ARPIT GARG

iarpitgarg@gmail.com | (505)-389-9495 | github.com/kirarpit

PUBLICATIONS

[1] **Arpit Garg**, Yazied A. Hasan, Adam Yañez and Lydia Tapia, "Defensive Escort Teams via Multi-Agent Deep Reinforcement Learning," In *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, under review, 2020

[2] **Arpit Garg**, Hao-Tien Lewis Chiang, Satomi Sugaya, Aleksandra Faust and Lydia Tapia, "Comparison of Deep Reinforcement Learning Policies to Formal Methods for Moving Obstacle Avoidance," In *Proceedings of IEEE International Conference on Intelligent Robots and Systems (IROS)*, Macau, China, Nov. 2019

PROJECTS

Cooperative Escorts via Multi-Agent Deep Reinforcement Learning (RL), Tapia Lab June - Sept 2019

- Trained multiple agents with partial observations to coordinate and safely escort a payload through environments with interacting moving obstacles (Video: https://youtu.be/SoYesKti4VA)
- Used convolutional neural networks (CNN) to infer geometry of the objects from raw LiDAR observations
- Simulated social force model for obstacle motion dynamics and wrote libraries for collision detection
- Implemented a distributed training system in Python using RLlib that scaled the training from 1 worker to 100 workers and 4 GPUs thus reducing the training time by 26x and collecting 100M samples in 24hrs

Custom Pseudo-C Compiler, UNM Jan - April 2018

- Built a large-scale standalone program, written in C++, that compiles a legal LOBO-C (a custom C like language) code block and generates MIPS assembly code
- The entire code is written in C++ and can be found at https://github.com/kirarpit/compiler-construction

Moving Obstacle Avoidance with Deep RL, Tapia Lab Nov 2018 - March 2019

- Implemented and trained a deep RL actor-critic A3C policy for robot navigation that outperformed a previous state-of-the-art method
- Performed a comprehensive comparison between the trained policy and a formal optimal method to gain insights on potential failure points of RL policy
- Used Distributed TensorFlow to implement and scale A3C thus increasing the training speed from 15steps/sec on a MacBook Pro to 1300steps/sec on 4 nodes with 16 CPU cores each

Solving Board Games Using Deep RL, Personal Project July - August 2018

- Implemented RL algorithms: Q-Learning (DQN), Policy Gradient, Actor-Critic, and AlphaGo Zero and solved the games Connect-4 and Tic-Tac-Toe on multiple board sizes using TensorFlow
- The code is written in Python and can be found at https://github.com/kirarpit/connect4

Web Push Notifications, MySmartPrice 2016

- Built and scaled a system that captured web-push notification subscriptions using JQuery, to be able to send notifications to over 650,000 daily active users and analyze their interactions with the website

EDUCATION

The University of New Mexico (UNM) New Mexico, USA
Master of Science in Computer Science; GPA 4.1/4.0 Jan 2018 – Dec 2019

Indian Institute of Technology, Roorkee Roorkee, India
Bachelor of Technology (B.Tech) in Civil Engineering July 2010 – May 2014

WORK EXPERIENCE

Research Assistant July 2018 - Present
Tapia Lab, UNM Albuquerque, USA

- Studied and implemented state-of-the-art Deep RL algorithms to tackle problems in robotic motion planning and systems and human-AI collaboration

Senior Product Engineer June 2014 – July 2017
MySmartPrice Web Technologies Pvt. Ltd. Hyderabad, India

- Contributed to the company's PHP codebase architecture and development, website security, and system administration as a back-end developer
- Designed and implemented business logic and data storage solutions at the back end

SKILLS

- Related Skill Sets: Deep RL, Motion Planning, Deep Learning, Back-end Web Development, Linux
- Programming Languages: Python, C++, PHP
- Software Packages – RLlib, TensorFlow, GNU Parallel, MySql, Nginx, Redis