpus comprises 28 studies. Methodological texts and conceptual marketing sources are used as complementary literature, distinct from this corpus.

The documentary synthesis relates the findings to the functions required by the proposal: acquiring and normalizing reviews, analyzing text, constructing indicators, and presenting queryable results. Inputs, procedures, outputs, and documented limitations are examined for each function. This organization will support integration decisions and distinguish contributions drawn from the literature from the adaptations proposed in this study.

The study covers the rationale, conceptual specification, and documentary comparison of the architecture. Results will be presented through a synthesis of prior studies, a description of the components, and a comparison of their functions with the reviewed solutions. Implementation, user testing, and measurement of the proposed system's performance belong to a subsequent stage.

Search Strategy and Selection Criteria

The review sought to identify how decision support systems turn customer opinions into management information and the role of language models in that process. It examined architectures and components, the transformation of aspect-level sentiment into indicators, prioritization methods, explanation generation, and evaluation. The Spanish-language search complemented these areas to identify studies published in that language. These questions guided study extraction and comparison; system evaluation was examined across their methods and results.

The search was conducted on October 25, 2026, and covered publications from January 1, 2016, through that date. Journal articles, conference papers, and reviews in English or Spanish were considered. OpenAlex was searched by title and abstract; IEEE Xplore by Document Title and Abstract; and ScienceDirect by title, abstract, and keywords. Google Scholar complemented the Spanish-language search. Date, language, and document-type filters were applied using the options available in each resource and completed during selection.

Eight generic search expressions were formulated and adapted to the syntax of each interface. S1 delimited MDSS; S2 connected decision support with customer opinions; S3 related ABSA to decisions and priorities; and S4 examined multicriteria methods. S5 and S6 addressed language models in decisions and review analysis, S7 focused on explainability and traceability, and S8 incorporated Spanish terms. The following expressions were used as the basis:

S1: ("marketing decision support system" OR "marketing decision support systems" OR "marketing decision support" OR "marketing decision-support")

S2: ("decision support" OR "decision support system" OR "decision support systems" OR "decision-making support") AND ("online reviews" OR "online review" OR "customer reviews" OR "customer review" OR "user-generated content" OR "customer feedback" OR "voice of the customer")

S3: ("aspect-based sentiment analysis" OR "aspect based sentiment analysis" OR "aspect-level sentiment" OR "aspect-based opinion mining" OR ABSA) AND ("decision support" OR "decision making" OR "decision-making" OR prioritization OR prioritisation OR "improvement priorities" OR "managerial implications")

S4: ("online reviews" OR "online review" OR "customer reviews" OR "product reviews" OR "aspect-based sentiment analysis" OR "opinion mining") AND (TOPSIS OR AHP OR VIKOR OR "multi-criteria decision" OR "multicriteria decision" OR "multi-attribute decision" OR MCDM OR MCDA)

S5: ("large language model" OR "large language models" OR LLM OR LLMs OR ChatGPT OR "generative AI" OR "generative artificial intelligence") AND ("decision support" OR "data-driven decision") AND (marketing OR "customer experience" OR "e-commerce" OR "business intelligence" OR "business analytics")

S6: ("large language model" OR "large language models" OR LLM OR LLMs OR GPT OR ChatGPT) AND ("customer reviews" OR "online reviews" OR "product reviews" OR "customer feedback") AND (summarization OR summarisation OR explanation OR explanations OR "decision support" OR "decision making" OR "managerial insights" OR "actionable insights")

S7: ("explainable" OR XAI OR interpretability OR traceability OR transparency) AND ("decision support system" OR "decision support systems") AND (marketing OR "e-commerce" OR "customer experience" OR "customer satisfaction" OR "customer reviews")

S8: ("sistema de soporte a decisiones" OR "sistemas de soporte a decisiones" OR "sistema de apoyo a la toma de decisiones" OR "soporte a la toma de decisiones" OR "toma de decisiones basada en datos") AND (marketing OR mercadotecnia OR "comercio electrónico" OR clientes OR reseñas)

In IEEE Xplore, terms were expanded in the title or abstract, whereas some expressions were divided and simplified in ScienceDirect to accommodate interface restrictions. The complementary Google Scholar search used: ("sistema de soporte a decisiones" OR "sistema de apoyo a la toma de decisiones") (marketing OR mercadotecnia) ("análisis de sentimientos" OR reseñas OR "inteligencia artificial"). The final versions of S4, S6, and S7 narrowed the exploratory searches to reviews, managerial support, and marketing contexts to reduce retrieval from other domains. The search log retains the expressions used in each source, the filters, and the results of each search.

Eligible studies addressed managerial or business decisions and used customer opinions or described a support architecture capable of integrating them. They had to document input management, analysis and aggregation, prioritization, explanation, or result presentation and describe their method and some evaluation, or constitute a relevant review or reference model. Open-access full text was required. Documents outside the domain, classification-only studies without a connection to decisions, systems aimed solely at consumers' purchasing choices, ineligible document types, and retracted publications were excluded.

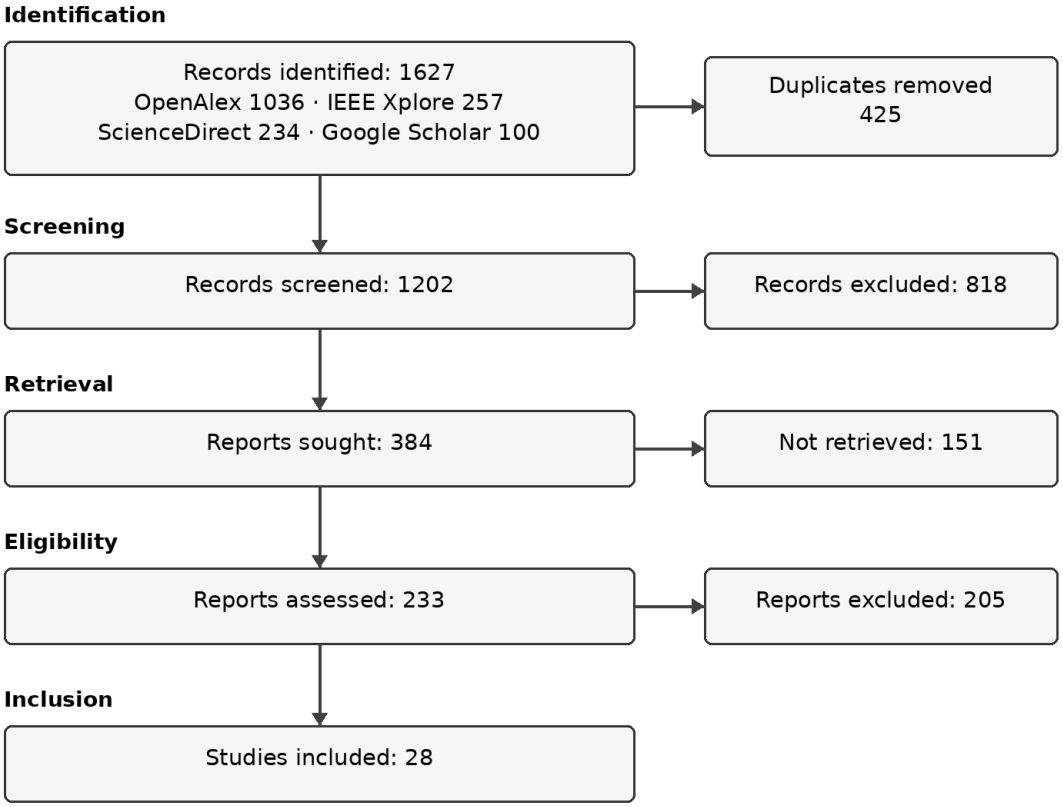
Duplicates were identified through matching DOIs or, when these were unavailable, normalized titles and years. Records with the same title and year but different DOIs were also examined to identify versions of the same document. Truncated title prefixes were compared in Google Scholar. Each record's provenance was retained to relate the consulted sources to the screened set.

