Karkala Shashank Hegde

@ khegde@usc.edu | 🔗 LinkedIn | 🔗 GitHub | 🌐 Website | 🎓 Scholar

Ph.D. Student with an interest in AI, Reinforcement Learning, and Deep Learning

Programming Skills

Languages:
Python (Tensorflow, PyBullet, Mujoco, Gym, PyTorch, Pandas, Numpy, Flask, Scikit-learn, Scipy, ROSpy),
MATLAB (Statistics and Machine Learning, Deep Learning, Signal Processing Toolboxes),
C++ (OpenAL, OpenCV)

Education

University of Southern California   Los Angeles, USA
PhD in Electrical and Computer Engineering (AI & Robotics); GPA: 3.94/4   *2021 – Present*
Minor Degree in Computer Science;   Advisor: Gaurav Sukhatme

University of Southern California   Los Angeles, USA
MS in Electrical and Computer Engineering (Machine Learning); GPA: 3.94/4   *2019 – 2021*

National Institute of Technology Karnataka   Surathkal, India
BTech in Electrical and Electronics Engineering (Signal Processing); GPA: 8.17/10   *2013 – 2017*
Thesis GPA: 9.5/10;   Advisor: Krishnan CMC

Work Experience

SalesDNA   Los Angeles, USA
Data Scientist Intern   *May 2021 – August 2021, Full-time*

- Built data pipelines for collection, cleaning, and real-time Markov decision process modeling of sales processes.
- Built model-free reinforcement learning algorithms to build AI strategies on this sales simulation.

Fidelity Investments: Asset Management Technology   Bangalore, India
Data Scientist   *July 2017 – July 2019, Full-time*

- Built a simulator using real trading data and trained an RL agent for portfolio construction in equity trading.
- Worked with the Equity Trading team to develop backend services with Java spring-boot, Python Flask, SQL, and Splunk.

Fidelity Investments: Asset Management Technology   Bangalore, India
Summer intern as Software Developer   *May 2016 – July 2016, Full-time*

- Worked with the fixed-income research team to build a complete end-to-end application using .NET and Excel VBA. Gained experience in the Microsoft Windows Presentation framework for building hard clients.

Publications

Hegde, S., Huang, Z., and Sukhatme, G.S., 2023. HyperPPO: A scalable method for finding small policies for robotic control. arXiv preprint arXiv:2309.16663.(Submitted to ICRA 2024)[site]

Hegde, S., Batra, S., Zentner, K.R. and Sukhatme, G.S., 2023. Generating Behaviorally Diverse Policies with Latent Diffusion Models. arXiv preprint arXiv:2305.18738. (Accepted at NeurIPS 2023)[site]

Hegde, S. and Sukhatme, G.S., 2023, May. Efficiently Learning Small Policies for Locomotion and Manipulation. In 2023 IEEE International Conference on Robotics and Automation (ICRA 2023) (pp. 5909-5915). IEEE.[site]

G. Salhotra, **S. Hegde**, SS. Batra, P. Englert, GS. Sukhatme (2022) *Guided Learning of Robust Hurdling Policies with Curricular Trajectory Optimization, Southern California Robotics Symposium*[site]

S. Hegde, Kanervisto, A., & Petrenko, A. (2021, August). Agents that listen: High-throughput reinforcement learning with multiple sensory systems. In 2021 IEEE Conference on Games (CoG) (pp. 1-5). IEEE. [site]

Hiteshi Sharma, R. Jain. and **S. Hegde** (2021) *Randomized Policy Learning for Continuous State and Action MDPs* (arxiv)
[pdf]

S. Hegde, V. Kumar, and A. Singh. (2018). *Risk aware portfolio construction using deep deterministic policy gradients*. IEEE Symposium Series on Computational Intelligence (SSCI) Bangalore, Nov. 2018. [pdf]

Singh, A., Kumar, V., and **Hegde, S.** (2018). *Reinforcement Learning: A Review in Finance*. Proceedings of Data Science Congress, Mumbai, May 2018. [pdf]

Severes, A., **Hegde, S.**, D'Souza, L. and Hegde, S. (2017). *Use of LED for enhanced lipid production in micro-algae based bio-fuels and predicting growth patterns*. Journal of Photochemistry and Photobiology B: Biology, Elsevier, Volume 170, Pages 235-240. [link]

RESEARCH EXPERIENCE

Research assistant -PhD   Los Angeles, USA
Robotic Embedded Systems Laboratory[link], USC   *Sept 2020 - present*

- Used CLIP language encoder with Latent Diffusion Models and Graph Hypernetworks for generative modeling in behavior space for robotic control.

- Develop and train sample efficient distributed learning methods for language-conditioned robotic control on SLURM.

- Create high-performing small Neural Networks on AWS EC2 instances for quadrotor flight control.

