

SWAMINATHAN GURUMURTHY

PhD Student At The Robotics Institute

(+1) 650-450-7962
gauthamsindia95@gmail.com
<https://github.com/swami1995>

EDUCATION

- Robotics Institute, Carnegie Mellon University** **September 2017 – December 2019**
- M.S in Robotics Research (**GPA: 4.09/4.33**)
- Indian Institute of Technology (BHU), Varanasi** **July 2013 – May 2017**
- B. Tech. in Electrical Engineering (**GPA: 8.99/10**)

SKILLS AND INTERESTS

- Languages: Python, C++, MATLAB
- Libraries and Packages: Pytorch, Tensorflow, Theano, OpenCV
- Areas of Interest: Deep Learning, Reinforcement Learning, Robotics, Visual SLAM, Computer Vision, NLP

PUBLICATIONS

Learning exploration Policies for Model Agnostic Meta-RL [[github](#)] [[paper](#)] **Spotlight/Poster to CoRL 2019**
Swaminathan Gurumurthy, Sumit Kumar, Katia Sycara

- We propose to explicitly model a separate exploration policy for the task distribution in Meta-RL given the requirements on sample efficiency. Having two different policies gives more flexibility during training and makes adaptation to any specific task easier. We show that using self-supervised or supervised learning objectives for adaptation stabilizes the training process and improves performance. (Oral at AMTL 2019 (ICML Workshop))

**Spotlight/Poster at WACV
2019**

3D Point Cloud Completion using Latent Optimization in GANs [[paper](#)]
Shubham Agarwal*, Swaminathan Gurumurthy* (equal contributions ordered alphabetically)

- In this project we address a fundamental problem with Neural Network based point cloud completion methods which reconstruct the entire structure rather than preserving the points already provided as input. These methods struggle when tested on unseen deformities. We address this problem by introducing a GAN based Latent optimization procedure to perform output constrained optimization using the regions provided in the input.

Community Regularization of Visually-Grounded Dialog [[github](#)] [[paper](#)] **Oral at AAMAS 2019**
Akshat Agarwal*, Swaminathan Gurumurthy*, Vasu Sharma*, Katia Sycara, Michael Lewis (equal contribution ordered alphabet)

- Here we aim to train 2 agents on the visual dialogue dataset where one agent is given access to an image and the other agent is tasked with guessing the contents of the image by establishing a dialogue with the first agent. The two agents are initially trained using supervision followed by Reinforce. In order to combat the resulting drift from natural language when training with Reinforce, we introduce a community regularization scheme of training a population of agents. (Oral also at ALA 2018 (FAIM Workshop))

Exploiting Data and Human Knowledge for Predicting Wildlife Poaching [[github](#)] [[paper](#)] **Oral/Poster - ACM COMPASS 2018**
Swaminathan Gurumurthy, Lantao Yu, Chenyan Zhang, Yongchao Jin, Weiping Li, Haidong Zhang, Fei Fang

- Using past data of traps/snares found in a wildlife Sanctuary, we predict the regions of high probability of traps/snares to guide the rangers to patrol those regions. We use novel frameworks of incorporating expert domain knowledge for the dynamic sampling of data points in order to tackle the imbalance in data. We further use these regions to produce optimal patrol routes for the rangers. This has now been deployed in a conservation area in China (Oral also at NeurIPS AISG 2018)

DeLiGAN: GANs for Diverse and Limited Data [[github](#)] [[paper](#)] **Poster at CVPR 2017**
Swaminathan Gurumurthy*, Ravi Kiran S.* and R. Venkatesh Babu

- In the paper, we try to explore the idea of finding high probability regions in the latent space of GANs by learning a latent space representation using learnable Mixture of Gaussians. This enables the GAN to model a multimodal distribution and stabilizes training as observed visually and by the intra-class variance measured using a modified inception score. Our modification is especially useful when the dataset is very small and diverse.

INTERNSHIPS

Differentiable-ilqr for Imitation learning

NURO

Jan 2020 - May 2020

- In this project I'm building a differentiable ilqr pipeline ground up in C++ without any autodiff packages. I will be using it for improving the imitation learning modules. Can't reveal details due to confidentiality agreements.

Off-on policy learning

Montreal Institute for Learning Algorithms

Aug-Nov 2018

- On policy methods are known to exhibit stable behavior and off-policy methods are known to be sample efficient. The goal here was to get the best of both worlds. We first developed a self-imitation based method to learn from a diverse set of exploratory policies which perform coordinated exploration. We then use a meta-learning objective to ensure that the off-policy updates to the policies are aligned with future on-policy updates. This leads to more stable training but fails to reach peak performance in most continuous control tasks we tested on.

Visual SLAM based SfM for Boreholes [\[github\]](#)

Australian Centre for Vision Technologies

May-July 2015

- Built a package to construct a sparse map and camera trajectory using SIFT features, fine-tuned using bundle adjustment and loop closure. It was tailored for boreholes and underground scenes with forward motion, where most of the current state of the art approaches like LSD SLAM, ORB SLAM and SVO struggled at both localization and mapping.

OTHER INTERESTING PROJECTS

Virtual assistants for human teams and Human-Machine teaming

Prof. Katia Sycara

Sep2019-Dec2019
Robotics Institute, CMU

- With this project, I was experimenting with a different approach towards research. This was a big project with 6-7 students working on different parts. My role was to help each one of them with their individual parts and also coordinate the different parts so that they can sync and integrate reasonably when they finish. The individual parts back then were: single agent meta-imitation; training populations of diverse teams; multi-agent/team meta-imitation; hierarchical imitation learning with sparse rewards; online test-time regret minimization for out-of-distribution observations; meta-imitation without the task labels; designing a semantic action space for collecting human data. And more..

Exploring interpretability in Atari Games for RL policies using Counterfactuals

Akshat Agarwal, Prof. Katia Sycara

May 2018 – July 2018
Robotics Institute, CMU

- In this project we aimed to understand what RL agents learn in simple games such as in Atari. We developed a GAN based method to find counterfactuals for the policies, i.e., we find small perturbations in the scene that can lead to changes in the agent action and use these to interpret agent behavior. GAN in this case is used to avoid adversarial examples and produce semantically meaningful perturbations.

Query Efficient Black Box Attacks in Neural Networks

Prof. Fei Fang
Prof. Martial Hebert

Nov 2017 – Feb 2018
Robotics Institute, CMU

- We test various methods to increase the sample efficiency of adversarial black box attacks on Neural nets. In one of the methods, we analyze the transferability of gradients and find that it has two components: Network specific components and Task specific components. The task specific component corresponds to the transferable properties of adversarial examples between architectures. Hence, we attempted to isolate this component and enhance the transfer properties. We then perform multiple queries on the black box network to obtain the architecture specific components using ES.

VOLUNTEERING ACTIVITIES

- **Teaching Assistant** : MA101: Mathematics (Calculus) : Prof. Tanmoy Som, IIT-BHU
: CSO101: Computer Programming : Prof. Anil Kumar Singh, IIT-BHU
Spring 2016
Fall 2016
- **Conference Reviewing** : AI For Social Good Workshop, NeurIPS 2018
: 3D-WiDGET, Workshop on Deep Generative Models for 3D understanding, CVPR 2019
- General Secretary - IEEE Student Branch IIT BHU during the session 2014-15.
- I organized a series of workshops on deep learning at IIT-(BHU) as part of our technical fest to expose students to DL.