Siyuan Guo

Curriculum Vitae

Correnstrasse 38
Tübingen, Germany
+49 1515 8778109
sguo26v@gmail.com
www.siyuanguo.com
siyuan.guo

Summary

Researcher at the intersection of the science of AI and AI for science, motivated by curiosity to understand intelligence. In 2025 she launched the Physics of Learning program; within a week of its arXiv release it prompted a fireside-chat invitation, multiple talk invitations and collaboration discussions with infrastructure support. Her work spans theory to application, most recently advancing pretrained foundation models on in-context causal inference for real-world scientific use cases. She thrives in cross-disciplinary teams and focuses on understanding intelligence and advancing science.

Education

2021–2025 **Ph.D. in Machine Learning**, *University of Cambridge & Max Planck Institute for Intelligent Systems & ELLIS*, United Kingdom and Germany
Thesis: Statistical Foundations for Causality and Machine Learning for Causal Inference
Supervisors: Prof. Ferenc Huszár and Prof. Bernhard Schölkopf
Expected Graduation Date: 2026.01

2020–2021 **Msc in Machine Learning**, *University College London*, UK, Distinction

2015–2019 **BA(Hons) and MMath in Mathematics**, *University of Cambridge*, UK, Distinction

Selected Publication

2025-09 **Siyuan Guo** and Bernhard Schölkopf. Physics of learning: A lagrangian perspective to different learning paradigms, *Arxiv Preprint*
TL;DR. We study the problem of building an efficient learning system; We show that learning follows the principle of least action and learning algorithms are derivable by seeking the stationary trajectories. Received multiple invitations to talks, fireside chats, and collaboration requests within one week of the Arxiv release.

2025-06 Jake Robertson*, Arik Reuter*, **Siyuan Guo**, Noah Hollmann, Frank Hutter, and Bernhard Schölkopf. Do-PFN: In-context learning for causal effect estimation. In NeurIPS 2025 *spotlight*, *Arxiv Version*
TL;DR. We pre-trained a foundation model on synthetic data for in-context causal inference. Demonstrate a modern deep learning based approach for practitioners to perform data science, with potential scientific and enterprise applications.

2022-03 **Siyuan Guo***, Viktor Tóth*, Bernhard Schölkopf, and Ferenc Huszár. Causal de Finetti: On the identification of invariant causal structure in exchangeable data. In NeurIPS 2023, *Arxiv Version*
TL;DR. We prove a fundamental representation theorem, Causal de Finetti, and derive core causal principles from standard Bayesian assumptions—replacing reliance on ad hoc, untestable postulates. The work has since been extended, leading to total 4 publications at NeurIPS and ICLR with 2 as **oral** and **spotlight** presentations.

Professional Experience

2024 **Research Scientist Intern**, *Meta FAIR (Fundamental AI Research)*, New York

2019–2020 **Quantitative Strategist**, *Goldman Sachs International*, London

Teaching

Spring 2025 Co-Instructor @ ETH Zurich
Seminar: Pattern Matching Beyond i.i.d. Data
Co-led with Prof. Bernhard Schölkopf, Dr. Simon Buchholtz, Dr. Antonio Orvieto.
Co-designed readings/assessments, held office hours, and evaluated student work.

July 2025 Invited Lecturer (Causality Day, opening lecture)
Cambridge Ellis Summer School on Probabilistic Machine Learning, 2025
Re-invited to deliver the opening lecture and curate the day's agenda and speaker lineup, following strong participant feedback in 2024.

July 2024 Invited Lecturer
Cambridge Ellis Summer School on Probabilistic Machine Learning , 2024
Lecture: *Causal de Finetti & Do Finetti: Bridging Causality and Probabilistic Models*.
Audience: ~ 100 graduate students, researchers and industry professionals. [Slides]

May 2024 Guest Lecturer
Causal Inference and AI Class @ Oregon State University
Lecture: *Causal de Finetti & Do Finetti: Bridging Causality and Probabilistic Models*.
Audience: graduate students.

Mentoring & Supervision

I take great pleasure in (co-)mentoring a few talented and highly motivated students. It has been a great privilege for me to constantly learn and get inspired by them.