- Experiment with audio-based communication between agents with multi-agent reinforcement learning for game AI

Research assistant - MS   Los Angeles, USA
Stochastic Systems & Learning Lab[link], Dynamic Robotics & Control Lab[link]   *Nov 2019 - May 2021*

- Build scale-able Reinforcement Learning policies using function approximators with lesser trainable parameters.

- Study and Apply state-of-the-art Imitation Learning techniques to self-driving vehicles and experiment on Hyper realistic simulations such as CARLA.

- Simulate and control a quadruped mini cheetah robot on Pybullet and Gazebo using stochastic control with policy gradient-based agents. Test the RL controller on the actual robot after integration with ROS.

- Experiment on different action spaces such as impedance, torque, and force control, and use hybrid learning methods with model predictive control to help faster learning. Use RLLib for distributed learning.

PROJECTS

Automatic paper tagging: Apply BERT sentence transformer to encode abstracts of hundreds of papers, and then find cosine similarity of the encoding with that of topic definitions to rank and tag them

Autonomous Vehicle Navigation: As a part of the Autonomous Vehicle lab, I worked on navigation, path planning, and simulation of an autonomous car to take part in IGVC 2021. I used Gazebo to build a simulation of the track, and implement path-finding algorithms such as A star.

Competetive and Co-operative Multi-Agent Reinforcement Learning[pdf]: As a part of my directed research with the Hardware Accelerated Learning group[link], I experiment with various multi-agent reinforcement learning algorithms. The goal of this project is to understand the state-of-the-art RL algorithms that work well in both competitive and cooperative environments.

Torque Transfer[code]: Use reinforcement learning and transfer learning to create robust AI agents. The AI agent should generalize to a variety of open-world self-driving simulations. After training an AI for a self-driving car simulation using Imitation learning and reinforcement learning, the learned policy was used as a pre-trained network for an AI agent in another self-driving simulation. The pre-trained model showed faster learning in the new simulation.

Emotion Transfer on speech using spectrogram images[code]: Use a conditional Generative Adversarial Neural Network to generate images on spectrograms of speech signals. By using cycle GANs we use style transfer on spectrograms of speech signals to embed emotion in them. The generated spectrogram is reconstructed back to speech using the Griffin-Lim algorithm.

Fashion compatibility prediction[code]: Use a Siamese Convolutional Neural Network to classify if two fashion objects are compatible. Then the pair-wise similarity scores are predicted to see if an outfit is compatible. To do this Google Tensorflow 2.0 was used and the models were trained on AWS p3.2xlarge instances (Tesla V100 GPUs)

Spoken Language classifier[code]: Implement a Gated Recurrent Unit-based Neural Network to classify the extracted MFCC features from speech audio. A streaming model classifies the language being spoken in real-time. Using this streaming model, we could analyze the probability of miss-classification at every instant of speech.

Prosthetic Voice (Thesis)[pdf]: Undergraduate Thesis: sEMG signal controlled speech production aid for speech-challenged individuals using Machine Learning. The signals were collected, filtered, pre-processed, and then fed to a classifier that would predict the hand action performed. The action would then be translated into speech.

Emotion Detection[pdf]: I was part of a three-member team that built a Machine Learning-driven emotion detector using variations in speech signals. We built a classifier using MFCC feature extraction and PCA on other features.

Teaching Experience

University of Southern California   Los Angeles, USA
Teaching Assistant   *2021 – 2023*
- CSCI567 – Machine Learning
- EE541 – A Computational Introduction to Deep Learning
- EE641 – Deep Learning Systems

Awards & Achievements

USC Annenberg Fellow: Awarded a Fellowship for my PhD.

Masters Student Honors Program[link]: Certificate for outstanding academic and research achievements.

The Data Open[pdf]: Was a finalist in the SoCal round of the Data Open Hackathon organized by Citadel.

Soda bottle classification contest[link]: Winner of image classification contest hosted by Deep Cognition.

Best Outgoing student in school, ranked in top 1% of All India Engineering exam.

Scholarship, National Level Science Talent Search Examination.

Talks And Presentations

Invited Speaker at Sears Deep Learning Center Program: Delivered a company-wide talk on SOTA applied Deep Reinforcement Learning. [link]

Guest lecture on Generative models for robotics: Presented a talk to a graduate Machine Learning class.

Career guidance talk: Presented to all class 12 students of Mangalore (a small town in Southern India) for career options in machine learning.

Service and Extracurricular

Reviewer: ICRA 2023, IROS 2023.

Organizer: USC robotics seminar.

Organizer: Machine Learning workshops at Fidelity Investments.

Secretary: Music club at National Institute of Technology Karnataka.

Band Manager and Lead Guitarist: Carnatic metal fusion band TATVAM.

Volunteer: SPIC MACAY Mangalore chapter.