BÖLÜM 14

FROM SEO TO ANSWER ENGINE OPTIMIZATION (AEO): GENERATIVE AI AND THE TRANSFORMATION OF SEARCH VISIBILITY

1. HİDAYET KARAMUK[1]

INTRODUCTION

In digital marketing, findability and visibility have long been strategically important because search often constitutes one of the earliest touchpoints in the consumer journey. The classic SEO literature shows that organic visibility is shaped not only by technical optimization but also by market demand and competitive dynamics, and therefore can be conceptualized as a marketing-relevant strategic problem rather than a purely technical one (Baye et al., 2016). Similarly, evidence that keyword choice functions as a long-term strategic investment—whose returns can accumulate over time—has further positioned SEO within broader marketing strategy and resource allocation decisions (Erdmann et al., 2022). Yet, as of 2025, the structure of search engine results pages (SERPs) and users' information-acquisition behaviors are undergoing a shift that

[1] Res. Asst., Kocaeli University, School of Business, Orcid: 0000-0002-7654-9337

directly challenges several assumptions embedded in this classic framework.

A central dimension of this shift is the increasing prominence of direct answers and rich SERP components. Featured snippets, knowledge panels, and direct-answer boxes can satisfy information needs without requiring users to click through to external pages, while also making content selection mechanisms more complex than rank alone (Strzelecki & Rutecka, 2020). In parallel, empirical work suggests that SERP features can substantially modulate organic click-through rates (CTR), implying that performance evaluation based solely on CTR may be insufficient in environments where visibility is partially realized on-SERP (Fubel et al., 2023). This limitation becomes even more pronounced in AI-summary contexts: analysis of March 2025 browsing data indicates that external link clicking declines when an AI summary is present, reinforcing the relevance of "zero-click" dynamics for marketing visibility and attribution (Pew Research Center, 2025).

Generative AI–enabled search experiences amplify this transformation by introducing an explicit "answer layer" that synthesizes information and guides discovery. Google's documentation on AI Overviews/AI Mode describes how responses may be supported by multiple related searches across subtopics ("query fan-out"), and notes that appearances in these AI features are included in Search Console reporting under the "Web" search type; importantly, it also states that no special optimization is required beyond foundational SEO practices (Google, n.d.-a). Thus, the core issue is not the disappearance of SEO, but the expansion of what "visibility" means—toward outcomes such as being included in the answer layer, being selected as a source, and being cited or summarized.

Within this context, Answer Engine Optimization (AEO) invites scholars and practitioners to rethink visibility management not merely as competition for rank, but as participation in the source-selection, synthesis, and verification practices of answer-producing systems. Information retrieval scholarship has long recognized that direct answers can improve convenience while potentially reducing user diligence and verification—captured as a "direct answer dilemma" (Potthast et al., 2021). More broadly, the characterization of search as moving "from finding to understanding" (i.e., exploratory search) provides a foundational lens for interpreting why answer-centered systems may reshape how users learn, compare, and decide (Marchionini, 2006). Recent work on conversational search further supports this perspective by emphasizing multi-turn interaction, context retention, and iterative clarification as defining features of contemporary information seeking (Mo et al., 2025). On the technical side, retrieval-augmented generation (RAG) has formalized architectures that aim to ground generated responses in retrieved evidence, underscoring the growing importance of provenance and sourceability—an emphasis that aligns directly with AEO's concern for being a trusted input to answer synthesis (Lewis et al., 2020).

Accordingly, this book chapter has three objectives: (i) to conceptualize the shift from SEO to AEO within the broader search–discovery–visibility literature, (ii) to examine how answer-centered architectures and changes in user information acquisition reshape visibility strategies, and (iii) to systematize AEO's core components alongside performance measurement approaches—while also addressing key limitations and risk areas. In doing so, the chapter also draws on search-engine guidance that stresses quality, accuracy, and contextual transparency, particularly in relation to scaled generative content production (Google, n.d.-e), as a relevant governance backdrop for AEO-oriented practice.

