tu.stephenl@gmail.com
https://stephentu.github.io

Stephen Tu

research interests

I study problems in the intersection of machine learning, optimization, and control theory. In particular, my Ph.D. work provides safety and performance guarantees for autonomous systems controlled using reinforcement learning.

education

Ph.D., EECS, University of California, Berkeley.
Advised by Prof. Benjamin Recht.
Thesis: *Sample Complexity Bounds for the Linear Quadratic Regulator.*

S.M., EECS, Massachusetts Institute of Technology.
Advised by Prof. Samuel Madden.
Thesis: *Fast Transactions for Multicore In-Memory Databases.*

B.A., Computer Science, University of California, Berkeley.
B.S., Mechanical Engineering, University of California, Berkeley.

preprints

Finite-time Analysis of Approximate Policy Iteration for the Linear Quadratic Regulator. arXiv 2019.
Karl Krauth*, **Stephen Tu***, and Benjamin Recht. (* denotes equal contribution.)

Certainty Equivalent Control of LQR is Efficient. arXiv 2019.
Horia Mania, **Stephen Tu**, and Benjamin Recht.

Learning Contracting Vector Fields For Stable Imitation Learning. arXiv 2018.
Vikas Sindhwani, **Stephen Tu**, and Mohi Khansari.

Non-Asymptotic Analysis of Robust Control from Coarse-Grained Identification. arXiv 2017.
Stephen Tu, Ross Boczar, Andrew Packard, and Benjamin Recht.

Large Scale Kernel Learning using Block Coordinate Descent. arXiv 2016.
Stephen Tu, Rebecca Roelofs, Shivaram Venkataraman, and Benjamin Recht.

publications

The Gap Between Model-Based and Model-Free Methods on the Linear Quadratic Regulator: An Asymptotic Viewpoint. COLT 2019.
Stephen Tu and Benjamin Recht.

On the Sample Complexity of the Linear Quadratic Regulator. FoCM.
Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and **Stephen Tu**.

Minimax Lower Bounds for $\mathcal{H}_\infty$-Norm Estimation. ACC 2019.
Stephen Tu*, Ross Boczar*, and Benjamin Recht. (* denotes equal contribution.)

Safely Learning to Control the Constrained Linear Quadratic Regulator. ACC 2019.
Sarah Dean, **Stephen Tu**, Nikolai Matni, and Benjamin Recht.

Regret Bounds for Robust Adaptive Control of the Linear Quadratic Regulator. NeurIPS 2018.
Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and **Stephen Tu**.

Least-Squares Temporal Difference Learning for the Linear Quadratic Regulator. ICML 2018.
Stephen Tu and Benjamin Recht.

Learning Without Mixing: Towards A Sharp Analysis of Linear System Identification. COLT 2018.
Max Simchowitz, Horia Mania, **Stephen Tu**, Michael I. Jordan, and Benjamin Recht.

On the Approximation of Toeplitz Operators for Nonparametric $\mathcal{H}_\infty$-norm Estimation. ACC 2018.
Stephen Tu, Ross Boczar, and Benjamin Recht.

Breaking Locality Accelerates Block Gauss-Seidel. ICML 2017.
Stephen Tu, Shivaram Venkataraman, Ashia C. Wilson, Alex Gittens, Michael I. Jordan, and Benjamin Recht.

Cyclades: Conflict-free Asynchronous Machine Learning. NeurIPS 2016.
Xinghao Pan, Maximilian Lam, **Stephen Tu**, Dimitris Papailiopoulos, Ce Zhang, Michael I. Jordan, Kannan Ramchandran, Christopher Ré, and Benjamin Recht.

Low-rank Solutions of Linear Matrix Equations via Procrustes Flow. ICML 2016.
Stephen Tu, Ross Boczar, Max Simchowitz, Mahdi Soltanolkotabi, and Benjamin Recht.

Machine Learning Classification over Encrypted Data. NDSS 2015.
Raphael Bost, Raluca Ada Popa, **Stephen Tu**, and Shafi Goldwasser.

Fast Databases with Fast Durability and Recovery through Multicore Parallelism. OSDI 2014.
Wenting Zheng, **Stephen Tu**, Eddie Kohler, and Barbara Liskov.

