

David I. Inouye

✉ dinouye@purdue.edu  <https://www.davidinouye.com>

› Research Vision and Topics

Develop trustworthy machine learning methods
that are robust to imperfect distributional and computational assumptions
using explainability, causality, and collaborative learning.

Topics: Trustworthy machine learning, Causal machine learning, Out-of-distribution robustness, Distribution shift, Robust collaborative learning, Explainable AI, Fairness

› Appointments

Assistant Professor (July 2019 - Present)
Elmore Family School of Electrical and Computer Engineering
Purdue University

Postdoctoral Researcher (Sep 2017 - July 2019)
Carnegie Mellon University
Supervisor: Prof. Pradeep Ravikumar

› Education

Ph.D. in Computer Science, The University of Texas at Austin (May 2017)
Advisers – Prof. Inderjit S. Dhillon
Prof. Pradeep Ravikumar (now at Carnegie Mellon University)
Dissertation Title – “Appropriate, Accessible and Appealing Probabilistic Graphical Models”

M.S. in Computer Science, The University of Texas at Austin, Austin, TX (May 2015)
B.S. in Electrical Engineering, Georgia Institute of Technology, Atlanta, GA (May 2012)
B.A. in Natural Sciences, Covenant College, Lookout Mountain, GA (May 2011)

› Research Funding

Co-PI (Lead PIs: Shreyas Sundaram, Shaoshuai Mou, Co-PIs: Christopher Brinton, James Goppert, Song Zhang, Michael Zoltowski, Senior Personnel: Dan DeLaurentis), Saab Inc. / Office of Naval Research, “Threat & Situational Understanding with Networked-Online Machine Intelligence (TSUNOMI)”, N00014-23-C-1016, June 2023 - January 2027, \$4.4M (Inouye: 10%).

Sole-PI, National Science Foundation (NSF), “RI: Small: Local and Forward-Oriented Deep Learning for Decentralized and Dynamic Environments”, IIS-2212097, October 2022 - September 2025, \$598,843.

Co-PI (PI: Saurabh Bagchi, Other Co-PIs: Prateek Mittal, Somali Chaterji, Mung Chiang), Army Research Laboratory, “Secure, Real-Time, Distributed Decision-Making for the Autonomous Battlefield”, W911NF-2020-221, August 2020 - August 2025, \$3.7M (shared among PI/Co-PIs).

Co-PI (PI: Saurabh Bagchi), Northrop Grumman Cybersecurity Research Consortium (NGCRC), “Secure, Real-Time Decision-Making for the Autonomous Battlefield,” September 2019-August 2020, \$125,000 (shared equally among PIs/Co-PIs).

› Manuscripts (‡ = students I advise, * = equal contribution)

Ruqi Bai[‡], Yao Ji, Zeyu Zhou, David I. Inouye. From Invariant Representations to Invariant Data: Provable Robustness to Spurious Correlations via Noisy Counterfactual Matching. Under Submission. [Preprint](#).

Surojit Ganguli[‡], Zeyu Zhou[‡], Christopher G. Brinton, David I. Inouye. Robust Collaborative Inference with Vertically Split Data Over Dynamic Device Environments. Under submission. [Preprint](#).

Wonwoong Cho[‡], Raymond A. Yeh, David I. Inouye. Causally Motivated Diffusion Sampling Frameworks for Harnessing Contextual Bias. Under submission.

David Burnett[‡], Surojit Ganguli[‡], Lance M. Kaplan, Devesh Upadhyay, David I. Inouye. Estimating and Evaluating Second-Order Uncertainty due to Missing Values at Inference Time. Under Submission.

Wonwoong Cho[‡], Yan-Ying Chen, Matthew Klenk, David I. Inouye, Yanxia Zhang. Att-Adapter: A Robust and Precise Domain-Specific Multi-Attributes T2I Diffusion Adapter via Conditional Variational Autoencoder. Under submission. [Preprint](#).

Brian Ko[‡], Ziyu Gong[‡], Jim Lim[‡], David I. Inouye. A Unified Framework for Comparing Distribution Matching Methods Across Trustworthy Machine Learning Tasks. In preparation.

