Takami Sato, Ph.D.

AI/ML Security Researcher

Daly City, California | takamisato1113@gmail.com | 949-266-7728
tkm2261.github.io | Google Scholar | GitHub | Kaggle | LinkedIn

Current positions: Machine Learning Engineer II, Uber Technologies, Inc.; Visiting Researcher, Keio University

U.S. visa status: H-1B | Google Scholar: 1,155 citations; h-index 16; i10-index 21 (accessed Oct. 30, 2026)

Research Profile

Security researcher working on **physical AI security**—the security of AI systems that perceive and act in the physical world, from autonomous vehicles to vision-language-action (VLA) robots. First/co-first-author work at USENIX Security, NDSS, and CVPR established physical-world attacks and defenses for autonomous-driving perception (camera, LiDAR, lane detection); current work extends this agenda to generative and embodied AI (CVPR 2024, ICLR 2025, CoRL 2026). Directly supervises graduate students at Waseda University and Keio University, largely within the JST CREST Physical AI Security Project (PI: Prof. Tatsuya Mori); supervised students' first-author papers have appeared at NeurIPS, CoRL, ACSAC, and IEEE Sensors Journal. Combines attack development, system-level evaluation, and defenses with production ML engineering.

Research interests: physical and embodied AI security; adversarial ML; sensor security; generative AI security; safety-aware evaluation; mathematical optimization and neural-network verification.

Education

Ph.D., Computer Science – University of California, Irvine Sep. 2024
Advisor: Prof. Qi Alfred Chen.

Dissertation: Exploring Novel Security Vulnerabilities and Their Safety Implications in Sensors and Perception for Autonomous Systems.

Tokyo Institute of Technology (now Institute of Science Tokyo) – Industrial Engineering and Management
M.Eng., Mar. 2013 | B.Eng., Mar. 2011 | Advisor: Prof. Kazuhide Nakata.
Research: mathematical optimization; undergraduate contribution to the SDPA open-source SDP solver.

Research Experience

Research Supervision – Physical AI Security 2023–Present
Visiting Researcher, Keio University; participating researcher, JST CREST (PI: Prof. Tatsuya Mori, Waseda University)

- Directly supervise graduate students in the Mori Lab (Waseda University) and the Yoshioka Lab (Keio University) end to end—problem formulation, attack/defense design, and experiments—and lead the structuring and writing of each paper; much of this work is part of the JST CREST Physical AI Security Project. Each resulting paper has the supervised student as first author (underlined in the Publication List).
- **PULSAR-Net** (R. Yoshida, Keio; NeurIPS 2026, [C1]): LiDAR point-cloud reconstruction under jamming, trained only on synthetic data; real-world vehicle reconstruction accuracy of 91.7% static / 73.2% driving (CenterPoint, IoU 0.4).
- **FLARE** / **ChromaGuard** (M. Watanabe, Keio; CoRL 2026, [C2]): illumination robustness and color-grounding failures in VLA policies; SmolVLA with ChromaGuard achieved 97.5% benign / 92.5% attacked success on real-robot color-dependent tasks.
- **Retroreflective patch attacks** (G. Tsuruoka, Waseda; ACSAC 2026, [C4]; VehicleSec 2024 WIP, [W1]): deployable, stealthy attacks on traffic-sign recognition.
- **LiDAR spoofing on the road** (R. Suzuki, Y. Hayakawa, Keio; IEEE Sensors Journal 2025, [J2, J3]): realizing and detecting high-speed / long-distance spoofing, and systematic security analysis of pulse-fingerprinting LiDARs.
- **Neural-network verification** (R. Ueda, Institute of Science Tokyo; arXiv 2025, [P1]): the interior-point vanishing problem in semidefinite relaxations.
- **Supervised-student awards:** MIRU 2026 Student Excellence Award (R. Yoshida, work leading to [C1]); CSS 2024 Student Paper Award (G. Tsuruoka, work leading to [C4]).

