SIDDHARTH MYSORE STHANESHWAR

SAN FRANCISCO BAY AREA

Website: sidmysore.github.io **LinkedIn:** www.linkedin.com/in/sidmysore **E-mail**: sidmys@bu.edu

EDUCATION

Ph.D. Computer Science – Boston University **[Aug '17 – Aug '22]**

Dissertation Title Conditioning Behavior Styles of Reinforcement Learning Policies
Research Advisor Kate Saenko
GPA 3.90

MSE Robotics – University of Pennsylvania **[Aug '14 – Dec '16]**

GPA 3.75

MEng Mechatronic Engineering – The University of Nottingham **[Sep '10 – Aug '14]**

Graduated in the First Class with Honors
Average Grade 80%
High Achievers Scholarship [2010 – 2014]
Dean's List of Top Performing Students in the Department of Electrical and Electronic Engineering [2012 – 2014]

TECHNICAL SKILLS

Programming (Primary) Python, PyTorch, TensorFlow
(Secondary) C, C++, C#, Java, Unity, OpenCV, MATLAB, ROS, OpenGL, Caffe

Computer Aided Design (CAD) SolidWorks, AutoCAD, PTC ProEngineer and Creo

Electronics Design Circuit Design, PCB design and population

Machining & Fabrication Laser Cutting, 3D Printing, Turning, Milling, Drilling, Tapping, Threading, Welding

WORK EXPERIENCE

Machine Learning Engineer in Sports Technology – Electronic Arts (EA SPORTS) **[Dec '23 - Present]**

Working with a specialized team within EA SPORTS to advance applications of ML in simulation, animation, character control.

A.I. Scientist III – Electronic Arts (Central Tech) **[Aug '22 – Dec '23]**

Researched, designed, and prototyped novel unsupervised machine learning mapping tools for automated player play-style analysis (patent application submitted).
Researched techniques for more generalized gameplay A.I. (won 'Tools and Platform' division at internal WhatIf technical fair)

A.I. Intern – Electronic Arts **[May '20 – Aug '20], [May '21 – Aug '21], [Sep '21 – Dec '21]**

Researched and published Reinforcement Learning tools for A.I. agents in games to learn and play with different behavior styles.
Developed novel actor-critic training architecture to enable a single actor network to successfully learn multiple play styles, thus significantly reducing runtime cost of using learned controllers.
Developed novel framework for blending reinforcement and imitation learning for play-styles.

Research Assistant – University of Pennsylvania **[Jun '16 – May '17]**

Assisted in design, fabrication, construction, and programming of custom-built quadrotor drones for autonomous flight research.
Researched novel computer vision-based techniques for fast and accurate product identification.

PUBLICATIONS

Ph.D. Dissertation

Mysore, S.; ***Conditioning Behavior Styles of Reinforcement Learning Policies***, Boston University Library, 2022

Publications

Mysore, S., Cheng, G., Zhao, Y., Saenko, K., Wu, M.; ***Multi-Critic Actor Learning: Teaching RL Policies to Act with Style***; *International Conference on Learning Representations, 2022*

Gongora, A.E., **Mysore, S.**, Li, B. et al. ; **Designing Composites with Target Effective Youngs Modulus using Reinforcement Learning**; *ACM Symposium on Computational Fabrication (SCF), 2021, DOI: https://doi.org/10.1145/3485114.3485123*

Mysore, S., Mabsout, B., Mancuso, R., Saenko, K.; **Regularizing Action Policies for Smooth Control with Reinforcement Learning**; *IEEE International Conference on Robotics and Automation, Xi'an, China, 2021*

Mysore, S., Mabsout, B., Mancuso, R., Saenko, K.; **Honey, I Shrunk the Actor: A Case Study on Preserving Performance with Smaller Actors in Actor-Critic RL**; *IEEE Conference on Games*, 2021

Mysore, S., Mabsout, B., Saenko, K., Mancuso, R.; ***How to Train your Quadrotor: A Framework for Consistently Smooth and Responsive Flight Control via Reinforcement Learning***; *ACM Trans. Cyber-Phys. Syst. 5, 4, Article 36 (October 2021), DOI: https://doi.org/10.1145/3466618*

Preprints

Mysore, S., Cheng, G., Zinno, F., Zhao, Y., Saenko, K.; ***Split-Critic Imitation Learning for Balancing Conflicting Imitation and Reinforcement Learning Objectives***; *2023*

