Journal of Digital Marketing and Communication

Research Article
Volume 6(2), 2026, 239–259
https://doi.org/10.53623/jdmc.v6i2.1433

The Fast Lane Keyword Methodology: A Structural Framework for Identifying and Prioritizing Zero-Data Long-Tail Keywords in Competitive Digital Markets

Pedro Pereira

Independent Researcher, Belo Horizonte, Minas Gerais, Brazil

Correspondence: pedrohbh@hotmail.com

SUBMITTED: 21 September 2026; REVISED: 21 October 2026; ACCEPTED: 23 October 2026

ABSTRACT: Keyword research tools report search volume, keyword difficulty (KD), and cost-per-click (CPC) only for queries that exceed their aggregation thresholds. Queries below these thresholds are routinely excluded. This study quantifies the proportion of queries with unavailable keyword metrics and evaluates the Fast Lane Keyword Methodology, a framework for prioritizing zero-data queries without relying on search volume. The Fast Lane Score combines four indicators, each scored from 0 to 3: query length, intent-modifier categories, search engine results page (SERP) weakness, and semantic parent match with a test publisher. Google autocomplete suggestions were collected for 29 seed keywords across three sectors (SEO tools, small-business consulting, and e-commerce) in the United States in October 2026, yielding 6,993 relevant queries. Keyword metrics were obtained from DataForSEO, and live SERPs were analyzed for each zero-data query. A total of 2,459 queries (35.2%) had no reported search volume, KD, or CPC, with proportions ranging from 14.7% to 70.0% across seed keywords. Zero-data queries were slightly longer than measurable queries (mean: 4.50 vs. 4.14 words) but did not contain more intent modifiers. Their top-ranking results were also more likely to lack pages directly matching the queries (72.8% vs. 58.1%). Based on the prespecified thresholds, 0.1% of zero-data queries were classified as Priority A, 9.5% as Priority B, and 45.9% as Priority C. In contrast, a conventional volume-based scoring method assigned a score of zero to all these queries. Approximately one-third of Google-suggested queries lack standard keyword metrics. The proposed structural scoring framework enables reproducible prioritization of these queries, providing small and medium-sized businesses with a systematic approach to identifying content opportunities overlooked by volume-based keyword research methods. However, the scoring thresholds require further calibration and validation against actual search-ranking outcomes.

KEYWORDS: Search engine optimization; query length; search intent; SERP analysis; small business content strategy

1. Introduction

Keyword research has evolved from selecting terms based solely on search volume to balancing search demand against competition, using metrics such as keyword difficulty, cost-per-click (CPC), and competitive density [1, 2]. Organic search visibility remains a major determinant of traffic to commercial websites [1], while query formulation influences which sources and search engine results page (SERP) features are displayed [3]. Therefore, selecting appropriate target queries remains one of the most consequential decisions in search engine marketing.

However, these metrics are generally available only for queries that occur frequently enough to be measured. Search demand is highly skewed, with a substantial proportion of queries occurring infrequently, many appearing only once, and user behavior differing between common and rare queries [4, 5]. Similar long-tail patterns have been observed in e-commerce search logs, where low-frequency queries generate sparse interaction signals that even large platforms struggle to interpret and address effectively [6, 7]. Keyword research tools that aggregate search data often report no search volume, keyword difficulty, or CPC for these infrequent queries. Conventional keyword research practices frequently interpret the absence of such metrics as an indication of limited opportunity, resulting in the exclusion of these queries during the initial screening process.

However, the consequences of this filtering approach are not necessarily neutral. In long-tail markets, estimation methods that exclude zero observations can produce biased demand estimates [8]. Furthermore, reductions in search costs have expanded the range of products purchased by consumers online, reinforcing the economic significance of long-tail markets [9, 10]. If zero-data queries represent genuine search demand, systematically excluding them may conceal potential opportunities from businesses that rely on conventional keyword research tools. Consequently, these queries may face less direct competition from businesses employing similar selection strategies.

This asymmetry is particularly relevant to small and medium-sized businesses (SMBs). Competing for high-volume queries often requires substantial domain authority, extensive backlink profiles, and significant content development budgets, which resource-constrained organizations may lack. Although SMBs recognize the importance of search engine optimization (SEO) for online visibility, limited resources often force them to treat SEO as a secondary business activity [11]. Moreover, existing SEO evaluation approaches have largely been developed for large, globally oriented websites rather than smaller, locally focused businesses [12]. From a resource-based perspective, sustainable competitive advantage may emerge from identifying and exploiting opportunities that competitors cannot easily recognize or replicate [13].

The search environment is also undergoing significant transformation. AI Overviews now appear in approximately half of real-user Google queries and are positioned above conventional organic search results [14]. Although users continue to interact with organic links through familiar browsing patterns [15], generative search engines select and cite information sources through mechanisms over which publishers have limited control [16, 17]. In addition, conversational search interfaces encourage users to formulate longer, more specific, and more structured queries than those typically used in traditional search environments [18, 19]. Such queries may be particularly likely to fall below the aggregation thresholds applied by conventional keyword research tools.

Previous studies have primarily examined long-tail queries from the search engine perspective, focusing on caching, information retrieval, and query rewriting [5–7]. Other research has investigated keyword selection and optimization strategies for queries with measurable search volumes [20]. However, to the author's knowledge, no published framework has specifically addressed the systematic evaluation of queries for which conventional keyword research tools report no measurable data. Furthermore, limited empirical evidence is available regarding the prevalence of these zero-data queries and the characteristics of their associated search results. These limitations highlight the need for an alternative keyword evaluation approach that does not depend exclusively on conventional search volume metrics.

This study proposes the Fast Lane Keyword Methodology, a four-signal scoring framework designed to prioritize zero-data queries independently of search volume, and evaluates its application using a dataset of 6,993 queries across three industry sectors. The study makes three primary contributions. First, it quantifies the proportion of zero-data queries generated through standard query-expansion methods and examines their structural characteristics and SERP features. Second, it introduces a transparent and reproducible scoring rubric, supported by explicit justifications for the selected indicators, weights, and classification thresholds. Third, it compares the Fast Lane Score with a conventional volume-based keyword opportunity score to identify differences in query prioritization and determine which potential opportunities are retained or excluded by each approach.

2. Literature Review

2.1. Long-tail queries and long-tail demand

The long-tail theory suggests that aggregated demand for niche products can rival the demand for a small number of popular products [10]. In online retail, reduced search costs result in a measurably less concentrated sales distribution compared with traditional retail channels [9]. Web search exhibits a similar pattern, with a substantial proportion of queries occurring infrequently and search behavior differing between high-frequency (head) and low-frequency (tail) queries [4]. The prevalence of queries that occur only once (singletons) also limits the efficiency of search-result caching [5]. Recent studies on e-commerce search confirm that long-tail query frequencies follow a Pareto distribution and generate sparse interaction signals [6]. Furthermore, existing query-rewriting methods often struggle to retrieve relevant results for these infrequent queries [7]. Collectively, this literature establishes that long-tail queries constitute a substantial proportion of search activity and present significant retrieval challenges. However, it does not adequately address the practical challenges faced by marketers in determining which long-tail queries warrant content investment when conventional keyword research tools provide no measurable data.

2.2. Keyword Metrics and the Limitations of Data Aggregation.

Keyword research tools rely on aggregated search query data to estimate metrics such as search volume, keyword difficulty, and cost-per-click (CPC). Queries that fall below a tool's minimum reporting threshold may be returned without measurable values. Research on keyword selection in search advertising demonstrates that keyword specificity interacts with search volume. Adding product-specific information to a keyword can influence click-through rates and advertising costs differently depending on the search volume associated with the product category [20]. More broadly, research on demand estimation indicates that excluding zero-sales

observations can introduce bias into demand estimates in long-tail markets [8]. A similar concern arises in keyword research when queries with unavailable search volume data are excluded, effectively treating the absence of measurement as evidence of absent demand. To the author's knowledge, no peer-reviewed study has systematically quantified the extent to which keyword research tools fail to report actual search activity. However, the present study does not attempt to measure this unreported activity directly. Instead, it quantifies the proportion of zero-data queries generated through a standardized keyword expansion procedure and examines the characteristics of their associated search engine results pages (SERPs).