Collaborative Research 2024–2025
Co-first-author and co-authored work

- Co-first-author work demonstrated moving-vehicle LiDAR spoofing with maximum tested speed 60 km/h and maximum tested range 110 m; co-authored SLAMSpoof localization-attack research (NDSS / ICRA 2025, [C7, C6]).
- Co-authored BALD research on environmental backdoors in embodied LLM decision-making, evaluated across virtual and real-world navigation, manipulation, and driving tasks (ICLR 2025, [C5]).

Graduate Research, UC Irvine 2018–2024
Adversarial ML and autonomous-system security

- Designed Dirty Road Patch attacks: at least 97.5% success and at most 0.903 s mean success time on 80 real-road-trace scenarios; 100% collision rates in the tested software-in-the-loop scenarios. Coordinated disclosure to 13 companies; 10 opened investigations (USENIX Security 2021, [C18]).
- Developed NDD / NDDA to study model-agnostic adversarial generation by diffusion models; evaluated physical traffic-sign attacks on a Tesla Model 3 (CVPR 2024, [C9]).
- Developed a study of nine LiDARs, including six next-generation systems, yielding 15 new findings on attack capabilities and defenses (NDSS 2024, [C10]).
- Designed E2E-LD / PSLD driving-oriented metrics and curated Comma2k19-LD to evaluate four lane-detection architectures (CVPR 2022, [C17]).
- Designed LaneGuard: 89% attack detection / 12% false positives across 80 attack and 11,558 benign scenarios; a separate 2 km highway trial at 120 km/h observed zero false positives (VehicleSec 2024, [C12]).

Research Funding

- **JST CREST**, *Security Assessment and Countermeasure Framework for AI-Driven Cyber-Physical Systems* (Physical AI Security Project), FY2023–. Principal Investigator (Research Director): Prof. Tatsuya Mori (Waseda University), with collaborating groups at Institute of Science Tokyo, UEC Tokyo, Osaka University, Keio University, and Ritsumeikan University. Research area: *Creation of System Software for Society 5.0 by Integrating Fundamental Theories and System Platform Technologies*. **Role:** participating researcher; supervise student research on physical-AI attacks and defenses in the Mori (Waseda) and Yoshioka (Keio) groups.

Publication List

Takami Sato is highlighted; * denotes co-first authorship; underlined authors are students whose research Takami Sato directly supervised. Acceptance rates are the overall rates reported by each venue. Links point to publisher or manuscript resources.

Conference Papers

[C1] Ryo Yoshida, **Takami Sato**, Wenlun Zhang, Yuki Hayakawa, Shota Nagai, Takahiro Kado, Taro Beppu, Ibuki Fujioka, Yunshan Zhong, and Kentaro Yoshioka. *Neural Reconstruction of LiDAR Point Clouds under Jamming Attacks via Full-Waveform Representation and Simultaneous Laser Sensing*. **NeurIPS**, 2026 (forthcoming). (acceptance rate: 25.7%)

[C2] Marino Watanabe, **Takami Sato**, and Kentaro Yoshioka. *Lights, Camera, Malfunction: When Illumination Robustness Leaves VLA Models Blind to Color*. **CoRL**, 2026 (forthcoming). (acceptance rate: 32.8%)

[C3] Shaoyuan Xie, Mohamad Habib Fakih, Junchi Lu, Fayzah Alshammari, Ningfei Wang, **Takami Sato**, Halima Bouzidi, Mohammad Abdullah Al Faruque, and Qi Alfred Chen. *FlyTrap: Physical Distance-Pulling Attack Towards Camera-based Autonomous Target Tracking Systems*. **NDSS**, 2026. (acceptance rate: 17.9%)

[C4] Go Tsuruoka, **Takami Sato**, Qi Alfred Chen, Kazuki Nomoto, Ryunosuke Kobayashi, Yuna Tanaka, and Tatsuya Mori. *Trapped by Their Own Light: Deployable and Stealth Retroreflective Patch Attacks on Traffic Sign Recognition Systems*. **ACSAC**, 2026 (forthcoming). (acceptance rate: 19.3%)

