OSWIN SO

EMAIL: oswinso@gmail.com **HOMEPAGE:** oswinso.xyz

RESEARCH INTEREST

Safe Autonomous Systems; Stochastic Optimal Control; Discrete Diffusion Models; Constrained Optimization; Machine Learning; Reinforcement Learning; Robotics

EDUCATION

Massachusetts Institute of Technology May 2024 -
Doctor of Philosophy in Aeronautics and Astronautics *Advisor:* Chuchu Fan

Massachusetts Institute of Technology Aug. 2022 - May 2024
Master of Science in Aeronautics and Astronautics, GPA: 5.0 *Advisor:* Chuchu Fan

- Thesis: Safe Nonlinear Control Under Control Constraints via Reachability, Optimal Control and Reinforcement Learning

Georgia Institute of Technology Aug. 2018 - May 2022
Bachelor of Science in Computer Science, GPA: 4.0 *Advisor:* Evangelos Theodorou

- Thesis: Distributionally Robust Optimization Techniques for Stochastic Optimal Control

PUBLICATIONS

* indicates equal contribution.

CONFERENCE

[C20] LASER: Latent Space Adjoint Matching for Support Constrained Entropy Regularized Offline RL
S. Zhang, **O. So**, E. Y. Yu, M. Cleaveland, P. Crowley-Dolen, C. Fan
Advances in Neural Information Processing Systems (NeurIPS), 2026. [pdf]

[C19] Causal Concept Explanations for Deep Neural Models
J. Rountree, P. Verma, **O. So**, C. Fan, J. Shah
Advances in Neural Information Processing Systems (NeurIPS), 2026.

[C18] Value Functions for Temporal Logic: Optimal Policies and Safety Filters
O. So*, W. Sharpless*, S. Herbert, C. Fan
IEEE Conference on Decision and Control (CDC), 2026. [pdf]

[C17] Learning-to-Optimize via Deep Unfolded Flows
[Spotlight, top 2%]
A. Saravanos, **O. So**, H. M. S. Ahmad, C. Fan
International Conference on Machine Learning (ICML), 2026.

[C16] Bellman Value Decomposition for Task Logic in Safe Optimal Control
W. Sharpless*, **O. So***, D. Hirsch, S. Herbert, C. Fan
Robotics: Science and Systems (RSS), 2026. [pdf]

[C15] Discrete Adjoint Matching
O. So, B. Karrer, C. Fan, Ricky T. Q. Chen, G.-H. Liu
International Conference on Learning Representations (ICLR), 2026. [pdf]

[C14] Solving Parameter-Robust Avoid Problems with Unknown Feasibility using Reinforcement Learning
O. So*, E. Y. Yu*, S. Zhang, M. Cleaveland, M. Black, C. Fan
International Conference on Learning Representations (ICLR), 2026. [pdf]

[C13] ReFORM: Reflected Flows for On-support Offline RL via Noise Manipulation
S. Zhang, **O. So**, H. M. S. Ahmad, E. Y. Yu, M. Cleaveland, M. Black, C. Fan
International Conference on Learning Representations (ICLR), 2026. [pdf]

[C12] Solving Multi-Agent Safe Optimal Control with Distributed Epigraph Form MARL
[Outstanding Student Paper Award, 0.6%]
S. Zhang*, **O. So***, M. Black, Z. Serlin, C. Fan
Robotics: Science and Systems (RSS), 2025. [pdf]

[C11] Safe Beyond the Horizon: Efficient Sampling-based MPC with Neural Control Barrier Functions
J. Yin*, **O. So***, C. Fan, P. Tsiotras
Robotics: Science and Systems (RSS), 2025. [pdf]

[C10] Discrete GCBF Proximal Policy Optimization for Multi-agent Safe Optimal Control
S. Zhang*, **O. So***, M. Black, C. Fan
International Conference on Learning Representations (ICLR), 2025. [pdf]

[C9] Solving Minimum-Cost Reach Avoid using Reinforcement Learning
O. So, C. Ge, C. Fan
Advances in Neural Information Processing Systems (NeurIPS), 2024. [pdf]

