

ARTICLES IN ENGLISH

INFORMATION AND DIGITAL TECHNOLOGIES

ARTIFICIAL INTELLIGENCE AND INTELLIGENT SYSTEMS

Статья поступила в редакцию: 03.08.2026

Статья принята к публикации: 02.09.2026

Статья опубликована: 28.09.2026

УДК 004.738.52

**Engineering methods for improving search relevance in e-commerce systems
based on Elasticsearch and OpenSearch**

Perelekhov Ivan

Senior Web Developer,

Oggetto Web LLC,

Taganrog, Russia

ORCID ID: 0009-0008-8448-3092

E-mail: ivanperelekhov@gmail.com

**Инженерные методы повышения релевантности поиска в e-commerce
системах на основе Elasticsearch и OpenSearch**

Перелехов Иван Владимирович

старший веб-разработчик, ООО Оджетто Веб

Российская Федерация, г. Таганрог

Abstract

The article presents an analytical review of engineering approaches to improving search relevance in e-commerce systems based on Elasticsearch and OpenSearch. The aim of the study is to systematise technical solutions applied at the stages of product-data preparation, index design, user-query processing, and result ranking. The review draws on recent scientific publications and industry studies addressing product-data quality, indexing, query interpretation and expansion, spelling correction, ranking mechanisms, user behaviour, and search-system performance. The reviewed approaches are grouped according to their engineering purpose, potential effect on search

Библиографическое описание: Перелехов И.В. Инженерные методы повышения релевантности поиска в e-commerce системах на основе Elasticsearch и OpenSearch // Universum: технические науки : электрон. научн. журн. 2026. 9(150). URL: <https://7universum.com/ru/tech/archive/item/23371>

For citation: Perelekhov I. Engineering methods for improving search relevance in e-commerce systems based on Elasticsearch and OpenSearch // Universum: Technical Sciences : electronic scientific journal. 2026. No. 9(150). Available at: <https://7universum.com/ru/tech/archive/item/23371>

quality, and principal implementation limitations. The reviewed sources indicate that search relevance can be influenced by attribute normalisation, appropriate field-type selection, analyser configuration, synonym management, structured parameter extraction, fuzzy matching, and controlled query relaxation. The review also highlights a recurring trade-off between recall and precision when query-expansion mechanisms are applied. Textual, attribute-based, behavioural, commercial, contextual, and infrastructure-related signals used in product ranking are systematised. The findings of the review allow relevance improvement to be considered as a set of interconnected engineering tasks performed at different stages of the search process while accounting for data-quality limitations, computational cost, and search-system performance.

Аннотация

В статье представлен аналитический обзор инженерных подходов к повышению релевантности поиска в e-commerce системах на основе Elasticsearch и OpenSearch. Цель исследования состоит в систематизации технических решений, применяемых на этапах подготовки товарных данных, проектирования индекса, обработки пользовательских запросов и ранжирования результатов. Материалом для обзора послужили современные научные публикации и отраслевые исследования, посвященные качеству товарных данных, индексированию, интерпретации и расширению запросов, коррекции опечаток, ранжированию, пользовательскому поведению и производительности поисковых систем. Рассмотренные подходы сгруппированы по их инженерному назначению, потенциальному влиянию на качество поиска и основным ограничениям применения. Показано, что на релевантность поиска могут влиять нормализация атрибутов, выбор типов полей, настройка анализаторов, управление синонимами, извлечение структурированных параметров, нечеткое сопоставление и контролируемое ослабление условий поиска. Обзор также выявляет характерный компромисс между полнотой и точностью выдачи при использовании механизмов расширения запросов. Систематизированы текстовые, атрибутные, поведенческие, коммерческие, контекстные и инфраструктурные сигналы, учитываемые при ранжировании. Результаты обзора позволяют представить повышение релевантности как совокупность взаимосвязанных инженерных задач, решаемых на разных этапах поискового процесса с учетом ограничений по качеству данных, вычислительным затратам и производительности системы.

Keywords: e-commerce, product search, relevance, Elasticsearch, OpenSearch, indexing, ranking.

Ключевые слова: e-commerce, товарный поиск, релевантность, Elasticsearch, OpenSearch, индексация, ранжирование.