2.3. Query Length, Specificity, and Conversational Search.

Longer queries generally contain more specific information, and variations in query wording can influence the search results retrieved by search engines [3, 21]. The growing adoption of conversational and large language model (LLM)-based search systems has also increased the relevance of longer and more complex queries. Prompts submitted to LLM-based assistants are frequently verbose and structured, often expressing broader or more detailed information needs than conventional search queries [18]. Conversational search systems are designed to process complex natural-language queries and maintain contextual information across multiple interactions [19]. Moreover, users of LLM-powered search systems exhibit querying behaviors that differ from those associated with traditional search engines [22]. Despite these developments, conventional search engines and conversational AI systems continue to serve distinct functions. Users frequently rely on search engines to browse and verify information, whereas chatbots are increasingly used to delegate tasks and obtain synthesized responses [23, 24]. Consequently, longer and more specific queries are likely to remain relevant to conventional search environments. Such queries may also be more likely to fall below the aggregation thresholds applied by keyword research tools, making them particularly relevant to alternative keyword evaluation approaches.

2.4. Search Intent.

Broder's taxonomy classifies search queries into three primary categories: informational, navigational, and transactional [25]. Automated classification of queries according to these categories has been extensively investigated [26], while topic-based query categorization has also been shown to improve information retrieval performance [27]. Recent studies suggest that generative artificial intelligence has transformed user expectations of search systems, highlighting the need for revised, user-centered taxonomies of search intent [28, 29]. Nevertheless, automated intent classification remains challenging, particularly for short or ambiguous queries, even when LLM-based approaches are employed [30]. To address these limitations, the Fast Lane Score employs a transparent, rule-based lexical approach rather than a machine-learning classifier. Specifically, it evaluates the presence of comparative, superlative, price-related, and audience-specific modifiers as observable indicators of query intent. This approach facilitates consistent application, minimizes dependence on complex classification models, and enables practitioners to reproduce the scoring procedure.

2.5. SERP Composition, User-Generated Content, and AI-Generated Answers.

The composition and presentation of search engine results pages influence how users access, evaluate, and interact with retrieved information. Different result presentation formats can affect search efficiency and user satisfaction depending on the nature of the search task [31].

Recent developments in generative search have further transformed SERP composition. AI Overviews and featured snippets displayed on the same results page may provide inconsistent information [32], while AI-generated answer boxes have become increasingly prominent SERP features [33]. Furthermore, sources cited in AI-generated responses overlap only partially with conventional organic search results [14, 34]. In health-related searches, AI Overviews have been observed to give greater prominence to user-generated multimedia content, including YouTube videos, compared with conventional organic results [34]. The effects of generative AI on online communities have also varied. For example, Stack Overflow experienced a decline in activity following the introduction of ChatGPT, whereas Reddit communities showed no comparable evidence of declining participation [35]. Search result quality remains another important concern. A year-long audit of Google, Bing, and DuckDuckGo revealed persistent difficulties in distinguishing high-quality product reviews from heavily optimized affiliate content [36]. Related research has proposed criteria for evaluating high-quality SEO content, including characteristics relevant to determining whether a ranking page adequately addresses a user's query [37]. However, to the author's knowledge, no peer-reviewed study has systematically quantified the extent to which forums and other user-generated content dominate top-ranking results for long-tail commercial queries. The present study addresses this gap by examining the prevalence of such content within its dataset.

2.6. Topical Authority and Semantic Relevance.

Modern information retrieval systems increasingly evaluate semantic relationships between queries and documents rather than relying exclusively on exact keyword matching. Dense retrieval models employ learned semantic representations to capture relationships between queries and textual content [38], while LLMs are increasingly integrated into information retrieval systems [39]. In SEO practice, these developments are frequently associated with the concept of "topical authority," which suggests that websites providing comprehensive coverage of a particular subject may be better positioned to rank for semantically related queries. Although this concept is widely discussed among SEO practitioners, its independent effect on search rankings has not been conclusively established through peer-reviewed empirical research. Accordingly, the Fast Lane Score adopts a narrowly defined and observable operational measure of topical relevance rather than assuming that topical authority directly determines ranking performance..

2.7. Composite Scoring and Weighting.

The Fast Lane Score is structured as an additive composite index that combines multiple indicators into a single evaluation measure. Weight selection and aggregation methods are among the most debated aspects of composite index construction [40, 41], particularly because subjective weighting decisions can introduce bias and influence the resulting rankings [42]. Previous research has demonstrated that unit-weighted linear composites can achieve predictive performance comparable to that of optimally weighted models across various applications [43, 44]. However, equal weighting is not universally optimal. For example, research involving partial least squares structural equation modeling (PLS-SEM) has shown that equally weighted composites may exhibit weaker out-of-sample predictive performance than composites using differentiated weights [45]. In the context of local-business SEO, an averaged composite indicator has also been employed and found to correlate with business performance [12]. These findings suggest that simple weighting approaches can provide a

transparent and practical foundation for composite scoring, although their suitability depends on the characteristics and objectives of the evaluation framework. Therefore, the Fast Lane Score adopts an explicitly defined weighting structure to support reproducibility and interpretability.

3. Materials and Methods

3.1. Origin of the Framework.

The framework was developed iteratively between February and March 2026, during which the author managed keyword strategies for two commercial keyword research tools. The initial approach employed a conventional keyword opportunity score, calculated as follows:

$$\text{Opportunity} = \text{Volume} \times (1 - \text{KD}/100) \times \text{Intent Match}$$

Intent weights were assigned as 1.0 for transactional or local queries, 0.6 for informational queries, and 0.3 for navigational queries. When applied to the complete set of queries generated by a keyword research tool, this formula assigns a score of zero to every zero-data query, regardless of its wording or SERP characteristics. The Fast Lane Score was therefore developed to prioritize queries that conventional volume-based scoring methods cannot adequately evaluate. Earlier versions of the framework employed different word-count categories and a tiered intent classification scale. The scoring rubric presented in this study was finalized before the validation study and remained unchanged after data collection. An earlier operational estimate suggested that approximately 5–15% of generated queries returned no keyword metrics. However, this estimate was based on informal observations across multiple campaigns and was not systematically documented.

3.2. Data Collection.

The study was conducted independently of NP Digital and its associated tools, using a DataForSEO account funded by the author. Data collection was performed on 21 October 2026, with the United States selected as the target market and English as the search language.

3.2.1. Verticals and Seed Keywords.

Three industry sectors were selected because they involve substantial participation by small and medium-sized businesses (SMBs) and face established competition from large online publishers: SEO and digital marketing tools, small-business consulting, and e-commerce. Ten seed keywords were initially identified for each sector, yielding 30 seed keywords in total (Appendix A). The seeds were selected to represent major product and service categories within each sector, emphasizing generic head terms rather than brand-specific keywords. One seed keyword, "rank tracker," was excluded after data collection because the term also refers to video-game ranking tools. Consequently, unrelated gaming suggestions satisfied the predefined relevance criteria and dominated the results for this seed.

Table 1. Seed Keywords By Vertical.