[C5] Ruochen Jiao, Shaoyuan Xie, Justin Yue, **Takami Sato**, Lixu Wang, Yixuan Wang, Qi Alfred Chen, and Qi Zhu. *Can We Trust Embodied Agents? Exploring Backdoor Attacks against Embodied LLM-based Decision-Making Systems*. **ICLR**, 2025. (acceptance rate: 32.1%)

[C6] Rokuto Nagata, Kenji Koide, Yuki Hayakawa, Ryo Suzuki, Kazuma Ikeda, Ozora Sako, Qi Alfred Chen, **Takami Sato**, and Kentaro Yoshioka. *SLAMSpoof: Practical LiDAR Spoofing Attacks on Localization Systems Guided by Scan Matching Vulnerability Analysis*. **ICRA**, 2025. (acceptance rate: 38.7%)

[C7] **Takami Sato***, Ryo Suzuki*, Yuki Hayakawa*, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, Ryo Yoshida, Qi Alfred Chen, and Kentaro Yoshioka. *On the Realism of LiDAR Spoofing Attacks against Autonomous Driving Vehicle at High Speed and Long Distance*. **NDSS**, 2025. (acceptance rate: 16.1%)

[C8] Ningfei Wang, Shaoyuan Xie, **Takami Sato**, Yunpeng Luo, Kaidi Xu, and Qi Alfred Chen. *Revisiting Physical-World Adversarial Attack on Traffic Sign Recognition: A Commercial Systems Perspective*. **NDSS**, 2025. (acceptance rate: 16.1%)

[C9] **Takami Sato**, Justin Yue, Nanze Chen, Ningfei Wang, and Qi Alfred Chen. *Intriguing Properties of Diffusion Models: An Empirical Study of the Natural Attack Capability in Text-to-Image Generative Models*. **CVPR**, 2024. (acceptance rate: 23.6%)

[C10] **Takami Sato***, Yuki Hayakawa*, Ryo Suzuki*, Yohsuke Shiiki*, Kentaro Yoshioka, and Qi Alfred Chen. *LiDAR Spoofing Meets the New-Gen: Capability Improvements, Broken Assumptions, and New Attack Strategies*. **NDSS**, 2024. (acceptance rate: 20.2%)

[C11] **Takami Sato***, Sri Hrushikesh Varma Bhupathiraju*, Michael Clifford, Takeshi Sugawara, Qi Alfred Chen, and Sara Rampazzi. *Invisible Reflections: Leveraging Infrared Laser Reflections to Target Traffic Sign Perception*. **NDSS**, 2024. (acceptance rate: 20.2%)

[C12] **Takami Sato**, Ningfei Wang, Yueqiang Cheng, and Qi Alfred Chen. *A Cross-Verification Approach with Publicly Available Map for Detecting Off-Road Attacks against Lane Detection Systems*. **VehicleSec**, 2024. (acceptance rate: 43.1%)

[C13] Sri Hrushikesh Varma Bhupathiraju, **Takami Sato**, Michael Clifford, Takeshi Sugawara, Qi Alfred Chen, and Sara Rampazzi. *On the Vulnerability of Traffic Light Recognition Systems to Laser Illumination Attacks*. **VehicleSec**, 2024. Short paper.

[C14] Ruochen Jiao, Xiangguo Liu, **Takami Sato**, Qi Alfred Chen, and Qi Zhu. *Semi-supervised Semantics-guided Adversarial Training for Robust Trajectory Prediction*. **ICCV**, 2023. (acceptance rate: 26.2%)

[C15] Ningfei Wang, Yunpeng Luo, **Takami Sato**, Kaidi Xu, and Qi Alfred Chen. *Does Physical Adversarial Example Really Matter to Autonomous Driving? Towards System-Level Effect of Adversarial Object Evasion Attack*. **ICCV**, 2023. (acceptance rate: 26.2%)

[C16] Ruochen Jiao, Juyang Bai, Xiangguo Liu, **Takami Sato**, Xiaowei Yuan, Qi Alfred Chen, and Qi Zhu. *Learning Representation for Anomaly Detection of Vehicle Trajectories*. **IROS**, 2023. (acceptance rate: 43%)