Introduction

Search in e-commerce systems is one of the key mechanisms of user interaction with a product catalog. Unlike general information retrieval, product search is aimed not only at identifying textual matches but also at aligning user intent with specific product characteristics,

category, attributes, availability, and commercial relevance. Therefore, the quality of search results is directly associated with navigation convenience, catalog browsing depth, bounce rate, and the probability of purchase.

The main challenge in improving product search accuracy is the diversity of product information, user input variability, synonymy, misspellings, vague query formulation, and inconsistencies in naming conventions.

The practical importance of product-data quality is illustrated by Salsify's 2025 consumer research conducted among 1,910 respondents in the USA and UK. The survey indicates that inconsistent or inaccurate product information can affect consumer trust and purchasing decisions (fig. 1).




Figure 1. Consumer response to inconsistent and inaccurate product information in e-commerce according to the Salsify survey [1]

These findings support the need to address product-data quality at the early stages of the search pipeline through data normalisation, attribute structuring, and consistent representation of product information before query-processing and ranking mechanisms are applied.

Despite the implementation of state-of-the-art search solutions such as Elasticsearch and OpenSearch, search effectiveness depends on engineering configuration, analysis of catalog structures and query flows, and adaptation of ranking rules to the characteristics of a particular platform.

Recent research on e-commerce search demonstrates a shift from isolated lexical matching techniques towards more complex approaches that integrate semantic processing, user behaviour, and relevance evaluation. Ren et al. [2] systematised information discovery mechanisms in e-commerce and showed that contemporary solutions combine information retrieval, recommender systems, natural language processing, and user interaction data. Nguyen et al. [3] focused on query rewriting and proposed a hybrid framework combining offline knowledge distillation with online reinforcement learning; experiments on the Amazon ESCI dataset demonstrated improvements in query relevance, diversity, and adaptability while addressing the latency and computational costs associated with large generative models. More recent work by Ou et al. [4] extends query rewriting towards alignment with human cognition and user preferences, reflecting the transition from generic query reformulation to preference-aware search optimisation. Jain et al. [5] addressed the evaluation stage by proposing the PROBES framework, which distinguishes exact, substitute, complement, and irrelevant results and identifies both precision and missed-recall errors; its evaluation across 100 product categories achieved an Actionable Error Rate of 76 %.

Overall, the reviewed studies address different components of e-commerce search, including product-data preparation, query processing, ranking, and relevance evaluation. Their combined consideration across the main stages of the product-search process is less common in the reviewed literature.

The aim of this analytical review is to systematise engineering approaches to improving search relevance in e-commerce systems based on Elasticsearch and OpenSearch and to examine their engineering roles and principal limitations at different stages of the product-search pipeline. The review objectives are to summarise the principal sources of relevance loss in heterogeneous product catalogues; examine approaches to data preprocessing, index design, and analyser configuration; systematise query-processing mechanisms and their limitations; and examine how textual, attribute-based, behavioural, commercial, contextual, and infrastructure-related signals are used in product ranking.

Materials and methods

The study was conducted as an analytical review of engineering approaches to improving product-search relevance in e-commerce systems based on Elasticsearch and OpenSearch. The source base included recent scientific publications and industry studies addressing product-data quality and preprocessing, search indexing, query interpretation and expansion, spelling correction, ranking and recommendation mechanisms, user behaviour, and search-system performance.

The reviewed approaches were organised according to the principal stages of the product-search pipeline: product-data preparation, index design, query processing, and ranking of retrieved items. Within each stage, the mechanisms described in the reviewed sources were grouped according to their engineering purpose, potential effect on search quality, and principal implementation limitations. Particular attention was given to the precision-recall trade-off, zero-result and irrelevant-result queries, filtering accuracy, computational cost, and other limitations associated with the use of individual search mechanisms.

The reviewed materials were then descriptively classified and summarised across the identified stages of the search pipeline. Search-quality indicators, including precision, recall, zero-result rate, click-through rate, conversion-related metrics, query latency, and computational cost, were considered as criteria for discussing the reviewed approaches rather than as empirically measured variables. The study did not include benchmarking or experimental comparison on a common dataset. Accordingly, its purpose was not to establish the quantitative superiority of individual methods, but to systematise their engineering functions, potential effects, and limitations in e-commerce search.

