Siddartha Devic

devic@usc.edu ◇ sid.devic.us ◇ Google Scholar ◇ @sid_devic ◇ US Citizen

Interests: My research aims to make machine learning and AI systems more *trustworthy* by investigating and evaluating their fairness, reliability, optimality, and safety. I am also interested in studying post-training methods for large language models (LLMs), especially for improving reliability and safety.

EDUCATION

University of Southern California *August 2021 - May 2026*
Ph.D. Computer Science GPA: 3.9/4
Thesis: Uncertainty and Trustworthiness in Modern Machine Learning Systems
Co-Advisors: Prof. Vatsal Sharan and Prof. Aleksandra Korolova
Awards: DOD National Defense Science and Engineering Graduate (NDSEG) Fellowship (2021-2024)
USC Viterbi Fellowship (2024-2025)

The University of Texas at Dallas *August 2017 - May 2021*
B.S. Mathematics, B.S. Computer Science Summa Cum Laude, GPA: 3.98/4
CS^2 Computer Science Honors Program, Collegium V Interdisciplinary Honors Program
Awards: Academic Excellence Scholarship (AES) covering full tuition plus stipend (2017-2021)

RESEARCH AND INDUSTRY EXPERIENCE

Apple *(Research Intern)* *May 2025 - September 2025*

- Uncertainty quantification, LLMs, and reasoning with Preetum Nakkiran and Aravind Gollakota.
- Post-training 7B LLMs using RLVR / SFT with vLLM, PEFT, verifiers, TRL.
- Cost-flexible model routing based on aleatoric and epistemic uncertainty decompositions.

University of California, Berkeley *(Visiting Student)* *August 2024 - December 2024*

- Graduate student visitor in "Modern Paradigms in Generalization" program at the Simons Institute.

Amazon *(Applied Scientist Intern)* *May 2024 - August 2024*

- Black-box classifier post-processing of language and vision models with Gaurush Hiranandani.
- Investigating adaptation of black-box classifiers to novel metrics/tasks via logit re-weighting.

Washington University in St. Louis *(Research Intern, NSF REU)* *May 2020 - September 2020*

- Reinforcement learning (RL) research with Prof. Brendan Juba; provable RL in large state spaces.

Johns Hopkins Applied Physics Labs *(Research Intern)* *May 2019 - August 2019*

- Machine learning research in the Machine Perception group with Vickram Rajendran, Eryk Banatt.
- Trained hundreds of neural networks in data labeling and active learning experiments.

SELECTED WORK

($\alpha\beta$) = ALPHABETICAL ORDER; * = EQUAL CONTRIBUTION

Trace Length is a Simple Uncertainty Signal in Reasoning Models
S. Devic, Charlotte Peale, Arwen Bradley, Sinead Williamson, Preetum Nakkiran, Aravind Gollakota
Preprint.

Stability and Multigroup Fairness in Ranking with Uncertain Predictions
S. Devic, Aleksandra Korolova, David Kempe, Vatsal Sharan.
ICML 2024. Also non-archival at Symposium on Foundations of Responsible Computing (FORC 2024).

When is Multicalibration Post-Processing Necessary?
Dutch Hansen*, S. Devic*, Preetum Nakkiran, Vatsal Sharan.
NeurIPS 2024.

OTHER PUBLICATIONS AND PREPRINTS

Are LLM Decisions Faithful to Verbal Confidence?
Jiawei Wang, Yanfei Zhou, S. Devic, Deqing Fu
Preprint.

From Calibration to Collaboration: LLM Uncertainty Quantification Should Be More Human-Centered
S. Devic, Tejas Srinivasan, Jesse Thomason, Willie Neiswanger, Vatsal Sharan.
Preprint.

Auditability and the Landscape of Distance to Multicalibration
($\alpha\beta$) Nathan Derhake, S. Devic, Dutch Hansen, Kuan Liu, Vatsal Sharan.
Innovations in Theoretical Computer Science (ITCS 2026).