Study Identification and Selection

Searches in OpenAlex, IEEE Xplore, ScienceDirect, and Google Scholar yielded 1627 records: 1036, 257, 234, and 100, respectively. Before screening, 425 duplicates were removed: 146 within-source matches, 269 between-source matches, and 10 duplicate versions. This left 1202 records for title and abstract screening.

Title and abstract screening excluded 818 records, after which 384 reports were sought. Of these, 151 could not be retrieved through open access, and 233 proceeded to documentary eligibility assessment. This stage examined complete abstracts and the full text of 44 preselected reports. A total of 205 reports were excluded, and 28 studies were included. Reasons comprised insufficient information to assess the contribution, indirect or overlapping evidence, a focus on consumer decisions, the absence of relevant data or methods, contexts outside marketing management, and sentiment classification without outputs for decision-makers. The complementary Spanish-language search added no studies to the final set. Figure 1 presents the selection process.

Figure 1. Study selection process reported with support from PRISMA 2020.



Source: prepared from the final selection record and the PRISMA 2020 reporting framework (Page et al., 2021). Identified records came from OpenAlex (1036), IEEE Xplore (257), ScienceDirect (234), and Google Scholar (100).

Two reviewers participated in study selection, considering the relevance of the problem, the data used, the method described, and the contribution to marketing decision support.

Extraction and Synthesis Procedure

Bibliographic information, methods, results, and limitations were extracted from the 28 included studies (Table 1). The comparison addressed input management, indicator construction, prioritization, and result interpretation, considering each study's task and evaluation conditions. The documentary comparison supported the specification of components and interfaces by clarifying what each stage receives, how it transforms the input, and what it delivers to the next stage. López Zerón et al. (2027) review deep learning architectures to support thermal time-series processing and decision-oriented outputs. In the MDSS, the specification is organized around continuity between reviews, aspect-level aggregates, and marketing reports.

Table 1. Synthesis of the studies included in the literature review with support from PRISMA 2020 (n = 28).

Authors (year)	Study title	Journal or source	Main contribution
Figueroa-Perez et al. (2019)	The Use of Marketing Decision Support Systems for New Product Design: A Review	International Journal of Computational Intelligence Systems	Reviews 61 publications on MDSS for product design, selected from 375 records. Classifies system types, decision support, and consumer satisfaction. Its 1998–2018 coverage predates current generative assistants.

Nave et al. (2018)	A decision support system framework to track consumer sentiments in social media	Journal of Hospitality Marketing & Management	Integrates text and sentiment mining, a dimensional model, and two dashboards to examine tourism reviews. The case connects opinions with managerial queries; its application focuses on one platform.
D'Aniello et al. (2022)	KnowMIS-ABSA: an overview and a reference model for applications of sentiment analysis and aspect-based sentiment analysis	Artificial Intelligence Review	Proposes KnowMIS-ABSA and distinguishes sentiment, affect, emotion, and opinion to guide analysis. Provides a conceptual basis for specifying outputs; these dimensions are neither interchangeable nor represented by a single metric.
Chornous et al. (2023)	A data science-based marketing decision support system for brand management	Innovative Marketing	Distributes monitoring, processing, pricing, portfolio, segmentation, and forecasting functions among nine data science agents. Provides an organizational precedent for MDSS; its agents do not necessarily correspond to LLM assistants.
Baier et al. (2025)	Collecting and Analyzing User-Generated Content for Decision Support in Marketing Management: An Overview of Methods and Use Cases	Schmalenbach Journal of Business Research	Reviews methods for collecting and analyzing user-generated content and their applications in marketing management. Connects data with business decisions while considering interpretability and ethical and legal conditions of use.
He and Chen (2026)	A Deep Learning-Integrated Big Data Analytics and Marketing Decision Support System	Procedia Computer Science	Integrates batch and stream processing with multimodal deep learning. Reports 18.5 ms at 40 200 requests/s in its configuration. Evaluation concerns business prediction and distributed infrastructure rather than conversational review analysis.
Lamrhari et al. (2019)	Business intelligence using the fuzzy-Kano model	Journal of Intelligence Studies in Business	Combines aspect extraction, fuzzy Kano, and SWOT to transform opinions into priorities. Reports 97.4 % accuracy and 92.4 % precision in extraction; these metrics do not evaluate the effects of marketing decisions.
Singh et al. (2022)	A Text Analytics Framework for Performance Assessment and Weakness Detection From Online Reviews	Journal of Global Information Management	Combines attribute-level sentiment, TOPSIS, and control charts to compare five car manufacturers and identify weaknesses. Connects opinions with diagnosis; review-sampling conditions limit chart interpretation.
McCloskey et al. (2024)	Natural language processing analysis of online reviews for small business: extracting insight from small corpora	Annals of Operations Research	Explores topic modeling and ChatGPT on reviews of a small family business. Obtains interpretable findings from a corpus of hundreds of comments. Its pilot nature and the need to verify interpretations limit generalization.
Hägele et al. (2025)	Aspect-based currency of customer reviews: A novel probability-based metric to pave the way for data quality-aware decision-making	Electronic Markets	Defines a probabilistic aspect-level currency metric based on changes in the object evaluated. Its hotel application shows that age and obsolescence are not equivalent; transfer requires identifying relevant product changes.
Shen et al. (2026)	Customer Requirements Analysis and Product Service	Journal of Theoretical and	Integrates reviews and question–answer content with NMF, aspect-based sentiment, and dual

