Search Whisperer: AI-Augmented Query Reformulation for Enhanced Information Sensemaking

[AUTHOR NAME], [Institution], [Country]

Search query formulation represents a critical yet often overlooked sensemaking challenge. Users frequently struggle to translate their information needs into effective queries, resulting in suboptimal search results and extended sensemaking cycles. We present *Search Whisperer*, an AI-powered tool that assists users in reformulating and expanding their search queries through intelligent keyword suggestions, domain-specific terminology recommendations, and automated Google dork generation. The system employs large language models to analyze query intent, detect query language, and provide contextually appropriate suggestions that preserve the specificity of detailed queries. We describe the system's design, its multi-language capabilities, and preliminary observations on how AI-assisted query reformulation affects the sensemaking process. Our work contributes to the growing discourse on human-AI collaboration in information-seeking tasks and raises questions about the changing nature of search literacy in an AI-augmented world.

Additional Key Words and Phrases: sensemaking, search, query reformulation, AI assistance, large language models, information seeking

ACM Reference Format:
[Author Name]. 2026. Search Whisperer: AI-Augmented Query Reformulation for Enhanced Information Sensemaking. 1, 1 (January 2026), 6 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

Sensemaking—the process of creating understanding from complex information spaces—fundamentally depends on effective information retrieval [4, 6]. Yet the gap between what users *want* to find and what they *query* for remains a persistent challenge in human-computer interaction. This "query formulation problem" is particularly acute when users venture into unfamiliar domains, where they may lack the vocabulary to express their information needs effectively.

Consider a user searching for information about infant mortality rates in a specific neighborhood. They might query "amsterdam west high baby deaths" when domain experts would search for "kraamsterfte epidemiologie amsterdam-west" (Dutch) or include terms like "perinatal mortality etiology socioeconomic factors." This vocabulary mismatch exemplifies what Belkin termed the "Anomalous State of Knowledge"—users cannot articulate what they don't know [2].

Recent advances in large language models (LLMs) present new opportunities to bridge this gap. We present *Search Whisperer*, an AI-powered query reformulation tool that assists users in the critical early stages of sensemaking. Rather than replacing human judgment, the system augments the user's query formulation capabilities by suggesting domain-appropriate terminology, generating advanced search operators, and adapting to the user's language—both literally and figuratively.

Author's address: [Author Name], [Institution], [City], [Country], [email].

2 RELATED WORK

2.1 Sensemaking and Search

Pirolli and Card's sensemaking model [4] describes information seeking as an iterative process of foraging and synthesis. Query formulation sits at the boundary between these phases: users must synthesize their current understanding to formulate queries, then forage for new information to refine that understanding. Russell et al. [6] demonstrated that the "cost structure" of sensemaking activities significantly influences behavior—when query reformulation is costly (in cognitive effort), users may settle for suboptimal results.

2.2 Query Expansion and Reformulation

Traditional query expansion techniques include relevance feedback [5], thesaurus-based expansion, and pseudo-relevance feedback. More recently, neural approaches have shown promise in understanding query intent [3]. However, these systems typically operate invisibly, expanding queries without user awareness or control. Our approach differs by making suggestions explicit, supporting user learning and agency.

2.3 AI-Assisted Information Seeking

The emergence of conversational AI has transformed information seeking, with systems like Bing Chat and Google's AI Overview providing direct answers. However, for complex sensemaking tasks, users often need to explore multiple sources critically rather than receive synthesized answers [7]. Search Whisperer occupies a middle ground: using AI to *prepare* users for traditional search rather than replacing it.

3 SYSTEM DESIGN

3.1 Architecture Overview

Search Whisperer is a web-based application that accepts natural language queries and provides multi-faceted assistance for query reformulation. The system is built on a Python/FastAPI backend with multiple LLM provider integrations (Claude, GPT-4, Groq/Llama) and a responsive JavaScript frontend. Figure 1 illustrates the system architecture.

3.2 Core Features

3.2.1 Query Quality Analysis. The system analyzes incoming queries across multiple dimensions: specificity, clarity, searchability, and domain appropriateness. Users receive a quality score (0-100) with an explanation of strengths and weaknesses, helping them understand *why* certain queries work better than others.

