AMAR VISWANATHAN

amarviswanathan@gmail.com • New York, NY • linkedin.com/in/amarviswanathan

PROFESSIONAL SUMMARY

Applied Science leader with **12+ years of experience** (8+ years post-PhD in industry) and **3 years of formal people management** - hiring, performance reviews, and goal-setting - leading teams of MLEs and Senior Scientists at Amazon, SAP, Verisk, and Siemens Research. Delivered **6 production launches** across B2C ads, enterprise AI, financial news, and insurance, driving **$26M+ in combined business and research impact**.

- **Tech Lead & Single Threaded Owner** of two major initiatives at Amazon Ads: (1) TACTIC, the first agentic system for targeting clause relevance and ranking optimization — driving ad campaign performance through automated promotion and demotion — and (2) AutoTrends, an autonomous trend-aware agentic system for detecting implicit and explicit trends in targeting; both launched Q1 2026 across 69M+ campaigns worldwide.
- **Single Threaded Owner - Distributed Optimization System (Amazon, NDA):** Developed a scalable distributed optimization system with DSPy-powered explainability; beta launched April 2026.
- **People leadership:** Formally managed teams of up to 16 across multiple companies; hired and grew ML teams, conducted structured 1:1s, set OKRs, and delivered performance reviews for MLEs and Senior Scientists.
- **Technical depth:** Deep expertise in RAG, LLM post-training, agentic AI, GraphRAG, Knowledge Graphs, NER/RE, embeddings, and vector retrieval. Production systems at scale: 69M+ ad campaigns, 22K enterprise conversations/day, 19M tweets/day.
- **Research credibility:** PhD (RPI), co-PI on $1.8M DARPA ASKE program, 15+ peer-reviewed publications and arXiv papers (ECIR, CIKM, ICIP), 1 granted US patent, 5+ patent applications in progress. Reviewer at NAACL, ACL, AAAI.
- Early adopter of coding agents (Amazon Kiro CLI, Amazon Q, Codex) for ML pipeline development, code generation, and rapid iteration — applied directly to TACTIC and AutoTrends codebases.

EDUCATION

PhD, Computer Science — Rensselaer Polytechnic Institute (RPI), Troy, NY *2011 – 2018*
Dissertation: Schema- and Data-aware Query Reformulation in Knowledge Graphs | Advisor: Dr. James A. Hendler | AAAI 2016 Doctoral Consortium Award

M.S., Computer Science — Rensselaer Polytechnic Institute (RPI), Troy, NY *2014 – 2016*

B.E., Computer Science — Anna University, Chennai, India *2003 – 2007*

PROFESSIONAL EXPERIENCE

Senior Applied Scientist | Amazon Ads & Special Projects | New York, NY *Apr 2025 – Present*

- **Tech Lead & Single Threaded Owner — TACTIC (Targeting Clause Tuning through Intelligent Control):** Designed and launched the first agentic targeting clause promotion and demotion system for 69M+ worldwide ad campaigns, orchestrated via an agentic architecture with chain-of-thought reasoning for explainability; contributed to an estimated $3M+ in ads revenue impact. Owned science roadmap and engaged cross-org stakeholders across Ads and Measurement orgs. 2+ patent applications in progress.
- **Tech Lead & Single Threaded Owner — AutoTrends:** Led development of an autonomous trend-aware LLM agentic system using a two-stage gated trend engine to detect implicit and explicit trends in targeting clauses for keywords. Operates in conjunction with TACTIC across 69M+ campaigns; delivered a 1.3× performance lift.
- **Single Threaded Owner — Distributed Optimization System (Amazon, NDA):** Developed a scalable

scalarized decentralized distributed optimization system integrated with a DSPy-powered explainability framework; orchestrated via Apache Airflow on EMR clusters; beta launched April 2026.

- Mentored junior scientists on agentic AI architectures and CoT evaluation frameworks, establishing team-wide best practices for autonomous targeting systems.

