Sreedhar Radhakrishnan

linkedin.com/in/sreedhar-radhakrishnan | sreedhar1895.github.io
sreedhar@andrew.cmu.edu | +1 (412) 499-1178

EDUCATION

CARNEGIE MELLON UNIVERSITY
MS in Information Networking
Machine Learning & Systems Track
GPA: 3.71/4.0 | Merit Scholar
May 2021 (Expected) | Pittsburgh, PA

PES UNIVERSITY
BTech in Computer Science
GPA: 9.23/10.0 | Merit Scholar
May 2018 | Bangalore, India

GRADUATE COURSES

MACHINE LEARNING COURSES
Machine Learning (10-601)
Deep Learning (11-785)
ML at Scale (17-700)
Deep Learning Research (14-708)
Applied Machine Learning (11-663)

SYSTEMS COURSES
Distributed Systems (14-736)
Computer Networks (14-740)
Operating Systems (14-712)
Data Software Engineering (17-601)

UNDERGRAD COURSES

Data Structures & Algorithms
Machine Learning & Big Data
Web Services & Databases
NLP & Information Retrieval
Digital Image Processing
Theory of Computation

SKILLS

PROGRAMMING LANGUAGES
• Java • Python • Go • C • PL/SQL
CLOUD COMPUTING AND BIG DATA
• AWS Kinesis • AWS DynamoDB
• AWS EKS • Apache Spark • Docker
MACHINE LEARNING TOOLS
• Spark MLlib • Numpy • Tensorflow
WEB TECHNOLOGIES AND DATABASES
• HTML • CSS • JavaScript • Vue.js • Flask
• MongoDB • MySQL • Go-kit

ACADEMIC AWARDS

CMU GRADUATE SCHOLARSHIP
PES UNDERGRADUATE SCHOLARSHIP
USC, VITERBI RESEARCH PROGRAM

WORK EXPERIENCE

ADOBE INC. | SDE Intern, Emerging Products Group
May 2020 – August 2020 | San Jose, CA

- Developed a Machine Learning Data Pipeline using Apache Spark to ingest sentence embedding (generated using BERT) from Slack Cloud onto a MongoDB NoSQL Database using Event Driven Cloud APIs.
- Developed a Recommendation System to provide solutions during deployment failures of Adobe Photoshop (10 million+ users).
- Collaborated with the Adobe Cloud team to integrate the design into the multi-cloud (AWS/Azure/GCP) platform.

GE | Software Development Engineer, Cloud and Big Data
August 2018 – June 2019 | Bangalore, India

- Applied Pub/Sub pattern and developed a Cloud-Native Data Pipeline to stream changes of over 5 million assets at the rate of 3000 records/second in real-time using Java, AWS Kinesis and DynamoDB.
- Implemented REST APIs using Go for large-scale retrieval and processing of aviation asset information from over 5 million records.

UNIVERSITY OF SOUTHERN CALIFORNIA
Research Intern - Deep Learning for Computer Vision
June 2017 – July 2017 | Los Angeles, CA

- Developed a Cycle GAN model for augmenting synthetic images to realistic urban scenes images. The images were used for training semantic segmentation models for tasks such as pedestrian detection.
- Publication Link: https://ieeexplore.ieee.org/document/8493745.

PROJECTS

MACHINE LEARNING AND BIG DATA ON THE CLOUD
CMU Machine Learning at Scale - Fall 2020 | Pittsburgh, PA

- Engineered high volume ETL Pipelines to ingest 150 GB of Taxi Data from HDFS to AWS S3 using PySpark and AWS EMR.
- Implemented Distributed Linear Regression using Spark MLlib to predict trip duration of taxi rides with an accuracy of 80%. and deployed the model using AWS Elastic Kubernetes Service.
- Link to Medium Article: https://link.medium.com/I0Eu5drJfcb
- Currently a Teaching Assistant for the Spring 2021 offering of the Machine Learning at Scale course by Dr. Heather Miller.

DISTRIBUTED FILE SYSTEM FROM SCRATCH
CMU Distributed Systems - Spring 2020 | Pittsburgh, PA

- Developed a Distributed File System which supports parallel CRUD operations and Data Replication with Read-Write locks.

DEEP LEARNING DRIVEN WEB SERVICE
PES University - Undergrad Thesis | Bangalore, India

- Implemented a GAN driven web service that generates and renders multiple car design images from a sketch. https://link.springer.com/chapter/10.1007/978-3-319-99740-7$_1$1.