Zeyu Zhou

zyzhou3@gmail.com (765) 418-0346 www.linkedin.com/in/zyzhou3 https://zyzhou3.github.io

Research Keywords

Robust and Trustworthy ML, Distribution Shift, Out-Of-Distribution Generalization, Causality, Generative AI

Educational Background

Purdue University, West Lafayette, IN May 2025

PhD in Electrical and Computer Engineering (ECE) Major GPA: 3.97/4.0

- Began PhD in Physics in 2019 and transferred to ECE in 2021

Wuhan University, Hubei, China June 2019

Bachelor of Science in Physics GPA: 3.86/4.0

Cornell University, Ithaca, NY December 2017

- Visiting student for one semester in the Department of Physics

Selected Publications

Zhou, Z., Liu, T., Bai, R., Gao J., Kocaoglu, M., & Inouye, D. I. (2024). Counterfactual Fairness by Combining Factual and Counterfactual Predictions. NeurIPS 2024.

Zhou, Z.*, Bai, R.*, Kulinski, S.*, Kocaoglu, M., & Inouye, D. I. (2024). Towards Characterizing Domain Counterfactuals For Invertible Latent Causal Models. ICLR 2024. (*equal contributions)

Zhou, Z., Azam, S. S., Brinton, C. G. & Inouye, D. I. (2023). Efficient Federated Domain Translation. ICLR 2023.

Zhou, Z., Gong, Z., Ravikumar, P. & Inouye, D. I. (2022). Iterative Alignment Flows. AISTATS 2022.

Work Experience

Machine Learning Quant Research Intern June 2023 – August 2023

***Bloomberg**, New York, NY*

- Developed recurrent neural network and transformer-based multivariate time series estimation models for real-world financial datasets characterized by significant noise, high missing rates, and potential absence of strong temporal signals; achieved a 40% improvement over existing approaches and a 20% improvement over traditional machine learning baselines.
- Conducted comprehensive research aimed at enhancing the model's capability of out-of-distribution generalization by leveraging domain knowledge including the option pricing formula under the Black-Scholes model.
- Modeled intraday pricing for a universe of equity options, many with sparse trading price data, serving as important input for next generation of improved Bloomberg volatility data products with tens of millions of annual revenue, along with a possible new stand-alone data product for trading and risk management.

Applied Scientist Intern June 2022 – September 2022

***Amazon**, Seattle, WA*

- Built Generative AI models to learn the data generating process from noisy and large real-world datasets at Amazon; tested and validated the scalability and robustness of the training framework.
- Provided a solid metric for evaluating downstream causal inference models for heterogeneous treatment effect (HTE) in the difficult real-world scenario where the ground truth data is unavailable and its distribution changes rapidly, by comparing the estimated HTE with that computed from the ground truth in synthetic datasets.
- Explored how to use deep generative models to capture the fixed effects in real-world datasets, which is essential for the downstream models to mitigate hidden confounding effects.
- Collaborated with a diverse team of economists, engineers, and applied scientists; earned trust from others, including senior stakeholders, through well-structured presentation, clear communication, and the delivery of convincing results.

Research Experience

Graduate Research Assistant June 2020 – Present

Advisor: Prof. David Inouye, Purdue University, West Lafayette, IN

Counterfactual Fairness

- Propose a method that is provably optimal in terms of predictive performance under the constraint of perfect counterfactual fairness, and characterize the inherent trade-off between counterfactual fairness and predictive performance.
- Derive bounds on the errors from estimated counterfactuals due to limited causal knowledge in practical scenarios, and introduce a novel objective function to mitigate the impact of these errors.
- Develop a post-processing algorithm that leverages estimated counterfactuals to mitigate biases in pretrained models, such as foundation models, without requiring retraining, while preserving the predictive performance.

Causality and Domain Counterfactual Generation

- Provide a practical yet theoretically grounded approach to estimating domain counterfactuals, which hypothesizes that causal variables are unobserved, and there is only access to observational data and domain labels.
- Derive a necessary and sufficient characterization of domain counterfactual equivalence; prove all invertible latent causal models can be split into disjoint equivalence classes with respect to their sparsity.
- Propose a novel model design of Generative AI models for domain counterfactual generation inspired by the theory; validate the methods through extensive simulated and image-based experiments.
- Use generated counterfactuals for downstream tasks, including improving fairness and out-of-distribution generalization.

Efficient Federated Domain Generalization

- Explore class conditional distribution shift in Federated Learning (FL), which is harder than the conventional non-IID setting of class imbalance or missing classes across clients.
- Propose a federated domain translation method that is easier to train than existing translation models and is more resource-efficient in terms of both communication and computation.
- Utilize the translation model to generate synthetic datasets for each client which could be useful for multiple downstream tasks; demonstrate that this methodology enables use of state-of-the-art domain generalization methods in a federated setting which enhances accuracy and robustness to increases in the synchronization period.

Unsupervised Domain Adaptation

- Build iterative distribution matching models by decomposing high dimensional distributions into orthogonal components and solving the 1D matching problems via closed-form solutions based on optimal transport theory; achieve faster and easier matching in comparison to most Generative Adversarial Networks (GANs) and flow-based methods.
- Achieve domain adaptation via direct distribution matching using invertible deep generative models.

Generative AI

- Build invertible deep generative models based on neural networks and iterative models; apply the models for tasks including density estimation, generative modeling, and distribution matching.
- Become proficient with the implementation of flow models, GANs, diffusion models, and VAEs using PyTorch.

Skills

Programming Skills: Python, MATLAB, LaTeX, SQL   *Deep Learning Framework*: PyTorch

Other tools: Transformers, PyTorch Geometric, Hadoop, Diffusers

Honors and Awards

- National Scholarship (award top 1% in the major)   2015 – 2016/ 2016 – 2017

Professional Activities

Talks

Machine Learning Reading Group, Purdue University, West Lafayette, IN

- AI for Science and Physics-inspired AI   November 2023
- Diffusion Models   October 2022

Conference Reviewer

- ICML2023, AISTATS 2024, ICLR 2024, ICML 2024, NeurIPS 2024, AAAI 2025