	Improvement Framework Using Multi-Source User-Generated Content and Dual Importance–Performance Analysis: A Case Study of Fresh E-Ecommerce	Applied Electronic Commerce Research	importance–performance analysis. Prioritizes requirements using PROMETHEE-II in fresh-food commerce. Distinguishes stated importance, derived importance, and performance.
Kyriakidis and Tsafarakis (2026)	MUSAsent: Integrating aspect-based sentiment into Multicriteria Satisfaction Analysis	Omega	Incorporates textual polarity as constraints in MUSA. Across thirteen datasets, Model I improves stability without reducing fit. Extending constraints to the global function requires addressing divergences between text and rating.
Maghsoudi et al. (2026)	A hybrid framework for notebook market analysis: Integrating social media sentiment mining with expert knowledge for feature prioritization	PLoS ONE	Analyzes 329 091 Twitter posts about seven laptop brands. Combines ABSA, fuzzy expert weights, and TOPSIS to prioritize attributes. Social data and expert weighting condition transfer to direct reviews.
Ng et al. (2026)	A blended sentiment analysis, large language model, and multi-criteria decision-making (BSLMM) approach to sustainable product innovation	Soft Computing	Combines BERT for attributes and sentiment, PCA for weighting, and fuzzy COPRAS for product comparison. Formalizes correlated criteria and uncertainty. The language model used is BERT; it does not constitute a generative query interface.
Gai et al. (2026)	Cruise service quality improvement: a quality function deployment approach with online reviews by large language models	Humanities and Social Sciences Communications	Uses LLMs to extract requirements and sentiment from cruise reviews; integrates Kano, QFD, and group consensus to prioritize improvements. The decision team's assessment forms part of the method alongside textual analysis.
You et al. (2026)	A Staged Framework for Mining User Requirements from Vehicle Online Reviews Based on Large Language Models and Sentiment Weighting	Systems	Uses LLMs to build a taxonomy and annotate requirements, BERT for large-scale categorization, and sentiment weighting for prioritization. The vehicle case separates semantic tasks, classification, and ranking; transfer requires domain adaptation.
Choi et al. (2026)	An omnichannel strategy development framework leveraging customer opinion divergence using large language models and explainable AI	Scientific Reports	Extracts and normalizes aspects using LLMs and analyzes cross-channel divergences with SHAP on 6.8 million reviews. Reports a 24.70-percentage-point normalization improvement under the specified comparison, rather than across the entire MDSS.
Lee and Song (2025)	From Unstructured Feedback to Structured Insight: An LLM-Driven Approach to Value Proposition Modeling	Electronics	Relates reviews of three watch generations to the Value Proposition Canvas. Combines extraction, embeddings, and clustering to generate profiles, gap density, and indicators with sensitivity analysis. Results alone do not establish business performance.
Teichert and Shah (2026)	From reviews to constructs: Using LLMs to model customer satisfaction in platform-based services	Journal of Retailing and Consumer Services	Generates ACSI-aligned indicators using LLMs and examines relationships with PLS-SEM in food-delivery services. Compares coding with human scales; partial construct coverage in reviews limits psychometric interpretation.

Xiao et al. (2025)	Do you actually need an LLM? Rethinking language models for customer reviews analysis	Artificial Intelligence Review	Compares small models, larger models, and a hybrid integration in polarity classification and product correlation. Documents performance–resource trade-offs; selection depends on the task and domain adaptation.
Praneeth et al. (2025)	Optimization of Customer Feedback Summarization Using Large Language Models (LLM) and Advanced Retrieval-Augmented Generation	IEEE Access	Combines hierarchical retrieval, RAG Fusion, and DSPy to answer feedback queries. Evaluates response relevance and consistency. The synthetic review dataset limits the scope of its results.
Dahish et al. (2025)	Enhancing phygital customer experience through generative AI: a social listening method for strategic retail decision-making	Journal of Strategic Marketing	Develops a social-listening method for retail decisions using LSTM and BERT. Relates customer content to management recommendations; the study's generative designation does not imply that all components are conversational assistants.
Özbay et al. (2026)	A Unified AI Framework for Turkish E-Commerce Review Analysis: Sentiment Classification, LLM-Based Summarization, and Fuzzy Evaluation	Applied Sciences	Integrates BERTurk, Kumru-2B summaries, and fuzzy evaluation on Turkish reviews. Reports a test macro-F1 of 0.9243 for classification. Qualitative summary evaluation is distinct from classifier performance.
Albesher and Alsanousi (2026)	Retrieval-augmented generation analysis of user reviews in Saudi mobile banking apps: a comparative user experience study of iOS and Android	PeerJ Computer Science	Examines 15 459 iOS and 230 734 Android reviews of banking applications. Combines sentiment and RAG with GPT-4o to identify usability problems. The banking context and translation condition transfer.
Cosentino et al. (2026)	Interpreting User Opinions: A Multidimensional Approach Leveraging Explainable AI and Generative Models	Machine Learning	Combines multidimensional classification, influential words, and explanation generation. Evaluates text, representations, and human assessments. Provides mechanisms to connect explanation with prediction without establishing causality between an opinion and a business outcome.
De Pretto et al. (2025)	AI for Knowledge-Driven Business Decisions: Capabilities and Constraints	European Conference on Knowledge Management	Compares ChatGPT and Perplexity with references and expert analysis in more than 210 experiments and approximately 900 prompts. Documents variation by task, context, and input format and the need for human validation.
Tselepatiotis et al. (2026)	A Large Language Model-Assisted Organizational Decision Support System Integrating Organizational Data with Management Frameworks	Information	Integrates indicators and semantic evidence with organizational frameworks. Evaluates three synthetic scenarios and a human assessment involving 16 participants. Functional consistency and explanation assessment do not constitute validation in a real organization.
Bhattacharya and Kumar (2026)	SupplyQual: an approach for using customer reviews for evaluating last-mile delivery performance	Applied Operations and Analytics	Integrates a knowledge graph and agents to query last-mile reviews. Compared with RAG, improves faithfulness by 6.3 points and context precision by 12.7, with lower answer correctness. It is a bounded proof of concept.

Source: authors' own work based on the 28 studies included in the review.

Thematic Synthesis of the Literature

The studies cover different parts of the process that transforms opinions into managerial information. Reviews and reference models define functions; applied cases connect textual analysis, indicators, and presentation; and recent studies incorporate generation and natural-language queries. Comparing these functions clarifies the scope of each precedent and supports the integration required by the MDSS.

Figueroa-Perez et al. (2019) classify MDSS for product design and show ways to organize decision support around models, data, and consumer knowledge. Chornous et al. (2023) assign brand-management tasks to nine agents based on data science methods. These studies present the MDSS as a combination of functions and resources extending beyond opinion classification. The multiagent organization in the second study does not imply that every function must be performed by an LLM.

