

Mukesh Mithrakumar

 [mukeshmithrakumar](#)

 mukesh@mukeshmithrakumar.com

 +1 605-690 7739

 [mukesh-mithrakumar](#)

Professional Experience

LEAD MACHINE LEARNING ENGINEER

84.51°

Oct 2024 – Present
Plymouth Meeting, PA

PLATFORM IMPACT & LEADERSHIP:

- Head engineering efforts to scale and productionize Kroger's personalization and search sciences, building the real-time platform that powers personalized user experiences for 80M+ households.
- Translated the personalization and search platform vision into actionable roadmaps, balancing performance, scalability, and cost while aligning product and engineering leadership on priorities, quarterly goals, and resourcing.
- Delivered: 10-15% uplift in cart adds and revenue, 2-3% overall sales lift, and ~5% increase in sales per search.

PERSONALIZATION & SEARCH:

- Implemented a multi-stage search system, combining lexical relevance, learned semantic embeddings, and real-time personalization signals to improve intent matching and result relevance across a 40M+ product catalog.
- Designed and deployed a two-stage real-time basket completion recommender system, leveraging user purchase history, session context, and item co-occurrence signals to drive incremental add-to-cart behavior.
- Built data parallel distributed training pipelines for recommendation and search models (two-tower, attention-based, xgboost, and JEPA), training on 1TB+ datasets using PyTorch, Kubeflow, Horovod, Ray, and NVIDIA Merlin.
- Reduced out-of-stock search results by building a real-time product availability microservice using Kafka, Beam, Dataflow and Bloom filters and integrating it across search, personalization and monetization, protecting \$250-\$500M in annual revenue.
- Built low-latency online inference architectures for search and personalization (sub-250ms P99) to support 1K+ QPS and 10B+ monthly personalized inferences using Vertex AI managed services, Cloud Run, Triton and TensorRT.

ML PLATFORM:

- Architected & built the MLOps foundation by templating model registration, deployment, and rollout strategies; building reusable CI/CD workflows and a centralized feature management and governance framework, enabling faster, safer model launches across multiple sciences.
- Led the design of end-to-end production data pipelines in collaboration with R&D, Data Engineering, and Infrastructure, spanning streaming systems (Kafka, Dataflow, Bigtable) and batch/ETL pipelines (PySpark, Dataproc, Rust).

SENIOR MACHINE LEARNING ENGINEER

IQVIA

Aug 2020 – Oct 2024
Plymouth Meeting, PA

PRODUCTS & BUSINESS IMPACT:

- Built and integrated IQVIA's AI assistant into the \$100 million+ Orchestrated Analytics products, enabling near real-time insights for 100k+ users.
- Delivered a suite of Next Best recommendation systems (Message, Content and Action) and a widget curation engine driving 5% click through rate, over 90% adoption rate and up to 15% uplift across 29 markets.
- Created a demand forecasting system achieving 95% accuracy for 30-day predictions, improving inventory efficiency by 25%.

ML SYSTEMS:

- Designed and implemented an AI-powered sales assistant using LLMs, Text-to-SQL, RAG, entity extraction, and knowledge graphs to provide automated insights and streamline complex data queries.
- Trained distributed deep learning models, including recommenders, CNNs, and transformer-based forecasters, and fine-tuned a small language model using Horovod, Dask, PyTorch, and TensorFlow.

ML PLATFORM:

- Established and maintained MLOps platforms on AWS, along with CI/CD infrastructure reducing deployment timelines from months to days by implementing automated workflows and establishing streamlined processes and best practices.
- Scaled ML pipelines using Kubeflow and Argo on Kubernetes, enabling reproducible training, orchestration, and deployment.
- Deployed models for real-time and batch serving using TensorFlow Serving and TorchServe, supporting 2 million weekly recommendations.

LEADERSHIP & MENTORING:

- Coordinated cross-functional collaboration across product, data science, research, DevOps, and engineering teams, managing efforts among 50+ members.
- Mentored 10+ ML engineers through 1:1s and workshops, improving project delivery by 40%.
- Shaped quarterly and annual roadmaps, delivering 15+ ML solutions annually.

Built and delivered multiple end-to-end machine learning products for early-stage customers, spanning computer vision, NLP, and generative systems, including cardboard quality detection for recycling, automated check processing for low-resource languages in banking, research paper search and summarization, and generative image tools for an animation studio.

Education

MS. IN COMPUTER SCIENCE

Arizona State University

Tempe, AZ

BS. IN PHYSICS.

South Dakota State University

Brookings, SD

Teaching

INSTRUCTOR

- Associate Instructor, Machine Learning and Machine Learning Operations. (FourthBrain). Summer 2021 & 2022.
- Lead Instructor, Data Science. (General Assembly). Fall 2020.

TEACHING ASSISTANT

- Machine Learning (FourthBrain). Fall 2021.
- Physics (PHY 2048 & PHY 2049). Fall 2015 – Spring 2017

MENTORING

- Data Science (Springboard). 2018-May 2024.
- Data Science (Thinkful). 2018-2020.

Publications

JOURNAL

- Y. Jin, P. Kharel, P. Lukashov, S. Valloppilly, B. Staten, J. Herran, I. Tatic, **M. Mithrakumar**, B. Bhusal, A. O'Connell, K. Yang, Y. Huh, R. Skomski and D. J. Sellmyer. "Magnetism and electronic structure of CoFeCrX (X = Si, Ge) Heusler alloys". Journal of Applied Physics 2016; (Vol.120, Issue 5).

CONFERENCE

- Tong Wu, Mateusz Buda, **Mukesh Mithrakumar**, Yunlong Wang, Srikanth Sankaran Iyer, Tanveer Ahmed Nasir. "Leveraging Language Model for Next Best Action in Promotion Campaigns to Augment HCP Engagement". Pharmaceutical Management Science Association 2023. (Poster).
- **M. Mithrakumar**, B. Bhusal, Y. Huh and P. Kharel. "Tuning Magnetic Properties of Co₂FeGe with Cr Substitution for Fe", South Dakota Academy of Science 2016. (Poster).

WORKSHOP

- **Mukesh Mithrakumar**. "Training a Distributed NER model for drug, disease and condition identification", Healthcare NLP Summit 2021. ([Pre-recorded talk](#)).
- **Mukesh Mithrakumar**. "Synthesis, Structure and Magnetism of Co₂Fe_{1-x}Cr_xGe (x=0,0.25,0.5) alloy", Undergraduate Research, Scholarship and Creative Activity Day 2016. (**Certificate of Excellence for Outstanding participation in Undergraduate Research**).

OPEN-SOURCE PROJECTS

- Mukesh Mithrakumar, [Deep Learning with TensorFlow 2.0 Book](#).
- Mukesh Mithrakumar, [TensorFlow Scientific Package](#).
- Mukesh Mithrakumar, [Keras RetinaNet](#).

COMPETITIONS

- Mukesh Mithrakumar. [Google AI Open Images – Object Detection](#). (Kaggle Top 100 – Bronze Medal). 2019.
- Mukesh Mithrakumar. [Google AI Open Images – Visual Relationship](#). (Kaggle Top 100 – Bronze Medal). 2019.

SELECTED BLOG POSTS

- Mukesh Mithrakumar. [How to tune a Decision Tree](#). Towards Data Science. 2019. (Over 300k Views).