[C17] **Takami Sato** and Qi Alfred Chen. *Towards Driving-Oriented Metric for Lane Detection Models*. **CVPR**, 2022. (acceptance rate: 25.3%)

[C18] **Takami Sato***, Junjie Shen*, Ningfei Wang, Yunhan Jack Jia, Xue Lin, and Qi Alfred Chen. *Dirty Road Can Attack: Security of Deep Learning based Automated Lane Centering under Physical-World Attack*. **USENIX Security**, 2021. (acceptance rate: 18.7%)

[C19] Ruochen Jiao, Hengyi Liang, **Takami Sato**, Junjie Shen, Qi Alfred Chen, and Qi Zhu. *End-to-end Uncertainty-based Mitigation of Adversarial Attacks to Automated Lane Centering*. **IV**, 2021. (acceptance rate: 49.3%)

Journal Articles

[J1] Sri Hrushikesh Varma Bhupathiraju, **Takami Sato**, Michael Clifford, Takeshi Sugawara, Qi Alfred Chen, and Sara Rampazzi. *To Go or Not to Go: Shedding Light on Traffic Light Signal Manipulation and Defense Strategies*. **ACM Transactions on Cyber-Physical Systems**, 2026.

[J2] Ryo Suzuki, **Takami Sato**, Yuki Hayakawa, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, Ryo Yoshida, Qi Alfred Chen, and Kentaro Yoshioka. *From Lab to Road: Realizing and Detecting LiDAR Spoofing Attacks Against Autonomous Vehicles at High Speed and Long Distance*. **IEEE Sensors Journal**, 25(13):25661–25681, 2025.

[J3] Yuki Hayakawa, **Takami Sato**, Ryo Suzuki, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, Ryo Yoshida, Qi Alfred Chen, and Kentaro Yoshioka. *Breaking the Shield: Systematic Security Analysis on Pulse Fingerprinting LiDAR Systems for Autonomous Driving*. **IEEE Sensors Journal**, 25(11):20523–20537, 2025.

Preprints

[P1] Ryota Ueda, **Takami Sato**, Ken Kobayashi, and Kazuhide Nakata. *Interior-Point Vanishing Problem in Semidefinite Relaxations for Neural Network Verification*. arXiv:2506.10269, 2025.

[P2] Junjie Shen, Ningfei Wang, Ziwen Wan, Yunpeng Luo, **Takami Sato**, Zhisheng Hu, Xinyang Zhang, Shengjian Guo, Zhenyu Zhong, Kang Li, Ziming Zhao, Chunming Qiao, and Qi Alfred Chen. *SoK: On the Semantic AI Security in Autonomous Driving*. arXiv:2203.05314, 2022.

[P3] **Takami Sato** and Qi Alfred Chen. *On Robustness of Lane Detection Models to Physical-World Adversarial Attacks in Autonomous Driving*. arXiv:2107.02488, 2021.

Selected Workshop / WIP Papers & Demos

[W1] Go Tsuruoka, **Takami Sato**, Qi Alfred Chen, Kazuki Nomoto, Ryunosuke Kobayashi, Yuna Tanaka, and Tatsuya Mori. *WIP: Adversarial Retroreflective Patches: A Novel Stealthy Attack on Traffic Sign Recognition at Night*. **VehicleSec**, 2024. WIP paper.

[W2] **Takami Sato**, Ryo Suzuki, Yuki Hayakawa, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, Qi Alfred Chen, and Kentaro Yoshioka. *Demo: Towards Practical LiDAR Spoofing Attack against Vehicles Driving at Cruising Speeds*. **VehicleSec**, 2024. Demo.

[W3] **Takami Sato***, Sri Hrushikesh Varma Bhupathiraju*, Michael Clifford, Takeshi Sugawara, Qi Alfred Chen, and Sara Rampazzi. *WIP: Infrared Laser Reflection Attack Against Traffic Sign Recognition Systems*. **VehicleSec**, 2023. **ETAS Best Short Paper Award; DENSO Best Demo Award**.