An External Fairness Evaluation of LinkedIn Talent Search
Tina Behzad, S. Devic, Aleksandra Korolova, David Kempe, Vatsal Sharan.
*AAAI 2026 (**Oral, top 8% of submitted papers in AI for Social Impacts Track**).*

An Efficient Plugin Method for Metric Optimization of Black-Box Models
S. Devic, N. Choudhary, Anirudh Srinivasan, Sahika Genc, Branislav Kveton, Gaurush Hiranandani.
Preprint.

Proper Learnability and the Role of Unlabeled Data
($\alpha\beta$) Julian Asilis, S. Devic, Shaddin Dughmi, Vatsal Sharan, Shang-Hua Teng.
Algorithmic Learning Theory (ALT) 2025.

Transductive Learning Is Compact
($\alpha\beta$) Julian Asilis, S. Devic, Shaddin Dughmi, Vatsal Sharan, Shang-Hua Teng.
NeurIPS 2024.

Open Problem: Can Local Regularization Learn All Multiclass Problems?
($\alpha\beta$) Julian Asilis, S. Devic, Shaddin Dughmi, Vatsal Sharan, Shang-Hua Teng.
COLT 2024 Open Problem.

Regularization and Optimal Multiclass Learning
($\alpha\beta$) Julian Asilis, S. Devic, Shaddin Dughmi, Vatsal Sharan, Shang-Hua Teng.
COLT 2024.

Fairness in Matching under Uncertainty
S. Devic, David Kempe, Vatsal Sharan, Aleksandra Korolova.
ICML 2023. Also non-archival at ACM conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO 23).

Polynomial Time Reinforcement Learning in Correlated FMDPs with Linear Value Functions
($\alpha\beta$) Zihao Deng, S. Devic, Brendan Juba.
AISTATS 2022. Also at NeurIPS 2021 Workshop on Ecological Theory of Reinforcement Learning.

Dynamic Bandwidth Allocation for PON Slicing with Performance-Guaranteed Online Convex Optimization
Genya Ishigaki, S. Devic, Riti Gour, Jason P. Jue.
IEEE GLOBECOM 2021.

Failout: Achieving Failure-Resilient Inference in Distributed Neural Networks
Ashkan Yousefpour, Brian Q Nguyen, S. Devic, Guanhua Wang, Aboudy Kreidieh, Hans Lobel, Alexandre M Bayen, Jason P Jue.
ICML 2020 Workshop on Federated Learning for User Privacy and Data Confidentiality.

DeepPR: Progressive Recovery for Interdependent VNFs with Deep Reinforcement Learning
Genya Ishigaki, S. Devic, Riti Gour, Jason P. Jue.
IEEE Journal on Selected Areas in Communications, 2020. Also appeared at IEEE GLOBECOM 2019.

AWARDS, ACADEMIC ACHIEVEMENTS, & ACADEMIC PROGRAMS ATTENDED

AAAI Student Travel Grant (2000$ Award) *2026*

DOD National Defense Science and Engineering Graduate (NDSEG) Fellowship *2021-2024*

USC Graduate School Research Award (1500$ Travel grant) *2023*

USC Center for AI in Society (CAIS) Poster Presentation Honorable Mention *2023*

UCLA IPAM Graduate Summer School on Algorithmic Fairness (Travel grant) *2022*

USC Viterbi Fellow (One year of funding for top students in entering PhD cohort) *2021*

NSF Graduate Research Fellowship Program (GRFP) Honorable Mention *2021*

Barry Goldwater Scholar Nomination (One of four STEM students representing UT Dallas) *2020*

Jonsson School Undergraduate Research Award (Awarded to ten engineering students) *2019-2020*

UT Dallas Undergraduate Research Scholar Award (Academic-year research grant) *2018-2019*

