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 – Aug. 2021*
  - My thesis focused on applied algorithms for deep reinforcement learning (RL). We studied approaches that utilize expert demonstrations for RL (imitation learning), address the issues of exploration and sparse environmental rewards, and improve sample efficiency with the transfer-RL and meta-RL paradigms.

- **University of Illinois, Urbana-Champaign**
  Master of Science in ECE; GPA: 3.92; Advisor: Dr. Josep Torrellas — *Aug. 2013 – Aug. 2016*
  - My Masters research was at the intersection of computer architecture and compilers. We proposed hardware-software co-design methods for efficient execution of graph-based and lock-free parallel programs.

- **Indian Institute of Technology, Kanpur**
  Bachelor of Technology (B.Tech) in Electrical Engineering; GPA: 9.5/10 — *Aug. 2007 – Aug. 2011*
  - My Bachelors research was on designing a hardware-level debugger for the AVR Atmega Micro-Controller using the JTAG protocol. We mapped the RTL-based debugger onto an FPGA fabric for speed and efficiency.

Publications (Please see website for code and paper links)

- **Imitation Learning from Observations under Transition Model Disparity**. *Tanmay Gangwani, Yuan Zhou, Jian Peng*; International Conference on Learning Representations, ICLR 2022. We present an algorithm for imitation learning when the expert and the imitator operate in MDPs with dissimilar transition dynamics. A greedy policy is learned in the imitator MDP using a similarity metric, and then deployed as a surrogate expert to further train an imitator that achieves the desirable long-horizon performance.

- **Hindsight Foresight Relabeling for Meta-Reinforcement Learning**. Michael Wan, Jian Peng, *Tanmay Gangwani*; International Conference on Learning Representations, ICLR 2022. Inspired by hindsight replay methods, we propose a trajectory relabeling approach that improves the sample efficiency of the meta-train phase of baseline off-policy meta-RL algorithms.

- **Learning Guidance Rewards with Trajectory-space Smoothing**. *Tanmay Gangwani, Yuan Zhou, Jian Peng*; Conference on Neural Information Processing Systems, NeurIPS 2020. Credit-assignment is challenging in environments with sparse, end-of-episode rewards. We derive surrogate rewards that provide dense supervision, are computed without auxiliary networks, and integrate easily with existing RL algorithms.

- **Harnessing Distribution Ratio Estimators for Learning Agents with Quality and Diversity**. *Tanmay Gangwani, Jian Peng, Yuan Zhou*; Conference on Robot Learning, CoRL 2020. We devise methods for training agents that achieve high task returns, while being behaviorally diverse. Our algorithm necessitates efficient density ratio estimation. We utilize noise-contrastive estimation and the DICE estimators for the same.

- **State-only Imitation with Transition Dynamics Mismatch**. *Tanmay Gangwani, Jian Peng*; International Conference on Learning Representations, ICLR 2020. We consider the imitation learning setting where the MDPs of the expert and the imitator differ in the specification of the transition dynamics function. Based on recent methods for adversarial inverse-RL, our algorithm addresses the dynamics shift by training the policy to imitate its own past trajectories, selected based on similarity to expert demonstrations in terms of state visitations.

- **Mutual Information Based Knowledge Transfer Under State-Action Dimension Mismatch**. Michael Wan, *Tanmay Gangwani, Jian Peng*; The Conference on Uncertainty in Artificial Intelligence, UAI 2020. We present an algorithm for knowledge-transfer in reinforcement learning where the teacher (expert) and the student (learner) agents may have arbitrarily different state- and action-spaces.

- **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 is 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.
- **Architectural Support for Relaxed Concurrent Priority Queueing in Chip Multiprocessors**. *Azin Heidarshenas*, Tanmay Gangwani*, Serif Yesil, Adam Morrison and Josep Torrellas*; International Conference on Supercomputing, ICS 2020. We propose a distributed hardware architecture for accelerating priority-queue-based applications on multicore processors. Our protocol affords scalability by dynamically adjusting the priority-queue relaxation in hardware.
- **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.** — Vancouver, BC
  Summer Internship — *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 offline-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 2019-2021, NeurIPS 2020, AISTATS 2021, ICML 2021, NeurIPS 2021, CoRL 2021
- **Conference Volunteer**. ICML 2020, ICML 2021
- **Teaching Assistant (@UIUC)**. Online Learning and Bandits, Applied Machine Learning, Introduction to Deep Learning, Compiler Construction, and Computer Architecture