

Hasan Khan

ML Engineer

Brooklyn, NY
hk3550@nyu.edu
+1 804 980 5375

Work

ML Engineer – Healthcare Risk Advisors

5/2023–Current

Deploying large language models for retrieval augmented generation (RAG), question answering and classification tasks on claims data. Particular focus on RAG systems, model fine tuning and domain adaptation on clinical and legal texts.

Data Scientist – LMI

8/2018–5/2021

Provided analytics and machine learning consulting services for U.S federal agencies. Highlighted projects:

Inventory Management – Built supervised regression models to optimize rare inventory lead time estimation for the Defense Logistics Agency. Random forests trained on 10+ years of inventory procurement data resulted in a 37% improvement of lead time predictions for rare procurements over baseline metrics, with an annual savings of \$45M on inventory holding costs.

Survey Analytics – Designed and developed a survey analytics web application for NASA Human Resources. Utilized BERT models for sentiment analysis and topic modeling on free response survey comments. Deployed application using Python Dash on Shiny.

Research

Improving Knowledge Distillation – Computational Linguistics, NYU

1/2023–5/2023

Research on improving the efficacy of Knowledge Distillation in LLMs by incorporating novel curriculum learning methods, where student models selectively rely on teacher distillation loss depending on datapoint difficulty. Conducted as part of a Computational Linguistics courses' semester project.

Neuronal Modelling – Meliza Lab, UVA

5/2016–5/2018

Modeled neurons in Python to simulate auditory cognition in Zebra Finches. Coded Recursive Bayesian Filters for biophysical parameter estimation. Classified auditory neurons into subtypes based on ion channel makeup.

Education

MS in Computer Science – New York University

8/2021–5/2023

BA in Computer Science – University of Virginia

8/2014–5/2018

BA in Cognitive Science – University of Virginia

8/2014–5/2018

Proficiencies

Languages: Python, Java, C++, SQL, HTML, Ada

Libraries: Langchain, LlamaIndex, Transformers, Sklearn, PyTorch, MLflow, SparkNLP

Platforms: Databricks, Azure, Git, Docker

Certifications: Databricks Generative AI Engineer Associate

Clearances: US Federal – Secret