Vertical	Seed keyword	Queries	Zero-data (%)
SEO and digital marketing tools	keyword research tool	237	83 (35.0)
SEO and digital marketing tools	seo audit tool	99	47 (47.5)
SEO and digital marketing tools	rank tracker	excluded	excluded
SEO and digital marketing tools	backlink checker	136	32 (23.5)
SEO and digital marketing tools	seo software for small business	2	0 (0.0)
SEO and digital marketing tools	content optimization tool	10	2 (20.0)
SEO and digital marketing tools	local seo tool	53	30 (56.6)
SEO and digital marketing tools	ai seo tool	98	65 (66.3)
SEO and digital marketing tools	website speed test	194	72 (37.1)
SEO and digital marketing tools	email marketing software	158	48 (30.4)
Small business consulting	small business consultant	139	67 (48.2)
Small business consulting	business plan writer	81	19 (23.5)
Small business consulting	small business bookkeeping	198	105 (53.0)
Small business consulting	llc formation service	49	14 (28.6)
Small business consulting	small business marketing agency	29	9 (31.0)
Small business consulting	hr outsourcing	90	27 (30.0)
Small business consulting	business coach	458	195 (42.6)
Small business consulting	payroll service	435	187 (43.0)
Small business consulting	small business grant	469	188 (40.1)
Small business consulting	fractional cfo	404	283 (70.0)
E-commerce	shopify store	498	171 (34.3)
E-commerce	dropshipping	677	112 (16.5)
E-commerce	print on demand	583	125 (21.4)
E-commerce	ecommerce website builder	116	57 (49.1)
E-commerce	product photography	490	149 (30.4)
E-commerce	amazon fba	470	69 (14.7)
E-commerce	etsy shop	502	173 (34.5)
E-commerce	abandoned cart email	122	47 (38.5)
E-commerce	subscription box business	51	20 (39.2)
E-commerce	ecommerce shipping	145	63 (43.4)

3.2.2. Query Generation.

Google autocomplete suggestions were collected for each seed keyword using 49 query probes, following the structure commonly employed by question-and-modifier keyword research tools. These probes comprised 11 question prefixes (how, what, why, where, when, who, which, can, will, are, and is), seven preposition and modifier suffixes (for, with, without, near, to, can, and is), five comparison suffixes (vs, versus, or, and, and like), and 26 alphabetical suffixes (a–z). This procedure generated 1,470 probes and yielded 9,036 unique suggestions after duplicate removal.

3.2.3. Relevance Filtering.

Google autocomplete may return suggestions that are unrelated to the intended topic of a seed keyword. For example, the seed "rank tracker" generated suggestions associated with video-game ranking systems. To address this issue, a relevance-filtering rule was established before keyword metrics were retrieved. Only queries containing every content word of the corresponding seed keyword were retained, with singular and plural forms treated as equivalent. This conservative filtering approach may exclude relevant paraphrases, such as "rank tracking for seo." Nevertheless, it provides a consistent and reproducible relevance criterion. The filtering procedure retained 7,185 queries. Following the exclusion of the ambiguous seed keyword, the final dataset comprised 6,993 queries derived from 29 seeds.

3.2.4. Keyword Metrics.

Search volume and cost-per-click (CPC) data were retrieved through the Google Ads endpoint, while keyword difficulty (KD) was obtained from DataForSEO Labs. The Google Ads endpoint accepts keywords containing no more than 80 characters and 10 words. Consequently, 33 longer queries could not be submitted and had no available search volume or CPC data from this source. This limitation represents an additional form of the measurement gap examined in the study. Section 4.5 presents a sensitivity analysis excluding these queries. A query was classified as zero-data when its search volume, KD, and CPC were all reported as zero or were unavailable.

3.2.5. SERP Data.

For each zero-data query, the top 10 organic search results were retrieved, including the domain, URL, and page title of each result. The estimated monthly organic traffic of every domain appearing among the top 10 results was also collected as a proxy for domain authority. In addition, the main textual content of the highest-ranking result was retrieved to determine its word count. To facilitate comparisons between zero-data and measurable queries, equivalent SERP data were collected for a simple random sample of 500 queries with measurable keyword metrics.

3.3. Measures.

The Fast Lane Keyword Methodology incorporates four complementary signals designed to evaluate the structural characteristics, commercial relevance, competitive conditions, and topical alignment of zero-data queries. Figure 1 illustrates the conceptual relationships between these signals and the proposed outcome. The conceptual model presented in Figure 1 distinguishes the observable indicators used in the scoring framework from the potential ranking outcomes that may result from their application. Although the four signals are combined into a composite opportunity score, the proposed relationship between this score and actual search-ranking performance is not empirically tested in the present study.

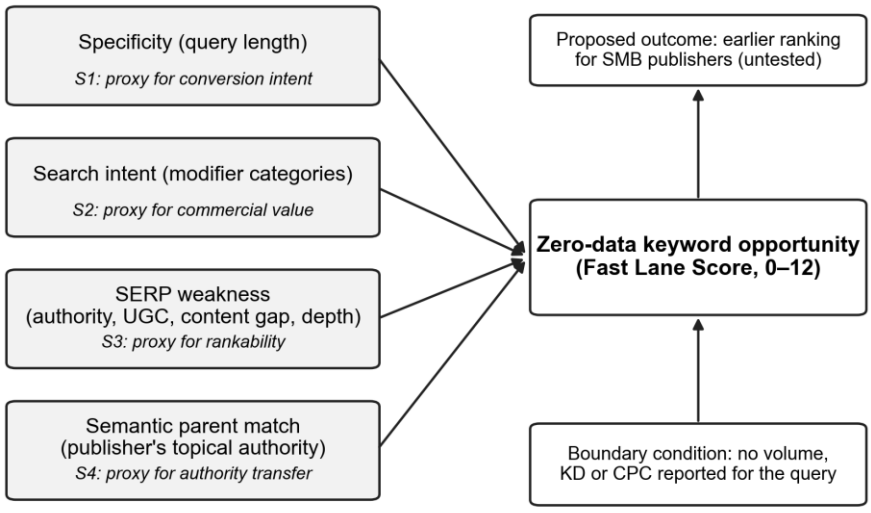


Figure 1. Conceptual Model of the Fast Lane Score. The Proposed Outcome is not Tested in This Study.

3.3.1. Signal 1.

Query length is used as a proxy for specificity, based on evidence that longer queries tend to express more detailed information needs [3, 21]. As presented in Table 2, queries containing 1–3 words receive 0 points, those containing 4–5 words receive 1 point, those containing 6–7 words receive 2 points, and those containing eight or more words receive 3 points. This scoring structure assumes that greater query specificity may indicate a more clearly defined information need, although longer queries do not necessarily demonstrate stronger commercial intent.

Table 2. Fast Lane Score: Signals, Scoring Criteria, and Interpretation.

Signal	Construct (proxy for)	Scoring criteria	Points
S1: Word count	Specificity and potential conversion intent	1–3 words = 0; 4–5 = 1; 6–7 = 2; ≥8 = 3	0–3
S2: Intent modifiers	Buyer stage and potential commercial value	+1 per modifier category (comparative, superlative, price/value, or audience), capped at 3	0–3
S3: SERP weakness	Competition and potential rankability	+0.75 per condition: no high-authority domain among the top five; ≥2 user-generated results among the top five; no direct query match in the top-five titles; highest-ranking result contains fewer than 1,000 words	0–3
S4: Semantic parent match	Topical relevance and potential authority transfer	+1 per condition: ranks for the parent keyword; ranks among the top 20 within the same topical cluster; ranks among the top 10 for the seed keyword	0–3
Fast Lane Score	Structural keyword opportunity	S1 + S2 + S3 + S4	0–12
Priority classification	Keyword prioritization	A: 9–12 (prioritize for immediate content development); B: 6 to less than 9 (review weaker signals before investment); C: 3 to less than 6 (monitor or consider selectively); Deprioritize: below 3	Four categories

Abbreviations: KD, keyword difficulty; CPC, cost-per-click; SERP, search engine results page.

3.3.2. Signal 2.

This signal evaluates lexical indicators associated with different stages of the buyer's decision-making process. As summarized in Table 2, one point is assigned for the presence of each of four modifier categories, with the total capped at 3 points. These categories comprise comparative modifiers (vs, versus, compare, alternative, or, instead of, replace, better than, and difference), superlative modifiers (best, top, cheapest, top rated, fastest, and easiest), price or value modifiers (free, cheap, affordable, price, pricing, cost, under $X, monthly, and lifetime), and audience qualifiers (for small business, for freelancers, for beginners, for agencies, and similar expressions). Multiple modifiers belonging to the same category contribute only one point. This category-based approach ensures consistency between the operational scoring procedure and the criteria presented in Table 1, resolving an inconsistency in an earlier version of the framework that counted individual modifier words rather than categories.