3.2.2 Intelligent Keyword Suggestions. Rather than simple synonym expansion, Search Whisperer generates contextually relevant keyword suggestions with explanations. Each suggestion includes a "why" annotation explaining its relevance:

> **perinatal mortality** — *Medical term that yields academic and clinical sources*
> **kraamsterfte etiologie** — *Dutch scientific terminology for studying causes*

3.2.3 Specificity Intelligence. A key design principle is matching suggestion depth to query depth. Short, exploratory queries ("climate change") receive broad suggestions for initial orientation. Detailed queries ("amsterdam west infant mortality socioeconomic factors 2020-2025") receive suggestions that *preserve*

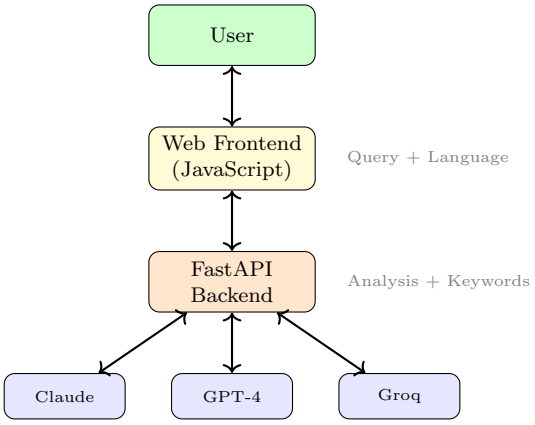

Fig. 1. System architecture showing the multi-provider LLM integration. The backend abstracts provider differences, allowing seamless failover.

all contextual elements while adding domain expertise. This prevents the common problem of AI systems "dumbing down" specific queries.

3.2.4 Google Dork Generation. For users comfortable with advanced search, the system generates Google dorks—specialized search operators that constrain results to specific domains, file types, or date ranges:

```
site:nih.gov "infant mortality"
  Amsterdam filetype:pdf 2022..2026
```

3.2.5 Multi-Language Support. The system automatically detects query language and adapts all outputs accordingly. When language is ambiguous (e.g., words shared between Dutch and German), it presents a clarification dialog. This is crucial for non-English users, who often receive English-biased suggestions from AI systems.

3.3 Interaction Modes

Search Whisperer offers three interaction modes reflecting different sensemaking stages (Figure 2):

- **Question Tab**: Initial query analysis and quick keyword suggestions for early exploration
- **Hotfix Tab**: Rapid generation of alternative phrasings and quick fixes for stuck searches
- **Full Repair Tab**: Comprehensive query reconstruction with domain expert "personas" (e.g., Academic Researcher, Investigative Journalist)

4 DESIGN RATIONALE

4.1 Transparency Over Automation

Unlike invisible query expansion, Search Whisperer makes all suggestions explicit. Users see *what* is suggested and *why*, supporting learning and informed decision-making. This aligns with principles of human-AI collaboration that preserve user agency [1].

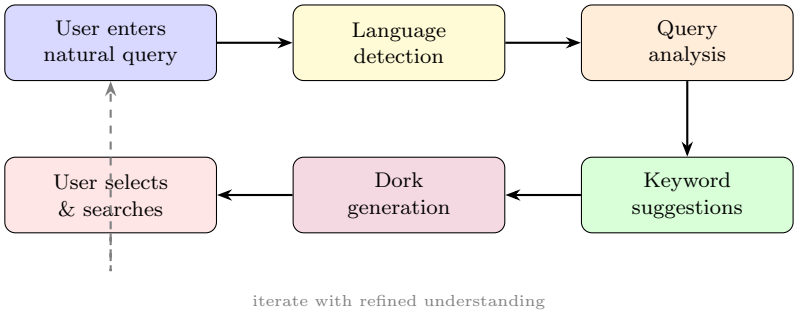

Fig. 2. User interaction flow showing the query reformulation cycle. Dashed arrow indicates iterative refinement as users develop domain vocabulary.

4.2 Language as First-Class Citizen

Many AI-assisted search tools assume English or provide perfunctory translation. We treat language as fundamental to sensemaking—a Dutch user researching Dutch topics should receive Dutch domain terminology, not translated English terms. The system includes Latin/Greek scientific terminology ("etiologie," "pathofysiologie") when appropriate, recognizing that academic discourse often transcends language boundaries.

Figure 3 illustrates how language detection influences the entire pipeline.

