Amlesh Sivanantham

samlesh@gmail.com • (408) 219-6474 • zamlz.org • zamlz

Education

08.2017 - 05.2019 — University of Southern California
M.S. in Computer Science (GPA: 3.78)
Concentration: Intelligent Robotics

09.2013 - 06.2017 — University of California, Santa Cruz
B.S. in Computer Engineering (Honors)
B.S. in Computer Science
Thesis: *Detecting Anomalies in Time-Series Data using Long Short-Term Memory Networks* - Advisor: Dr. Patrick Mantey

Skills

Languages: **Python**, **C++**, C, SQL, Java, JavaScript, Scheme, LaTeX, Bash, Verilog
Libraries: **TensorFlow**, PyTorch, OpenAI Gym, NumPy, Spacy, Matplotlib, Flask, DL4J

Work Experience

11.2021 - Present — **Advanced Software Engineer**
Multiply Labs

01.2021 - 10.2021 — **Big Data Developer**
Wells Fargo (Transitioned to Full-Time 11.2019)
Currently utilizing Big Data technologies like Hadoop and PySpark to create features for data analysis within the Enterprise Data Lake.

06.2019 - 01.2021 — **NLP Engineer**
Wells Fargo (Transitioned to Full-Time 11.2019)
Worked with Natural Language Processing and Deep Learning for A.I. projects within Wells Fargo.

09.2017 - 06.2019 — **Graduate Research Assistant**
University of Southern California - Robotic Embedded Systems Laboratory (RESL)
Performed graduate research in Deep Reinforcement Learning and its application to Robotics. Some of the areas I have worked on have been to teach neural networks to infer inverse dynamics of a system (system identification) and learning to integrate control theory with current deep reinforcement learning algorithms.

Publications

W1. V. Chockalingam, T. T. Sung, F. Behbahani, R. Gargeya, A. Sivanantham, and A. Malysheva. Extending World Models for Multi-Agent Reinforcement Learning in MALMO. In *Joint Proceedings of the AIIDE 2018 Workshops*. AIIDE, Dec 2018

Projects

Learning Inverse Dynamics of a System for Deep RL *(USC RESL)*

Instead of having an RL policy learn a mapping from states to actions, we had it learn a mapping from states to desired states. We also learnt an inverse dynamics model concurrently from data generated by the policy. We found that the policy's performance was marginally worse than the standard approach.

PPO with Curriculum Learning for Quadrotor Navigation *(USC RESL)*

Used Proximal Policy Optimization (PPO) with curriculum learning to train a policy to control a quadrotor in a simple OpenGL quadrotor simulator we wrote. We were able to solve the task when we used perfect state information, but when we changed the state to RGB image data from a camera and IMU information, it was unable to learn the task.