3.3.3. Signal 3.

This signal assesses the competitive characteristics of organic search results using four observable conditions, each contributing 0.75 points, for a maximum of 3 points, as detailed in Table 2. The first condition is satisfied when none of the top five organic results originates from a high-authority domain, operationally defined as a domain receiving an estimated 1,000,000 or more organic visits per month. The second condition is satisfied when at least two

of the top five results consist of forum discussions, question-and-answer pages, or other user-generated content. The third condition is satisfied when none of the top five result titles contains at least 80% of the query's content words, indicating the absence of a direct lexical match. The fourth condition is satisfied when the highest-ranking result contains fewer than 1,000 words of main textual content. An earlier version of the rubric assessed authority using a third-party domain-authority metric. In the present framework, estimated monthly organic traffic replaces that metric to improve reproducibility through the use of a consistent data source. The four conditions summarized in Table 2 collectively provide a standardized measure of SERP weakness, with higher scores indicating potentially less competitive search environments.

3.3.4. Signal 4.

This signal evaluates the topical alignment between a target query and an established publisher within the corresponding industry sector. A test publisher was selected for each vertical using a rule established before scoring. Specifically, the domain appearing most frequently among the top 10 organic results for that vertical's zero-data queries was selected, excluding user-generated content platforms and high-authority domains. This procedure identifies an established, SMB-scale publisher that already competes for relevant queries while minimizing subjective intervention in publisher selection (Table 3).

Table 3. Test publishers.

Vertical	Test publisher	Appearances in zero-data top 10	Estimated monthly organic visits
SEO and digital marketing tools	semrush.com	65	946,988
Small business consulting	revenuebase.ai	115	43,792
E-commerce	shopify.com	210	669,160

The signal awards 1 point for each of three conditions: (a) the publisher ranks for a keyword containing the seed keyword, indicating existing parent-keyword content; (b) the publisher ranks among the top 20 results for a keyword sharing at least two content words with the target query, indicating membership in the same topical cluster; and (c) the publisher ranks among the top 10 results for the seed keyword itself, indicating an established pillar topic. Each condition is evaluated using the publisher's ranked-keyword data. Higher scores indicate stronger existing relationships between the publisher's ranking profile and the target query, although they do not establish a causal relationship between topical authority and subsequent ranking performance. he priority classifications presented in Table 1 provide an operational basis for allocating content development resources. Priority A identifies queries that satisfy several favorable conditions within the framework, whereas Priority B identifies potentially attractive queries requiring additional assessment. Priority C represents lower-priority opportunities that may warrant monitoring, while queries scoring below 3 are deprioritized. These classifications represent proposed decision rules rather than empirically validated performance thresholds.

3.4. Procedure

The methodological procedure consists of sequential stages, beginning with seed keyword selection and ending with query scoring and prioritization. As illustrated in Figure 2, the process integrates autocomplete-based query expansion, relevance filtering, keyword metric retrieval, zero-data classification, SERP analysis, and composite scoring. The figure highlights

the distinction between conventional keyword evaluation and the proposed Fast Lane approach. Under conventional volume-based scoring, queries without measurable search volume receive a score of zero and are effectively excluded from further prioritization. In contrast, the Fast Lane procedure retains these queries and evaluates them using observable structural, lexical, competitive, and topical characteristics. All data collection and scoring procedures were automated using a single script. Each API response was stored to preserve the original retrieved information and support reproducibility. The script, seed keyword list, and scored dataset are intended to be deposited in an open-access Zenodo repository following manuscript acceptance. The corresponding digital object identifier (DOI) will be added to this section once the repository record becomes available.

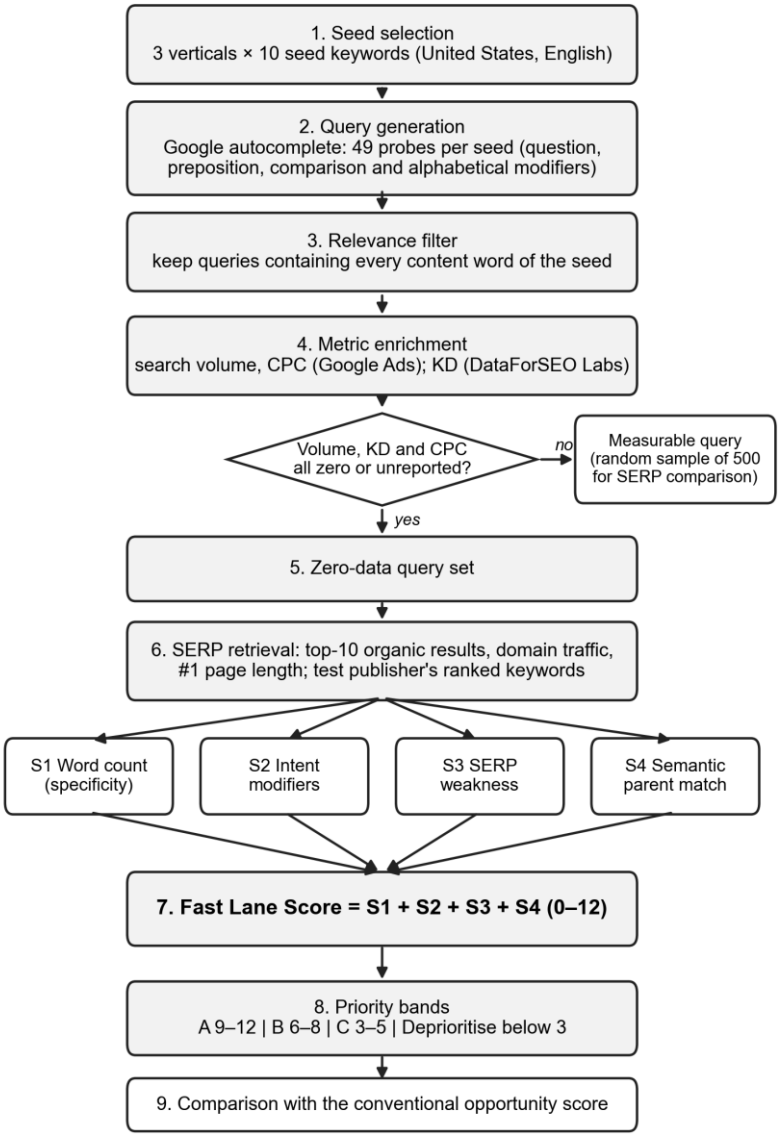


Figure 2. Methodological Flowchart of the Fast Lane Procedure.

3.5. Justification of weights and thresholds

3.5.1. Equal Weighting.

Each of the four signals contributes a maximum of 3 points to the composite Fast Lane Score. Because empirical ranking outcomes were unavailable for estimating differentiated weights,

equal weighting was adopted as a transparent and reproducible initial approach. Previous studies have demonstrated that unit-weighted linear composites can achieve predictive performance comparable to that of optimally weighted models across various applications [43, 44]. Equal weighting also reduces reliance on potentially biased subjective judgments when determining the relative importance of individual indicators [42]. Nevertheless, equal weighting is not universally optimal, and differentiated weights may improve predictive performance under certain conditions [45].

3.5.2. Three-Point Scale for Each Signal.

Each signal is evaluated using a standardized scale ranging from 0 to 3 points. Signals 1, 2, and 4 each have four possible scoring levels, whereas Signal 3 comprises four conditions, each contributing 0.75 points. Standardizing the maximum contribution of each signal prevents any individual indicator from disproportionately influencing the composite score solely because it contains more scoring categories or conditions. This approach also improves interpretability by maintaining a consistent scoring range across the four dimensions. However, the use of a common numerical scale does not necessarily imply that the indicators have equivalent predictive importance or contribute equally to actual search-ranking performance.

3.5.3. Priority Thresholds.

The composite Fast Lane Score ranges from 0 to 12, with higher values indicating stronger structural opportunities based on the proposed scoring criteria. Queries scoring 9 or above are classified as Priority A, reflecting favorable performance across several dimensions. Scores ranging from 6 to less than 9 are assigned to Priority B, indicating potentially attractive opportunities with one or more weaker dimensions. Queries scoring from 3 to less than 6 are classified as Priority C, while those scoring below 3 are deprioritized. These thresholds were established before the validation study to provide a practical, transparent, and interpretable classification framework. Section 4.3 evaluates the resulting priority categories by examining differences in query structure and SERP characteristics across the scoring bands. However, these comparisons do not establish whether the proposed thresholds can reliably predict subsequent search-ranking performance.

