Ajay Jain

Berkeley, CA · US Citizen
ajayj@berkeley.edu · ajayjain.net · (408) 252-9390

Education

University of California, Berkeley — BERKELEY
Ph.D. in Computer Science — Jun 2019 – May 2025 (expected)

- Advised by Prof. Pieter Abbeel and Prof. Joseph E. Gonzalez.
- Interests in unsupervised learning, generative models, combinatorial optimization, compilers.
- Awarded NSF Graduate Research Fellowship.

Massachusetts Institute of Technology — CAMBRIDGE
B.S. in Computer Science and Engineering (Course 6-3) — Aug 2016 – May 2019

- Advised by Prof. Saman Amarasinghe.
- Activities: Machine Intelligence Community, HackMIT.
- *Undergraduate GPA:* 5.00/5.00

Experience

Berkeley Artificial Intelligence Research — BERKELEY
Research Assistant, Robot Learning Lab — Jun 2019 – Aug 2019

- Extended convolutional autoregressive models for inpainting, ensembling (in submission).
- Proposed error correcting sparse graphical memory for visual navigation (in submission).
- ML practitioners simplify their models to fit within limited high-bandwidth GPU memory, but compute is abundant. Published Checkmate training system for optimal memory/recomputation tradeoff at Conference on Machine Learning and Systems (MLSys) 2020.

Uber ATG — TORONTO
Research Intern, Toronto R&D team — Jun 2018 – Jan 2019

- Advised by Prof. Raquel Urtasun.
- Accurate forecasts of vehicle & pedestrian behaviors are critical for safe self-driving. Published Discrete Residual Flow, a deep, tractable generative model that predicts future actor behavior at Conference on Robot Learning (CoRL) 2019. Our trajectory joint distribution is efficiently marginalizable, allowing uncertainty-aware, cost-based planning, and multi-modal.

MIT CSAIL — CAMBRIDGE
Research Assistant, Compilers @ MIT group — Mar 2018 – May 2019

- Published compiler autovectorization technique at Compiler Construction (CC) 2019. Our LLVM pass automatically upgrades hand-vectorized code to higher vector widths and new architectures for performance & portability.
- How to learn with combinatorial output constraints? Presented learned automatic code scheduler, POCSNet, at ML for Systems @ ISCA 2019.

Facebook — MENLO PARK
Software Engineering Intern, Applied Machine Learning team — May 2017 – Sep 2017

- Trained fast facial expression recognition models for core FB app mobile videos.

Kensho Technologies CAMBRIDGE
Software Engineering Intern Jan 2017 – Feb 2017

- Worked on named entity recognition and knowledge graph construction problems. Added regions to graph database to match hundreds of thousands of entity instances in news articles.

Juniper Networks SUNNYVALE
Software Engineering Intern Jun 2016 – Aug 2016

- Built support ticket routing system with many-class text classification pipeline.

Papers

Conference publications

[*] Denotes equal contribution

MLSys 2020 Paras Jain[*], Ajay Jain[*], Aniruddha Nrusimha, Amir Gholami, Pieter Abbeel, Kurt Keutzer, Ion Stoica, Joseph E. Gonzalez. Checkmate: Breaking the Memory Wall with Optimal Tensor Rematerialization. *Conference on Machine Learning and Systems*, 2020.

CoRL 2019 Ajay Jain[*], Sergio Casas Romero[*], Renjie Liao[*], Yuwen Xiong[*], Song Feng, Sean Segal, Raquel Urtasun. Discrete Residual Flow for Probabilistic Pedestrian Behavior Prediction. *Conference on Robot Learning*, 2019.

CC 2019 Charith Mendis[*], Ajay Jain[*], Paras Jain and Saman Amarasinghe. Revec: Program Rejuvenation through Revectorization. *International Conference on Compiler Construction*, 2019.

OCEANS 2017 C Mirabito, DN Subramani, T Lolla, PJ Haley, A Jain, PFJ Lermusiaux, C Li, DKP Yue, Y Liu, FS Hover, N Pulsone, J Edwards, KE Railey, G Shaw. Autonomy for Surface Ship Interception. *IEEE OCEANS–Aberdeen*, 2017.

Workshops

ISCA 2019 Ajay Jain, Saman Amarasinghe. Learning Automatic Schedulers with Projective Reparameterization. *ML for Sys. at Intl. Symposium on Computer Architecture*, 2019.

ICML 2019 Kavya Ravichandran, Ajay Jain, Alexander Rakhlin. Using Effective Dimension to Analyze Feature Transformations in Deep Neural Networks. *Identifying and Understanding DL Phenomena at International Conference on Machine Learning*, 2019.

NeurIPS 2018 Paras Jain, Xiangxi Mo, Ajay Jain, Harikaran Subbaraj, Rehan Durrani, Alexey Tumanov, Joseph Gonzalez, Ion Stoica. Dynamic Space-Time Scheduling for GPU Inference. *LearningSys Workshop at Neural Information Processing Systems*, 2018.

Preprints

2020 Michael Laskin[*], Scott Emmons[*], Ajay Jain[*], Thanard Kurutach, Pieter Abbeel, Deepak Pathak. Sparse Graphical Memory for Robust Planning. *In submission*, 2020.

2020 Ajay Jain, Pieter Abbeel, Deepak Pathak. Locally Masked Convolution for Autoregressive Models. *In submission*, 2020.

2019 Paras Jain, Xiangxi Mo, Ajay Jain, Alexey Tumanov, Joseph E Gonzalez, Ion Stoica. The OoO VLIW JIT Compiler for GPU Inference. *arXiv*, 2019.

2018 Anand Srinivasan, Ajay Jain, Parnian Barekatain. An Analysis of the Delayed Gradients Problem in Asynchronous SGD. 2018.