Ruqi Bai[‡], James Z. Hare, Nicholas R. Waytowich, David I. Inouye. StarCraft Motion: A Motion Forecasting Dataset with Ground Truth Unit-Level Goals and Strategic Behavior. Under submission.

› Publications

Ziyu Gong^{*‡}, Jim Lim^{*‡}, David I. Inouye. Improving Distribution Matching via Score-Based Priors and Structural Regularization. In *International Conference on Machine Learning (ICML)*. 2025. [Preprint](#). (26.9% acceptance)

Zeyu Zhou[‡], Tianci Liu, Ruqi Bai[‡], Jing Gao, Murat Kocaoglu, David I. Inouye. Counterfactual Fairness by Combining Factual and Counterfactual Predictions. In *Neural Information Processing Systems (NeurIPS)*. 2024. (25.8% acceptance)

Wonwoong Cho[‡], Hareesh Ravi, Midhun Harikumar, Vinh Khuc, Krishna Kumar Singh, Jingwan Lu, David I. Inouye, Ajinkya Kale. Enhanced Controllability of Diffusion Models via Feature Disentanglement and Realism-Enhanced Sampling Methods. In *European Conference on Computer Vision (ECCV)*. 2024. (27.9% acceptance)

Mai Elkady[‡], Hong Minh Thu Bui, Bruno Ribeiro, David I. Inouye. Vertical Validation: Evaluating Implicit Generative Models for Graphs on Thin Support Regions. In *Uncertainty in Artificial Intelligence (UAI)*. 2024. (27% acceptance)

Ziyu Gong[‡], Ben Usman, Han Zhao, David I. Inouye. Towards Practical Non-Adversarial Distribution Matching. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2024. (27.6% acceptance)

- Ruqi Bai[‡], Saurabh Bagchi, David I. Inouye. Benchmarking Algorithms for Domain Generalization in Federated Learning. In *International Conference on Learning Representations (ICLR)*. 2024. **(Spotlight, 5% acceptance)**
- Zeyu Zhou^{*‡}, Ruqi Bai^{*‡}, Sean Kulinski^{*‡}, Murat Kocaoglu, David I. Inouye. Towards Characterizing Domain Counterfactuals For Invertible Latent Causal Models. In *International Conference on Learning Representations (ICLR)*. 2024. (31% acceptance)
- Sean Kulinski[‡], David I. Inouye. Towards Explaining Distribution Shifts. In *International Conference on Machine Learning (ICML)*. 2023. (27.9% acceptance)
- Sean Kulinski[‡], Nicholas R. Waytowich, James Z. Hare, David I. Inouye. StarCraftImage: A Dataset For Prototyping Spatial Reasoning Methods For Multi-Agent Environments. In *IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR)*. 2023. (25.8% acceptance)
- Zeyu Zhou[‡], Sheikh S. Azam, Christopher Brinton, David I. Inouye. Efficient Federated Domain Translation. In *International Conference on Learning Representations (ICLR)*. 2023. (32.0% acceptance)
- Wonwoong Cho^{*‡}, Ziyu Gong^{*‡}, David I. Inouye. Cooperative Distribution Alignment via JSD Upper Bound. In *Neural Information Processing Systems (NeurIPS)*. 2022 (25.6% acceptance).
- Mai Elkady^{*‡}, Jim Lim^{*‡}, David I. Inouye. Discrete Tree Flows via Tree-Structured Permutations. In *International Conference on Machine Learning (ICML)*. 2022. (21.9% acceptance)
- Zeyu Zhou[‡], Ziyu Gong[‡], Pradeep Ravikumar, David I. Inouye. Iterative Alignment Flows. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 2022. (29.2% acceptance)
- Rui Wang[‡], Xiaoqian Wang, David I. Inouye. Shapley Explanation Networks. In *International Conference on Learning Representations (ICLR)*. 2021. (28.7% acceptance)
- Sean Kulinski[‡], Saurabh Bagchi, David I. Inouye. Feature Shift Detection: Localizing Which Features Have Shifted via Conditional Distribution Tests. In *Neural Information Processing Systems (NeurIPS)*. 2020. (20% acceptance)
- David I. Inouye, Liu Leqi, Joon Sik Kim, Bryon Aragam, Pradeep Ravikumar. Automated Dependence Plots. In *Uncertainty in Artificial Intelligence (UAI)*. 2020. (27.6% acceptance)
- Chih-Kuan Yeh, Cheng-Yu Hsieh, Arun S. Suggala, David I. Inouye, Pradeep Ravikumar. On the (In)fidelity and Sensitivity for Explanations. In *Neural Information Processing Systems (NeurIPS)*. 2019. (21.6% acceptance)
- David I. Inouye and Pradeep Ravikumar. Deep Density Destructors. In *International Conference on Machine Learning (ICML)* 35, 2018. Code available: <https://github.com/davidinouye/destructive-deep-learning>. (oral presentation)
- David I. Inouye, E. Yang, G. I. Allen, Pradeep Ravikumar. A Review of Multivariate Distributions for Count Data Derived from the Poisson Distribution. *WIRES Computational Statistics*, 9: e1398, 2017. Code available.