Mabsout, B., Roozkhosh, S., **Mysore, S.**, Saenko, K., Mancuso, R.; ***The SwaNNFlight System: On-the-Fly Sim-to-Real Adaptation via Anchored Learning***; *2022*

Mysore, S., Platt, R., Saenko, K.; ***Reward-guided Curriculum for Robust Reinforcement Learning***; *2019*

Workshops

Mysore, S., Platt, R., Saenko, K.; ***Reward-guided Curriculum for Learning Robust Action Policies***;
Workshop on Multi-task and Lifelong Reinforcement Learning at ICML 2019

Academic Research Experience

Improving Practical Robustness and Cross-Domain Policy Transfer in Reinforcement Learning **[Aug '17 – Aug '22]**

Research Advisor Prof. Kate Saenko; Boston University

Investigating techniques to improve the practical utility of learned control policies and to bridge the domain gap when applying them subject to distributional shift, with primary applications in robotics and videogames.

Developed automated bandit-based learning scheme for developing Reinforcement Learning control policies that are robust to perturbations in actor environments (***published at**: MTLRL Workshop ICML '19*).

Developed an intuitive training regularization scheme for generally improving smoothness of learned policies without additional runtime computation while achieving significant reduction in power consumption *(**published at**: ICRA '21).*

Studied and characterized the impact of Actor-Critic deep RL architectures on network architectures and demonstrated that careful tuning of network sizes allow for highly performant yet compact policy networks *(**published at**: CoG '21).*

Developed tools for single controllers to learn multiple tasks or behavior styles – continuation and extension of Summer 2020 internship work done with Electronic Arts *(**published at**: ICLR '22)*.

Developed techniques for blending Reinforcement and Imitation Learning while avoiding conflicts in the optimization criteria.

Neural Network based Flight Control for High Performance Racing Drones **[Apr '19 – Aug '22]**

Collaboration With Prof. Renato Mancuso; Boston University *(Project webpage: https://wfk.io/neuroflight/)*

Designed novel training regime to improve the accuracy and stability of reinforcement learning-based control output in deploying neural network flight-controller models trained in simulation to real drone hardware *(**published at**: TCPS '21, ICRA '21)*

Developed extension for online learning and hot-swapping neural network graph for adjusting control on the fly. (*under review*)

Reinforcement Learning Aided Design **[Mar '18 – Aug '21]**

Collaboration With Profs. Emily Whiting, Keith Brown, Elise Morgan & Wojciech Matusik (MIT); Boston University

Developed a Reinforcement Learning based pipeline to automatically adjust designs of composites to satisfy specific physical properties with 95% success rate and significant runtime speed-up (***published at**: SCF '21*).

Autonomous Flight Guided by Event-Based Camera **[Jul '16 – Apr '17]**

Supervising Professor Prof. Kostas Daniilidis; University of Pennsylvania, GRASP Lab

Developed vision-based quadrotor flight control with using event-based (iniVation DVS) cameras, as part of a group effort.

Designed and built custom quadrotors to serve as testing platforms.

Specific responsibilities: State estimation and controller design; Hardware design, construction and management.

Product ID and Retrieval from Large Catalogues **[Jun '16 – Nov '16]**

Supervising Professor Prof. Kostas Daniilidis; University of Pennsylvania, GRASP Lab

Developed logVLAD - a framework for efficient multi-class product labeling and localization in natural images of shelves in stores, using logarithmically scaled VLAD image encodings, with high precision, better response to feature burstiness, and good energy distribution (*work contributed to Master's Thesis*).

Low Cost Manipulation (LoCoMa) **[Mar '16 – May '17]**

Supervising Professor Prof. Mark Yim; University of Pennsylvania, GRASP Lab – Modular robotics Lab (ModLab)

Helped prototype a novel, low cost, 0-DoF end-effector and manipulation scheme.

Camera Localization **[May '15 – Aug '15]**

Supervising Professor Prof. Jianbo Shi; University of Pennsylvania, GRASP Lab

Implemented real-time image-based camera localization in known 3D spaces using feature matching, visual odometry and SfM.

Investigated application of deep-learning in characterizing the camera motion between sequential images

Developed method and software to interface with existing Vicon hardware infrastructure to track of multiple objects in real-time.

