

Tanmay Gangwani

Homepage : <https://tgangwani.github.io>

Email : gangwan2@illinois.edu

Github : <https://github.com/tgangwani>

EDUCATION

University of Illinois, Urbana-Champaign

- *Ph.D. in Computer Science; Advisor: Dr. Jian Peng* *Aug. 2016 –*
 - I'm interested in machine learning, especially Reinforcement Learning. My research is mainly focused on designing algorithms which efficiently leverage expert demonstrations for RL (imitation learning), address the exploration challenge in complex environment, and use generative modeling methods for model-based RL.

University of Illinois, Urbana-Champaign

- *Master of Science in ECE; GPA: 3.93; Advisor: Dr. Josep Torrellas* *Aug. 2013 – Aug. 2016*
 - My Masters research was at the intersection of Computer Architecture, Compilers and Systems. We proposed hardware-software co-design methods for efficient execution of graph-based and lock-free parallel programs.

PUBLICATIONS (PLEASE SEE WEBSITE FOR CODE AND PAPER LINKS)

- [[Under review](#)]. New paper on imitation learning under submission at ICLR 2020. Details and link omitted to preserve anonymity.
- **Learning Belief Representations for Imitation Learning in POMDPs.** *Tanmay Gangwani, Joel Lehman, Qiang Liu, Jian Peng*; The Conference on Uncertainty in Artificial Intelligence, UAI 2019. We propose a framework for imitation of expert demonstrations via adversarial learning, in partially observable environments. We highlight the importance of learning rich representations from incomplete data.
- **Learning Self-Imitating Diverse Policies.** *Tanmay Gangwani, Qiang Liu, Jian Peng*; International Conference on Learning Representations, ICLR 2019. Policy optimization for deep-RL by exploiting past good experiences of the agent. A discriminator is trained to produce shaped rewards, which provide a strong learning signal, especially for sparse-reward environments. Stein-variational policy gradient with Jensen-Shannon kernel used for population-based exploration.
- **Policy Optimization by Genetic Distillation.** *Tanmay Gangwani, Jian Peng*; International Conference on Learning Representations, ICLR 2018. Policy optimization algorithm for continuous control (MuJoCo) tasks in deep-RL, inspired by evolutionary computing approaches such as genetic algorithms. We use imitation learning to do crossover between two policies parameterized by deep neural networks. Policies in the ensemble also go through mutation and selection operators.
- **Distributed and Secure ML using Self-tallying Multi-party Aggregation.** *Yunhui Long**, *Tanmay Gangwani**, *Haris Mughees and Carl Gunter*; NeurIPS'18 workshop on Privacy Preserving Machine Learning (PPML). We train ML models in a distributed setting, while providing input data confidentiality to each data contributor using constructions based on cryptographic primitives like zero-knowledge proofs.
- **CASPAR: Breaking Serialization in Lock-Free Multicore Synchronization.** *Tanmay Gangwani, Adam Morrison, and Josep Torrellas*; Architectural Support for Programming Languages and Operating Systems, ASPLOS 2016. Hardware support to improve performance of lock-free parallel algorithms which use x86 compare-and-swap atomics.

INDUSTRY EXPERIENCE

D-Wave Systems Inc.

- *Summer Internship* *Vancouver, BC*
May 2019 - Aug 2019
 - **Model-based RL for Industrial Process Control.** Worked on model-based RL methods for optimizing control in industrial-process applications such as chemical and manufacturing plants. We investigated techniques for learning robust neural network models of the plant dynamics, which are accurate in long-horizon predictions, and use them for learning policies in the batch-RL setting. Mentored by Jason Rolfe and William Macready.

- **Uber AI Labs** San Francisco, CA
Summer Internship *May 2018 - Aug 2018*
 - **Imitation Learning in POMDPs.** Proposed methods for imitation learning in partially-observable RL environment using adversarial methods. Work culminated in a research paper (published at UAI). Mentored by Joel Lehman and Kenneth Stanley.
- **Advanced Micro Devices** Bellevue, WA
Summer Internship *May 2015 - Aug 2015*
 - **Hardware-software Co-optimization.** Explored techniques for improving the efficiency of irregular, unbalanced OpenCL code segments on GPUs. Added new hardware modules to the AMD GEM5 (C++) simulator.
- **Intel Corporation** Bangalore, India
Software Engineer (Full-time) *Aug 2011 - Aug 2013*
 - **High Performance Computing.** Involved in the development and testing of Xeon Phi (Knights Corner, Knights Landing) accelerator chips, as a member of the HPC group at Intel. Also worked on performance-analysis and code-optimization for scientific applications such as NAMD and LBM.

ACADEMIC DUTIES

- **Conference Reviewer.** ICLR 2019, ICLR 2020
- **Teaching Assistant (@UIUC).** Online Learning and Bandits, Applied Machine Learning, Introduction to Deep Learning, Compiler Construction, and Computer Architecture