- David I. Inouye, Pradeep Ravikumar, Inderjit S. Dhillon. Square Root Graphical Models: Multivariate Generalizations of Univariate Exponential Families that Permit Positive Dependencies. In *International Conference on Machine Learning (ICML)* 33, 2016. (oral presentation)
- David I. Inouye, Pradeep Ravikumar, Inderjit S. Dhillon. Fixed-Length Poisson MRF: Adding Dependencies to the Multinomial. In *Neural Information Processing Systems (NeurIPS)* 28, 2015. Code available.
- David I. Inouye, Pradeep Ravikumar, Inderjit S. Dhillon. Capturing Semantically Meaningful Word Dependencies with an Admixture of Poisson MRFs. In *Neural Information Processing Systems (NeurIPS)* 27, 2014. Code available.
- David I. Inouye, Pradeep Ravikumar, Inderjit S. Dhillon. Admixture of Poisson MRFs: A Topic Model with Word Dependencies. In *International Conference on Machine Learning (ICML)* 31, 2014. (oral presentation)
- Joshua Inouye, David I. Inouye, Silvia Blemker. Towards Undistorted and Noise-Free Speech in an MRI Scanner: Correlation Subtraction followed by Spectral Noise Gating. *Journal of the Acoustical Society of America (JASA)*, 135(3), pp. 1019-1022, January 2014.
- Beaux Sharifi, David I. Inouye, Jugal K. Kalita. Summarization of Twitter Microblogs. *The Computer Journal*, 57(3), pp. 378-402, 2014.
- David I. Inouye & Jugal K. Kalita. Comparing Twitter Summarization Algorithms for Multiple Post Summaries. In *IEEE International Conf. on Social Computing*, pp. 298-306, 2011. (9.8% acceptance rate, oral presentation)

› Workshop Publications

- Avi Amalanshu[‡], Yash Sirvi, David I. Inouye. Decoupling Vertical Federated Learning using Local Self-Supervision. In *NeurIPS 2024 Workshop on Self-Supervised Learning - Theory and Practice*. 2024. <https://openreview.net/forum?id=nFdhJ1Imrp>
- Zeyu Zhou[‡], Ruqi Bai[‡], David I. Inouye. Improving Practical Counterfactual Fairness with Limited Causal Knowledge. In *ICLR 2024 Workshop on Navigating and Addressing Data Problems for Foundation Models (DPFM)*. 2024. <https://openreview.net/forum?id=Tz9v3zln4Y>
- Sean Kulinski^{*‡}, Zeyu Zhou^{*‡}, Ruqi Bai^{*‡}, Murat Kocaoglu, David I. Inouye. Towards Characterizing Domain Counterfactuals for Invertible Latent Causal Models. *NeurIPS 2023 Workshop on Causal Representation Learning*. 2023.
- Surojit Ganguli[‡], Avi Amalanshu[‡], Amritanshu Ranjan[‡], David I. Inouye. Internet Learning: Preliminary Steps Towards Highly Fault-Tolerant Learning on Device Networks. *ICML 2023 Workshop on Localized Learning*. 2023. <https://openreview.net/forum?id=75PEgr4xTl>
- Sean Kulinski[‡], David I. Inouye. Towards Explaining Image-Based Distribution Shifts. *CVPR Workshop on Vision Datasets Understanding*. 2022. (64% acceptance rate)
- Wonwoong Cho^{*‡}, Ziyu Gong^{*‡}, David I. Inouye. Why be adversarial? Let's cooperate!: Cooperative Dataset Alignment via JSD Upper Bound. Accepted to *ICML 2021 Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models (INNF+, spotlight talk, 14/40)*.