Principal Scientist | SAP Labs | Palo Alto, CA *Mar 2024 – Mar 2025*

- **Shipped SAP Consultant Capability (SCC) to General Availability (Q4 2024)** — a flagship enterprise RAG product driving $22M in business impact, serving 22K multi-turn conversations/day at conversation precision of 0.79.
- **Applied LLM post-training, agentic architectures, and GraphRAG techniques** to reduce hallucinations and improve answer precision across the SCC pipeline, enabling robust handling of complex multi-turn enterprise queries.
- **Formally managed a cross-functional team as Principal Scientist:** grew direct reports from 2 to 8 (including interns); instrumental in broader org hiring of PMs, SWEs, and MLEs, scaling the total team to 16. Conducted structured 1:1s, set individual OKRs, and delivered performance reviews.
- **Hired and raised the bar:** led technical interviews and hiring panels, significantly raising applied ML rigor across the SCC product team.
- Architected the end-to-end KG-RAG pipeline; selected and fine-tuned retrieval model (BGE-EN-ICL) and generation model (Claude 3.5 Sonnet); deployed multi-node Milvus DB for vector retrieval; optimized inference via vLLM.
- Developed a strong evaluation framework for RAG incorporating domain experts and RAGAS-based LLM-as-a-judge methodology (patent application in progress).
- Drove research roadmap for encoder pretraining and SLM improvements, influencing SAP's broader Joule conversational AI strategy.
- Published at ECIR 2025 and SIGIR; 3+ patent applications in progress.

Senior Research Scientist | Dataminr | New York, NY *Jul 2022 – Feb 2024*

- **Launched RAG-based fact retrieval system** into Dataminr's B2B product using E5-Mistral embeddings and fine-tuned LLaMA-7B; achieved 82% factual accuracy on production data processing 19M tweets/day.
- Designed and deployed hallucination detection and correction system for summarization pipeline (mBART + GPT), correcting 802K hallucinated entities per day at 66% accuracy.
- Built end-to-end knowledge graph construction pipeline (Neo4j) processing 19M tweets/day; fine-tuned BERT for NER, LUKE for relation extraction (precision 70.2); orchestrated via Apache Airflow and Databricks.

Machine Learning Research Manager | Verisk AI | New York, NY *Jun 2021 – July 2022*

- **Formally managed a team of 5 MLEs** — conducted hiring, structured 1:1s, set quarterly OKRs, and delivered performance reviews; grew team capability in applied NLP and knowledge graphs.
- **Built and led a government grants research initiative from scratch:** secured $250,000 in DARPA funding by assembling a team and leading proposal submissions — a net-new revenue stream for Verisk.
- **Shipped 2 B2B launches for the insurance industry:** an abstractive summarization system for large-scale multimodal event news (ROUGE-F1 $>$ 35%, deployed on Apache Airflow) and an event knowledge graph ('The Machine') with 1M+ entities and 80 property edges — driving a $1M+ business need for Verisk clients.
- Deployed knowledge graph as both RDF and Property Graph representations (Amazon Neptune, Neo4j); managed full data pipeline via Apache Airflow.

Acting Senior Key Expert – KG / Staff Research Scientist | Siemens Corporate Technology | Princeton, NJ
Aug 2018 – May 2021

- **Co-Principal Investigator on $1.8M DARPA ASKE program** (Deep Code Curator); led 15 members across industry and academic partners — set milestones, ran project reviews, and coordinated cross-org deliverables across two successful program phases.
- **Managed 2 research scientists and 2 software engineers** on multimodal KG and visual QA projects; provided performance input for annual reviews and set individual research goals.
- Built end-to-end BiLSTM-CRF extraction pipeline converting 960 scientific papers into queryable RDF knowledge graphs; architected KB construction and oversaw image extraction pipeline.
- Developed KG-enabled caption generation and Visual Question Answering systems using DBpedia/YAGO, encoder-decoder networks, trained on COCO/Pascal VOC datasets.
- Promoted from Staff Research Scientist to Acting Senior Key Expert within one year based on DARPA program leadership impact.

Senior Systems Engineer | Infosys Technologies | Bengaluru, India *Oct 2007 – Jul 2011*

- Built iSEE (Infosys Sentiment Extraction Engine) for customer opinion mining; resulted in a granted US patent (2013).
- Developed Semantic Wiki integrations and BPEL-based workflow automation for enterprise clients; published at WWW 2009.