Nave et al. (2018) connect opinion mining, dimensional storage, and dashboards, while Baier et al. (2025) extend this relationship to different content types and marketing decisions. He and Chen (2026) address infrastructure for heterogeneous data and batch and stream processing. Their results concern a distributed business prediction architecture and do not establish millisecond response times as a requirement for the proposed system. Together, these studies support separating acquisition, organization, analysis, and querying while accounting for each application's domain. Preparing inputs involves examining their currency and aspect consistency. Hägele et al. (2025) show that, following changes to a product or service, the same review may remain current for one attribute and become obsolete for another. Choi et al. (2026) group equivalent consumer expressions through aspect normalization. Both precedents connect data preparation with interpretation of the resulting outputs. Constructing management indicators requires defining the meaning of analytical outputs. D'Aniello et al. (2022) distinguish sentiment, affect, emotion, and opinion; Lee and Song (2025) generate profiles and measures associated with a value proposition; and Teichert and Shah (2026) operationalize satisfaction constructs using an explicit framework and comparison with human coding. The relationship with marketing therefore depends on defining the indicator and connecting it to the phenomenon being interpreted.

Shen et al. (2026) distinguish stated importance, derived importance, and performance when analyzing requirements through reviews and question–answer content. Kyriakidis and Tsafarakis (2026) introduce sentiment as constraints in MUSA and compare variants across thirteen datasets. Model I's greater stability without loss of fit relates to that formulation and the evaluated data. An overall rating and its text may diverge; treating them as equivalent signals requires examining their relationship.

Domain and scale influence the analysis. McCloskey et al. (2024) explore the combination of topics and an LLM in a small corpus. You et al. (2026) separate category construction, annotation, and requirement classification into different stages. Xiao et al. (2025) compare models across tasks and describe trade-offs between resources and performance. These studies support selecting each component by its function rather than assigning all MDSS processes to a single model. Improvement prioritization takes different forms. Lamrhari et al. (2019) integrate opinion analysis, fuzzy Kano, and SWOT; Singh et al. (2022) use TOPSIS and control charts to identify weaknesses; and Maghsoudi et al. (2026) add expert weighting to the analysis of social-media posts about laptops. These procedures establish rules for converting assessments into relative positions or priorities beyond mention counts. Frequency, polarity, and importance are distinct variables, so combining them requires explicit criteria and assumptions.

Ng et al. (2026) apply PCA to correlated criteria and compare alternatives using fuzzy COPRAS. Gai et al. (2026) integrate LLM-extracted requirements with Kano, QFD, and decision-team consensus. Prioritization options differ in the data and human participation they require. This evidence distinguishes opinion aggregation from multicriteria prioritization and clarifies the criteria and data needed for each procedure.

Choi et al. (2026) connect cross-channel opinion divergence with retail strategies through aspect normalization and SHAP. Dahish et al. (2025) relate social listening to decisions about physical and digital experiences. These domains broaden the review's context; however, the initial MDSS proposal focuses on direct product reviews. Social-media evidence is retained as a precedent for future extensions rather than incorporated as a current input source. For result presentation and interpretation, Özbay et al. (2026) combine sentiment

classification, generative summarization, and fuzzy evaluation. The macro-F1 of 0.9243 concerns the classifier on the test set; summaries are assessed qualitatively. This distinction must be retained when comparing systems, because a classification metric does not demonstrate explanatory faithfulness or the usefulness of a managerial recommendation. Albesher and Alsanousi (2026) integrate sentiment, RAG, and the presentation of usability problems in a multilingual banking context.

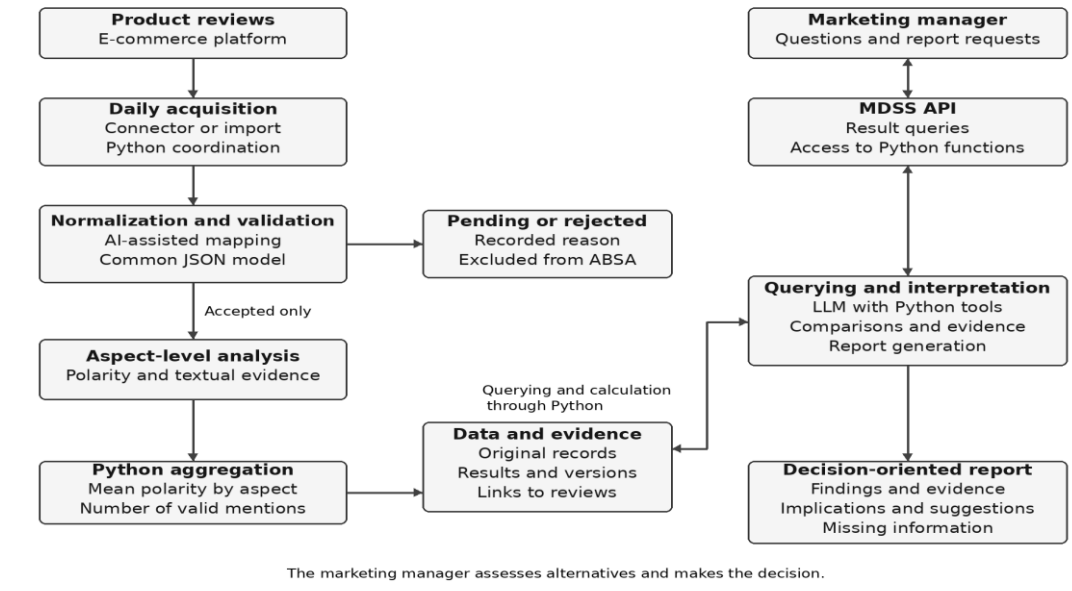
Praneeth et al. (2025) query synthetic feedback through hierarchical retrieval, RAG Fusion, and DSPy. Bhattacharya and Kumar (2026) answer last-mile questions using a knowledge graph and agents. Compared with RAG, they report gains of 6.3 percentage points in faithfulness and 12.7 in context precision, alongside lower answer correctness. This trade-off distinguishes retrieval of relevant evidence from a complete answer to a question, although the two properties are related. Cosentino et al. (2026) complement this perspective by connecting explanations with words that influence classification.

Across more than 210 experiments, De Pretto et al. (2025) document variation in assistants according to task, context, and file format. Tselepatiotis et al. (2026) evaluate a DSS through synthetic scenarios, result consistency, and human assessment with 16 participants. Their work provides criteria for examining assisted interpretation without demonstrating the proposed MDSS's effectiveness in real decisions. These evaluations serve as comparative evidence and as references for subsequent work. The review supports examining review provenance and currency, the meaning of analytical outputs, indicator-construction rules, and the relationship between responses and evidence together. The synthesis thus identifies the connections that the MDSS must specify for direct e-commerce reviews. The following section defines system inputs, responsibilities, and outputs based on these documented functions. Evaluation of the proposed integration belongs to subsequent work.

