Heydar Soudani

Applied Researcher · LLM Agents · Post-Training & RL · Retrieval · Uncertainty Quantification

+31 6 10607261 · heydar.soudani@gmail.com · LinkedIn · heydarsoudani.github.io · github.com/HeydarSoudani · Google Scholar

Based in the Netherlands · Wrapping up PhD (expected 2027) · Open to full-time roles

Summary

Final-year PhD candidate in Computer Science at Radboud University training, building and evaluating LLM systems with two focuses: **LLM agents** (deep research agents, agentic search, tool use) and **uncertainty quantification**. At **Thomson Reuters Labs**, built a multi-agent search system for legal research and a controller for deep research agents based on trajectory-level signals, improving both the search trajectory and final task performance (EMNLP 2026); at **KPN**, brought uncertainty estimation to a customer-facing RAG chatbot. Hands-on with LLM fine-tuning (full, LoRA), RL post-training (RLVR), retrieval and reranking, synthetic data generation and benchmark construction. Published at ACL, EMNLP, COLING, SIGIR, SIGIR-AP, ICTIR and ACM Computing Surveys.

Experience

Thomson Reuters Labs *Nov 2025 – May 2026*

Data Scientist Intern (Applied AI Research) Zug, Switzerland (Remote)

- **Multi-agent search for legal research**: replaced a single-step retrieval pipeline with a modular **multi-agent search system** (planner, query decomposer, retriever, reranker and summarizer agents) to improve task completion. Worked with the **product team** on **prompt engineering** and **agent orchestration**.
- **Stagnant reasoning** (*EMNLP 2026*): through intermediate analysis of 7 deep research agents on 3 datasets, found a failure mode: recall plateaus after a few search calls while agents keep searching. Built **RAAC**, a controller that uses 4 unsupervised trajectory signals to make the agent continue, intervene or stop, and backed it with extensive **trajectory-level analysis** of agent search behavior.
- **Query reformulation & reranking study**: benchmarked decomposition and reformulation methods (HyDE, Query2doc, Rewrite-Retrieve-Read, ThinkQE) with 8 rerankers (MonoT5, RankLLaMA, Rank1, Rank-R1, Qwen3, RankGPT, RankZephyr) and **rank fusion** (RRF, interleaving, bandits). Found that one-shot methods are capped by the LLM's own knowledge, which motivated agentic search.
- **Infrastructure**: **AWS** (SageMaker multi-GPU jobs, Bedrock), **Langfuse** trajectory tracing, Claude and GPT APIs, and open-weight Qwen models served with **vLLM**.

KPN *May 2025 – Oct 2025*

Data Scientist Intern Amsterdam, the Netherlands

- **Uncertainty for a production chatbot**: applied the axiomatic uncertainty method (*ACL 2025*) to KPN's customer-facing **RAG chatbot**, which answers from a **knowledge graph** and text, to score its confidence in each answer and flag unreliable responses.

Radboud University *Feb 2023 – Present*

PhD Candidate, Trustworthy Generative Information Access Nijmegen, the Netherlands

- **Uncertainty for search agents** (*SIGIR 2026*): built **R2C**, a **multi-sampling** uncertainty method that perturbs an agent's reasoning-and-search trajectory and scores how consistent its answers stay. Across 5 agents: +5% AUROC. Tested the uncertainty scores on downstream tasks: **abstention** (about 5% better F1) and **model selection** (about 7% higher exact match).
- **Total Recall QA** (*SIGIR 2026*): co-built a **verifiable benchmark** for deep research agents where answering a query requires finding **all** relevant documents (total recall), with about 2.3K queries over a 57.7M-document corpus. It remains challenging for state-of-the-art agents.
- **Multimodal uncertainty** (*ICTIR 2026*): co-developed **LeMUQ**, a **fine-tuned RoBERTa** scorer based on **input perturbation**: it reads a VLM's token probabilities with the image or retrieved context removed, improving AUROC by 0.038 on average.
- **Fine-tuning vs. RAG** (*SIGIR-AP 2024*): to find whether RAG or fine-tuning better injects **long-tail, less popular knowledge** into LLMs, compared **full fine-tuning**, **QLoRA**, synthetic QA generation and RAG on 12 LMs (80M–11B). RAG won on long-tail entities, and the proposed **Stimulus RAG** beat fine-tuning with no training cost.