3.6. Data Analysis.

Descriptive statistics were calculated for the complete dataset and separately for each industry sector. Zero-data and measurable queries were compared in terms of word count, intent-modifier categories, and SERP characteristics. Mann–Whitney U tests were employed to compare ordinal measures, while chi-square (χ^2) tests were used to assess differences in categorical proportions. The distribution of Fast Lane Scores was examined using a histogram, as illustrated in Figure 4, with the number of queries in each priority category reported separately for each industry sector. Figure 4 provides a visual representation of the overall score distribution and the relative concentration of queries across the proposed priority categories.

Three sensitivity analyses were conducted to assess the robustness of the findings. First, the high-authority domain threshold used in Signal 3 was adjusted from the baseline of 1,000,000 estimated monthly organic visits to 500,000 and 5,000,000 visits. Second, the analysis was repeated after excluding queries that exceeded the Google Ads submission limits. Third, the previously excluded seed keyword was reintegrated into the dataset to evaluate its

influence on the overall results. A statistical significance level of $p < 0.05$ was applied throughout the analysis.

4. Results and Discussion

4.1. The Size of the Zero-Data Category.

Of the 6,993 queries in the dataset, 2,459 (35.2%) had no reported search volume, keyword difficulty (KD), or cost-per-click (CPC) data. An additional 4,473 queries had reported search volumes, while 61 had only KD values. Even among queries with measurable search volumes, the values were concentrated near the reporting threshold. The median monthly search volume was 10, the lowest value reported by Google Ads, and 62.9% of these queries had a monthly search volume of 20 or fewer. The proportion of zero-data queries varied substantially across seed keywords, as illustrated in Figure 3. Among the 27 seed keywords that generated at least 20 queries, the proportion ranged from 14.7% for "amazon fba" to 70.0% for "fractional cfo", with a median of 37.1%. Two narrowly defined seed keywords generated insufficient suggestions for meaningful comparison: "seo software for small business" (2 queries) and "content optimization tool" (10 queries). Significant differences were also observed across industry verticals ($\chi^2(2) = 244.66$, $p < 0.001$). Small-business consulting exhibited the highest proportion of zero-data queries at 46.5% (1,094 of 2,352), followed by SEO and digital marketing tools at 38.4% (379 of 987) and e-commerce at 27.0% (986 of 3,654). These findings indicate that the prevalence of zero-data queries varies considerably across both seed keywords and industry verticals.

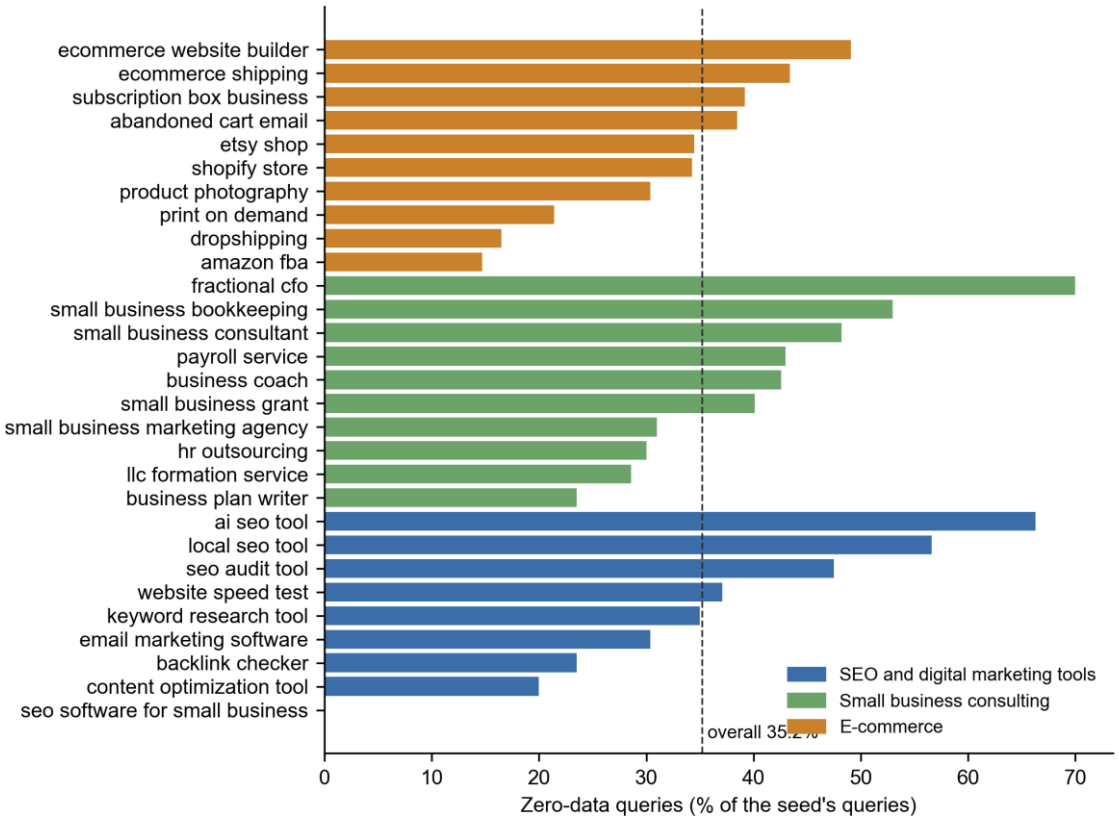


Figure 3. Share of Zero-Data Queries by Seed Keyword and Vertical. the Dashed Line Marks the Overall Share (35.2%).

The observed proportion of zero-data queries was more than twice the upper limit of the preliminary 5–15% estimate presented in an earlier version of this study (Section 3.1). This finding is consistent with previous research on long-tail search behavior, which indicates that most distinct search queries occur infrequently [4, 5]. Consequently, keyword research tools that report search volumes only above specific thresholds may fail to provide measurable data for a substantial proportion of expanded queries. As demonstrated in Figure 3, the zero-data category represents a substantial component of keyword datasets rather than an isolated occurrence. Applying search-volume-based filtering to a keyword list generated through a standard expansion procedure would exclude approximately one-third of the queries before further analysis. Moreover, the extent of this exclusion varies considerably across topics, potentially introducing systematic bias into keyword selection and subsequent analyses.

4.2. Zero-data Queries Compared with Measurable Queries.

Table 2 compares zero-data queries with measurable ones. Zero-data queries were slightly longer (mean 4.50 words against 4.14; Mann–Whitney $U = 6{,}407{,}236.5$, $p < 0.001$), and more often had eight or more words (2.7% against 1.0%), but the median was four words in both groups. They did not differ in intent modifiers: 13.3% of zero-data queries and 12.7% of measurable queries contained at least one modifier category ($\chi^2 = 0.48$, $p = 0.49$), and fewer than 1% of either group contained two or more.

Table 2. Zero-data and Measurable Queries Compared.

Characteristic	Zero-data (n = 2,459)	Measurable (n = 4,534)	Test
Mean words (median)	4.50 (4)	4.14 (4)	$U = 6{,}407{,}236.5$, $p < 0.001$
8 or more words	2.7%	1.0%	
At least one intent-modifier category	13.3%	12.7%	$\chi^2 = 0.48$, $p = 0.49$
Two or more categories	0.8%	0.7%	$\chi^2 = 0.15$, $p = 0.70$
SERP characteristics (top five)	*n = 2,459*	*random sample, n = 499*	
No high-authority domain	20.6%	20.6%	$\chi^2 = 0.00$, $p = 1.00$
Two or more user-generated results	25.7%	32.3%	$\chi^2 = 8.66$, $p = 0.003$
No title directly matching the query	72.8%	58.1%	$\chi^2 = 42.38$, $p < 0.001$
#1 result under 1,000 words	59.3%	60.3%	$\chi^2 = 0.15$, $p = 0.69$
Mean S3 score (0–3)	1.34	1.29	$U = 637{,}391.5$, $p = 0.14$

Note: SERP data for measurable queries come from a simple random sample of 500 (one returned no organic results). High-authority domain: estimated organic traffic of 1,000,000 or more visits a month.