RESTRUCTURING THE SEARCH ENVIRONMENT WITH GENERATIVE ARTIFICIAL INTELLIGENCE

The integration of generative artificial intelligence (GenAI) into search is transforming search from a system primarily designed to retrieve webpages through indexing and ranking into an information tool that increasingly provides answer generation and manages research workflows. The technical foundation of this shift lies in the Transformer architecture, which enabled large language models (LLMs) to be trained at scale (Vaswani et al., 2017). Transformer-based models expanded search interaction beyond a keyword-matching logic by delivering contextual language representations, capturing long-range dependencies, and providing a general-purpose generation capability reusable across a wide range of tasks. In this sense, they represent a critical breakpoint that reconfigures how users formulate queries and how systems respond.

A second driver accelerating the value of LLMs in search has been advances in instruction following and alignment with user intent. Scaling-oriented models such as GPT-3 demonstrated strong few-shot learning performance, opening broad application space for search-adjacent tasks including question answering, summarization, and text generation (Brown et al., 2020). Subsequently, alignment approaches using human feedback (e.g., reinforcement learning from human feedback, RLHF) enabled models to respond more consistently to user instructions, making the query → direct answer flow more viable in real-world settings (Ouyang et al., 2022). Taken together, these developments made it technically more feasible for search engines not only to present "the best links," but also to interpret user goals and synthesize information into answer-form outputs.

However, generation capability alone is insufficient in a domain like search, where expectations for accuracy and timeliness are high. For this reason, search systems have increasingly moved toward architectures that connect generative models to external information sources. Retrieval-Augmented Generation (RAG) aims to "ground" the model's output in content retrieved from an external document store, supporting more accurate and more traceable outputs through provenance (Lewis et al., 2020). Within the search ecosystem, such architectures have become a foundational element of "answer engine" designs because they address both the need for up-to-date information and the need to indicate what sources an answer is based on.

Industry implementations also make this restructuring visible at the interface level. Google has framed this shift through generative AI–enabled search experiences that prioritize synthesis and answer-like summaries, highlighting changes in how users discover information (Reid, 2024). Similarly, Microsoft has introduced generative search experiences in Bing and later Copilot Search, where generative summaries/answers appear in an upper layer of the search results while providing cited links and traditional results for verification and deeper exploration (Bing, 2024, 2025).

Overall, GenAI is restructuring the search environment at three levels. First, the interaction level: user queries become longer, more contextual, and more goal-oriented, while systems shift toward supporting multi-step research, planning, and sensemaking workflows. Second, the output-format level: the experience moves from a "link list" orientation toward a composite of "answer + supporting sources + discovery suggestions." Third, the ecosystem level: the generative answer layer reorganizes marketing-relevant dynamics such as publisher traffic flows, advertising placements, and competition for visibility. This shift makes it necessary to re-examine core assumptions in the SEO literature through the lens of AEO. The next section therefore elaborates the conceptual framing of this transformation and its relationship to the search–discovery–visibility literature.

CONCEPTUAL SHIFT FROM SEO TO ANSWER ENGINE OPTIMIZATION (AEO) IN THE SEARCH, DISCOVERY, AND VISIBILITY LITERATURE

In the digital marketing literature, search engines have traditionally been positioned as a high-intent marketplace in which visibility is translated into attention and, ultimately, measurable traffic outcomes. Early SEO research demonstrates that organic traffic is not merely a function of technical compliance, but also reflects market demand and competition, indicating that SEO is inseparable from broader marketing strategy and competitive positioning (Baye et al., 2016). Likewise, the strategic value of keyword selection has been framed as a long-term investment decision, reinforcing that visibility building is an ongoing resource allocation problem rather than a one-off optimization task (Erdmann et al., 2022).

However, search systems have been undergoing two interrelated transformations that complicate the classic "rank → click" logic. The first concerns the SERP's interface evolution through direct answers and rich results. Direct-answer boxes and featured snippets can satisfy informational needs on the SERP, and the mechanisms of selection into these formats are not always reducible to rank position alone (Strzelecki & Rutecka, 2020). Correspondingly, large-scale evidence indicates that SERP features can substantially modulate organic CTR, implying that traditional performance indicators may understate informational impact when attention is absorbed by on-SERP components (Fubel et al., 2023). This dynamic becomes more salient in AI-summary environments: browsing-data analysis suggests that users are less likely to click external links when an AI summary appears, supporting the argument that "zero-click" satisfaction can meaningfully reshape visibility outcomes (Pew Research Center, 2025).