PUBLICATIONS & PATENTS

2025

- Min, Mathew, Pan, Bansal, Keshavarzi, **Amar Viswanathan Kannan**. Efficient Knowledge Graph Construction and Retrieval from Unstructured Text for Large-Scale RAG Systems. *arXiv 2507* (2025).
- Min, Bansal, Pan, Keshavarzi, Mathew, **Amar Viswanathan Kannan**. Towards Practical GraphRAG: Efficient Knowledge Graph Construction and Hybrid Retrieval at Scale. *arXiv:2507.03226* (2025).
- **Amar Viswanathan Kannan** & Sasaki. Combining Knowledge Graphs and Retrieval Augmented Generation for Enterprise Resource Planning. *ECIR 2025.*

2020

- **Amar Viswanathan Kannan**, Fradkin, Akrotirianakis, Kulahcioglu, Canedo, Roy, Yu, Malawade, Al Faruque. Multimodal Knowledge Graph for Deep Learning Papers and Code. *CIKM 2020.*
- Roy, Akrotirianakis, **Amar Viswanathan Kannan**, Fradkin, Canedo, Koneripalli, Kulahcioglu. Diag2Graph: Representing Deep Learning Diagrams in Research Papers as Knowledge Graphs. *ICIP 2020.*

2019

- **Amar Viswanathan Kannan**, Akrotirianakis, Roy. Deep Code Curation (DCC) — A Project to Build Multi-modal Knowledge Graphs from Deep Learning Papers. *AKBC/SLKB Workshop 2019.*

2018

- **Amar Viswanathan Kannan**, de Mel, Hendler. Feature-based Reformulation of Entities in Triple Pattern Queries. *arXiv:1807.01801* (2018).
- **Amar Viswanathan Kannan**. Schema- and Data-aware Query Reformulation in Knowledge Graphs. *PhD Thesis, RPI* (2018).

2017

- **Amar Viswanathan Kannan**, Michaelis, de Mel, Hendler. In-context Query Reformulation for Failing SPARQL Queries. *SPIE 2017.*

- Rashid, **Amar Viswanathan Kannan**, Gross, Kendall, McGuinness. Leveraging Semantics for Large-Scale Knowledge Graph Evaluation. *WebSci 2017*.

2016

- **Amar Viswanathan Kannan**. Pragmatic Querying in Heterogeneous Knowledge Graphs. *AAAI 2016 Doctoral Consortium*.
- **Amar Viswanathan Kannan**, De Mel, Hendler. Pragmatics and Discourse in Knowledge Graphs. *AAAI 2016 Workshop on Symbiotic Cognitive Systems*.

2009–2013

- Erickson, **Amar Viswanathan Kannan**, Shinavier, Shi, Hendler. Open Government Data: A Data Analytics Approach. *IEEE Intelligent Systems* (2013).
- **Amar Viswanathan Kannan**, Venkatesh, Vasudevan, Balakrishnan, Shastri. Suggestion Mining from Customer Reviews. (2011).
- Hussain, Balakrishnan, **Amar Viswanathan Kannan**. Semantic Wiki Aided Business Process Specification. *WWW 2009*.

Patents

- Balakrishnan, Vasudevan, **Amar Viswanathan Kannan**, Raghunathan, Ravindran. Methods for Analyzing User Opinions and Devices Thereof. *US Patent App. 13/946,832* (2013).

AWARDS & RECOGNITION

- Co-PI, DARPA ASKE Program — $1.8M award, 2018–2020
- AAAI 2016 SIGAI Doctoral Consortium Award
- NSF Award, U.S. Semantic Technologies Symposium (US2TS 2018)
- Finalist, Three Minute Thesis (3MT), RPI Graduate Research Symposium 2016

PROFESSIONAL SERVICE

PC Member / Reviewer: NAACL 2026, 2025, ACL 2026, 2025, 2024, AAAI 2022, CogSci 2015–2020, Big Data Journal • Conference Chair: CIKM 2020 (Demos & Posters) • Organizer: WebSci 2017 • Book Reviewer: Manning Publications (Knowledge Graphs)