Mai Elkady^{*‡}, Jim Lim^{*‡}, David I. Inouye. Discrete Tree Flows via Tree-Structured Permutations. Accepted to *ICML 2021 Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models (INNF+)*.

Manush Bhatt[‡], David I. Inouye. Learning Flows by Parts. *NeurIPS 2020 Workshop Beyond Backpropagation*. 2020.

David I. Inouye, Pradeep Ravikumar. WeakFlow: Iterative Invertible Distribution Transformations via Weak Destructive Flows. *ICML 2020 Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models*. 2020.

David I. Inouye, Pradeep Ravikumar. Deep Point Process Destructors. *NeurIPS Workshop on Learning with Temporal Point Processes*. 2019.

David I. Inouye, Pradeep Ravikumar, Pradipto Das, Ankur Datta. Hyperparameter Selection under Localized Label Noise via Corrupt Validation. Presented at *Learning with Limited Labeled Data NeurIPS Workshop*, 2017.

<https://www.davidinouye.com/publication/inouye-2017-hyperparameter/inouye-2017-hyperparameter.pdf>

› Meeting Organization

“Localized Learning: Decentralized Model Updates via Non-Global Objectives”. Workshop at the International Conference on Machine Learning (ICML) 2023. Co-organizers: Mengye Ren, Mateusz Malinowski, Michael Eickenberg, Gao Huang, Eugene Belilovsky. (26.3% acceptance rate)

› Invited Talks

Georgia Institute of Technology. Atlanta, GA. January 29, 2025. “Domain Counterfactuals for Explainability, Fairness, and Domain Generalization”. Machine Learning at Georgia Tech (ML@GT) Seminar.

Carnegie Mellon University. Pittsburgh, PA. October 29, 2024. “Domain Counterfactuals for Explainability, Fairness, and Domain Generalization”. Machine Learning / Duolingo Seminar.

University of Michigan. Ann Arbor, MI. October 11, 2024. “Domain Counterfactuals for Explainability, Fairness, and Domain Generalization”. Communications and Signal Processing (CSP) Seminar.

University of Wisconsin-Madison. Madison, WI. September 4, 2024. “Domain Counterfactuals for Explainability, Fairness, and Domain Generalization”. Systems, Information, Learning and Optimization (SILO) Seminar.

UCLA. Los Angeles, CA. March 14, 2024. “Domain Counterfactuals for Trustworthy ML via Sparse Interventions”. Computer Science Department Seminar.

UC San Diego. San Diego, CA. March 13, 2024. “Domain Counterfactuals for Trustworthy ML via Sparse Interventions”.

UC Berkeley. Berkeley, CA. March 11, 2024. “Domain Counterfactuals for Trustworthy ML via Sparse Interventions”.

Stanford University. Palo Alto, CA. March 8, 2024. “Domain Counterfactuals for Trustworthy ML via Sparse Interventions”.

TOC4Fairness Seminar supported by the Simons Collaboration on the Theory of Algorithmic Fairness. Virtual. November 15, 2023. “Towards Trustworthy Machine Learning via Distribution Matching”.

The University of Texas at Austin. Austin, TX. November 3, 2023. “Towards Trustworthy Machine Learning via Distribution Matching”.

Texas A&M University (TAMU). College Station, TX. November 1, 2023. “Towards Trustworthy Machine Learning via Distribution Matching”.