The second transformation concerns how discovery and information acquisition are increasingly mediated by answer-like synthesis rather than document lists. This introduces a key tension: while direct answers and summary layers improve convenience, they may also reduce diligence and verification behavior—captured in the information retrieval literature as the "direct answer dilemma" (Potthast et al., 2021). Related work suggests that answer-like components can influence credibility judgments and perceptions of authority, further indicating that visibility is not only about being seen but also about being trusted and accepted as a source within the answer layer (Bink et al., 2022; Strzelecki & Rutecka, 2020).

In this context, AEO can be understood as a conceptual extension of SEO that shifts the focal outcome from ranking-based clicks to answer-layer inclusion, source selection, and citation-worthiness. Google's documentation for AI features in Search positions these experiences as an additional surface of visibility, indicating that they are reported in Search Console under the "Web" search type and that no special optimization is required beyond foundational SEO principles (Google, n.d.-a). At the same time, Google's guidance on generative AI content stresses that scaled content creation must add value and maintain quality, underscoring that visibility in AI-mediated environments is tied to accuracy, usefulness, and contextual transparency rather than volume alone (Google, n.d.-e). These ideas also align with quality evaluation discussions that emphasize experience, expertise, authoritativeness, and trust as a lens for assessing content credibility in search contexts (Google, 2022; Google, 2025).

Taken together, these developments provide a foundation for explaining why the SEO-to-AEO transition is not merely terminological. It reflects a structural shift in how visibility is produced and consumed: from "ranked list exposure leading to clicks" toward "answer-layer presence, source selection, and trust-mediated influence," where persuasion and decision-making can increasingly occur before any click happens (Pew Research Center, 2025; Potthast et al., 2021).

ANSWER-CENTERED SEARCH ARCHITECTURES AND THE TRANSFORMATION OF USER INFORMATION ACQUISITION

With generative AI, search is shifting from "listing relevant documents" toward producing answers while linking to evidence. At the architectural core of this shift is retrieval-augmented generation (RAG), which combines a retrieval component (fetching relevant passages or documents) with a generator that composes an answer grounded—at least in principle—in retrieved evidence rather than relying solely on parametric memory (Lewis et al., 2020). This move is significant because it reframes search output as a synthesized response and positions sourceability and provenance as central requirements for credibility and user decision-making.

In practice, answer-centered systems add an answer layer on top of the traditional results page. Google's documentation on AI features in Search indicates that these experiences may surface relevant links and can rely on multi-step processes such as "query fan-out," issuing multiple related searches across subtopics to support response construction (Google, n.d.-a). This reflects a design shift from "single query → ranked list" to "multi-subquery exploration → integrated answer," which changes how users encounter information and how content is selected for inclusion in the search interface.

Similar dynamics are visible in Microsoft's search ecosystem. Bing's generative search and later Copilot Search emphasize blended experiences where a synthesized response is presented together with links and traditional results to support verification and deeper exploration (Bing, 2024, 2025). This interface logic suggests that visibility increasingly depends not only on ranking but on whether content is selected as an input to synthesis and displayed as a supporting source within the answer layer.

These architectural changes also reshape user information acquisition. In the information retrieval literature, exploratory search is characterized as moving “from finding to understanding,” involving iterative querying, comparison, and sensemaking (Marchionini, 2006). Answer-centered systems amplify this trajectory by promoting multi-turn interaction: users ask follow-up questions, refine intent conversationally, and expect systems to maintain context and produce structured syntheses. Work on conversational search and conversational information seeking similarly foregrounds clarification, reformulation, and iterative information gain across turns as defining features of contemporary search behavior (Mo et al., 2025; Schneider et al., 2024; Zamani et al., 2022).

From a marketing and visibility perspective, one major implication is the potential weakening of click-based discovery as more consumption happens on the SERP itself. Using March 2025 browsing data, Pew Research Center reports that when an AI summary appears, users are less likely to click external links, supporting the argument that “zero-click” satisfaction may become more prevalent in AI-mediated search contexts (Pew Research Center, 2025). In such environments, success is increasingly tied to being represented within the answer layer (e.g., through citation, mention, or selection as a supporting link) rather than only to being clicked.

