Tejas Lodaya

https://tejaslodaya.me

Email : tejas.lodaya1@gmail.com

Mobile : +91-8892-246-348

ML Engineer with hands-on experience in Python, R, Scala, and Spark. Experience of working in diverse teams and on use cases in the travel, food, retail and SaaS industry. My work lies at the intersection of machine learning and software engineering to provide end-to-end scalable ML solutions

Education

- **PES Institute of Technology** — Bangalore, India
  Bachelor of Engineering in Computer Science; GPA: 9.00/10.0 (Top 5%) — *2013 – 2017*

Experience

- **Swiggy** — Bangalore, India
  Machine Learning Engineer — *May 2019 - Present*
  - Served as an away-team to deploy payment routing model that picks the best gateway based on historic success rate, bin number and other customer-related features
  - Building a platform to create deep learning pipelines starting from data ingestion to transformation using Apache Beam. To orchestrate pipelines and to serve them at sub-100ms latencies, kuberenetes is used
  - Built dependency graphs using airflow to achieve checkpointing and robust alerting for offline batch models. To streamline monitoring of features and real-time models, grafana and prometheus were wired
  - To boost unit economics per order and to drive inorganic growth among restaurants, implemented demand shaping in search, ads and listing
  - Achieved 10x speedups in production crons by converting row-wise operations into columnar using vectorization
- **Freshworks** — Chennai, India
  Machine Learning Engineer — *Oct 2018 - May 2019*
  - Migrated legacy codebase from open-source redis cluster to enterprise redis labs, reducing burn-rate by $250k per year. Implemented memcached to further increase key-fetch rate and reduce latency
  - Architected database model for storing normalized term frequency and document frequency across articles, achieving O(1) read and write speeds
  - Built APIs for exposing tf-idf ranking model to end customers through chatbot. Also integrated diverse use-cases like smalltalk, open-domain question answering, gibberish detector and custom intent detection engine
  - Worked on language agnostic spell-correct microservice achieving average search complexity of O(1), at the cost of pre-calculation time and storage space of n deletions
- **Noodle.ai** — Bangalore, India
  Associate Data Scientist - Product — *Apr 2018 - Sep 2018*
  - Built a novel time series clustering framework on variational recurrent auto-encoder using pytorch (on GPU) and extended to anomaly detection, serving as a major component in predictive maintaince and energy optimization
  - Developed an incremental learning framework using global-local ensemble model, where global serves as a long term model and local serves as a short term model
  - Developed a proprietary ensemble modelling technique consisting of various models such as arima, xgboost, croston, prophet to capture heterogeneity of various timeseries
  - Built an inhouse EDA tool which automatically munges data to plot the features,performs statistical tests and summarizes the result
- **Noodle.ai** — Bangalore, India
  Associate AI Engineer — *Jan 2017 - Apr 2018*
  - Built and orchestrated demand forecasting ecosystem for real-time consumption (using R). Wrote DAGs using Airflow as the workflow schedule system to run batch jobs
  - Worked on scaling compute by employing SIMD on N cores using different parallel backends like doSnow, doParallel in R
  - Built automated training and prediction backend batch jobs

Open source projects

- **Variational Recurrent Autoencoder**: Unsupervised, feature-based timeseries clustering algorithm in pytorch
- **Troop**: Simple library to perform chunkwise processing on data.frame across multiple cores of a single machine using SNOW clusters with a low memory footprint