[C8] How to Train Your Neural Control Barrier Function: Learning Safety Filters for Complex Input-Constrained Systems
O. So, Z. Serlin, M. Mann, J. Gonzales, K. Rutledge, N. Roy, C. Fan
IEEE International Conference on Robotics and Automation (ICRA), 2024. [pdf]

[C7] Solving Stabilize-Avoid Optimal Control via Epigraph Form and Deep Reinforcement Learning
O. So, C. Fan
Robotics: Science and Systems (RSS), 2023. [pdf]

[C6] MPOGames: Efficient Multimodal Partially Observable Dynamic Games
O. So, P. Drews, T. Balch, V. Dimitrov, G. Rosman, E. A. Theodorou
IEEE International Conference on Robotics and Automation (ICRA), 2023. [pdf]

[C5] Deep Generalized Schrodinger Bridge
G.-H. Liu, T. Chen*, **O. So***, E. A. Theodorou
Advances in Neural Information Processing Systems (NeurIPS), 2022. [pdf]

[C4] Decentralized Safe Multi-agent Stochastic Optimal Control using Deep FBSDEs and ADMM
M. A. Pereira, A. D. Saravanos, **O. So**, E. A. Theodorou
Robotics: Science and Systems (RSS), 2022. [pdf]

[C3] Maximum Entropy Differential Dynamic Programming
O. So, Z. Wang, E. A. Theodorou
IEEE International Conference on Robotics and Automation (ICRA), 2022. [pdf]

[C2] Variational Inference MPC using Tsallis Divergence
Z. Wang*, **O. So***, J. Gibson, B. Vlahov, M. S. Gandhi, G.-H. Liu, E. A. Theodorou
Robotics: Science and Systems (RSS), 2021. [pdf]

[C1] Adaptive Risk Sensitive Model Predictive Control with Stochastic Search
Z. Wang, **O. So**, K. Lee, E. A. Theodorou
Learning for Dynamics and Control (L4DC), 2021. [pdf]

JOURNAL

[J8] Safe and Scalable Multi-Drone Payload Transport Via CBF-Based Reinforcement Learning With Zero-Shot Sim-to-Real Transfer
J. Choi, **O. So**, S. Zhang, C. Taylor, C. Fan
IEEE Robotics and Automation Letters, 2026. [pdf]

[J7] Second-order constrained dynamic optimization
Y. Aoyama, **O. So**, A. D. Saravanos, E. A. Theodorou
The International Journal of Robotics Research, 2026. [pdf]

[J6] Safety on the Fly: Constructing Robust Safety Filters via Policy Control Barrier Functions at Runtime
L. Knoedler*, **O. So***, Y. Ji, M. Black, Z. Serlin, P. Tsiotras, J. Alonso-Mora, C. Fan
IEEE Robotics and Automation Letters (RAL), 2025. [pdf]

[J5] Learning plasma dynamics and robust rampdown trajectories with predict-first experiments at TCV
A. M. Wang, A. Pau, C. Rea, **O. So**, C. Dawson, O. Sauter, M. D. Boyer, A. Vu, C. Galperti, C. Fan, A. Merle, Y. Poels, C. Venturini, F. Felici, S. Marchioni, TCV Team
Nature Communications, 2025. [pdf]

[J4] Active ramp-down control and trajectory design for tokamaks with neural differential equations and reinforcement learning
A. M. Wang, C. Rea, **O. So**, C. Dawson, D. T. Garnier, C. Fan
Communications Physics, 2025. [pdf]

[J3] GCBF+: A Neural Graph Control Barrier Function Framework for Distributed Safe Multi-Agent Control
S. Zhang*, **O. So***, K. Garg, C. Fan
IEEE Transactions on Robotics (TRO), 2025. [pdf]

[J2] Comment(s) on Control barrier functions for stochastic systems
O. So, C. Fan
Automatica, 2024. [pdf]

[J1] Learning safe control for multi-robot systems: Methods, verification, and open challenges
K. Garg, S. Zhang, **O. So**, C. Dawson, C. Fan
Annual Reviews in Control, 2024. [link]

WORKSHOP PAPERS

[W2] WISH: Margin-Free Latent Control Barrier Functions via World-Model Imagination
[Spotlight; top 8.3%]
E. Y. Yu*, **O. So***, A. Saravanos, M. Cleaveland, C. Fan
The Science of Physical AI Safety Workshop, Conference on Robot Learning (CoRL), 2026.