Conceptual Mdss Architecture

The architecture combines daily review acquisition, normalization and validation, aspect-level analysis, Python aggregation, and LLM-assisted querying and interpretation. A decision-oriented report is the main output, accompanied by the data and evidence supporting each finding. Python is proposed to coordinate operations, validate inputs, and calculate results; JSON defines the common exchange model, and an API exposes services to external applications. Figure 2 shows the components, their exchanges, and the connection between reviews and the report for the marketing manager.

Figure 2. Conceptual MDSS architecture for daily review analysis and marketing decision support.



Source: authors' own work. Arrows represent the processing flow and query exchanges; storage retains records and their links to results.

Automated Daily Acquisition

The proposal adopts daily acquisition of direct reviews for configured products, coordinated through a scheduled Python process. Each connector would use the platform's official API when available and when authorized access exists. Importing exported files is considered a complementary route. The initial run would retrieve available reviews within the specified period; subsequent runs would seek new records and, when supported by the source, changes to existing ones. Provenance, acquisition date, execution status, and the continuation point after interruptions would be recorded. This acquisition interval is independent of the publication period used in queries.

Input Normalization and Validation

Each received record would be retained with its normalized representation. An AI component would assist in proposing mappings between source fields and the common model; once approved, the mapping would be reused in subsequent runs. Before admitting a record to analysis, Python would check its structure, types, and required fields. This assistance is proposed for field mapping and does not constitute evidence of universal compatibility between platforms.

The input JSON contract combines original content, product identification, provenance, and processing metadata (Table 2). Product identification is interpreted in relation to its platform, and the normalized record retains its connection to the original source.

Table 2. Normalized JSON record structure and validation rules.

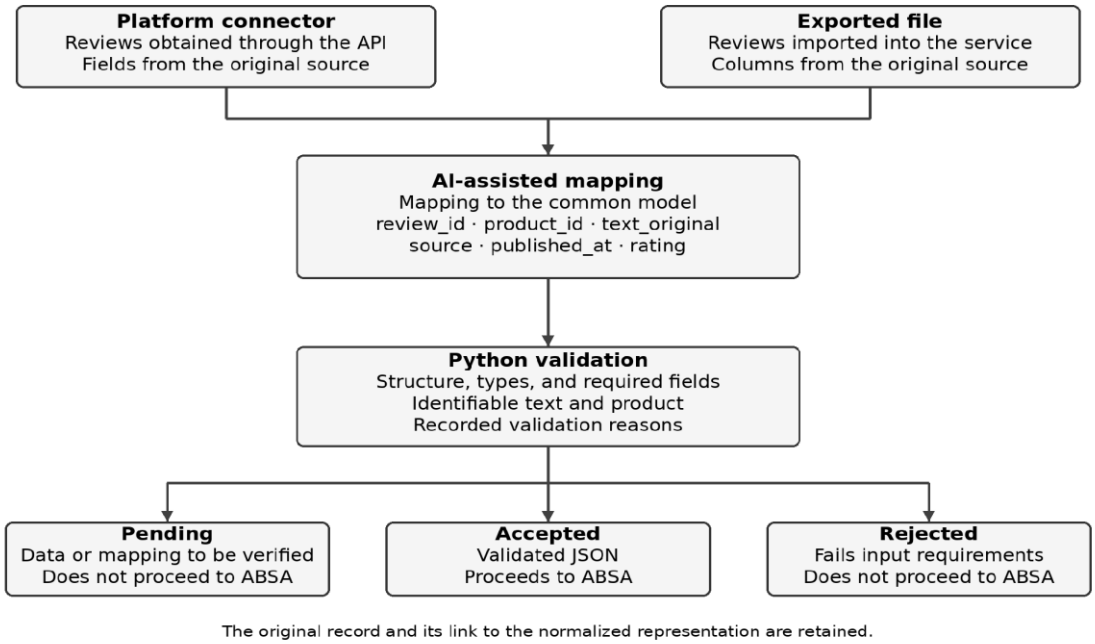
Field	JSON type	Rule and function
review_id	string	Internal identifier assigned by the service; required.
source_review_id	string or null	Platform identifier, when available.
product_id	string	Product identified within the platform; required.
source	string	Source platform; required.
text_original	string	Received nonempty text; required for analysis.
language	string or null	Recorded language; null when unknown.
published_at	string or null	Publication date in ISO 8601 format; null if unavailable.
acquired_at	string	Acquisition date and time in ISO 8601 format, assigned by the service.
rating	object or null	Rating with its original value and scale; null if unavailable.
source_reference	string	Review URL or reference to the imported file and record; required.
schema_version	string	Version of the common model applied by the service.
validation_status	string	*accepted*, *pending*, or *rejected*.
validation_errors	array	Validation reasons; an empty list when no issues are present.

Source: authors' own work. Fields and their rules form part of the proposed contract.

Admission to ABSA depends on validation status. *accepted* identifies records that pass the checks; *pending* refers to inputs requiring correction or verification; and *rejected* indicates failure to meet requirements. Only *accepted* records proceed to analysis. Reviews without publication dates may form part of the overall aggregate but are excluded from comparisons between publication periods.

Figure 3 illustrates how inputs converge into the common model and how their admission to analysis is determined. The flow retains record provenance and identifies validation status before forwarding the record to ABSA.

Figure 3. Normalization of heterogeneous inputs and validation of the common model.



Source: authors' own work. Two input routes and selected common-model fields are shown; both retain a reference to the original record.

Aspect-Level Analysis and Evidence Preservation

The ABSA analytical function receives *accepted* reviews and returns aspect mentions with positive, negative, or neutral polarity. Each mention's output would include review_id, product_id, mention_id, aspect, polarity, and evidence. A review could yield several mentions, each retained separately with its original text span. Python would verify response structure and correspondence between the evidence and the received text. The absence of identified aspects would be distinguished from a processing error: the latter would be recorded and excluded from aggregation without assigning it a neutral label.

The conceptual distinctions in D'Aniello et al. (2022) support separating textual expression, aspect, and assessment. Grouping mentions into equivalent categories would require a common catalog, whose version would accompany the results. Choi et al. (2026) provide an aspect-normalization precedent for examining this semantic correspondence. The MDSS specification leaves the internal architecture and the model used for ABSA open.