School of Engineering Dean's List (Top 10% within engineering) *4 of 5 Semesters*

UT Dallas Academic Excellence Scholarship (Full undergraduate tuition + stipend) *2017 - 2021*

Anson L. Clark Undergraduate Research Scholar (Participant & advisor) *Summers 2017, 2018*

POSTERS & TALKS

Stability and Multigroup Fairness in Ranking with Uncertain Predictions *(Poster, ICML 2024)*
(Talk, FORC 2024)
(Talk, MILA × Vector Institute DEFirst Seminar 2024)

Regularization and Optimal Multiclass Learning *(Talk, Harvard CS theory seminar (TGINF), 2023)*
(Talk, USC ML and Lunch seminar 2023)

Fairness in Matching under Uncertainty *(Poster, ICML 2023)*
(Poster, EAAMO 2023)

Learning Quickly in MDPs with Many States *(USC Theory Lunch 2022)*

Gradient Descent and Clustering in Hyperbolic Space *(Slides, Report, Graduate Course Project 2020)*

Online PR with Bounded Regret *(Poster, UTD Undergraduate Research Contest 2020)*

Point Packing in Hypercubes *(Slides, UTD Mathematics Problem Solving Group 2019)*

ALICE for Deep Active Learning *(Talk, Johns Hopkins Applied Physics Labs 2019)*

Failure-Resilient Distributed Deep Learning Inference *(Poster, Huawei Future Networks Summit 2019)*

Convex Functions for Reinforcement Learning *(Poster, UTD Undergraduate Research Contest, 2019)*

Robust Optimization with Applications in Networking *(Slides, UTD Graduate Seminar, 2019)*

A Reinforcement Learning Based Approach to Networking *(Slides, UTD Graduate Seminar, 2019)*

MENTORSHIP

Alex Reyes Aranda *(Undergraduate)* *May 2025 - Oct. 2025*

- Phenomenological deep learning, USC CURVE undergraduate research program. Project on the impact of label noise on feature selection in neural networks. Co-mentored with Julian Asilis.

Kuan Liu, Nathan Derhake *(Undergraduates)* *Oct. 2024 - Dec. 2025*

- Multicalibration auditing, USC CURVE undergraduate research program. Our work on multicalibration auditing led to a publication at ITCS 2026.
- Worked with both Kuan and Nathan on their PhD applications for 2025-2026 cycle.

Dutch Hansen *(Undergraduate)* *Jan 2023 - May 2025*

- Empirical aspects of multicalibration, USC CURVE undergraduate research program. Our work on multicalibration resulted in publications at NeurIPS 2024 and ITCS 2026.
- I also helped Dutch apply for PhD programs and fellowships. Dutch has grown into a wonderful researcher, and is starting his PhD in Fall 2025 at the University of Washington!

Anish Jayant *(Undergraduate)* *Jan 2023 - May 2023*

- Aspects of multicalibration, Freshman undergraduate at USC interested in theory.

Jayron Martinez *(High School Senior)* *May 2022 - Dec 2022*

- Bias in Machine Learning Algorithms. Part of the USC SHINE high school program.
- After the program, I spent an additional 6 months working with Jayron helping him apply to 10+ universities for his undergraduate degree. He is now an undergraduate at USC!
- Continued mentorship, meeting every few months during his undergraduate degree at USC.

STUDENT ACTIVITIES

VGSA Department Senator Elected to represent the USC Computer Science department as part of the graduate student association. Organized events and career opportunities for Masters and PhD students. *5 hours/week, 2021-2022 Academic year.* [VGSA]

ACM UTD Chapter President Led the largest CS organization at UTD (70 officers, 700+ Members). I proposed and established a $30k perpetual endowed scholarship with club funds. Coordinated student-based semester long technical projects, mentorship programs, a 750+ person hackathon, funding for student startups, industry talks, and more. *8-10 hours/week, March 2018 - Dec. 2020.* [ACM]