Results and discussion

Problems of search relevance in product catalogs

Based on the sources considered in this review, four recurring groups of factors affecting relevance in e-commerce product search can be distinguished: product-data heterogeneity, variability of user queries, semantic mismatch between user language and product descriptions, and commercial constraints. These factors occur at different stages of the search process and are

addressed by different engineering mechanisms described in the literature. In the e-commerce environment, the same product may be described through its brand, model, category, functional purpose, technical specification, color, size, material, or consumer use case. Therefore, a user query is rarely a simple set of keywords. It represents a compressed expression of demand that must be matched with both structured and unstructured elements of a product card [6]. If such queries are processed insufficiently, the search system may return results that formally match the query but are practically irrelevant, potentially reducing the usefulness of search results for the user.

One recurring source of relevance problems is the heterogeneity of product data. Information from different sources, such as suppliers, marketplace sellers, inventory management systems, and external databases, can result in inconsistencies in attribute names, measurement units, value formats, categories, and the completeness of product descriptions. These differences create difficulties when creating mappings, analyzers, and normalization rules for Elasticsearch and OpenSearch. If attributes are not separated into distinct fields and normalised during indexing, a larger proportion of product information may have to be processed as unstructured text. In this way, the effectiveness of filters, faceted browsing, and ranking would be negatively impacted, especially for complex searches using structured query elements. Thus, product-data heterogeneity represents a relevance problem that originates before query processing and ranking are performed. However, even a well-structured product catalog does not eliminate relevance loss, since search quality also depends on how the system interprets user input.

The second problem category is connected to the variability of user queries. Users may enter specific names, common categories, descriptions, abbreviations, professional terminology, jargon, transcribed words, and even misspelled entries. Given that the search window still continues to be one of the key gateways to a product catalog, it is necessary for the system to understand both standardized names of products and their unstandardized variations. This problem is confirmed by Baymard Institute's Search UX 2026 benchmark, which analyzed more than 10,000 assessments across over 170 e-commerce websites and applications. According to the benchmark, 56 % of platforms did not provide sufficient support for users' search needs, while approximately half of usability test participants preferred to start product search through the search bar rather than through navigation menus [7]. Thus, errors in query processing affect one of the central scenarios of user interaction with a product catalog.

A separate problem is synonymy and the semantic gap between user language and product-card terminology. For example, a customer may search for "running sneakers," while relevant products are categorized as "training shoes," or enter "laptop for design," although suitable items are described through processor, graphics card, and display specifications. In such cases, simple token matching does not reflect actual relevance. Approaches described for addressing this problem include synonym dictionaries, morphological analysis, query expansion, field boosting for categories and attributes, and controlled use of fuzzy search. Without appropriate configuration, relevant items may be excluded from the retrieved set, while unrelated items may receive relatively high scores because of accidental term matches.

Last but not least, in e-commerce, relevance is constrained by the problem that search results that are correct might be useless at the same time. In particular, a product may fully match the user's query but still have a low rating, be unavailable, or be unsuitable for purchase. Conversely, a product with lower lexical similarity may be more appropriate for the user's actual purchase intent

than an item with a stronger textual match. Thus, relevance in e-commerce extends beyond lexical matching and may require a combination of textual, attribute-based, behavioural, and commercial criteria. Therefore, the problem of relevance for product catalogs arises in the confluence of data quality, query variation, semantics, and business aspects. Accordingly, approaches to improving search quality are not limited to ranking adjustments at the final stage and may also involve product-data preparation, structuring of searchable fields, and index design.

Indexing and data preprocessing as the basis of relevant search results

The reviewed approaches indicate that relevance problems observed at the ranking stage may originate earlier, during product-data preparation and index construction. For this reason, preprocessing and indexing were examined as interconnected components of the search pipeline rather than as isolated technical operations. In e-commerce search, indexing is not only the technical loading of product cards into Elasticsearch or OpenSearch but also an engineering stage of data modelling that can influence the accuracy of search, filtering, and ranking [8]. The product card generally holds heterogeneous data that includes the name of the product, its brand, category, description, price, model, size, color, specifications, and availability. The lack of normalization and separation of these data into appropriate field types can lead to incorrect matching, product exclusion, or irrelevant product placement in search results.