Polarity Aggregation in Python

The only sentiment aggregate defined is the arithmetic mean of polarity by aspect. Each validly classified mention would be assigned −1 for negative polarity, 0 for neutral, and +1 for positive. The mean for an aspect a, calculated over a defined set of reviews for a product, is expressed in Equation (1):

$$\bar{s}_{\mathrm{a}} = \frac{1}{n_{\mathrm{a}}} \sum_{i=1}^{n_{\mathrm{a}}} s_{\mathrm{a},\mathrm{i}} \quad (1)$$

In Equation (1), n_a denotes the number of validly classified mentions of the aspect, and $s_{a,i}$ is the encoded polarity of mention i. The set would be defined by product and, if requested, publication period. The mean would range from −1 to +1 and be displayed alongside n_a, which counts mentions rather than reviews. When n_a

= 0, an unavailable value would be returned. A zero mean may reflect neutral mentions or opposing opinions that cancel each other out, so individual mentions and evidence would be retained. This coding summarizes polarity; it does not represent a validated scale of emotional intensity or overall satisfaction.

After daily processing, Python would calculate aggregates and record the dataset used, the category version, and the excluded statuses. Each query would retrieve results from an identified version, so subsequent updates would not change the reference for an existing report. The design includes neither indicator selection nor a composite priority index.

Interactive Querying and Decision-Oriented Reporting

The query layer would connect the marketing manager's questions with review analysis. The LLM would organize the request and ask the Python tools for results, comparisons between products or periods, and evidence needed to respond. Python would validate parameters and perform calculations on identified datasets using compatible categories and the same polarity coding. The response would bring together findings, observed differences, and mention counts to relate them to possible actions and the information needed to assess those actions.

Praneeth et al. (2025) and Bhattacharya and Kumar (2026) use evidence retrieval, while Cosentino et al. (2026) connect classification with explanation. These precedents support examining correspondence between responses and sources. In the proposed design, the LLM would access defined operations, and each interaction would record the query, its parameters, tool results, and the evidence used. Generation rules would require preserving calculated values and identifying insufficient information. Compliance with these rules and response faithfulness would be evaluated later.

Each report finding would be presented through five elements: the observation, supporting evidence, its marketing implication, a possible action, and the information still needed to decide. Means and counts would remain available as support for the explanation. Suggestions would be identified as such and linked to findings; comparisons of alternatives would distinguish this support from assumptions about costs, feasibility, or expected outcomes. Unavailable data would be identified as missing information. The design does not presume that a recommendation will improve business performance.

The report expresses the design's orientation toward data-driven decision-making (DDM) by connecting opinions and calculations with alternatives that are understandable to the marketing manager. Through their discussion of Marketing 7.0, Carneiro (2026) and Negricea and Purcarea (2026) provide a context for relating automation, consumer understanding, and human interpretation. Translating that context into system responsibilities is a conceptual decision of the proposal. Users retain responsibility for assessing alternatives and making the final decision.

The layer is specified with GPT-6.1 Sol as the proposed primary configuration and Claude Sonnet 5.5 as an alternative (Table 3). Requirements include tool use and structured outputs documented by the providers (OpenAI, n.d.-a, n.d.-b; Anthropic, n.d.-b, n.d.-c). Either model would occupy the same interface: it would receive the question and its scope, request Python operations, and produce the report from the retrieved results and text spans. Changing providers would require adapting calls to their API while retaining the MDSS contracts.

Table 3. Proposed configurations for the query and interpretation layer.

Characteristic	Primary configuration	Alternative configuration
Model	GPT-6.1 Sol	Claude Sonnet 5.5
API identifier	gpt-6.1-sol	claude-sonnet-5-5
Proposed role	Querying, interpretation, and reporting	Same functions through another adapter

External tools	Documented function calling	Documented tool use
Structured output	Documented	Documented
Context window	1 050 000 tokens	1 000 000 tokens
Maximum output	128 000 tokens	128 000 tokens
Input per million tokens	USD 2.00	USD 2.00
Output per million tokens	USD 10.00	USD 10.00
Conditions of use	Paid API; quotas depend on account tier	Paid API; limits depend on account and provider
Selection status	Proposed reference configuration	Proposed alternative; subsequent evaluation

Source: technical documentation from OpenAI (n.d.-a, n.d.-b) and Anthropic (n.d.-a, n.d.-b, n.d.-c). Base API rates, excluding caching or batch discounts. Different rates apply to GPT-6.1 Sol requests exceeding 272 000 input tokens. Limits and prices may change.

Table 3 describes features and conditions for integration and usage; report faithfulness would be examined using domain evidence. The proposal involves sending the results and text spans needed for each query and recording the model, parameters, calls, and tokens used. Usage would depend on the context sent, response length, and retries, alongside Python hosting and storage. Queries would run on request, independently of daily acquisition. This comparison complements the report specification and its orientation toward data-driven decisions.

API Services and Output Contract

API services would allow users to query results by product and period, ask questions, request reports, and check daily processing status. Each question or report request would specify its scope, and the service would record the data version used. Product or period selection would delimit the query without changing mean polarity as the sentiment aggregate. Reports would be produced on request using the latest available processed dataset, without needing to coincide with daily acquisition. The API response organizes the report and its analytical support into the blocks described in Table 4. This organization presents findings and their implications before the queryable calculation and evidence details. Each block serves a specific function in the response.

Table 4. Output JSON structure for decision-oriented queries and reports.

Block	JSON type	Content and function
scope	object	Products, period, update date, and version of the dataset used. Defines the response scope.
results	array	Product, aspect, mean polarity, and number of valid mentions. Retains calculations returned by Python.
interpretation	array	Findings and their marketing implications, linked to results.
evidence	array	Review and mention identifiers, with the text spans supporting the findings.
suggestions	array	Action alternatives, their relationship with evidence, and assumptions requiring assessment.
limitations	array	Restrictions of the analyzed dataset and information still needed for interpretation or decision-making.

Source: authors' own work. Fields and their rules form part of the proposed contract.

RESULTS AND DISCUSSION

The study produces an MDSS architecture specification connecting direct e-commerce reviews with reports that support marketing decisions. The review of 28 studies supported the separation of acquisition, normalization, aspect-level analysis, aggregation, and querying. This organization is specified through daily updates, input and output JSON contracts, validation statuses, and API operations. The contribution lies in connecting these functions and preserving the relationship between each interpretation and the data and reviews used.

Comparison with Nave et al. (2018) positions the design among systems that organize opinions for managerial queries. Their work integrates text mining, a dimensional model, and dashboards; the proposal incorporates natural-language queries that request calculations and evidence from Python tools. Chornous et al. (2023) distribute brand management among nine data science agents. The MDSS proposed here focuses on product reviews and specifies responsibilities for textual analysis, calculation, and interpretation. The differences reflect the service's purpose and organization.

Özbay et al. (2026) integrate classification, generative summarization, and fuzzy evaluation for e-commerce reviews. The proposed MDSS defines mean polarity as a fixed aggregate and uses it in queries, comparisons, and a report containing implications and suggestions. Ng et al. (2026) and Gai et al. (2026) incorporate weighting and multicriteria methods into prioritization. In this architecture, comparing opinions and evidence supports the assessment of alternatives, while the marketing manager assesses their importance and feasibility and decides on action. The mean describes observed opinions; deciding also requires interpreting the product's context.

