

Mukesh Mithrakumar



[mukeshmithrakumar](#)



mukesh@mukeshmithrakumar.com



[+1 605-690 7739](tel:+16056907739)



[mukesh-mithrakumar](#)

Professional Experience

LEAD MACHINE LEARNING ENGINEER, PERSONALIZATION & LOYALTY STRATEGY

Oct 2024 – Present

84.51*

Plymouth Meeting, PA

- Leading the design and implementation of a real-time, two-stage recommender system for next-item prediction while modernizing legacy batch processing infrastructure.
- Built distributed data processing infrastructure with Kubeflow and PySpark on Dataproc for feature processing at scale.
- Designed and implemented an end-to-end ML feature platform using Vertex AI Feature Store and Google Big Table enabling real-time feature serving. Built CI/CD pipelines that reduced feature setup time from hours to minutes.

SENIOR MACHINE LEARNING ENGINEER, NEXT BEST TEAM

Aug 2020 – Oct 2024

IQVIA

Plymouth Meeting, PA

Led the development of enterprise-scale AI/ML products, serving over YYYYk sales representatives globally and driving more than \$XX million in annual revenue from pharmaceutical client sales.

KEY PRODUCTS & ACHIEVEMENTS:

- Delivered and integrated IQVIA's AI assistant into the \$XXX million+ Orchestrated Analytics products, enabling near real-time insights for YYYY+ users.
- Delivered a Next Best Action recommendation system and widget curation engine with >90% adoption rate and up to XX% uplift across 29 markets.
- Created a demand forecasting system achieving 95% accuracy for 30-day predictions, improving inventory efficiency by 25%.

TECHNICAL LEADERSHIP:

- 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.
- Established and maintained MLOps platforms on both AWS and GCP, along with CI/CD infrastructure reducing deployment timelines from months to days by implementing automated workflows and establishing streamlined processes and best practices.
- Implemented, trained, and fine-tuned distributed deep learning and large language models, including recommenders, CNNs, and transformer-based forecasters, using frameworks such as Horovod, Dask, PyTorch, TensorFlow, NVIDIA Merlin, and Keras.
- Scaled ML pipelines using Kubeflow on Kubernetes and Vertex AI, deploying models for real-time and batch serving with Vertex AI Endpoints, TorchServe, TensorFlow Serving, and NVIDIA Triton, delivering X million weekly recommendations.

STRATEGIC LEADERSHIP & MENTORING:

- Led cross-functional collaboration across product, front-end, back-end, data science, research, DevOps, and SRE teams, coordinating efforts among 50+ members.
- Mentored 10+ ML engineers through 1:1s and workshops, improving project delivery by 40%.
- Shaped technical strategy through quarterly roadmap planning and agile implementation.

FOUNDER AND CTO

Aug 2018 - Aug 2020

ASTRUM AI SOLUTIONS

San Francisco, CA

Developed and launched innovative Machine Learning products for early-stage startups, enhancing their market competitiveness.

- Built a Web Application for Research Paper Search, Summarization, and Chatbot using JavaScript, node.js, Python, AWS S3, API Gateway, ElasticBeanstalk, Lambda functions, SageMaker for hosting the BERT summarizer, and AWS Lex.
- Successfully led the development of a Multi-lingual Optical Character Recognition (OCR) Mobile App enabling automated check deposits for low resource languages (Tamil & Sinhala), enhancing access to banking services for underserved populations.
- Developed and deployed an object detection model on AWS to detect various grades of cardboard for a recycling company.

Education

Arizona State University

MS. IN COMPUTER SCIENCE.

Tempe, AZ

South Dakota State University

BS. IN PHYSICS.

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. Lukashev, S. Valloppilly, B. Staten, J. Herran, I. Tadic, **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).

Other

SKILLS

- Programming Languages- Python, GO, Java, JavaScript, Shell, Git, SQL.
- Big Data & Streaming: Apache Spark, PySpark, Hadoop, Apache Kafka, Dataproc
- Deep Learning Frameworks- TensorFlow, NVIDIA Merlin, PyTorch, Keras, Pandas, OpenCV.
- Distributed Deep Learning Frameworks- Horovod, Dask.
- Machine Learning Frameworks – SciKit-Learn, XGBoost, NumPy, SciPy.
- Database Management- MongoDB, Google BigQuery, db2.
- Workflow Management- Apache Airflow, Kubeflow, Vertex AI.
- Cloud Computing- Amazon Web Services, docker, Kubernetes, Google Cloud Platform, Dataproc.
- Web Development- Flask, HTML, CSS, jQuery, Node.js, Streamlit.
- Profiling Framework- Grafana, Jaeger, Prometheus.
- Product Management- Jira, Confluence.
- Software Skills- Jenkins, Lean, Agile, Scrum, Test Driven Development, CI/CD, Software Design Patterns, MLOps, SonarQube.