[W4] **Takami Sato***, Yuki Hayakawa*, Ryo Suzuki*, Yohsuke Shiiki*, Kentaro Yoshioka, and Qi Alfred Chen. *WIP: Practical Removal Attacks on LiDAR-based Object Detection in Autonomous Driving*. **VehicleSec**, 2023. **ETAS Best Short Paper Award Runner-Up**.

[W5] **Takami Sato**, Junjie Shen, Ningfei Wang, Yunhan Jack Jia, Xue Lin, and Qi Alfred Chen. *Demo: Security of Deep Learning based Automated Lane Centering under Physical-World Attack*. **IEEE Security and Privacy Workshops (SPW)**, 2021, p. 244. Demo abstract.

[W6] Christopher DiPalma, Ningfei Wang, **Takami Sato**, and Qi Alfred Chen. *Security of Camera-based Perception for Autonomous Driving under Adversarial Attack*. **IEEE Security and Privacy Workshops (SPW)**, 2021, p. 243. Abstract.

Selected Poster Abstracts

[A1] **Takami Sato**, Yuki Hayakawa, Ryo Suzuki, Yohsuke Shiiki, Kentaro Yoshioka, and Qi Alfred Chen. *Towards Large-Scale Measurement Study on LiDAR Spoofing Attacks against Object Detection*. **ACM CCS**, 2022. Poster abstract.

[A2] **Takami Sato**, Junjie Shen, Ningfei Wang, Yunhan Jack Jia, Xue Lin, and Qi Alfred Chen. *Security of Deep Learning based Lane Keeping Assistance System under Physical-World Adversarial Attack*. **NDSS**, 2020. Poster abstract; **Best Technical Poster Award**.

Supervised-Student Workshop Papers, Demos & Posters

[S1] Ryo Yoshida, **Takami Sato**, Wenlun Zhang, Yuki Hayakawa, Shota Nagai, and Kentaro Yoshioka. *Demo: Real-World Reconstruction of LiDAR Perception under Jamming Attacks*. **VehicleSec**, 2026. Demo.

[S2] Ryo Yoshida, **Takami Sato**, Yuki Hayakawa, Ryo Suzuki, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, and Kentaro Yoshioka. *Poster: Discovering Sensor-Fusion-Vulnerabilities in Autonomous Driving Systems against LiDAR Attacks*. **NDSS**, 2025. Poster.

[S3] Ryo Suzuki, **Takami Sato**, Yuki Hayakawa, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, Qi Alfred Chen, and Kentaro Yoshioka. *WIP: Towards Practical LiDAR Spoofing Attack against Vehicles Driving at Cruising Speeds*. **VehicleSec**, 2024. WIP paper.

[S4] Yuki Hayakawa, **Takami Sato**, Ryo Suzuki, Kazuma Ikeda, Ozora Sako, Rokuto Nagata, Qi Alfred Chen, and Kentaro Yoshioka. *WIP: An Adaptive High Frequency Removal Attack to Bypass Pulse Fingerprinting in New-Gen LiDARs*. **VehicleSec**, 2024. WIP paper.

[S5] Ozora Sako, **Takami Sato**, Yuki Hayakawa, Ryo Suzuki, Kazuma Ikeda, Rokuto Nagata, Qi Alfred Chen, and Kentaro Yoshioka. *Poster: Intensity-Aware Chosen Pattern Injection LiDAR Spoofing Attack*. **NDSS**, 2024. Poster.

Selected Domestic Papers (Japan)

Presented in Japanese; titles translated.

[D1] Go Tsuruoka, **Takami Sato**, Ryo Iijima, and Tatsuya Mori. *Black Hole Attack: A Hijacking Attack on Autonomous Driving Using Occupancy Prediction and Its Countermeasures*. **Computer Security Symposium (CSS)**, 2026.