2024– **Moritz Miller**, *M.Sc. in Statistics @ ETH Zurich*
Supervising Master's Thesis and proposed research direction.
Project: Understanding how transformers perform counterfactual reasoning.
Outcome: Accepted at NeurIPS 2025. (ArXiv)

2024– **Chao Chen**, *M.Sc. Computer Science, Machine Intelligence Track @ ETH Zurich*
Supervising Master's Thesis and proposed research direction.
Project: Understanding the representation and composition of linguistic structure.
Outcome: Under submission for ACL 2026.

2024– **Anna Kerekes**, *Ph.D. Computer Science @ ETH Zurich*
Mentoring junior Ph.D. and proposed the research direction.
Project: Statistical Foundations of (Causal) Structure in Time-Series.
Outcome: Ongoing.

2024 **Bálint Mucsányi**, *M.Sc. Machine Learning @ University of Tübingen*
co-mentored with Patrik Reizinger.
Project: Investigate how policy diversity affects identifiability of learned representations in RL.
Outcome: Preprint available (arXiv); under review at ICLR 2026.

2024 **Szilvia Ujváry**, *Ph.D. Machine Learning @ University of Cambridge*
co-mentored with Prof. Ferenc Huszár.
Project: Algorithmic Information Extension of Causal de Finetti.
See Ferenc's blog post on the very interesting research proposal.

Full Publication List

[1] **Siyuan Guo** and Bernhard Schölkopf. Physics of learning: A lagrangian perspective to different learning paradigms. *arxiv preprint arXiv:2509.21049*, 2025.

[2] Jake Robertson*, Arik Reuter*, **Siyuan Guo**, Noah Hollmann, Frank Hutter, and Bernhard Schölkopf. Do-pfn: In-context learning for causal effect estimation. In *NeurIPS*, 2025. *joint first. NeurIPS **spotlight** with acceptance rate **3.19%**.

[3] Moritz Miller, Bernhard Schölkopf, and **Siyuan Guo**. Counterfactual reasoning: an analysis of in-context emergence. In *NeurIPS*, 2025.

[4] Patrik Reizinger*, **Siyuan Guo***, Ferenc Huszár, Bernhard Schölkopf†, and Wieland Brendel†. Identifiable exchangeable mechanisms for causal structure and representation learning. In *ICLR*, 2025. *equal contribution. † equal supervision. ICLR **spotlight** with acceptance rate **5.1%**.

[5] Patrik Reizinger*, Bálint Mucsányi*, **Siyuan Guo***, Benjamin Eysenbach, Bernhard Schölkopf, and Wieland Brendel. Skill learning via policy diversity yields identifiable representations for reinforcement learning. *arXiv preprint arXiv:2507.14748*, 2025.

[6] Charles Arnal, Vivien Cabannes, Niklas Nolte, **Siyuan Guo**, Leon Bottou, and Bernhard Schölkopf. Exploring transformers' inductive biases with structured latent processes. 2025.

[7] **Siyuan Guo**, Chi Zhang, Karthika Mohan, Ferenc Huszár*, and Bernhard Schölkopf*. Do finetti: On causal effects for exchangeable data. In *NeurIPS*, 2024. *joint senior. NeurIPS **oral** with acceptance rate **0.46%**.

[8] **Siyuan Guo**, Jonas Wildberger, and Bernhard Schölkopf. Out-of-variable generalization for discriminative models. In *ICLR*, May 2024.

[9] Aniket Didolkar, Anirudh Goyal, Nan Rosemary Ke, **Siyuan Guo**, Michael Valko, Timothy Lillicrap, Danilo Rezende, Yoshua Bengio, Michael Mozer, and Sanjeev Arora. Metacognitive capabilities of llms: An exploration in mathematical problem solving. In *NeurIPS*, 2024.

[10] **Siyuan Guo**, Aniket Didolkar, Nan Rosemary Ke, Anirudh Goyal, Ferenc Huszár†, and Bernhard Schölkopf†. Learning beyond pattern matching? assaying mathematical understanding in llms. *arXiv preprint arXiv:2405.15485*, 2024. † joint senior authors.