Data preprocessing can influence how accurately catalogue attributes are matched with user queries. In e-commerce systems, this stage is especially important because queries often combine several product features rather than reproduce a full product name. If these features are not represented in the index as searchable and filterable fields, the system becomes overly dependent on textual matching. The main methods of preparing product data for search indexing are presented in table 1.

Table 1.

Methods of preparing product data for search indexing

Data preparation method	Engineering role	Potential effect on search relevance
Text cleaning	Removes markup, noise, service symbols, and redundant fragments.	Can reduce accidental matches and irrelevant results.
Case and spelling normalization	Unifies variants of brands, categories, and attributes.	Can reduce fragmentation of search results and facets.
Unit standardization	Converts values to a single measurement format.	Can improve numeric search and attribute-based filtering.
Attribute extraction	Separates color, size, material, model, and compatibility into fields.	Can improve filtering and category search.
Brand and model normalization	Links alternative spellings and product-line variants.	Can improve consistency of exact search and aggregations.
Product card deduplication	Detects duplicate or near-duplicate product records.	Can reduce information noise in search results.
Data enrichment	Adds tags, categories, synonyms, compatibility, and behavioral features.	Can expand search context and support intent matching.

After data preprocessing, the next step is to define how each product element should be represented in the index. At this stage, the same logical entity may need to be stored differently depending on the search purpose. The name of the product will be used for searching, whereas the brand will be used for matching. The price is required for sorting and filtering by ranges, and the specifications are used in faceted navigation. The use of inappropriate field types may lead to structured data being treated as regular text or vice versa. Therefore, mappings should take into account both the internal structure of the product card and the characteristics used for attribute extraction and query matching [9, 10]. The main indexed field types and their functions in product search are presented in table 2.

Table 2.**Indexed field types and their role in product search**

Field group	Data examples	Preferred indexing method	Search function
Full-text fields	Product name, description, short description.	Text fields with language analyzers.	Natural-language product search.
Exact values	Brand, SKU, model, product ID.	Keyword fields with normalizer.	Exact matching, aggregation, filtering.
Category data	Category, subcategory, product type.	Keyword fields or hierarchical structures.	Category search and faceted navigation.
Numeric attributes	Price, screen size, memory, weight, power.	Integer, float, scaled_float.	Range filtering and sorting.
Nested attributes	Product specifications, variation parameters.	Nested or object fields.	Search by complex attribute combinations.
Location and stock data	Delivery region, warehouse, availability.	Keyword, boolean, geo fields if needed.	Regional availability and stock-aware search.
Behavioral signals	Views, clicks, purchases, add-to-cart events.	Numeric or aggregated fields.	Additional ranking based on user behavior.

Configuring analyzers is the next level of index engineering optimization. An analyzer determines how the source text is transformed into searchable tokens. It usually includes character filters, a tokenizer, and token filters. In product search, this makes it possible to convert words to lowercase, remove insignificant symbols, process morphological forms, and apply stemming, lemmatization, n-grams, or synonym dictionaries.

Multi-field usage becomes particularly beneficial when one and the same field is indexable by means of multiple formats. For instance, the name of a product will be available in a full-text field, a keyword field, and an n-gram field for implementing search queries for the purpose of completion. Multi-fields with normalisers for keyword fields can be useful for values that require normalisation without tokenisation, such as brand names, SKU numbers, product codes, categories, colour designations, size indications, and inventory status.

This is particularly important for large marketplaces, where product data is supplied by multiple sellers and providers. In such environments, incomplete attributes, inconsistent categories, and different naming conventions become a systemic issue. Elasticsearch and OpenSearch may not fully compensate for low-quality source content if product data is not validated and normalized

before indexing. Therefore, indexing should be integrated into a broader data management process, while well-prepared data and a correctly designed index reduce the load on ranking mechanisms and make search more stable as the catalog grows.

Technical mechanisms for processing user queries

The reviewed query-processing mechanisms illustrate a recurring trade-off between retrieval precision and recall. Mechanisms intended to broaden matching or compensate for query variation include query expansion, spelling correction, fuzzy search, and query relaxation [11, 12]. Recent e-commerce studies also consider learned query rewriting as an approach to adapting user input before product retrieval [13]. These approaches can reduce the number of zero-result queries but may simultaneously increase the probability of retrieving irrelevant products. In e-commerce, this is especially important because queries often contain typos, abbreviations, incomplete product names, colloquial terms, mixed brand and attribute references, and implicit purchase intent. Therefore, query processing must balance recall and precision: excessive expansion may return irrelevant results, while overly strict matching may lead to zero-result searches and lost conversions. The main query processing mechanisms are presented in table 3.