Finally, answer-centered architectures sharpen long-standing concerns about the risks of direct answers. While these systems can increase convenience, they may also reduce user diligence and verification, creating a tension between speed and epistemic caution. This has been discussed as a “direct answer dilemma” in the information retrieval literature (Potthast et al., 2021). Accordingly, answer-centered search should be understood not merely as an interface update but as a socio-technical transformation that reshapes how users interpret authority, credibility, and evidentiary signals, with direct consequences for AEO-oriented visibility strategies.

ANSWER ENGINE OPTIMIZATION (AEO) CORE COMPONENTS: CONTENT ARCHITECTURE, STRUCTURED MARKUP, AND AUTHORITY SIGNALS

In answer-centered search experiences (e.g., generative summaries and answer layers), visibility can no longer be reduced to "ranking highly." Instead, AEO can be framed as increasing the likelihood that content becomes selectable, quotable, and citable within the answer layer. Conceptually, AEO can be organized around three components: (i) content architecture, (ii) structured markup, and (iii) authority signals.

Content architecture (answerability & sourceability). AEO-oriented content should be structured to match user intent and to provide a clear, bounded response that can be extracted at the passage level. Google's "helpful, reliable, people-first content" guidance emphasizes that content should primarily benefit people rather than be created to manipulate ranking systems (Google, n.d.-b). Within this framing, content architecture typically prioritizes (a) explicit question/goal signaling via headings, (b) an early concise answer, and (c) follow-up justification and evidence with transparent scope (definitions, assumptions, and exceptions). The goal is to make content modular and traceable—qualities that support passage selection, summarization, and citation behaviors in answer-layer systems.

Structured markup (structured data & semantics). Answer engines benefit from machine-readable semantic cues. Google notes that structured data can help systems understand page content and provides general guidelines and policies for structured data use (Google, n.d.-d). At a practical level, Google's introduction to structured data markup further clarifies how structured data is implemented and interpreted in Search (Google, n.d.-f). For AEO, structured markup matters because it can clarify content type and relationships (e.g., article, organization, author, FAQ), supporting more reliable extraction and presentation. In Q&A contexts, Google's documentation distinguishes appropriate use of FAQPage versus QAPage, noting that QAPage is intended for pages where users can submit answers while FAQPage is designed for publisher-authored question–answer sets (Google, n.d.-c; Google, n.d.-g). These practices align with schema.org vocabularies that define the FAQPage and QAPage types as standardized semantic representations (Schema.org, n.d.-a; Schema.org, n.d.-b).

Authority signals (credibility, E-E-A-T, and transparency). Answer-layer inclusion is also shaped by whether content is treated as a credible source. Google has highlighted the expansion from E-A-T to E-E-A-T by adding "Experience" as an evaluative lens, reinforcing the role of credibility cues in quality assessment (Google, 2022). Google's search quality evaluation documentation further signals that quality judgments consider factors related to trust and reputation, even though the rater guidelines themselves are not a direct description of ranking algorithms (Google, 2023, 2025). Accordingly, AEO-relevant authority signals include verifiable authorship and ownership, editorial policies, update and correction practices, and explicit citations—elements that support user-facing verifiability and trust in answer-mediated environments.

Taken together, AEO's core components can be summarized as answerable content architecture, machine-readable semantic markup, and institutionalized credibility/authority signaling. This triad also sets the conceptual basis for the measurement and risk discussion in the subsequent section.

MEASURING AEO PERFORMANCE: METRICS, LIMITATIONS, AND RISK AREAS

In answer-centered search, evaluating AEO solely through rankings and organic click-through rate (CTR) is increasingly insufficient. A key reason is the growing prevalence of on-SERP satisfaction, where users consume answers in the answer layer without clicking out. Using March 2025 browsing data, Pew Research Center reports that traditional-result clicks occurred in 8% of visits when an AI summary was present versus 15% when it was absent; clicks on links inside the AI summary occurred in about 1% of such visits (Pew Research Center, 2025). This suggests that a substantial portion of "visibility outcomes" may materialize before any click.