[D2] Ryo Yoshida, **Takami Sato**, Wenlun Zhang, Shota Nagai, Yuki Hayakawa, Yunshan Zhong, Takahiro Kado, Ibuki Fujioka, Taro Beppu, and Kentaro Yoshioka. *Neural Reconstruction of LiDAR Point Clouds under Jamming Attacks via Full-Waveform Representation and Simultaneous Laser Sensing*. **MIRU**, 2026. **MIRU Student Excellence Award.**

[D3] Go Tsuruoka, **Takami Sato**, Ryo Iijima, Kazuki Nomoto, Yuna Tanaka, and Tatsuya Mori. *Is Next-Generation Autonomous Driving Safe? Attacks on Occupancy Prediction and Their System-Level Impact*. **Symposium on Cryptography and Information Security (SCIS)**, 2026.

[D4] Go Tsuruoka, **Takami Sato**, Qi Alfred Chen, Kazuki Nomoto, Ryunosuke Kobayashi, Yuna Tanaka, and Tatsuya Mori. *Proposal and Evaluation of Adversarial Patch Attacks Exploiting Reflected Headlight Light*. **Computer Security Symposium (CSS)**, 2024. **CSS Student Paper Award.**

Professional Experience

Uber Technologies, Inc. Sep. 2024–Present
Machine Learning Engineer II

- Developed the Driver Cancel prediction model, now used throughout Latin America, delivering +0.11% Gross Bookings and +0.21% Trips.
- Led driver-embedding integration into matching models, increasing Gross Bookings by 0.38%. Separately, engineered matching-graph features for driver-cancel prediction, increasing Gross Bookings by 5% and revenue by 7.5% for a US-based product.
- Architected daily incremental training pipelines and optimized GPU utilization to improve ETA-surprise prediction.
- Productionized Transformers, DNNs, and XGBoost for wait-time and session-value prediction.
- Owned prediction-system operations, including experiment monitoring, real-time alerts, continuous deployment, and automated rollback.

Qualcomm Incorporated Jun.–Sep. 2024
Ph.D. Research Intern – AI and Automotive Security Research

- Evaluated state-of-the-art attacks and defenses for Snapdragon Ride.

Uber Technologies, Inc. Jun.–Sep. 2023
Ph.D. Research Software Engineer Intern

- Built Transformer-based driver-activity and incentive-payout forecasts over a 168-hour horizon; improved city-level weekly payout MAPE against an XGBoost baseline.
- Architected an AUCC-based evaluation pipeline and conducted counterfactual analysis of incentive campaigns, accounting for missing observations in holdout groups.

Toyota InfoTech Labs Jun.–Sep. 2022
Security Researcher Intern
- Investigated infrared-laser attacks on traffic-sign perception; co-first-author research published at VehicleSec 2023 and NDSS 2024 ([W3, C11]).

NIO Inc. Mar.–Sep. 2021
Security Researcher Intern
- Investigated Level-2 lane-centering safety and defenses; resulting LaneGuard research published at VehicleSec 2024 ([C12]).

Independent Consulting 2017–2018
Machine Learning Consultant
- Provided machine learning consulting and mentored practitioners in ML and Kaggle competitions.

Recruit Communications Co., Ltd. Sep. 2016–Jun. 2017
Machine Learning Engineer
- Developed customer-retention models and prediction APIs deployed on AWS and GCP.

BrainPad Inc. Apr. 2013–Jun. 2016
Data Scientist / Software Engineer / Team Manager
- Delivered machine learning and optimization applications, including demand forecasting and vehicle-routing systems; owned analysis, architecture, implementation, code review, and monitoring.
- Managed five engineers and data scientists while remaining hands-on.

Research Service
- **Program committees:** USENIX Security 2026, 2027; AAAI 2026; EuroS&P 2024; CPSIoTSec 2023, 2024; IEEE MSN 2023.
- **ML / vision reviewing:** AISTATS (2025), ICML (2025), ICLR (2025); NeurIPS (2024–2026); ICCV (2023, 2025); ECCV (2022, 2024, 2026); CVPR (2022–2024, 2026).
- **External / sub-reviewing:** 50+ papers for IEEE S&P, USENIX Security, ACM CCS, NDSS, AsiaCCS, INFOCOM, ACSAC, TDSC, ICCPS, BMVC, and SPML.

