

EDUCATION

- George Mason University, Virginia** August 2016 – May 2018
Master of Science, Computer Science
Relevant coursework: Theory & Applications of Data Mining, Data Mining on Multimedia Data, Pattern Recognition, Applied Statistics, Mining Massive Datasets, Computer Vision
- Fast.AI International Fellow – Deep Learning** October 2017 – December 2017
- University of Mumbai, India** May 2010 – June 2014
Bachelor of Engineering, Computer Engineering

SKILLS

Languages: Python (3 years), Java (4+ years), SQL(4+ years), C++ (< 1 year), R (< 1 year)
Tools: AWS(S3, EC2, EMR), Git, CI/CD (Travis, CircleCI), Redshift, Hadoop, Spark, Tableau, IPython, Weka, Docker
Techniques: Clustering, classification, natural language processing
Libraries: scikit-learn, Numpy, Pandas, Scipy, GraphLab, matplotlib, seaborn, gensim
Beginner level knowledge of Torch and Keras.

PROJECTS

Impressionist Artist Identification Using Deep Learning: Classifying Impressionist paintings from Wikiart based on the genre & artist using convolutional neural networks (CNNs) in Torch.
Urban Sound Classification: Implemented a research paper to classify environmental sounds from the UrbanSound8K dataset using CNNs in Keras.
Music Mood Classification Using the Million Song Dataset: Used audio features and machine learning for mood analysis. Techniques: segment aggregation, Random Forest, Support Vector Machines
Scipy: I've made a few contributions to the open source library. The closed PRs are available here: bit.do/scipy-prs
Raven: An open-source CLI tool to manage your Spotify music library.

WORK EXPERIENCE

- GIS Programmer**, George Mason University, *Fairfax, VA* August 2016 – May 2018
- Developed George Mason University's OpenGeoportal using Java, Apache Solr, Python.
 - Lead workshops on ArcGIS, QGIS, Python and CartoDB to train students and faculty.
 - Assisted students with framing research questions and designing projects based around GIS and data analysis.
- Software Engineer Intern- Data Science**, Udacity, *Mountain View, CA* June 2017 – August 2017
- Improved organizational access to data by building a Slack bot with Flask and SQLAlchemy.
 - Built an ETL system to migrate platform events to Amplitude for real-time analytics.
 - Communicated insights about Udacity products to PMs after analyzing 300,000+ events.
 - Built data pipelines using Apache Airflow.
- Senior Software Engineer**, Capgemini, *Mumbai, India* June 2014 – July 2016
- Developed features for forecasting, asset management and resource allocation in a web application using C# and ASP .NET.
 - Built dashboards & reporting systems enabling management to access up-to-date data and aid decision making.
 - Improved the performance of SQL stored procedures by 0.5-2x.
 - Improved perceived page load performance on dashboards by 30%.
 - Reduced page load speed by 6s for frequently accessed pages using caching mechanisms.
 - Mentored and onboarded 2 junior engineers by training them on the technical and business aspects of the project.
 - Developed a Java/J2EE product lifecycle management application for a Fortune 500 engineering client.