

PRATYUSH KAR

Graduation date: May 2019

✉ pkar@cs.utexas.edu · ☎ (+1) 512-496-5335 · [in pratyushkar](https://www.linkedin.com/in/pratyushkar) · [p-kar](https://github.com/p-kar)

🎓 EDUCATION

The University of Texas at Austin, USA Aug, 2017 – Present

M.S. in Computer Science **GPA: 4.0/4.0**

Birla Institute of Technology & Science, Pilani, India July, 2013 – May, 2017

B.E. (Hons.) in Computer Science **CGPA: 9.4/10**

👛 WORK EXPERIENCE

Nvidia Inc., Santa Clara, California May, 2018 – Aug, 2018

Deep Learning Intern

- Implemented the YOLOv3 object detection pipeline in TensorRT for ~70% faster inference (code open source on [Github](#))
- Using attention transfer, trained and successfully pruned YOLOv3 architecture by ~40% with <1% drop in mean average precision

Qualcomm India Pvt. Ltd., Hyderabad, India May, 2016 – July, 2016

Software Engineering Intern

- Developed a parser and command sequencer (in C++ and Python) for running commands present in the log files to simulate a voice call on the Hexagon DSP simulator

🔗 PROJECTS

Style Transfer in Text using WGANs Mar, 2018 – May, 2018

Advisor: Prof. Ray Mooney

- Used adversarial training to overcome distribution mismatch problem arising due to MLE training of seq2seq models
- Implemented Wasserstein GANs trained with cyclic consistency loss for text style transfer between formal and informal English

Gossip-based Broadcast for Cluster Membership Apr, 2018 – May, 2018

Advisor: Prof. Vijay Chidambaram

- Implemented a rumor mongering broadcast solution for Rapid, a scalable distributed membership service
- Improved network utilization by ~12% in disseminating failures during complex network failures (tested on ~1000 VMs)

Distributed Key-Value Store with Eventual Consistency Feb, 2018 – Mar, 2018

Advisor: Prof. Vijay Chidambaram

- Implemented a Distributed Key-Value store with monotonic reads and read-your-writes session guarantees in Go. The application was able to respond to ~600 requests/sec

Actor-Critic Methods with Selective Transfer for ViZDoom Nov, 2017 – Dec, 2017

Advisor: Prof. Risto Miikkulainen

- Implemented a synchronous variant of the A3C algorithm with an attention network (in PyTorch) to play the first-person shooter game, Doom, with only the video frame input
- Improved the performance of the vanilla model by selectively transferring knowledge from simpler tasks to a complex battle scenario

⚙️ SKILLS

- Languages: C++, Python, Go, Java
- Frameworks: PyTorch, Docker, Spark, Hadoop, AWS
- Courses: Neural Networks, NLP, Deep Learning, Distributed Computing, Virtualization