Table 3.

Query processing mechanisms and their potential engineering effects in Elasticsearch and OpenSearch

Query processing mechanism	Technical implementation	Potential engineering effect	Potential risk
Query normalization	Lowercasing, whitespace removal, symbol unification, ASCII-folding.	Can improve consistency of query-index matching.	May remove meaningful symbols in model names or SKUs.
Morphological processing	Language analyzers, stemming, lemmatization, token filters.	Can expand matching across grammatical forms.	May merge semantically close but distinct terms.
Synonym handling	Synonym token filters and controlled synonym dictionaries.	May improve recall for alternative product formulations.	May over-expand queries and mix product categories.
Typo correction	Fuzzy query, edit distance, suggesters.	Can reduce zero-result searches caused by input errors.	May return irrelevant items for short or ambiguous queries.
Autocomplete	Edge n-gram, completion suggester, search-as-you-type fields.	Supports incomplete input and accelerates search.	May increase index size and generate irrelevant suggestions.
Structured attribute parsing	Query parsing and extraction of brand, model, size, color, and capacity.	Converts raw input into text, term, and range conditions.	May produce parsing errors for non-standard product names.
Field boosting	Multi-match query, field boosting, dis_max.	Improves ranking of matches in high-value fields.	May penalize relevant items with incomplete product cards.
Zero-result handling	Relaxation logic, fallback queries, partial filter removal.	Preserves the search scenario when exact matches are unavailable.	May show products that do not satisfy critical constraints.

The mechanisms presented above show that user query processing in e-commerce search cannot be reduced to simply passing the input string to full-text search. In practical implementations, a query may undergo several processing steps, including normalisation, linguistic processing, typo handling, structured-attribute extraction, synonym-based expansion, and field weighting. These mechanisms can be organised as a managed pipeline and evaluated using search logs, frequent queries, zero-result cases, and user-behaviour metrics.

Ranking search results using textual, behavioral, and business factors

For analytical purposes, the ranking mechanisms discussed in the reviewed sources can be grouped into six categories of signals: textual, attribute-based, behavioural, commercial, contextual, and infrastructure-related. Their functions differ: some determine query-product correspondence, whereas others reflect product utility, user behaviour, commercial conditions, personal context, or the computational constraints of the search system. In product search, ranking may account not only for lexical similarity to the query but also for alignment with user intent [14]. In e-commerce contexts, the practical usefulness of a retrieved item may additionally depend on its availability, required attributes, price, rating, and completeness of the product description. In Elasticsearch and OpenSearch, this logic can be implemented through field boosting, multi-match queries, `dis_max`, function score, script score, rescore, and combinations of full-text conditions with filters. Accordingly, textual relevance may be supplemented with attribute-based, behavioural, commercial, and contextual signals, while field weights should be adjusted according to product category and search scenario (table 4).

Table 4.

Signal groups used in product search ranking

Signal group	Feature examples	Technical implementation	Potential ranking role
Textual signals	Product name, description, category, brand, model.	BM25, multi-match, <code>dis_max</code> , field boosting.	Represent basic textual proximity between the query and product card.
Attribute signals	Color, size, memory capacity, material, technical parameters.	Term, range, nested queries, filters, facets.	Can refine matching by specific product characteristics.
Behavioral signals	Clicks, views, add-to-cart events, purchases.	Aggregated fields, feature store, function score, ML models.	Can reflect observed patterns of user interaction with products.
Commercial signals	Margin, stock availability, delivery time, price, rating.	Script score, <code>rank_feature</code> , business rules.	Can incorporate commercial characteristics into ranking.
Contextual signals	Region, device, user history, seasonality.	Personalized boost rules, separate indices, hybrid ranking.	Can account for the current user or request context.
Infrastructure signals	Response time, cluster load, index size, computation cost.	Rescore window, caching, shard routing, script limits.	Can help manage trade-offs between ranking complexity and system performance.