Invited Talks
- **March 2026 – Stanford University, 38th Japan-America Student Summit (JASS):** Autonomous Driving Security and Life in the Academic Security Community (Japanese).
- **September 2024 – FIT 2024 Top Conference Session, Japan:** LiDAR Spoofing Meets the New-Gen (NDSS 2024).
- **November 2021 – Google Cloud ML Summit:** BigQuery Techniques for Kagglers (Japanese).

Selected Media Coverage
Coverage of co-authored research.
- **Keio University & JST press release** (2025): disabling autonomous-driving LiDAR from long distance while driving ([C7]); also covered by Nikkei, Nikkei xTECH, and Jidounten Lab.
- **UC Irvine News** (2025): commercial traffic-sign recognition and vehicle-safety deficiencies ([C8]).
- **UC Irvine News** (2024): road-object spoofing and vanishing attacks on new-generation LiDARs ([C10]).
- **Keio University & JST press release** (2024): first comprehensive security study of LiDAR sensors for autonomous driving ([C10]); also covered by Nikkei.
- **Automotive News** (2024): LiDAR spoofing and object-removal vulnerabilities ([C10]).
- **Nikkei xTECH** (2024): LiDAR object-removal attacks and the HFR attack ([C10]).
- **IoT World Today** and **Electro Optics** (2024): LiDAR spoofing vulnerabilities ([C10]).
- **ITmedia NEWS** (2024): invisible infrared-laser attacks on traffic-sign recognition ([C11]).
- **UC Irvine ICS News** (2023): named Public Impact Fellow; **UC Irvine ICS News** (2020): team won 1st place at the Baidu AutoDriving CTF.

Selected Research Artifacts
Public resources associated with co-authored research.
- **NDD / NDDA project and physical-world demonstrations** – generative adversarial examples ([C9]).
- **BALD evaluation code** – embodied LLM backdoors ([C5]).
- **SLAMSpoof code** – LiDAR localization attacks ([C6]).
- **Comma2k19-LD dataset and evaluation code** – driving-oriented lane-detection metrics ([C17]).
- **Dirty Road Patch attack code and demo videos** – physical-world lane-centering attacks ([C18]).

Selected Awards & Honors

- **ECCV Outstanding Reviewer** (2024).
- **VehicleSec 2023:** DENSO Best Demo and ETAS Best Short Paper Award ([W3]); ETAS Best Short Paper Award Runner-Up ([W4]).
- **Graduate Dean's Dissertation Fellowship**, UC Irvine (2023); **Public Impact Fellowship**, UC Irvine (2022); **Beall Family Foundation Graduate Student Entrepreneur Award** (2021).
- **NDSS Best Technical Poster Award** (2020); **Baidu AutoDriving Security CTF**, 1st place (team, 1/24, 2020).
- **Kaggle Competitions Grandmaster** (highest competition tier, 2021).
- **Kaggle Days x Z by HP World Championship:** 1st place, Final (Barcelona, 2022); 2nd place (San Francisco, 2/82, 2022).
- **Kaggle 2nd-place finishes:** Foursquare Location Matching (2/1,079, 2022); PetFinder Pawpularity (2/3,537, 2022); chaii Hindi/Tamil Question Answering (2/943, 2021); Shopee Product Matching (2/2,426, 2021).
- **Kaggle gold medals:** Tweet Sentiment Extraction (5/2,227, 2020); solo: Instacart Market Basket Analysis (15/2,622, 2017), TalkingData AdTracking Fraud Detection (11/3,946, 2018).
- **JASSO scholarship repayment exemption** for outstanding academic achievement (2013).

Technical Expertise & Community

Research methods: adversarial ML; physical-world attack evaluation; generative and embodied AI security; sensor spoofing; counterfactual evaluation; mathematical optimization.

Engineering: Python, C++; Transformers, diffusion models, XGBoost; production ML pipelines; sparse linear algebra and semidefinite-programming implementations.

Community: Organizer, Kaggler-ja (Japanese Kaggle community, 10,000+ members); mentoring through independent ML consulting.