- **Uncertainty in RAG** (*ACL 2025*): for the **retrieve-then-generate** RAG setup, defined **5 axioms** for how uncertainty should change under retrieval, showed existing methods meet only 2, and derived a **calibration function** that brought RAG AUROC to or above the no-retrieval level.
- **Synthetic conversational data** (*ACM Computing Surveys 2026*): led a survey of **synthetic multi-turn dialogue generation** for open-domain, task-oriented and information-seeking systems. Proposed a general framework covering seed data creation, utterance generation and **quality filtering**.
- **Tutorials**: Data Augmentation for Conversational AI, 2 editions (*CIKM 2023*, *WebConf 2024*).

Atrovan *Dec 2017 – Jun 2022*
Web Developer (part-time) Tehran, Iran

- **Products**: built and maintained **4 production web products** for industrial and individual users: smart metering, fleet management, home automation and building management.
- **Engineering**: **React.js**, JavaScript, **REST APIs**, **PostgreSQL**, **Docker**, **Kubernetes** and Git, in a 10+ person **Scrum** team, contributing to **system design**.

Selected Publications

[1] **EMNLP 2026**. When Deep Research Agents Stagnate: Enhancing Reasoning with Retrieval-Aware Agent Control
[2] **SIGIR 2026**. Uncertainty Quantification for Retrieval-Augmented Reasoning
[3] **SIGIR 2026**. Total Recall QA: A Verifiable Evaluation Suite for Deep Research Agents
[4] **ICTIR 2026**. Uncertainty Quantification for Multimodal Retrieval Augmented Generation
[5] **ACM Computing Surveys 2026**. A Survey on Recent Advances in Conversational Data Generation
[6] **ACL 2025**. Why Uncertainty Estimation Methods Fall Short in RAG: An Axiomatic Analysis
[7] **SIGIR-AP 2024**. Fine Tuning vs. Retrieval Augmented Generation for Less Popular Knowledge
[8] **COLING 2022**. Persian Natural Language Inference: A Meta-learning Approach

Technical Skills

Agents & multi-agent systems: deep research agents (ReAct, Search-R1, Tongyi-DR), multi-agent search, query decomposition & reformulation, LangChain / LangGraph, prompt engineering
Retrieval & ranking: dense & sparse retrieval (BM25, SPLADE, E5, Qwen3-emb), LLM rerankers, rank fusion, FAISS
LLM training & post-training: SFT, full fine-tuning, LoRA / QLoRA, RL post-training (GRPO, PPO, DPO, RLVR), reward design, synthetic data generation
Uncertainty, reliability & evaluation: uncertainty quantification (multi-sampling, input perturbation), calibration, LLM-as-judge, benchmark construction, retrieval & trajectory analysis
LLM inference & models: vLLM, LiteLLM, OpenRouter, vision-language models
Languages: Python, SQL, Bash, JavaScript, C/C++
ML / DL: PyTorch, Hugging Face (Transformers, PEFT, Datasets), scikit-learn, TensorFlow, NumPy, pandas
MLOps & cloud: AWS (SageMaker, Bedrock), HPC / Slurm, Langfuse, W&B, Docker, Kubernetes
Software engineering: Git, Linux, REST APIs, PostgreSQL, system design, Claude Code, Cursor

Education

PhD Computer Science — Radboud University, 2023–2027 (expected)
MSc Computer Engineering — Sharif University of Technology, 2022
BSc Electrical Engineering — Polytechnic of Tehran, 2019
BSc Computer Engineering (minor) — Polytechnic of Tehran, 2018

References

Dr. Faegheh Hasibi — Assistant Professor, Radboud University, Nijmegen · f.hasibi@cs.ru.nl
Dr. Hamed Zamani — Associate Professor, University of Massachusetts Amherst · zamani@cs.umass.edu
Dr. Navid Rekabsaz — Research Scientist, Grubel (formerly Thomson Reuters Labs) · navid.rekabsaz@gmail.com