[W1] Data-driven discovery of non-Newtonian astronomy via learning non-Euclidean Hamiltonian
O. So, G. Li, E. A. Theodorou, M. Tao
Machine Learning and the Physical Sciences Workshop, NeurIPS, 2022. [pdf]

IN SUBMISSION

[S4] CONTRL: Contrastive Reinforcement Learning for Preference-Aligned Autonomous Driving
Z. Xu*, **O. So***, V. Chanana, J. Kang, K. S. Refaat
In Submission.

[S3] LEAP-CBF: A Safety Filter for Uncertain Systems with Least-Effort Adversarial Potentials
O. So*, E. Y. Yu*, C. Fan
In Submission.

[S2] WISH: Margin-Free Latent Control Barrier Functions via World-Model Imagination
E. Y. Yu*, **O. So***, A. Saravanos, M. Cleaveland, C. Fan
In Submission.

[S1] Flow-Skill: Learning Skill based Hierarchical Policies via Flow Models for Offline Goal Conditioned Reinforcement Learning
H. M. S. Ahmad, S. Zhang, **O. So**, A. Saravanos, C. Fan, W. Li
In Submission.

EXPERIENCES

Waymo (Sunnyvale, CA) June 2026 - September 2026
Software Engineer Intern
- Project: Learning reward models for fine-tuning VLAs in autonomous driving
- Mentor: Khaled Refaat

Meta (FAIR) (New York City, NY) May 2025 - August 2025
Research Scientist Intern
- Project: Discrete Adjoint Matching
- Mentor: Guan-Horng Liu, Ricky T. Q. Chen

Toyota Research Institute (Cambridge, MA) May 2022 - July 2022
Research Intern
- Project: Efficient Multimodal Partially Observable Dynamic Games
- Mentor: Paul Drews

Aurora Innovation (Pittsburgh, PA) May 2021 - July 2021
Motion Planning Intern
- Project: Cost function learning for autonomous driving.
- Mentor: Paul Vernaza

Amazon Robotics (Boulder, CO) May 2020 - July 2020
Software Development Engineer Intern

- Project: LiDAR scan matching for improving docking of autonomous robots.

Greenzie (Atlanta, GA) May 2019 - July 2019
Robotics Engineer Intern

- Project: Communication and control of autonomous lawn mowing robots.

HONORS AND AWARDS

Fellowships

MathWorks Fellow 2024
Yao T. Li (1938) Fellowship 2022

Awards

Outstanding Paper Award Finalist, RSS 2025
Outstanding Student Paper Award, RSS 2025

INVITED TALKS

2. "Policy Optimization under Specifications with Reinforcement Learning." MAE 248: Safety for Autonomous Systems, UCSD (Virtual), 2025. [link]

1. "Solving Stabilize-Avoid Optimal Control via Epigraph Form and Deep Reinforcement Learning." Safe RL Seminar, Virtual, 2023.

ACADEMIC SERVICES

- **Journal Reviewer**: Artificial Intelligence, Automatica, IEEE Transactions on Robotics (TRO), IEEE Transactions on Automatic Control (TAC), IEEE Control Systems Letters (L-CSS), IEEE Robotics and Automation Letters (RA-L).
- **Conference Reviewer**: RSS, ICRA, IROS, NeurIPS, ICLR, CORL, IASEAI, L4DC

SELECTED PRESS COVERAGE

- "AI Steers Systems Safely through Uncertain Conditions", by Muhammad Rohail, *Quantum Zeitgeist*, 2026. [link]
- "New prediction model could improve the reliability of fusion power plants", by Jennifer Chu, *MIT News*, 2025. [link]
- "MIT engineers help multirobot systems stay in the safety zone", by Jennifer Chu, *MIT News*, 2025. [link]
- "A step toward safe and reliable autopilots for flying", by Adam Zewe, *MIT News*, 2024. [link]