

MATTHEW SOTOUDEH

masotoudeh@ucdavis.edu · (408) 832-5833 · <https://matthewsot.github.io/>

EDUCATION

University of California, Davis

Davis, CA

BS Computer Science, BS Mathematics; Regents Scholar;

GPA: 4.0

Grad: Jun 2021

Lynbrook High School

San Jose, CA

GPA: 3.9

Grad: Jun 2017

EXPERIENCE

Davis Automated Reasoning Group

Davis, CA

Student Researcher

December 2018 — Present

- Working on a variety of research topics related to automated reasoning, verification, and artificial intelligence under Professor Aditya Thakur with the Davis Automated Reasoning Group (DARG).
- “Bounded Model Checking of Deep Neural Network Controllers” won second place at the International Symposium on Principles of Programming Languages (POPL) 2020 Student Research Competition.
- “Computing Linear Restrictions of Neural Networks” (<https://arxiv.org/abs/1908.06214>) accepted at the Conference on Neural Information Processing Systems (NeurIPS) 2019.
- “A Symbolic Neural Network Representation and its Application to Understanding, Verifying, and Patching Networks” (<https://arxiv.org/abs/1908.06223>) currently undergoing peer review.

Intel AIPG, Office of the CTO

San Diego, CA

Research Intern

Jun 2018 — Sep 2018

- Developed a novel, fully-automated compiler system for heterogeneous, parallel systems.
- Fully automatic lowering of arbitrary linear algebra computations onto fixed-function accelerator instruction sets using a novel two-operand version of TVM IR, efficient sub-graph isomorphism solver, and feedback-driven transformation-space search.
- Unified interface to compiler heuristics enables rapid, reproducible testing of new heuristics (including ML-driven models).
- Tests on an upcoming deep learning architecture can achieve up to 3-5X faster execution times than hand-optimized kernel libraries.
- Work has already influenced other compiler and software teams across the company.
- “ISA Mapper: A Compute and Hardware Agnostic Deep Learning Compiler” (<https://doi.org/10.1145/3310273.3321559>) presented and published as a full paper at the ACM International Conference on Computing Frontiers (CF), 2019.

Intel Labs

Santa Clara, CA

Research Intern

Jul 2016 — Jan 2018

- Developed state-of-the-art ML parameter compression method. Up to 60% accuracy improvement over existing work after compressing models up to 1000×.
- Optimized compressed matrix-multiplication routine achieves 15× faster inference than MKL.
- Wrote a unified library that simplifies implementation of multiple compression methods in TensorFlow by expressing each method as a weight-generating sub-graph.
- “C3-Flow: Compute Compression Co-Design for Deep Neural Networks” (<https://doi.org/10.1145/3316781>) presented and published as a full paper at the Design Automation Conference (DAC), 2019.

Develop Summer Academy

Co-Founder & Instructor

San Jose, CA
Feb 2016 — August 2017

- Taught over 50 local middle school students programming, leadership, and other life skills
- Developed courses, marketed the camp, handled logistics, and taught classes
- Recognized by our school, district, and House Representative Ro Khana
- Over \$30,000 in revenue over two summers.

Action

Co-Founder & Chief Software Engineer

San Jose, CA
Nov 2014 — Jan 2016

- Improved the meeting follow-up experience for 1,000s of meetings at Google, Microsoft, UC Berkeley, and others. Featured on the Chrome Web Store.

INTERESTS

I am broadly interested in understanding and explaining computational processes. My recent research has focused on exploring how ideas from the fields of programming languages and automated reasoning can help design safer and more efficient deep learning systems. Longer-term, I am interested in the development of new automated reasoning primitives for underspecified program analysis and manipulation problems. I intend to continue my research while sharing what I have learned through regular teaching.

PROJECTS

See my open-source projects at <https://github.com/matthewsot>

AWARDS

2nd Place at the International Symposium on Principles of Programming Languages (POPL) 2020 Student Research Competition

Jan 2020

Recognized for my work on Bounded Model Checking of Deep Neural Network Controllers. The award was accompanied by a financial award of \$800.

University of California, Davis Computer Science Department Undergraduate Travel Award

Jan 2020

I won an award of \$300 from the UC Davis computer science department to attend the Symposium on Principles of Programming Languages (POPL) 2020 Student Research Competition, where I presented my work on Bounded Model Checking of Deep Neural Network Controllers (and won second place).

NeurIPS Travel Award

Nov 2019

I won an award of \$1,000 from the NeurIPS Foundation to attend the 2020 International Conference on Neural Information Processing Systems, where I presented my work on Computing Linear Restrictions of Deep Neural Networks.

University of California, Davis Undergraduate Research Center Travel Award

May 2019

I won the competitive UC Davis Undergraduate Research travel award (\$500) to support my travel to the International Conference on Computing Frontiers, where I presented my work on a compiler system for modern Deep Learning hardware.

Rep. Ro Khanna Congressional Award

I was officially recognized by House Representative Ro Khanna for my role in founding the innovative Develop Summer Academy.

1st Place FBLA State & National Competitions

Jun 2014, Apr 2015

Won first place against hundreds of teams across the country for two e-business websites I built and presented.

1st Place Application at CodeDay SV

Oct 2014

An early version of our meeting efficiency add-on Action beat dozens of other teams at the CodeDay SV hackathon.

Nokia Developer Wiki 1st Place Winner

Jan 2013

I won two Nokia Developer Wiki competitions for instructional articles describing how to use new Windows APIs.