Toyota Technological Institute at Chicago (TTIC). Chicago, IL. September 29, 2023. “Towards Trustworthy ML via Distribution Alignment”.

› Internal Talks

“Will AI Write All Our Software?”. ECE Stories of Success Panel. Virtual. March 8, 2023.

“Unifying and Advancing the Science of Deep Distribution Alignment”. Center for Innovation in Control, Optimization, and Networks (ICON) weekly seminar. West Lafayette, IN. October 21, 2022.

“Unifying and Advancing the Science of Deep Distribution Alignment”. Purdue ECE Seminar Series. Virtual. August 30, 2022.

› Awards

ICLR 2024 Spotlight Paper (5% acceptance rate)

Rising Star Faculty Award 2021, Center for Resilient Infrastructures, Systems, and Processes (CRISP)

NSF Graduate Research Fellowship

ICLR 2022 Highlighted Reviewer (Top 9% of ICLR reviewers)

NeurIPS 2021 Outstanding Reviewer Award (Top 8% of NeurIPS reviewers who were judged to be instrumental to the review process based on Area Chair and author feedback)

Eagle Scout; Georgia Tech PURA Research Award; Elks Club Scholarship

› Students

Ruqi Bai (PhD)	01/2020 - Present
Zeyu Zhou (PhD)	05/2020 - Present
Mai Elkady (PhD in CS, co-advised with Petros Drineas)	08/2020 - Present
Ziyu Gong (PhD)	08/2020 - Present
David Burnett (PhD)	10/2020 - Present
Hyung (Jim) Lim (PhD)	12/2020 - Present
Wonwoong Cho (PhD)	01/2021 - Present
Surojit Ganguli (PhD)	01/2022 - Present
Yi-Chung Chen (PhD, co-advised with Jing Gao)	05/2024 - Present
Sean Kulinski (PhD in ECE, Next: MosiacML/DataBricks)	01/2020 - 12/2023

Vijay Prasad (MS in ECE, Next: MITRE)	05/2021 - 08/2022
Manush Bhatt (MS in CS, Next: SoFi)	05/2020 - 08/2021
Kyung Min (Brian) Ko (Year after graduation)	02/2024 - Present
Avi Amalanshu (Undergraduate, REU/SURF)	05/2023 - Present
Connie Kang (Undergraduate, REU/SURF)	05/2023 - 08/2023
Amritanshu Ranjan (SURF, summer after grad., Next: Walmart Global Tech)	05/2023 - 08/2023
Nikita Ravi (Undergraduate, Next: Software Engineer @ Salesforce)	01/2022 - 05/2022
Xinji (Jimmy) Jiang (Undergraduate)	01/2022 - 05/2022
Rui Wang (Undergraduate, Next: Startup in China, deferred UW admission)	01/2020 - 06/2021

› Teaching Experience

Taught ECE 47300 (Spring 2024): Introduction to Artificial Intelligence undergraduate course in Purdue ECE. Updated lecture content including Reinforcement Learning algorithms and a new assignment. (164 undergraduate students)

Taught ECE 57000 (Fall 2023): Artificial Intelligence graduate course in Purdue ECE. Extended template for reports to focus on paper storyline, claim-evidence structure, and implementation details. (293 students: 138 undergraduate / 155 graduate, 66 online / 227 in-person)

Taught ECE 47300 (Spring 2023): Introduction to Artificial Intelligence undergraduate course in Purdue ECE. Updated lecture content including natural language processing and Markov decision processes. (124 students: 123 undergraduates, 1 graduate)

Taught ECE 57000 (Fall 2022): Artificial Intelligence graduate course in Purdue ECE. Updated lecture content and organization from previous year including a new assignment and the new topics of diffusion models and distribution alignment. (254 students: 212 in person, 42 online; 158 graduate, 96 undergraduate)

Taught ECE 57000 (Fall 2021): Artificial Intelligence graduate course in Purdue ECE. Updated lecture content and organization from previous year. (183 students: 91 graduate, 92 undergraduate)