Behavioural and commercial signals may become particularly relevant when product search extends beyond lexical matching. If products have similar textual relevance, higher positions may be assigned to items with stronger evidence of user interest, such as clicks, add-to-cart events, or purchases. Short-term interaction history can also be incorporated into ranking to account for the

buyer's current search intent [15]. However, these signals require careful interpretation, since high click-through rates may reflect visual appeal, low price, or temporary promotion rather than actual satisfaction. These observations indicate that behavioural signals can be interpreted together with commercial and outcome-related metrics, such as conversion rate, revenue per search, purchase rate, and returns to search results.

Hybrid ranking architectures represent one of the approaches discussed for combining lexical retrieval with additional ranking signals. Elasticsearch and OpenSearch can be used for candidate retrieval, filtering, and primary scoring, while an external machine-learning component performs additional reranking of a limited set of retrieved products. Such architectures can be used to combine lexical and attribute-based matching with behavioural signals, user context, product availability, and other features that are difficult to incorporate into a static scoring formula.

Taken together, the reviewed approaches show that product-search ranking can incorporate textual and attribute-based matching together with behavioural, commercial, and contextual signals. In Elasticsearch and OpenSearch, built-in scoring mechanisms may also be combined with analytics, feature engineering, machine-learning components, and experimental evaluation where these are appropriate to the search scenario.

For the purposes of this review, the discussed approaches can be summarised according to the main stages of the product-search process, each addressing different engineering tasks and limitations. A summary of the reviewed methods, their potential effect on search quality, and their principal risks is presented in table 5.

Table 5.**Summary of engineering methods for improving product-search relevance**

Search stage	Main methods	Potential effect	Main limitations
Product data preparation	Normalisation, cleaning, attribute extraction, deduplication.	Potentially more consistent data and more accurate filtering.	Dependence on source-data quality.
Index design	Text, keyword, numeric and nested fields; analysers; multi-fields.	Potential improvement in query-product matching and filtering.	Larger and more complex indices.
User query processing	Synonyms, typo correction, attribute extraction, query relaxation.	Potentially higher recall and fewer zero-result queries.	Possible loss of precision.
Result ranking	BM25, field boosting, function score, behavioural and commercial signals.	Potential improvement in the ordering of retrieved products.	Weight calibration and ranking bias.

Thus, individual mechanisms should not be considered in isolation. Their practical effects may depend on decisions made at the data, index, query-processing, and ranking levels, as well as on trade-offs between search quality and computational cost.

The sources considered in this review suggest that no single Elasticsearch or OpenSearch mechanism is sufficient to address all aspects of product-search relevance. Data normalisation and index design can support more accurate retrieval, whereas synonym expansion, fuzzy matching, and query relaxation may increase recall at the cost of additional irrelevant results. Behavioural and commercial signals can provide additional information for ranking but require careful weighting to avoid excessive preference for popular or promoted products. At scale, these mechanisms may also

introduce technical constraints: extensive use of n-grams and enriched fields increases index size, complex scoring and reranking may increase query latency, and behavioural features require regular updating and additional data-processing infrastructure. As an analytical review, the study does not provide experimental validation or a quantitative comparison of the discussed mechanisms on a common dataset. Accordingly, the effects described above should be interpreted as literature-based engineering tendencies and potential trade-offs rather than as experimentally verified performance gains.

Conclusion

The review indicates that search relevance in Elasticsearch- and OpenSearch-based e-commerce systems can be addressed at several interconnected stages of the search pipeline. At the data and indexing level, the key engineering measures include normalisation of product attributes and measurement units, separation of searchable and filterable characteristics into appropriate field types, and the use of suitable analysers and structured fields. These measures can reduce dependence on unstructured textual matching and support more structured filtering and ranking.

At the query-processing and ranking stages, relevance depends on the controlled combination of query normalisation, synonym handling, typo correction, structured attribute extraction, field boosting, and fallback logic for zero-result queries. The reviewed mechanisms illustrate a recurring precision-recall trade-off: broader matching may improve recall but can also increase the number of irrelevant results, whereas stricter filtering may improve precision while increasing the risk of excluding potentially relevant products. Final ranking can therefore combine textual matching with attribute-based, behavioural, commercial, and contextual signals using mechanisms such as multi-match, dis_max, function score, rank_feature, and rescore.