The SERPs of the two groups were similar on authority and page depth. The clearest difference was the content gap: in 72.8% of zero-data SERPs no top-five title addressed the query directly, against 58.1% for measurable queries. User-generated results were common in both groups but slightly less common for zero-data queries. Reddit alone appeared in the top five of 40.3% of zero-data SERPs. These results qualify the premise that zero-data queries are a distinct, more specific kind of query. In this dataset they look much like measurable queries in wording, and Google ranks similar kinds of sites for both. What sets them apart is that search results answer them less directly, which is the gap a publisher could fill. This fits evidence that specific wording changes what search engines return [3] and that results for commercial queries are often a poor match for what users ask [36]. The prominence of Reddit and other forums is consistent with research showing that user-generated content holds a large place in current results [34, 35], although in this dataset it was not more prominent for zero-data queries.

4.3. Distribution of Fast Lane Scores and Priority Bands.

All 2,459 zero-data queries were scored. Fast Lane Scores ranged from 0 to 9.5, with a mean of 3.45 (SD 1.73) and a median of 3.25 (Figure 4). Only 2 queries (0.1%) reached Priority A, 234 (9.5%) were Priority B, 1,129 (45.9%) Priority C and 1,094 (44.5%) were deprioritised. The bands differed by vertical. Priority B made up 18.7% of the SEO-tools queries and 15.0% of the e-commerce queries, but only 1.4% of the consulting queries, where 60.4% were deprioritised. Most of that difference comes from Signal 4, which depends on the test publisher: mean S4 was 1.59 against semrush.com, 1.40 against shopify.com and 0.42 against revenuebase.ai.

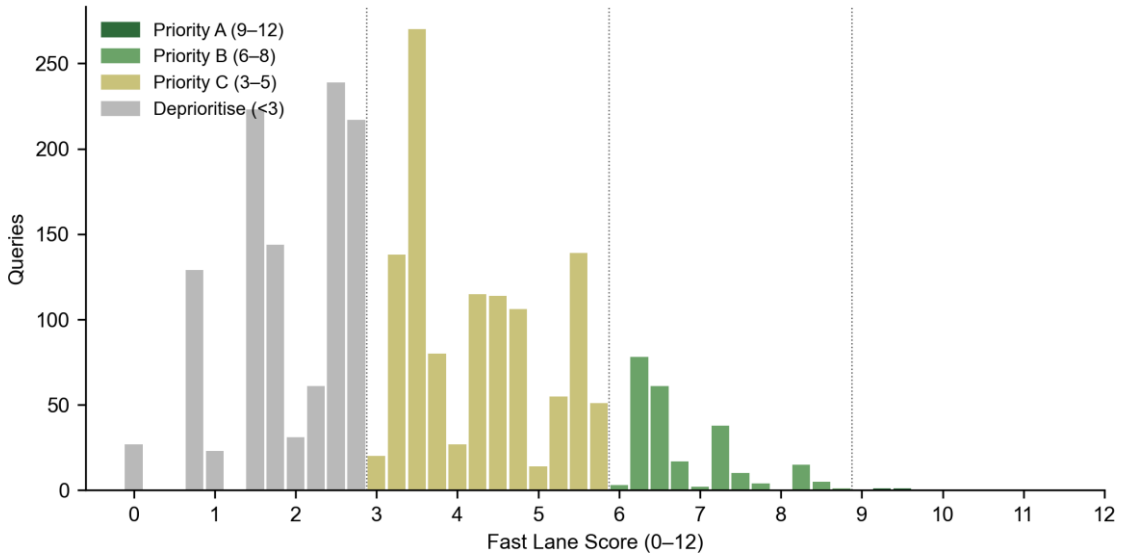


Figure 4. Distribution of Fast Lane Scores Across the 2,459 Zero-Data Queries, Coloured by Priority Band.

Signal by signal, the mean scores were 0.98 for S1, 0.14 for S2, 1.34 for S3 and 0.99 for S4. Signal 2 contributed least. Of the zero-data queries, 86.7% contained no intent modifier and none contained more than two. Moving up the bands, every signal rises: deprioritised queries averaged 0.61 on S1, 0.03 on S2, 1.09 on S3 and 0.19 on S4, while Priority B queries averaged 1.72, 0.46, 1.83 and 2.77. The bands therefore separate queries that are longer, face weaker SERPs and sit closer to the publisher's existing topics, as intended. Without Signal 4, only 6 queries would score 6 or more. The near-empty Priority A band is the main result of this section. The thresholds were set before data collection, on the assumption that the strongest opportunities combine several intent modifiers with a long, specific query. Queries of that kind exist (Section 4.4), but autocomplete rarely produces them: it completes short prefixes into short phrases, and only 0.8% of its zero-data suggestions carried two or more modifier categories. Under the fixed thresholds the bands are therefore conservative. In practice the framework is better read as a ranking within a list than as an absolute standard: the top decile of this dataset (Priority B and above) contains 236 queries a publisher could act on.

4.4. Worked examples.

The three queries used as examples in the earlier version were re-scored against live SERPs on 21 October 2026 with the corrected rubric (Table 3). Signal 4 was scored against the SEO-tools

test publisher, with "ubersuggest" as the seed. These queries were written by hand and are not part of the dataset.

Table 3. Re-scoring of the Three Example Queries (United States, 21 October 2026).

Query	Words	S1	S2 (categories)	S3 (conditions met)	S4	Score	Band
how does ubersuggest lifetime pricing compare to monthly seo tools	10	3	2 (comparative, price)	2.25 (UGC, no direct match, thin #1)	1	8.25	B
best seo tools under $50 per month for small businesses	10	3	3 (superlative, price, audience)	0.75 (no direct match)	1	7.75	B
is ubersuggest good enough to replace semrush for freelancers	9	3	2 (comparative, audience)	0.75 (no direct match)	0	5.75	C

S1, word count; S2, intent-modifier categories; S3, SERP weakness; S4, semantic parent match; UGC, two or more user-generated results in the top five.

None of the three reaches Priority A. The earlier version reported scores of 10 to 11.25 for these queries, all in Priority A; the difference comes from correcting the scoring errors the reviewers identified and from scoring Signals 3 and 4 against live data rather than by inspection. The examples illustrate how the rubric is applied. They also show that hand-written queries carry far more intent modifiers than autocomplete suggestions (two or three categories each, against a dataset mean of 0.14), which is the reason the Priority A band is nearly empty in Section 4.3.

4.5. Fast Lane Score Compared with Conventional Keyword Selection.

The conventional opportunity score assigned a value of zero to all 2,459 zero-data queries, as it must by construction, and a positive value to 4,430 of the 4,534 measurable queries. A volume-based workflow would therefore have discarded 35.2% of the dataset before ranking, including all 236 queries in Priority A and B. The Fast Lane Score gives those queries an order. It cannot say whether they carry more or less demand than measurable queries; it says which of them combine specificity, a weak SERP and a good fit with the publisher's existing content. Table 4 compares the approaches.

Table 4. The Fast Lane Methodology Compared with Conventional Keyword-Selection Methods.