Taught ECE 20875 (Spring 2021): Python for Data Science undergraduate course in Purdue ECE. (363 undergraduate students: 282 residential, 81 online)

Taught ECE 57000 (Fall 2020): Artificial Intelligence graduate course in Purdue ECE. Updated lecture content and organization from previous year. Developed project checkpoints with peer review to keep students on track. Adapted lecture style to hybrid in-person and online learning setting by leveraging iPad for slide notation or derivations. (118 students: 65 graduate, 53 undergraduate)

Taught ECE 20875 (Spring 2020): Python for Data Science undergraduate course in Purdue ECE. Helped update lecture content and organization from the previous semester with Prof. Chris Brinton. Adapted lectures and exams to the online learning setting mid-year including learning new online technologies. (238 undergraduate students)

Taught ECE 57000 (Fall 2019): Artificial Intelligence graduate course in Purdue ECE. Updated content to focus on unsupervised learning and probabilistic models. (76 students, mix of graduate and undergraduate students)

Attended Effective College Teaching workshop hosted by Purdue University in Fall 2019.

Participated in Fall 2019 College of Engineering Diversity & Inclusion Workshop.

Presented two guest lectures at Carnegie Mellon University for “Introduction to Machine Learning” (10-701) (2018).

Participated in teaching and learning summit held by the Eberly Center Teaching Excellence & Educational Innovation at Carnegie Mellon University (2018).

Worked as a teaching assistant (TA) at The University of Texas at Austin for “Information Retrieval and Web Search” (CS 371R) taught by Prof. Raymond Mooney (2015). Responsibilities included designing rubric for homework assignments, holding office hours, and grading.

Participated in teaching seminars hosted by Sanger Learning Center and the Center for Teaching and Learning at The University of Texas at Austin that included a workshop on “Mastering TA-to-Student Communication” and a microteach seminar to practice teaching strategies and receive constructive feedback (2015).

› Service

Area Chair for NeurIPS 2025, ICML 2025, NeurIPS 2024

NSF Reviewer 2020, NSF Panel 2019

Reviewer awards

ICML 2022 Top Reviewer (Top 10% of reviewers)

ICLR 2022 Highlighted Reviewer (Top 9% of reviewers)

NeurIPS 2021 Outstanding Reviewer Award (Top 8% of reviewers)

Reviewer

ICLR (2025, 2023, 2022, 2021)

ICML (2025AC, 2024, 2022, 2021, 2020, 2019)

NeurIPS (2025AC, 2024AC, 2023, 2022, 2021, 2020, 2019, 2018)

AISTATS (2024, 2023, 2022, 2019, 2017, 2016)

CVPR (2023, 2021) ICCV (2025)

Others: ICML 2025 Workshops, UAI (2025), AAI (2018, 2015), JMLR (2023, 2022), TMLR (2024, 2023, 2022), TPAMI, TNNLS, TKDE, WWWJ, Bayesian Analysis, MLHC (2017, 2016),

ICDM (2015), Machine Learning, Annals of Applied Statistics

Faculty Search Committee (Fall 2020, Spring 2021, Fall 2023, Spring 2024), Graduate Committee (Fall 2020, Spring 2021, Fall 2021, Fall 2022, Spring 2023), Graduate Admissions Committee (Spring 2020)

Organized Purdue ECE’s Machine Learning Reading Group (Fall 2021-present)

Panelist at Purdue’s Engineering Academic Careers Club event called “Obtaining a Faculty Position.”

Taught “Hour of Code” <<https://code.org/learn>> to middle school and high school students (2016).

Led English conversation group for international students in research lab (2015).

Served on the Masters in CS Admissions Committee at The University of Texas at Austin by reviewing candidate applications (2013).

Organized field day for underprivileged students as part of Bright Futures Atlanta, an inner-city academy and after-school program for underprivileged children (2011). Tutored underprivileged high school student from Bright Futures Atlanta (2012).

Served as President of the Ultimate Frisbee club at Covenant College (2008-2010).

Worked with underprivileged children and helped tsunami relief efforts with a group of students (2008).