Anti-Caching: A New Approach to Swapping in Main Memory OLTP Database Systems. VLDB 2014.
Justin DeBrabant, Andrew Pavlo, **Stephen Tu**, Michael Stonebraker, and Stan Zdonik.

Speedy Transactions in Multicore In-Memory Databases. SOSP 2013.
Stephen Tu, Wenting Zheng, Eddie Kohler, Barbara Liskov, and Samuel Madden.

Processing Analytical Queries over Encrypted Data. VLDB 2013.
Stephen Tu, M. Frans Kaashoek, Samuel Madden, and Nickolai Zeldovich.

The HipHop Compiler for PHP. OOPSLA 2012.
Haiping Zhao, Iain Proctor, Minghui Yang, Xin Qi, Mark Williams, Guilherme Ottoni, Charlie Gao, Andrew Paroski, Scott MacVicar, Jason Evans, and **Stephen Tu**.

The Case for PIQL: A Performance Insightful Query Language. SoCC 2010.
Michael Armbrust, Nick Lanham, **Stephen Tu**, Armando Fox, Michael Franklin, and David Patterson.

PIQL: A Performance Insightful Query Language For Interactive Applications. SIGMOD 2010 Demo.
Michael Armbrust, **Stephen Tu**, Armando Fox, Michael Franklin, David Patterson, Nick Lanham, Beth Trushkowsky, and Jesse Trutna.

industry

Software Engineering Intern. Google Brain. Summer 2017.
Worked on projects related to trajectory optimization and learning Lyapunov functions from data.
Hosted by Vikas Sindhwani.

Developer. Data-microscopes team, Qadium. Summer 2014.
Wrote the first implementation of `data-microscopes`, a Bayesian non-parametric library for Python.
Project page: https://datamicroscopes.github.io/

Software Engineering Intern. HPHP team, Facebook. 4/2011–8/2011.
Implemented various performance improvements in Facebook's PHP source-to-source translator.

Software Engineering Intern. Datacenters team, Facebook. 1/2011–4/2011.
Worked on deploying a row level consistency checker for Facebook's distributed MySQL deployment.

Software Engineering Intern. Intuit. Summer 2009.
Built tools for encoding tax specifications in XPath.

Lead Programmer. RSSP-IT, UC Berkeley. 2008-2010.
Maintained various internal tools for managing the residential dorm network at UC Berkeley.

invited talks

Amazon, Palo Alto. Summer 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Facebook AI Research, Menlo Park. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Robot Locomotion Group, Massachusetts Institute of Technology. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Machine Learning Seminar, University of Washington. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Google DeepMind, London. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Google Brain Robotics, Mountain View. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Computer Science Colloquium, Princeton University. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Microsoft Research, New York City. Spring 2019.
Title: *Safe and Reliable Reinforcement Learning for Continuous Control.*

Stanford Information Systems Laboratory. Spring 2018.
Title: *Finite Sample Guarantees for Control of an Unknown Linear Dynamical System.*

Simons Institute. Fall 2017.
Title: *A Lyapunov Analysis for Accelerated Block Gauss-Seidel.* (Lightning talk)

SIAM Conference on Optimization. Spring 2017.
Title: *Convergence and Geometry of Non-convex Matrix Sensing.*

Microsoft Research, Redmond. Summer 2013.
Title: *Speedy Transactions in Multicore In-Memory Databases.*

teaching

Graduate Student Instructor. CS 189–*Introduction to Machine Learning*, UC Berkeley. Fall 2018.

Graduate Student Instructor. CS 189–*Introduction to Machine Learning*, UC Berkeley. Fall 2016.
Section notes: https://people.eecs.berkeley.edu/~stephentu/cs189-fa16.

awards

2018: Google PhD Fellowship in Machine Learning.

2011: CRA Undergraduate Research Award Honorable Mention.

service

Reviewer for OSDI 2014, NeurIPS (2016, 2018, 2019), ICML 2019, ACC 2019, AISTATS 2019, COLT 2019, JMLR, NeurIPS 2018 Workshop (*Machine Learning for Intelligent Transportation Systems*), and ICML 2018 Workshop (*Modern Trends in Nonconvex Optimization for Machine Learning*),