

Vedang Waradpande

+1 (848) 437-1822

vedang.waradpande@gmail.com

[vedangw.github.io](https://github.com/vedangw)

[LinkedIn](#)

[GitHub](#)

Machine Learning Engineer and Researcher with 3+ yrs. of exp. specializing in NLP, GenAI, Data Science, Big Data, MLOps, & Backend

EXPERIENCE

Machine Learning Engineer II, *TikTok* – Mountain View, CA, USA

Aug 2024 – Present

- Developing ML workflows for creator monetization integrity on the TikTok platform involving Transformer-based image content embeddings, traditional ML models, and social graph clustering.
- Optimized human-in-the-loop content moderation efficiency by developing ML-based heuristics for isolating high-confidence spam cases, directly allowing for 27% more cases to be covered and reducing daily capital loss.
- Improved our core spam classifier pipeline through feature selection, conditional hyperparameter optimization based on data shift, and fixing retraining frequency, reducing training job execution time while maintaining model performance.
- Leading a project across several risk surfaces to design a business-oriented cluster quality metric for our graph clustering workflow.

Machine Learning Researcher (RTS III), *Rutgers University* – Newark, NJ, USA

Jun 2023 – Aug 2024

- Worked on TB drug discovery by developing an ML-based screening pipeline for large-scale molecular datasets using NLP, graph-based, and tree-based ML models for property prediction, leading to the discovery of 12 promising drug candidates.
- Predicted compound toxicity by training traditional (XGBoost, LightGBM, CatBoost, and Random Forest) and state-of-the-art deep learning models (CNN, LSTM, Graph Neural Networks, Transformers, BERT) with a test accuracy of over 95%.
- Built distributed model training and hyperparameter tuning pipelines using Ray Tune and PySpark in combination with PyTorch, Hyperopt, and Scikit-Learn, implementing multi-GPU training and inference.

Data Science Intern, *Faire* – San Francisco, CA, USA

May 2022 – Aug 2022

- Deployed an ML model (XGBoost) for predicting users with no payment intent to reduce underwriting risk and proposed actions to take on users based on the model score. Directly addressed 40% defaults and showed a 25% increase in commission.
- Developed an end-to-end pipeline using SQL for preprocessing & feature selection and Python for model training & analysis.
- Designed an optimization framework to select the best heuristic and thresholds to aggregate model results for users across orders.
- Designed an A/B test for the model by identifying primary, secondary, and guardrail metrics, and devising decision rules.

Software Engineer; Data Scientist, *Razorpay* – Bengaluru, KA, India

Jun 2019 – Aug 2021

E-Commerce Fraud Detection Team

- Pioneered the development of critical components within Thirdwatch, a fraud detection platform, encompassing core business logic, ML models, the Data Science platform, and supporting backend components and infrastructure (used the AWS stack).
- Created and deployed an ensemble XGBoost + CNN model to predict if an address is incomplete for delivery with 94% AUC.
- Trained XGBoost models for predicting fraud e-commerce orders with 75% AUC and deployed using Docker and AWS EC2.
- Implemented & scaled pipelines for model training, deployment, monitoring, and data ingestion using MLOps tools and Airflow.

Payments Data Team

- Designed and implemented a payment method recommender system for users using historical time-series data achieving a 4% boost in user retention during checkout.
- Developed a model testing pipeline using Apache Spark to test the model on extensive data with over 100 million samples.
- Implemented a fast real-time prediction REST API for the model using Scala and Akka HTTP with a latency of less than 300ms.

SKILLS

Expertise: Machine Learning, Deep Learning, Natural Language Processing, Computer Vision, Large Language Models (LLMs), Retrieval Augmented Generation (RAG), Recommender Systems, Fraud Detection, Data Science, Data Analysis, Optimization, REST APIs

Languages: Python, SQL, Scala, Go, C++, Bash (Linux), R

ML Tools: PyTorch, Tensorflow, Keras, Scikit-learn, Pandas, NumPy, MLLib, Hugging Face, NLTK, XGBoost, LightGBM, Matplotlib, OpenCV

Backend, Data & ML Infra: MySQL, Airflow, PySpark, Ray, MLFlow, Flask, Docker, Snowflake, MapReduce, AWS (S3, EC2, Sagemaker)

EDUCATION

M.S. in Computer Science - Rutgers University, New Brunswick

Sep 2021 – May 2023

B.E. (Honors) in Computer Science - Birla Institute of Technology and Science, Pilani

Aug 2015 – Jul 2019

PROJECTS

Retrieval Augmented Generation using LLMs for Question Answering, Fact-Checking, and Entity Linking [\[Code\]](#) [\[Paper\]](#)

Visual Question Answering with Generative Image & Language Models [\[Code\]](#) [\[Paper\]](#)

PUBLICATIONS

- **Waradpande, V.A.;** Prakash, P.V.S.; Jhaveri, N; Agarwal S “Predicting Completeness of Unstructured Shipping Addresses Using Ensemble Models” **[SIGIR eCom '21]**