The comparison also considers traceability. Praneeth et al. (2025) use augmented retrieval for feedback queries, and Bhattacharya and Kumar (2026) organize last-mile information in a graph queried by agents. In the proposal, contracts link product, review, and mention identifiers with the text span, calculated result, and dataset version. This establishes what information must accompany a response and which records allow its basis to be reconstructed. Interpretive faithfulness is thus distinguished from classification accuracy and the relevance of a business suggestion.

The connection with data-driven decision-making is expressed in the report: each finding is accompanied by evidence, a marketing implication, an action alternative, and missing information. Metrics support an understandable explanation and remain available for querying. The treatment of consumer experience and human participation relates the design to the Marketing 7.0 context discussed by Carneiro (2026) and Negricea and Purcarea (2026). This context informs the service's purpose, while the technical contribution lies in the interfaces and rules connecting analysis with reporting.

The documentary comparison positions the architecture functionally against prior studies whose metrics concern their own tasks, data, and configurations. The specification provides a basis for subsequently examining mapping correspondence, aspect-level classification, calculation consistency, and report faithfulness. Evaluation would also need to consider daily updates, integration errors, and the usefulness of responses to the marketing manager. The configuration comparison specifies interfaces and economic conditions for the query layer; rates are usage references rather than measurements of MDSS costs.

CONCLUSION

The literature review informed an MDSS architecture that transforms direct product reviews into information for marketing decisions. The design defines a flow from daily acquisition and input validation through aspect-level analysis, mean polarity, and report generation. JSON contracts and the API specify information exchanges and the responsibilities of components coordinated in Python.

The contribution consists of connecting each finding with the calculations, mentions, and text spans that support it. The LLM layer is connected to query and comparison tools to structure explanations and possible

actions. The report presents these alternatives alongside their assumptions and missing information, leaving assessment and decision-making to the user. This organization connects data-based support with the consumer understanding and human judgment that contextualize the proposal through Marketing 7.0.

Comparison with prior studies supports the distribution of functions and clarifies the role of aggregation and assisted interpretation. Subsequent development would need to examine the complete flow using reviews from the application domain, considering mapping quality, classification, traceability, report faithfulness, and operating conditions. Extending the design to other consumer sources is a line of future work that would require reviewing the contracts and data comparability.

REFERENCES

1. Albesher, A. S., & Alsanousi, B. (2026). Retrieval-augmented generation analysis of user reviews in Saudi mobile banking apps: A comparative user experience study of iOS and Android. PeerJ Computer Science, 12, e3565. https://doi.org/10.7717/peerj-cs.3565
2. Anthropic. (n.d.-a). Claude Sonnet 5.5. Claude Platform Docs. https://platform.claude.com/docs/en/models/sonnet-5-5/overview
3. Anthropic. (n.d.-b). Structured outputs. Claude Platform Docs. https://platform.claude.com/docs/en/build-with-claude/structured-outputs
4. Anthropic. (n.d.-c). Tool use with Claude. Claude Platform Docs. https://platform.claude.com/docs/en/agents-and-tools/tool-use/overview
5. Baier, D., Decker, R., & Asenova, Y. (2025). Collecting and analyzing user-generated content for decision support in marketing management: An overview of methods and use cases. Schmalenbach Journal of Business Research, 77(3), 419–455. https://doi.org/10.1007/s41471-025-00208-7
6. Bernal Torres, C. A. (2010). Metodología de la investigación: Administración, economía, humanidades y ciencias sociales (3.ª ed.). Pearson Educación.
7. Bhattacharya, M., & Kumar, M. (2026). SupplyQual: An approach for using customer reviews for evaluating last-mile delivery performance. Applied Operations and Analytics, 2(1), 2743465. https://doi.org/10.1080/29966892.2026.2743465
8. Carneiro, M. R. (2026). A nova miopia no Marketing 7.0: Chuteiras rosas e caixa-preta na era da IA autônoma. En Administração e Gestão Organizacional: Estratégias, Processos e Decisões 3 (1.ª ed., pp. 1–15). Atena Editora. https://doi.org/10.22533/at.ed.851132605081
9. Choi, B., Son, Y., & Choi, Y. (2026). An omnichannel strategy development framework leveraging customer opinion divergence using large language models and explainable AI. Scientific Reports. https://doi.org/10.1038/s41598-026-71125-w
10. Chornous, G., Fareniuk, Y., Giedraitis, V. R., Ulvidienė, E., & Kharlamova, G. (2023). A data science-based marketing decision support system for brand management. Innovative Marketing, 19(2), 38–50. https://doi.org/10.21511/im.19(2).2023.04
11. Cosentino, C., Gunduz Cure, M., Marozzo, F., & Ozturk Birim, S. (2026). Interpreting User Opinions: A Multidimensional Approach Leveraging Explainable AI and Generative Models. Machine Learning, 115(6), 140. https://doi.org/10.1007/s10994-026-07037-7
12. Dahish, Z., Miah, S. J., Pandit, A., & Roy, S. K. (2025). Enhancing phygital customer experience through generative AI: A social listening method for strategic retail decision-making. Journal of Strategic Marketing, 1–21. https://doi.org/10.1080/0965254X.2025.2540267
13. D'Aniello, G., Gaeta, M., & La Rocca, I. (2022). KnowMIS-ABSA: An overview and a reference model for applications of sentiment analysis and aspect-based sentiment analysis. Artificial Intelligence Review, 55(7), 5543–5574. https://doi.org/10.1007/s10462-021-10134-9
14. De Pretto, G., Bolisani, E., & Nakash, M. (2025). AI for Knowledge-Driven Business Decisions: Capabilities and Constraints. European Conference on Knowledge Management, 26(1), 234–243. https://doi.org/10.34190/eckm.26.1.3764
15. Figueroa-Perez, J. F., Leyva-Lopez, J. C., Santillan, L. C., Pérez Contreras, E. O., & Sánchez, P. J. (2019). The use of marketing decision support systems for new product design: A review. International Journal of Computational Intelligence Systems, 12(2), 761–774. https://doi.org/10.2991/ijcis.d.190624.001