[11] **Siyuan Guo***, Viktor Tóth*, Bernhard Schölkopf, and Ferenc Huszár. Causal de finetti: On the identification of invariant causal structure in exchangeable data. In *NeurIPS*, 2023. *joint first.

[12] Ishan Kumar, Zhijing Jin, Ehsan Mokhtarian, **Siyuan Guo**, Yuen Chen, Mrinmaya Sachan, and Bernhard Schölkopf. CausalCite: A causal formulation of paper citations. In *Findings of the Association for Computational Linguistics (ACL)*. Association for Computational Linguistics, August 2024.

[13] Limor Gultchin, **Siyuan Guo**, Alan Malek, Silvia Chiappa, and Ricardo Silva. Pragmatic fairness: Developing policies with outcome disparity control. In Francesco

Locatello and Vanessa Didelez, editors, *Proceedings of the Third Conference on Causal Learning and Reasoning*, volume 236 of *Proceedings of Machine Learning Research*, pages 243–264. PMLR, 01–03 Apr 2024.

[14] Jonas Wildberger, **Siyuan Guo**, Arnab Bhattacharyya, and Bernhard Schölkopf. On the interventional kullback-leibler divergence. In *Proceedings of the Second Conference on Causal Learning and Reasoning (CLeaR)*, volume 213 of *Proceedings of Machine Learning Research*, pages 328–349. PMLR, April 2023.

[15] Andrei Paleyes*, **Siyuan Guo***, Bernhard Schölkopf, and Neil Lawrence. Dataflow graphs as complete causal graphs. In *2nd International Conference on AI Engineering - Software Engineering for AI (CAIN)*, pages 7–12. IEEE, April 2023. *joint first.

[16] **Siyuan Guo**, Soo Ling Lim, and Peter J. Bentley. Teams frightened of failure fail more: Modelling reward sensitivity in teamwork. In *2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 109–116, 2020.

Invited Talks / Contributed Talks & Seminars

Nov 2025 **Invited Talk**, *Physics of Learning & Do-PFN*, kausable.ai, Heidelberg, Germany

Oct 2025 **Invited Talk**, *Physics of Learning*, Friday Talks @ Tübingen, Tübingen, Germany

Oct 2025 **Invited Fireside Chat**, *Physics of Learning*, Ploutos, Stream available [Link]

Oct 2025 **Invited Talk**, *Physics of Learning*, UCL Gatsby Computational Neuroscience Unit, Slides [Link]

April 2023 **Contributed Speaker**, *Causal representation learning workshop at Max Planck Institute for Intelligent Systems*, Germany
Contributed Talk on *Moving beyond I.I.D.: Causal de Finetti and OOV Generalization*

Jan 2023 **Invited Speaker**, *Causal Reading Group in Swarma*, China
Invited Talk on *Causal de Finetti* with more than 1000 attendees.

Dec 2022 **Invited Speaker**, *1st Causal Methods in Environmental Science Workshop*, UK
Invited Talk on *Causal de Finetti: On the Identification of Invariant Causal Structure in Exchangeable Data*. [Link]

Nov 2022 **Department Seminar**, *DeepMind ELLIS CSML Seminar Series, UCL*, UK
Invited Talk on *Causal de Finetti*. [Link]

May 2022 **Department Seminar**, *Computational Statistics & Machine Learning Group (OxC-SML) in University of Oxford*, UK

Professional Services

Organizing: Causal Representation Learning workshop [Link]
Causal Digital Twins at Ellis unconference

Reviewing: NeurIPS 2024, NeurIPS 2025, ICLR 2026, Journal of Causal Inference

Media, Awards and Fellowships

- Received MIT EECS Rising Star. [News Article]
- Received 2025 G-research PhD Prize University of Cambridge [News Article]

- Won 2024 MPI-IS Outstanding Female Doctoral Student Prize [News Article]
- Received Cambridge-Tübingen PhD Fellowship in Machine Learning [Link]
- Received Premium Research Studentship [Link]
- Awarded Dean's List at University College London
- Awarded Openshaw Prize, Exhibition & Foundation Scholarship at University of Cambridge for performances in mathematical Tripos [Link]