Dimension	Volume-based selection	Volume × difficulty (opportunity score)	Fast Lane Methodology
Data required	Search volume	Volume, KD (± CPC, intent)	Query text, live SERP, publisher's ranking data
Dependence on search volume	Total	Total (score = 0 when volume = 0)	None
Competition assessment	None	Tool's KD estimate	Direct SERP inspection (authority, UGC share, content gap, depth)
Intent evaluation	None	Coarse weight	Modifier categories in the query
Applicability to SMBs	Low: favours head terms held by established domains	Moderate	High: targets queries large publishers do not track
Advantages	Simple; widely available	Balances demand and difficulty	Surfaces queries other methods cannot see; reproducible
Limitations	Ignores competition and the long tail	Cannot score zero-data queries	Demand is unmeasured by definition; SERP data needs collecting; weights not yet calibrated on outcomes

The band counts were stable when the authority threshold of Signal 3 was moved. At 500,000 visits a month the counts were A 2, B 225, C 1,127 and deprioritised 1,105; at 5,000,000 they were 2, 245, 1,145 and 1,067. Excluding the 33 queries that Google Ads could not accept changed the counts by at most 15. Including the excluded seed "rank tracker" raised the zero-data share to 35.5% (2,554 of 7,185) and, because the selection rule then chose a gaming-statistics site as the SEO-tools test publisher, lowered Priority B to 184 queries. That case shows that Signal 4 is only as valid as the choice of publisher.

4.6. Implications, Limitations and Future Research.

The long-tail account holds that rare items together make up a large share of demand [9, 10]. This study shows that a large share of the search tail is also invisible to the tools practitioners use: about a third of the queries a standard expansion produces, and up to 70% for some topics. Treating those queries as worthless repeats, in keyword research, the bias that comes from discarding zero observations in demand estimation [8]. The comparison in Section 4.2 adds that the missing queries are not, as a group, more specific than the measured ones. The measurement gap is better explained by thresholds in the tools than by anything unusual about the queries.

For SMBs with limited SEO resources [11, 12], the results suggest three things. First, a volume filter removes a large and topic-dependent part of the opportunity set, so its cost should be checked rather than assumed. Second, the most common weakness in zero-data SERPs is not low authority but the absence of a page that answers the query directly, which is a gap a small publisher can fill. Third, the Fast Lane Score is most useful as a way to rank a zero-data list; a publisher would start with the top decile rather than wait for queries that meet every threshold. How AI answers will change the value of these queries remains open [14, 16].

First, demand for zero-data queries is unmeasured by definition; the study shows that the queries exist (Google suggests them) and what their SERPs look like, not how much traffic they would bring. Second, the weights are equal by design and not yet calibrated on outcomes [45], and the band thresholds proved conservative for autocomplete data. Third, the data cover one market, one language and one day. Fourth, Signals 3 and 4 use proxies (estimated traffic for authority, title match for direct answers) that a manual review might score differently. Two of the three Signal 4 conditions depend only on the seed, so the signal varies mainly between seeds, and its values depend on the choice of test publisher. Fifth, autocomplete reflects Google's own filtering, so the zero-data share depends on the expansion method; the relevance rule also removed some relevant paraphrases, and one ambiguous seed was excluded after collection.

A longitudinal test is the natural next step: publish content for a sample of queries from each band and compare time to rank, impressions and clicks. That design would also allow the weights and thresholds to be estimated from outcomes rather than set in advance. A second step is to test the framework on query sets that autocomplete under-represents, such as the longer queries users type into conversational interfaces [18, 19].

5. Conclusions

A missing number in a keyword tool is a statement about the tool, not about the query. In this dataset, about a third of the queries Google itself suggested carried no volume, difficulty or CPC, and those queries were neither shorter nor less commercial than the ones that did. Any

method that ranks by volume discards them without examining them. The Fast Lane Score shows that they can be ranked by evidence that is available for every query, namely its wording, its live results and the publisher's existing rankings. The strict thresholds proposed for the top band were rarely met, so the framework's present value lies in ordering a list rather than in certifying single queries, and its claim to predict ranking outcomes still has to be tested. The opportunity it points to is largest for publishers that cannot compete for head terms, because the gap it measures is the one that volume-based competitors do not look at.

Acknowledgments

This research received no external funding. Data collection costs were paid by the author.

Author Contribution

Conceptualization, methodology, data collection, data analysis and writing: P.P.

Competing Interest

The author developed the original version of this framework while employed by NP Digital, which owns the keyword research tools AnswerThePublic and Ubersuggest, until July 2026. The validation study reported here was designed, self-funded and conducted independently after that employment ended, using a commercial data provider unaffiliated with NP Digital. The author has no current financial relationship with NP Digital and no financial interest in the commercial adoption of the methodology. An earlier version of this manuscript was posted as a preprint on SSRN (https://doi.org/10.2139/ssrn.6458202).

References

[1] Baye, M.R.; De los Santos, B.; Wildenbeest, M.R. (2016). Search engine optimization: What drives organic traffic to retail sites? *Journal of Economics & Management Strategy*, *25*, 6–31. https://doi.org/10.1111/jems.12141.

[2] Killoran, J.B. (2013). How to use search engine optimization techniques to increase website visibility. *IEEE Transactions on Professional Communication*, *56*, 50–66. https://doi.org/10.1109/TPC.2012.2237255.

[3] van Hoof, M.; Trilling, D.; Moeller, J.; Meppelink, C.S. (2024). It matters how you google it? Using agent-based testing to assess the impact of user choices in search queries and algorithmic personalization on political Google Search results. *Journal of Computer-Mediated Communication*, *29*, zmae020. https://doi.org/10.1093/jcmc/zmae020.

[4] Downey, D.; Dumais, S.; Horvitz, E. (2007). Heads and tails: Studies of web search with common and rare queries. In Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '07); pp. 847–848. https://doi.org/10.1145/1277741.1277939.

[5] Skobeltsyn, G.; Junqueira, F.; Plachouras, V.; Baeza-Yates, R. (2008). ResIn: A combination of results caching and index pruning for high-performance web search engines. In Proceedings of the 31st Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '08); pp. 131–138. https://doi.org/10.1145/1390334.1390359.

[6] Kekuda, A.; Zhang, Y.; Udayashankar, A. (2024). Embedding based retrieval for long tail search queries in ecommerce. In Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24); pp. 771–774. https://doi.org/10.1145/3640457.3688039.

[7] Peng, W.; Li, G.; Jiang, Y.; Wang, Z.; Ou, D.; Zeng, X.; Xu, D.; Xu, T.; Chen, E. (2024). Large language model based long-tail query rewriting in Taobao search. In Companion Proceedings of

the ACM Web Conference 2024 (WWW '24); pp. 20–28. https://doi.org/10.1145/3589335.3648298.

[8] Adam, H.; He, P.; Zheng, F. (2024). Machine learning for demand estimation in long tail markets. *Management Science*, *70*, 5040–5065. https://doi.org/10.1287/mnsc.2023.4893.

[9] Brynjolfsson, E.; Hu, Y.J.; Simester, D. (2011). Goodbye Pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science*, *57*, 1373–1386. https://doi.org/10.1287/mnsc.1110.1371.

[10] Anderson, C. (2006). The Long Tail: Why the Future of Business Is Selling Less of More; Hyperion: New York, NY, USA.

[11] Schultheiß, S.; Lewandowski, D. (2025). Becoming visible with limited resources: Non-profit journalists' perspectives on search engine optimization. *PLOS ONE*, *20*, e0322573. https://doi.org/10.1371/journal.pone.0322573.

[12] Singh, S.K.; Gupta, S.; Kumar, S.; Gupta, B.B.; Alhalabi, W.; Arya, V.; Zhang, J. (2024). A novel cumulative indicator score using indicator averaging for optimizing local business websites of enterprise systems. *Enterprise Information Systems*, *18*, 2301658. https://doi.org/10.1080/17517575.2023.2301658.

[13] Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, *17*, 99–120. https://doi.org/10.1177/014920639101700108.

[14] Grossman, R.; Liu, S.; Chen, M.K.; Smith, M.; Borcea, C.; Chen, Y. (2026). How generative AI disrupts search: An empirical study of Google Search, Gemini, and AI Overviews. In Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26); pp. 448–459. https://doi.org/10.1145/3805712.3809667.

[15] Allawati, S.; McKay, D.; Sanderson, M.; Thomas, P.; Trippas, J. (2026). An eye tracking study: Are AI Overviews changing search behavior? In Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26); pp. 29–40. https://doi.org/10.1145/3805712.3809530.

[16] Aggarwal, P.; Murahari, V.; Rajpurohit, T.; Kalyan, A.; Narasimhan, K.; Deshpande, A. (2024). GEO: Generative engine optimization. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '24); pp. 5–16. https://doi.org/10.1145/3637528.3671900.