Academic Service

Reviewer for IEEE Robotics and Automation Letters (RAL) **[RAL '23]**

Reviewer for IEEE International Conference on Robot Systems (IROS) **[IROS '22, '23]**

Reviewer for IEEE International Conference on Robotics and Automation (ICRA) **[ICRA '22]**

Reviewer for Conference on Neural Information Processing Systems (NeurIPS) **[NeurIPS '21, '22, '23]**

Reviewer for Conference on Robot Learning (CoRL) **[CoRL '21, '22, '23]**

Reviewer for IEEE Conference on Games (CoG) **[CoG '21, '22, '23, '24]**

Coordinator for Image and Video Computing (IVC) group meetings and website at BU **[Aug '18 – May '20]**

Manage planning duties for IVC's weekly meetings and also manage the group's website

Technical advisor for AI4ALL Summer Program **[May '19 – Aug '19]**

Trained undergraduates involved in the AI4All program at Boston University to develop and execute Machine Learning code

Reviewer for the Machine Learning Journal **[Jun '19]**

Teaching

Boston University

Guest Lecturer – CS 542 Machine Learning (Fall '21)

Teaching Assistant – CS 542 Machine Learning (Spring '22)

Grader – CS 581 Computational Fabrication (Spring '22), CS 480/680 Introduction to Computer Graphics (Fall '19, Fall '20), CS 542 Machine Learning (Fall '18, Spring '20), CS 440 Artificial Intelligence (Fall '20)

University of Pennsylvania

Teaching Assistant – CIS 581 Computer Vision and Computational Photography (Fall '15)

Invited Talks

Electronic Arts ML For Animation **[Feb '23]**

Electronic Arts Research Deep Dive **[Jul '21]**

Boston University AI Research (AIR) Initiative **[Apr '21]**

Electronic Arts ML Special Interest Group **[Aug '20]**

Boston University Image and Video Computing (IVC) Student Talk **[Aug '20]**

Technical Reports & Selected Coursework

University of Pennsylvania

Master's Thesis – logVLAD: A Novel Pipeline for Image Retrieval (see notes on Product ID research for details)

Stereo Visual Odometry
Implemented and compared different approaches to stereo visual odometry against the KITTI vision benchmark suite. Implementations computed odometry by solving the 3D-3D affine Procrustes problem, by solving the 3D-2D Perspective-n-Point (PnP) problem, and by using optical flow.

Learning Path-planning
Extracted feature-maps from a satellite-view map using Gaussian Mixture Models (GMMs) trained to recognize colors, and used them to build a cost-map over which Dijkstra's and A* algorithms were applied to determine the best traversable path between arbitrary start and goal set on the map.

Gesture Recognition
Applied Hidden Markov Models towards learning to recognize hand gestures utilizing inputs from an Inertial Measurement Unit (IMU) attached to the arm.

Autonomous Quad-rotor Flight
Implemented programs to facilitate autonomous path-planning and flight control in real-time on research- and consumer-grade quad-rotor drones.

Reconstructing 3D Scenes from Image Sequences
Applied concepts of Structure from Motion (SfM) and epipolar geometry to build sparse 3D scene reconstructions from a sequence of images.

Face Replacement in Images
Developed a program that attempts to seamlessly replace faces detected in an image with some other face, with skin-tone blending and masking to account for face rotations.

Light-Painting with a Robot Arm
Modeled DH-parameters and Inverse Kinematics of a robot arm to plan and execute the 'painting' of a picture with an LED mounted at the arm's tool-tip, captured on a long-exposure image.

Machine Learning applied to Pricing Estimation
Utilized several machine learning techniques attempts to estimate a price range for real-estate listings using bags of words provided to describe each listing.

The University of Nottingham

4th-Year Group Design Project – *Automated Search and Rescue Robot with 3D Environment Mapping*
Supervising Professor Prof. Kevin Lee
Specific area of focus: Object Detection and Recognition for robot guidance.
Constructed a mobile robotic platform which used computer-vision-guided navigation and arm control for collecting rescue pre-determined rescue targets.

Language Skills

English (1st Language) – Proficient; TOEFL Score 120 (Tested Oct '13)

German – Certified to the A2 level (Aug '11); Studied at the B1 level (Sep '13 – Apr '14)

Kannada – Basic conversational ability