Based on the review, product-data preparation, index design, query transformation, ranking, and operational monitoring can be considered as interconnected stages of the product-search process. The performance of such search configurations can be assessed using retrieval, behavioural, commercial, and infrastructure indicators, including precision, recall, zero-result rate, click-through rate, conversion-related metrics, query latency, and computational cost. The study is limited by its analytical design and does not include an experimental comparison of the identified mechanisms on a common dataset, which defines a direction for further research.

References

1. Salsify. 2025 Consumer Research [Electronic resource]. – URL: <https://www.salsify.com/resources/report/2025-consumer-research> (дата обращения: 04.08.2026).
2. Ren Z., He X., Yin D. de Rijke M. Information discovery in e-commerce // Foundations and Trends in Information Retrieval. – 2024. – Vol. 18, No. 4–5. – P. 417–690. – DOI: 10.1561/15000000097.
3. Nguyen D.A., Mohan R.K., Yang S., Akash P.S., Chang K.C.-C. MiniELM: A lightweight and adaptive query rewriting framework for e-commerce search optimization // Findings of the Association for Computational Linguistics: ACL. – 2025. – P. 6952–6964. – DOI: 10.18653/v1/2025.findings-acl.363.

4. Ou R., Wang K., Shao J., Zhang T., Huo C. Aligning query rewriting with human cognition and preference in e-commerce search // *Proceedings of the ACM Web Conference*. – 2026. – P. 7933–7942. – DOI: 10.1145/3774904.3792806.
5. Jain S. DSouza C.A., Rana J.B., Joshi A., Yenigalla P. PROBES: Performance and relevance observation for better search // *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 5: Industry Track)*. – 2026. – Vol. 5. – P. 625–635. – DOI: 10.18653/v1/2026.eacl-industry.48.
6. Zhenyu L., Haoran C., Mingkai W. Personalizing search, navigation, and product discovery experiences in B2C e-commerce using machine learning models // *Chronicle of Emerging Scientific Paradigms*. – 2025. – No. 10.
7. Baymard Institute. Ecommerce Search UX 2026: 8 Search “Query Types” UX Best Practices [Electronic resource]. – URL: <https://baymard.com/blog/ecommerce-search-query-types> (дата обращения: 04.08.2026).
8. Sharma C., Kaur A., Datta P., Gulzar Y. Optimizing eCommerce data: effective approaches for data collection, cleansing, and preprocessing // *Strategic Innovations of AI and ML for E-Commerce Data Security*. IGI Global Scientific Publishing. – 2025.
9. Loughnane R., Liu J., Chen Z., Wang Z., Giroux J., Du T., Schroeder B., Sun W. Explicit attribute extraction in e-commerce search // *Proceedings of the Seventh Workshop on e-Commerce and NLP @ LREC-COLING*. – 2024. – P. 125–135.
10. Moulton M., Ng Y.-K. Boolean interpretation, matching, and ranking of natural language queries in product selection systems // *Discover Computing*. – 2024. – Vol. 27. – DOI: 10.1007/s10791-024-09432-x.
11. Allahim A., Cherif A., Imine A. Semantic approaches for query expansion: taxonomy, challenges, and future research directions // *PeerJ Computer Science*. – 2025. – P. 1–28. – DOI: 10.7717/peerj-cs.2664.
12. Güzel Turhan C. Leveraging large language models for spelling correction in Turkish // *PeerJ Computer Science*. – 2025. – P. 1–26. – DOI: 10.7717/peerj-cs.2889.
13. Dai A., Zhu Z., Hu H., Tang G., Liu L., Xu S. Enhancing E-Commerce Query Rewriting: A Large Language Model Approach with Domain-Specific Pre-Training and Reinforcement Learning // *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*. – 2024. – P. 4439–4445. – DOI: 10.1145/3627673.3680109.
14. Saadany H., Bhosale S., Agrawal S., Kanojia D., Orasan C., Wu Z. Centrality-aware product retrieval and ranking // *Proceedings of the*. – 2024. – P. 215–224. – DOI: 10.18653/v1/2024.emnlp-industry.17.
15. Sheng T., Muthiah S., Islam A., Feng J. Progressive refinement of e-commerce search ranking based on short-term activities of the buyer // *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*. – 2025. – P. 4370–4374. – DOI: 10.1145/3726302.3731966.