[17] Strauss, I.; Yang, J.; O'Reilly, T.; Rosenblat, S.; Moure, I. (2026). The attribution crisis in LLM search results: Estimating ecosystem exploitation. *Data & Policy*, *8*, e15. https://doi.org/10.1017/dap.2026.10064.

[18] Trippas, J.R.; Al Lawati, S.F.D.; Mackenzie, J.; Gallagher, L. (2024). What do users really ask large language models? An initial log analysis of Google Bard interactions in the wild. In Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '24); pp. 2703–2707. https://doi.org/10.1145/3626772.3657914.

[19] Mo, F.; Mao, K.; Zhao, Z.; Qian, H.; Chen, H.; Cheng, Y.; Li, X.; Zhu, Y.; Dou, Z.; Nie, J.-Y. (2025). A survey of conversational search. *ACM Transactions on Information Systems*, *43*, 1–50. https://doi.org/10.1145/3759453.

[20] Ampler, N.; Lehmann-Zschunke, N.; Olbrich, R. (2025). How to design keywords in search engine advertising: A multi-group comparison based on the search volume of the product type. *Review of Marketing Science*, *23*, 173–206. https://doi.org/10.1515/roms-2024-0020.

[21] Jansen, B.J.; Spink, A. (2006). How are we searching the World Wide Web? A comparison of nine search engine transaction logs. *Information Processing & Management*, *42*, 248–263. https://doi.org/10.1016/j.ipm.2004.10.007.

[22] Sharma, N.; Liao, Q.V.; Xiao, Z. (2024). Generative echo chamber? Effect of LLM-powered search systems on diverse information seeking. In Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24); pp. 1–17. https://doi.org/10.1145/3613904.3642459.

[23] Lichtenegger, E.; Urman, A.; Hannak, A. (2026). A comparative study of users' information-seeking practices across search engines and generative AI chatbots. In Proceedings of the 49th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '26); pp. 1094–1105. https://doi.org/10.1145/3805712.3809525.

[24] Zhou, T.; Li, S. (2026). Understanding user switch of information seeking: From search engines to generative AI. *Journal of Librarianship and Information Science*, *58*, 696–708. https://doi.org/10.1177/09610006241244800.

[25] Broder, A. (2002). A taxonomy of web search. *ACM SIGIR Forum*, *36*, 3–10. https://doi.org/10.1145/792550.792552.

[26] Jansen, B.J.; Booth, D.L.; Spink, A. (2008). Determining the informational, navigational, and transactional intent of Web queries. *Information Processing & Management*, *44*, 1251–1266. https://doi.org/10.1016/j.ipm.2007.07.015.

[27] Liu, F.; Yu, C.; Meng, W. (2002). Personalized web search by mapping user queries to categories. In Proceedings of the Eleventh International Conference on Information and Knowledge Management (CIKM '02); pp. 558–565. https://doi.org/10.1145/584792.584884.

[28] Lichtenegger, E.; Urman, A.; Hannak, A. (2026). A new taxonomy of web search: A user-centered framework for search intent in the AI era. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26); pp. 1–16. https://doi.org/10.1145/3772318.3791050.

[29] Shah, C.; White, R.; Andersen, R.; Buscher, G.; Counts, S.; Das, S.; Montazer, A.; Manivannan, S.; Neville, J.; Rangan, N.; Safavi, T.; Suri, S.; Wan, M.; Wang, L.; Yang, L. (2025). Using large language models to generate, validate, and apply user intent taxonomies. *ACM Transactions on the Web*, *19*, 1–29. https://doi.org/10.1145/3732294.

[30] Alexander, D.; de Vries, A.P. (2025). In a few words: Comparing weak supervision and LLMs for short query intent classification. In Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25); pp. 2977–2981. https://doi.org/10.1145/3726302.3730213.

[31] Taieb-Maimon, M.; Harush, H. (2025). Web search engine results page viewing formats for different search tasks. *International Journal of Human–Computer Interaction*, *41*, 6350–6365. https://doi.org/10.1080/10447318.2024.2376358.

[32] Hu, D.; Baumann, J.; Urman, A.; Lichtenegger, E.; Forsberg, R.; Hannák, A.; Wilson, C. (2026). Auditing Google's AI Overviews and featured snippets: A case study on baby care and pregnancy. *Proceedings of the International AAAI Conference on Web and Social Media*, *20*, 1044–1062. https://doi.org/10.1609/icwsm.v20i1.42681.

[33] Donabauer, P.; Elsweiler, D. (2025). Exploring the impact of warnings on user perception towards AI-generated content in search results. In Proceedings of the 34th ACM International Conference on Information and Knowledge Management (CIKM '25); pp. 605–614. https://doi.org/10.1145/3746252.3761089.

[34] Sünkler, S.; Lewandowski, D.; Schultheiß, S.; Koop, O. (2026). Information diversity and authority shifts in Google's AI Overviews: An audit of health queries. In Companion Publication of the 18th ACM Web Science Conference (WebSci Companion '26); pp. 24–28. https://doi.org/10.1145/3795513.3807429.

[35] Burtch, G.; Lee, D.; Chen, Z. (2024). The consequences of generative AI for online knowledge communities. *Scientific Reports*, *14*, 10413. https://doi.org/10.1038/s41598-024-61221-0.

[36] Bevendorff, J.; Wiegmann, M.; Potthast, M.; Stein, B. (2024). Is Google getting worse? A longitudinal investigation of SEO spam in search engines. In Advances in Information Retrieval (ECIR 2024); Lecture Notes in Computer Science; Springer: Cham, Switzerland; pp. 56–71. https://doi.org/10.1007/978-3-031-56063-7_4.

[37] Sezer Çırakoğlu, F.; Koşaner, Ö. (2024). Linguistic insights into high-quality content for SEO: Optimizing high-quality content with Grice's conversational maxims. *Telematics and Informatics Reports*, *16*, 100169. https://doi.org/10.1016/j.teler.2024.100169.

[38] Zhao, W.X.; Liu, J.; Ren, R.; Wen, J.-R. (2024). Dense text retrieval based on pretrained language models: A survey. *ACM Transactions on Information Systems*, *42*, 1–60. https://doi.org/10.1145/3637870.

[39] Zhu, Y.; Yuan, H.; Wang, S.; Liu, J.; Liu, W.; Deng, C.; Chen, H.; Liu, Z.; Dou, Z.; Wen, J.-R. (2026). Large language models for information retrieval: A survey. *ACM Transactions on Information Systems*, *44*, 1–54. https://doi.org/10.1145/3748304.

[40] Greco, S.; Ishizaka, A.; Tasiou, M.; Torrisi, G. (2019). On the methodological framework of composite indices: A review of the issues of weighting, aggregation, and robustness. *Social Indicators Research*, *141*, 61–94. https://doi.org/10.1007/s11205-017-1832-9.

[41] OECD; European Commission Joint Research Centre. (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide; OECD Publishing: Paris, France. https://doi.org/10.1787/9789264043466-en.

[42] Pereira Libório, M.; Ekel, P.I.; Laudares, S.; Martins, C.A.P.S. (2024). A goal-based weighting for composite indicators constructed through Ordered Weighted Averaging (OWA) operator. *MethodsX*, *12*, 102575. https://doi.org/10.1016/j.mex.2024.102575.

[43] Dawes, R.M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist*, *34*, 571–582. https://doi.org/10.1037/0003-066X.34.7.571.

[44] Wainer, H. (1976). Estimating coefficients in linear models: It don't make no nevermind. *Psychological Bulletin*, *83*, 213–217. https://doi.org/10.1037/0033-2909.83.2.213.

[45] Hair, J.F.; Sharma, P.N.; Sarstedt, M.; Ringle, C.M.; Liengaard, B.D. (2024). The shortcomings of equal weights estimation and the composite equivalence index in PLS-SEM. *European Journal of Marketing*, *58*, 30–55. https://doi.org/10.1108/EJM-04-2023-0307.



© 2026 by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/).