16. Gai, T., Wu, J., Xing, Y., Liu, Y., & Cao, M. (2026). Cruise service quality improvement: A quality function deployment approach with online reviews by large language models. Humanities and Social Sciences Communications, 13(1), 619. https://doi.org/10.1057/s41599-026-06941-6
17. Gooljar, V., Issa, T., Hardin-Ramanan, S., & Abu-Salih, B. (2024). Sentiment-based predictive models for online purchases in the era of marketing 5.0: A systematic review. Journal of Big Data, 11, Article 107. https://doi.org/10.1186/s40537-024-00947-0
18. Hägele, L., Klier, M., Moestue, L., & Obermeier, A. (2025). Aspect-based currency of customer reviews: A novel probability-based metric to pave the way for data quality-aware decision-making. Electronic Markets, 35, Article 10. https://doi.org/10.1007/s12525-025-00760-4
19. He, Z., & Chen, Q. (2026). A Deep Learning-Integrated Big Data Analytics and Marketing Decision Support System. Procedia Computer Science, 291, 155–165. https://doi.org/10.1016/j.procs.2026.09.035
20. Hernández Sampieri, R., Fernández Collado, C., & Baptista Lucio, M. del P. (2014). Metodología de la investigación (6.ª ed.). McGraw-Hill Education.
21. Kyriakidis, A., & Tsafarakis, S. (2026). MUSAsent: Integrating aspect-based sentiment into Multicriteria Satisfaction Analysis. Omega, 143, 103552. https://doi.org/10.1016/j.omega.2026.103552
22. Lamrhari, S., Elghazi, H., & El Faker, A. (2019). Business intelligence using the fuzzy-Kano model. Journal of Intelligence Studies in Business, 9(2), 43–58. https://doi.org/10.37380/jisib.v9i2.468
23. Lee, J., & Song, C. H. (2025). From Unstructured Feedback to Structured Insight: An LLM-Driven Approach to Value Proposition Modeling. Electronics, 14(22), 4407. https://doi.org/10.3390/electronics14224407
24. López Zerón, T. O., Sánchez Ortega, A., Quezada Quiterio, M., Mohedano Torres, E. de J., Marrón Hernández, U., & Jiménez Zerón, G. (2027). A multi-output CNN1D–GRU–MLP architecture for thermal time-to-risk classification in cold chain logistics. International Journal of Combinatorial Optimization Problems and Informatics, 18(1), 237–255. https://doi.org/10.61467/2007.1558.2027.v18i1.1465
25. Maghsoudi, M., Bakhtiari, M., & Bakhtiari, H. (2026). A hybrid framework for notebook market analysis: Integrating social media sentiment mining with expert knowledge for feature prioritization. PLOS One, 21(2), e0342067. https://doi.org/10.1371/journal.pone.0342067
26. McCloskey, B. J., LaCasse, P. M., & Cox, B. A. (2024). Natural language processing analysis of online reviews for small business: Extracting insight from small corpora. Annals of Operations Research, 341(1), 295–312. https://doi.org/10.1007/s10479-023-05816-2
27. Nave, M., Rita, P., & Guerreiro, J. (2018). A decision support system framework to track consumer sentiments in social media. Journal of Hospitality Marketing & Management, 27(6), 693–710. https://doi.org/10.1080/19368623.2018.1435327
28. Negricea, C., & Purcarea, I. M. (2026). CMOs' focus on continual growth and progress to strategically reinvent marketing's operating model. Holistic Marketing Management Journal, 16(3), 9–16. https://ideas.repec.org/a/hmm/journl/v16y2026i3p09-16.html
29. Ng, C. Y., Tang, Y. M., & Ho, G. T. S. (2026). A blended sentiment analysis, large language model, and multi-criteria decision-making (BSLMM) approach to sustainable product innovation. Soft Computing, 30(7), 4685–4701. https://doi.org/10.1007/s00500-026-11402-y
30. OpenAI. (n.d.-a). Function calling. OpenAI API. https://developers.openai.com/api/docs/guides/function-calling
31. OpenAI. (n.d.-b). GPT-6.1 Sol Model. OpenAI API. https://developers.openai.com/api/docs/models/gpt-6.1-sol
32. Özbay, E., Altunbey Özbay, F., & Özer, A. B. (2026). A unified AI framework for Turkish e-commerce review analysis: Sentiment classification, LLM-based summarization, and fuzzy evaluation. Applied Sciences, 16(12), Article 5849. https://doi.org/10.3390/app16125849
33. Page, M. J., Moher, D., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., Shamseer, L., Tetzlaff, J. M., Akl, E. A., Brennan, S. E., Chou, R., Glanville, J., Grimshaw, J. M., Hróbjartsson, A., Lalu, M. M., Li, T., Loder, E. W., Mayo-Wilson, E., McDonald, S., … McKenzie, J. E. (2021). PRISMA 2020 explanation and elaboration: Updated guidance and exemplars for reporting systematic reviews. BMJ, 372, n160. https://doi.org/10.1136/bmj.n160

34. Praneeth, B., Mohana, Nattem, E. C., Jetti, K., Kavyashree, B. K., Rakshitha, D., Ramakanth Kumar, P., & Sreelakshmi, K. (2025). Optimization of Customer Feedback Summarization Using Large Language Models (LLM) and Advanced Retrieval-Augmented Generation. IEEE Access, 13, 124319–124332. https://doi.org/10.1109/ACCESS.2025.3588337
35. Shen, Z., Zhao, C., & Li, Y. (2026). Customer Requirements Analysis and Product Service Improvement Framework Using Multi-Source User-Generated Content and Dual Importance–Performance Analysis: A Case Study of Fresh E-Ecommerce. Journal of Theoretical and Applied Electronic Commerce Research, 21(1), 19. https://doi.org/10.3390/jtaer21010019
36. Singh, A., Jenamani, M., Thakkar, J., & Dwivedi, Y. K. (2022). A Text Analytics Framework for Performance Assessment and Weakness Detection From Online Reviews. Journal of Global Information Management, 30(8), 1–26. https://doi.org/10.4018/JGIM.304069
37. Teichert, T., & Shah, A. M. (2026). From reviews to constructs: Using LLMs to model customer satisfaction in platform-based services. Journal of Retailing and Consumer Services, 88, 104539. https://doi.org/10.1016/j.jretconser.2025.104539
38. Tselepatiotis, M., Kariofyllas, C., Kontogianni, A., & Alepis, E. (2026). A Large Language Model-Assisted Organizational Decision Support System Integrating Organizational Data with Management Frameworks. Information, 17(10), 971. https://doi.org/10.3390/info17100971
39. Xiao, Y., Li, Y., Chen, S., Barker, H., & Rad, R. (2025). Do you actually need an LLM? Rethinking language models for customer reviews analysis. Artificial Intelligence Review, 58(10), 332. https://doi.org/10.1007/s10462-025-11308-5
40. You, Z., Peng, S., Zhang, W., & Tan, H. (2026). A Staged Framework for Mining User Requirements from Vehicle Online Reviews Based on Large Language Models and Sentiment Weighting. Systems, 14(6), 645. https://doi.org/10.3390/systems14060645