A Multi-Layer Measurement Framework

AEO performance can be operationalized more robustly by combining multiple layers of metrics rather than relying on a single indicator set. A first layer concerns exposure, capturing whether content is included in the answer layer (e.g., as a supporting link within AI Overviews/AI Mode) and how such appearances are reflected in reporting. Google notes that AI-feature appearances are included in Search Console reporting under the "Web" search type and that no special optimizations are required beyond foundational SEO practices (Google, n.d.-a). Exposure should also account for visibility in SERP features such as direct answers/featured snippets, since being selected into these formats can reflect a mechanism that is not reducible to classic ranking alone (Strzelecki & Rutecka, 2020).

A second layer focuses on engagement, especially forms of interaction that occur without clicks. In answer-centered search, SERP dwell, follow-up queries, and query chains become informative signals of how users progress through information acquisition. In this context, Google's description of "query fan-out"—issuing multiple related searches across subtopics to develop a response—suggests that engagement may unfold as a process rather than a single click event (Google, n.d.-a). Relatedly, CTR requires careful interpretation: because SERP features can substantially modulate organic CTR, a decline in CTR does not necessarily imply a proportional decline in informational impact (Fubel et al., 2023).

A third layer captures value outcomes. Answer layers may reduce outbound clicks while still accelerating decisions and shifting conversions to other touchpoints. Accordingly, AEO evaluation should combine Search Console with analytics and track multi-channel outcomes such as assisted conversions, brand search lift, and direct traffic, rather than treating organic sessions as the sole indicator of performance (Google, n.d.-a).

A fourth layer concerns trust and quality. In interfaces such as Copilot Search, prominent citations and verification affordances elevate "being cited" or selected as a source into a meaningful outcome, distinct from being clicked (Bing, 2025). At the same time, the answer layer can amplify quality risks—wrong context, outdated claims, or incorrect brand associations—requiring systematic auditing and governance. Guidance on using generative AI content emphasizes the importance of adding value, maintaining quality, and providing appropriate transparency, which is directly relevant for monitoring and mitigating misattribution risks (Google, n.d.-e).

These layers also foreground key limitations and risk areas. As zero-click behavior increases, classic KPIs (sessions, CTR) may under-capture informational influence and brand impact (Pew Research Center, 2025). Direct answers can improve convenience while potentially reducing diligence and verification, aligning with the "direct answer dilemma" discussed in the information retrieval literature (Potthast et al., 2021). Finally, policy and scaled-generation risks become more salient when content is mass-produced without adding value, underscoring the need for accuracy, quality control, and contextual transparency as part of AEO governance (Google, n.d.-e).

REFERENCES

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*(1), 6–31. https://doi.org/10.1111/jems.12141

Bing. (2024, July 24). *Introducing Bing generative search*. Bing Search Blog. Retrieved September 9, 2025, from https://blogs.bing.com/search/July-2024/generativesearch

Bing. (2025, April 4). *Introducing Copilot Search in Bing.* Bing Search Blog. Retrieved September 10, 2025, from https://blogs.bing.com/search/April-2025/Introducing-Copilot-Search-in-Bing

Bink, M., Zimmerman, S., & Elsweiler, D. (2022). Featured snippets and their influence on users' credibility judgements. In *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval (CHIIR '22)* (pp. 113–122). Association for Computing Machinery. https://doi.org/10.1145/3498366.3505766

Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., ... Amodei, D. (2020). *Language models are few-shot learners*. arXiv. https://doi.org/10.48550/arXiv.2005.14165

Erdmann, A., Arilla, R., & Ponzoa, J. M. (2022). Search engine optimization: The long-term strategy of keyword choice. *Journal of Business Research*, *144*, 650–662. https://doi.org/10.1016/j.jbusres.2022.01.065

Fubel, E., Groll, N. M., Gundlach, P., Han, Q., & Kaiser, M. (2023). *Beyond rankings: Exploring the impact of SERP features on organic click-through rates*. arXiv. https://doi.org/10.48550/arXiv.2306.01785

Google. (2022, December 15). *Our latest update to the quality rater guidelines: E-A-T gets an extra E for experience*. Google Search Central. Retrieved September 12, 2025, from https://developers.google.com/search/blog/2022/12/google-raters-guidelines-e-e-a-t

Google. (2023, November). *Search quality rater guidelines: An overview* [PDF]. Google. Retrieved September 13, 2025, from https://services.google.com/fh/files/misc/hsw-sqrg.pdf

