HARSHIT MEHTA

Senior Applied Scientist

+15129987271 • harshit.mehta@utexas.edu • linkedin • Redmond, WA

Summary

Experienced Applied Scientist with a robust background in machine learning, natural language processing (NLP), and generative AI. Adept at delivering scalable, data-driven solutions and leading cross-functional teams. Consistently achieves measurable business outcomes through innovation and technical excellence.

Experience

Microsoft - Viva Engage
Redmond, WA

Senior Applied Scientist
01/2024

- Spearheaded the rollout of generative AI features on the Viva Engage platform, including theme extraction, writing assistance, and content moderation. Most of these features improved active user engagement metric, while content moderation improved bad content filtering effectiveness. Leveraged different prompting technique like few-shot, Cot, Tot, LLM as a judge, ReAct etc.
- Partnered with 10+ teams across 5 organizations within Microsoft to evaluate the risks of generative AI features across different harms like Jailbreak, copyright, ungrounded, sexual, violence etc. This effort led to 30 + features rollout during FY 2024.
- Conducted A/B experiment to test bivariate logic for impression discounting in Recommender system, online metrics like home feed active engagement and messages views improved by 2% and 5% respectively. This increased MAU by 200K.
- Developed an offline evaluation framework for testing generative AI features for (quality, performance and Responsible AI metrics). This speed up iterations/experimentations and reduce dependency on manual reviews.
- Conducted multiple RecSys AB experiments to test two-tower architecture, Learn to Rank, and different feature engineering methods.

Microsoft - Azure Storage
Redmond, WA

Senior Data & Applied Scientist
08/2022 - 12/2023

- Led telemetry pipeline development in collaboration with multiple teams, optimizing data center storage efficiency and achieving $200M in cost savings within six months across the U.S. region.
- Designed Azure customer usage dashboards to support internal stakeholders in short- and long-term forecasting of data center demands.

Dell Technologies
Austin, TX

Consultant, Data Science
04/2019 - 08/2022

- Led the development of a pre-trained/fine-tuned multi-classification transformer model for customer survey intent prediction, enhancing CSAT score tracking accuracy by 50%.
- Developed LSTM /1D CNN multi-classification models to resolve customer issues, improving accuracy by 20% in six months. Resulted in $9M dollars saving.
- Led the development of autoencoder based anomaly detection pipelines to identify anomalous log messages.
- Designed experimentation frameworks to assess the impact of in-house NLP models.
- Mentored Data Science Interns across different teams on NLP tasks (Entity Recognition, Sentiment Analysis and Question Answering).

EZ
Austin, TX

Data Scientist
06/2018 - 03/2019

- Developed SQL queries to extract data from an SQL Server and implemented Isolation Forest and autoencoder models for anomaly detection.
- Used LDA for topic modeling. Used coherence score and perplexity for deciding optimal number of topics.
- Applied sequential model-based optimization for hyperparameter tuning, leading to the successful deployment of a LightGBM binary classifier. Saved 500K dollars over six months.

Skills

Machine learning:

Neural Network, NLP, Linear Regression, Logistic Regression, Decision Trees, SVM, Naïve Bayes, Clustering, K Nearest Neighbors, Transfer Learning, Extreme Classification, Transformers, Retriever, Ranker, LLM, Prompting, Generative AI

Programming languages: Python, Pyspark, Rstudio, SQL, KQL, SCOPE

NN Frameworks: Pytorch, Tensorflow, Huggingface

Education

University of Texas at Austin

Masters in Operation Research and Statistics | 3.85/4.00 | 08/2016 - 05/2018

Delhi Technological University

Bachelors in Engineering | 3.90/4.00 | 08/2009 - 06/2013

Key Achievements

- Collaborated across multiple organizations within Microsoft to develop an AI-powered tool for Street Grace, an NGO dedicated to protecting children from sexual exploitation.
- Co-author of two research papers with 500 + citations each.
- Prize Winner of multiple hackathons:
  Microsoft (1)
  Dell Technologies (2)
  Kaggle/bitgrit (4).

Projects

Build Multi-class Model to predict effectiveness of argumentative essays — Kaggle

https://www.kaggle.com/competitions/feedback-prize-effectiveness/leaderboard

Achieved 3rd rank among 1500+ teams.

- Built custom T5 model for text augmentation.
- Used Span MLM for pre-training and custom-built Span Classification model using LLM models.
- Used Multi-tasking, Adversarial Weight Perturbation, and mask augmentation for fine-tuning.

Build multi-linguistic recommender engine — Kaggle

https://www.kaggle.com/competitions/learning-equality-curriculum-recommendations/leaderboard

Achieved 9th rank among 1000 + teams

- Built retriever + reranker pipeline. Used Inverse Cloze Task for pre-training.
- Built custom dataloader to create language wise batches.
- Used the ArcFace approach to generate embeddings and passed them to the reranker for filtering irrelevant topics.

Classification model for Health Care Intervention Program — Curriculum project

Built Random forest classifier which reduced cost of intervention program by $4.3 Million.

Build profile matching algorithm — Bitgrit

https://bitgrit.net/competition/4

Achieved 4th rank among 2000 participants

- Built Multi-Classification model to predict swipe status for two users at a time.
- Used null importance and permutation importance techniques for feature selection.

Find me online

Kaggle
https://www.kaggle.com/harshit92

Google Scholar
https://scholar.google.com/citations?user=tWThbZ4AAAAJ&hl=en&authuser=1

GitHub
https://github.com/hmehta92

Languages

English Proficient ●●●●○

Hindi Native ●●●●●