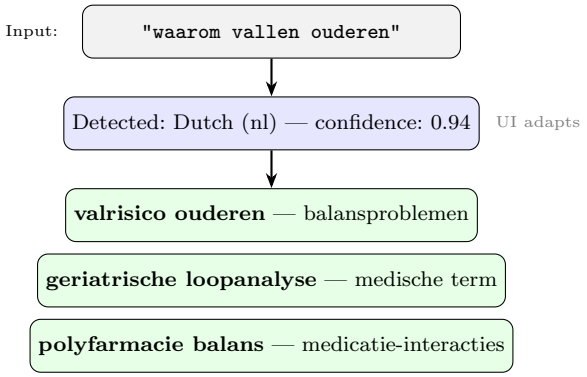

Fig. 3. Language-aware keyword generation. Dutch query receives Dutch domain terminology with Dutch explanations.

4.3 Supporting Search Literacy

By explaining suggestions and demonstrating advanced search techniques, the system aims to improve users' search literacy over time. The goal is not dependency on AI assistance, but rather scaffolded learning that transfers to unassisted search.

5 PRELIMINARY OBSERVATIONS

While formal evaluation is ongoing, informal use has revealed several patterns:

Vocabulary Discovery: Users frequently express surprise at domain-specific terms they hadn't considered. A user researching "why old people fall" discovered "fall risk assessment geriatric," "gait analysis elderly," and "polypharmacy balance"—vocabulary that opened new information spaces.

Specificity Preservation: Early versions over-generalized detailed queries. The current "Specificity Intelligence" approach maintains user-specified constraints, which users find respectful of their existing knowledge.

Language Comfort: Non-English users report feeling "understood" when the system responds in their language, reducing the cognitive load of mental translation during sensemaking.

Dork Skepticism: Many users are unfamiliar with advanced search operators and initially skeptical. Once they see results, adoption increases significantly.

6 DISCUSSION AND WORKSHOP QUESTIONS

Our work raises several questions for the sensemaking community:

Changing Search Literacy: As AI assistance becomes ubiquitous, what does "search literacy" mean? Should we teach query formulation, or teach effective prompting of AI assistants?

Trust and Verification: Users cannot easily verify that AI-suggested keywords are appropriate for their domain. How do we design for appropriate trust calibration?

Sensemaking Loop Effects: Does AI-assisted query reformulation accelerate sensemaking, or does it create new dependencies that slow independent exploration?

Language and Culture: How should AI search tools handle the inherent biases toward English and Western knowledge structures in their training data?

Expert vs. Novice Needs: Domain experts need different assistance than novices. How can systems adapt to user expertise levels?

7 CONCLUSION

Search Whisperer demonstrates how AI can augment human sensemaking by addressing the query formulation bottleneck. By providing transparent, explained suggestions in the user's language, the system aims to enhance rather than replace human judgment in the information-seeking process. We look forward to discussing these ideas and learning from the diverse perspectives at the workshop.

ACKNOWLEDGMENTS

[Acknowledgments removed for review]

REFERENCES

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kiber, and Eric Horvitz. 2019. Guidelines for Human-AI Interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (2019), 1–13. https://doi.org/10.1145/3290605.3300233

[2] Nicholas J. Belkin. 1980. Anomalous States of Knowledge as a Basis for Information Retrieval. *Canadian Journal of Information Science* 5, 1 (1980), 133–143.

[3] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. *arXiv preprint arXiv:1901.04085* (2019).

[4] Peter Pirolli and Stuart Card. 2005. The Sensemaking Process and Leverage Points for Analyst Technology as Identified Through Cognitive Task Analysis. *Proceedings of the International Conference on Intelligence Analysis* 5 (2005), 2–4.

[5] Joseph John Rocchio. 1971. Relevance Feedback in Information Retrieval. In *The SMART Retrieval System: Experiments in Automatic Document Processing*, Gerard Salton (Ed.). Prentice-Hall, 313–323.

[6] Daniel M. Russell, Mark J. Stefik, Peter Pirolli, and Stuart K. Card. 1993. The Cost Structure of Sensemaking. In *Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems (CHI '93)*. ACM, New York, NY, USA, 269–276. https://doi.org/10.1145/169059.169209

[7] Chirag Shah and Emily M. Bender. 2022. Situating Search. *Proceedings of the 2022 Conference on Human Information Interaction and Retrieval* (2022), 221–232. https://doi.org/10.1145/3498366.3505816