Google. (2025, September 11). *Search quality evaluator guidelines* [PDF]. Google. Retrieved September 13, 2025, from https://guidelines.raterhub.com/searchqualityevaluatorguidelines.pdf

Google. (n.d.-a). *AI features and your website*. Google Search Central. Retrieved September 11, 2025, from https://developers.google.com/search/docs/appearance/ai-features

Google. (n.d.-b). *Creating helpful, reliable, people-first content.* Google Search Central. Retrieved September 19, 2025, from https://developers.google.com/search/docs/fundamentals/creating-helpful-content

Google. (n.d.-c). *FAQ (FAQPage) structured data*. Google Search Central. Retrieved September 19, 2025, from https://developers.google.com/search/docs/appearance/structured-data/faqpage

Google. (n.d.-d). *General structured data guidelines*. Google Search Central. Retrieved September 19, 2025, from https://developers.google.com/search/docs/appearance/structured-data/sd-policies

Google. (n.d.-e). *Google Search's guidance on using generative AI content on your website*. Google Search Central. Retrieved September 19, 2025, from https://developers.google.com/search/docs/fundamentals/using-gen-ai-content

Google. (n.d.-f). *Introduction to structured data markup in Google Search.* Google Search Central. Retrieved September 14, 2025, from https://developers.google.com/search/docs/appearance/structured-data/intro-structured-data

Google. (n.d.-g). *Q&A (QAPage) structured data*. Google Search Central. Retrieved September 13, 2025, from https://developers.google.com/search/docs/appearance/structured-data/qapage

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems, 33*. https://arxiv.org/abs/2005.11401

Marchionini, G. (2006). Exploratory search: From finding to understanding. *Communications of the ACM, 49*(4), 41–46. https://doi.org/10.1145/1121949.1121979

Mo, F., Mao, K., Zhao, Z., Qian, H., Chen, H., Cheng, Y., ... & Nie, J. Y. (2025). A survey of conversational search. *ACM*

Transactions on Information Systems, *43*(6), 1-50. https://doi.org/10.1145/3759453

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., & Christiano, P. (2022). *Training language models to follow instructions with human feedback* (arXiv:2203.02155). arXiv. https://doi.org/10.48550/arXiv.2203.02155

Pew Research Center. (2025, July 22). *Google users are less likely to click on links when an AI summary appears in the results*. Retrieved September 13, 2025, from https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/

Potthast, M., Hagen, M., & Stein, B. (2021). The dilemma of the direct answer. *ACM SIGIR Forum, 54*(1), Article 14, 1–12. https://doi.org/10.1145/3451964.3451978

Reid, E. (2024, May 14). *Generative AI in Search: Let Google do the searching for you*. Google. Retrieved September 15, 2025, from https://blog.google/products/search/generative-ai-search/

Schema.org. (n.d.-a). *FAQPage.* Retrieved September 11, 2025, from https://schema.org/FAQPage

Schema.org. (n.d.-b). *QAPage*. Retrieved September 15, 2025, from https://schema.org/QAPage

Schneider, P., Poelman, W., Rovatsos, M., & Matthes, F. (2024). Engineering conversational search systems: A review of applications, architectures, and functional components. In *Proceedings of the 6th Workshop on NLP for Conversational AI (NLP4ConvAI 2024)* (pp. 73–88). Association for Computational Linguistics. https://aclanthology.org/2024.nlp4convai-1.5.pdf

Strzelecki, A., & Rutecka, P. (2020). Featured snippets results in Google web search: An exploratory study. In Á. Rocha, J. L. Reis, M. K. Peter, & Z. Bogdanović (Eds.), *Marketing and smart technologies: Proceedings of ICMarkTech 2019* (pp. 9–18). Springer. https://doi.org/10.1007/978-981-15-1564-4_2

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). *Attention is all you need* (arXiv:1706.03762). arXiv. https://doi.org/10.48550/arXiv.1706.03762

Zamani, H., Trippas, J. R., Dalton, J., & Radlinski, F. (2022). *Conversational information seeking* (arXiv:2201.08808). arXiv. https://doi.org/10.48550/arXiv.2201.08808