Viraj Navkal

Machine Learning Mathematician

virajnavkal@gmail.com
(978) 394-5445
Redmond, WA

github.com/vnavkal
linkedin.com/in/vnavkal

Profile

- Machine learning engineer with strong mathematical background
- Developed core IP for two startups in different industries
- 6 years' experience as both an individual contributor and a technical lead

Skills

- Technologies:
  - Fluent with Tensorflow, numpy, pandas, scikit-learn
  - Open-source contributions to Tensorflow, scikit-image, patsy, celery
  - Experience with multiple data storage formats, both structured (e.g. PostgreSQL) and unstructured (e.g. protocol buffer, Redis)
  - Familiar with web frameworks: Ruby on Rails, Django
  - Comfortable with technologies for managing environments (Docker, Kubernetes), versioning code (git), and testing code (pytest, Jenkins)
  - Other languages: fluent with Ruby, some experience with Scala/Spark, Java, C++, emacs lisp
- Topics: Deep learning, statistical modeling, applied mathematics, computer vision, asynchronous programming, functional programming

Experience

Helm.ai (autonomous driving) | Research Scientist & Engineer

2018 - present

- I am currently the technical lead for the perception team, which is responsible for training computer vision models for the task of autonomous driving. The team consists of me and five other PhDs.
- Developed new, task-agnostic semi-supervised learning techniques for comprehending video data. My work played an important role in securing several large R&D contracts.
- Worked on many of the standard problems in computer vision, including semantic segmentation, instance segmentation, object detection, monocular depth estimation, tracking
- Designed and implemented data structures for saving model evaluation results efficiently
- Implemented important functionality in the model training and validation pipelines, enabling other researchers to train and validate models efficiently
- Helped set up and maintain a continuous integration test suite to check the exact numerical behavior of model training code, to guard against regressions
- Mentored several junior and mid-level researchers

Upstart (consumer lending) | Data Scientist

2013 - 2018

- As the first data scientist, I oversaw the loan platform's growth from $0 to over $1B in investments and the data science team's growth from 0 to 10 members
- Adapted XGBoost and Tensorflow to support the custom modeling framework that our team developed

- Managed the interface between the Django pricing application and the Ruby on Rails web application, using RabbitMQ to prioritize requests and ensure connection reliability
- Wrote an extensible, asynchronous utility to extract data from the web application's PostgreSQL databases and transform it into a single table on which statistical models could be trained
- Built financial models to estimate the distributions of future portfolio cashflows
- Received excellent performance reviews, available upon request

Education

University of California, Los Angeles / Ph.D., Mathematics

2007 - 2013

- GPA: 3.98
- Thesis topic: Commutative algebra (a generalization of linear algebra)

University of Pennsylvania / B.A., Mathematics and Physics

2003 - 2007

- GPA: 3.93
- Academic honors:
    - One of only two students in my class to receive the Phi Beta Kappa award twice
    - Member of the Benjamin Franklin Scholars honors society

Volunteer work

- Wrote backend for energy data analysis app (www.energizeapps.com) for the Energize Andover organization