

LIU ZIYIN

Email: liu.ziyin.p@gmail.com / ziyinl@mit.edu

Website: <https://www.mit.edu/~ziyinl/>

Citizenship: China

Office: Research Laboratory of Electronics, 26-209, 77 Massachusetts Ave, Cambridge, MA 02139

RESEARCH INTEREST

- **Physics-Inspired AI Foundations:** Identifying the principles derived from physics and mathematics that underpin modern artificial intelligence systems.
- **Empirics of Artificial Neural Networks:** Investigating the behaviors, like phase transitions and emergence of capabilities, in artificial neural networks and large language models.
- **Biological and Artificial Network Connections:** Studying the similarities and dissimilarity between how artificial networks and biological systems form and utilize latent representations.
- **Efficient Algorithm:** Applying theory to devise efficient algorithms for differentiable relaxation problems.

EDUCATION

The University of Tokyo, Japan

Sep. 2023

Doctoral Advisor: Masahito Ueda

Masters (2018-2020) and PhD (2020-2023) in Physics

Doctoral Thesis: *Symmetry breaking in deep learning*

Carnegie Mellon University, Pittsburgh, USA

Aug. 2018

Bachelor of Science in Physics;

Bachelor of Science in Mathematics (Dual Degree)

Minor in Computer Science

EMPLOYMENT

Postdoctoral Fellow, Research Laboratory of Electronics, MIT

Mar. 2024 – Now

Collaborator: Isaac Chuang (MIT Physics, EECS)

Postdoctoral Fellow, Physics & Informatics Laboratories, NTT Research

Mar. 2024 – Now

Collaborator: Ryan Hamerly (MIT, NTT Research)

JSPS Postdoctoral Fellow, Department of Physics, UTokyo

Oct. 2023 – Mar. 2024

PROFESSIONAL ACTIVITY

Visiting Researcher, Center for Brain Science, Harvard University

Mar. 2022 – Dec. 2022

Host: Prof. Venkatesh Murthy

Research Intern, Physics & Informatics Labs, NTT Research

Mar. 2022 – Dec. 2022

Host: Dr. Hidenori Tanaka

Quantitative Finance Research Intern, Preferred Networks

Jul. 2020 – Oct. 2021

Research Intern, High-Dimensional Statistics Unit, RIKEN AIP

Jun. 2018

Host: Prof. Makoto Yamada

Research Assistant, Language Technology Institute, CMU

2017 – 2018

Host: Prof. Louis Philippe Morency

Research Assistant, Pittsburgh Supercomputing Center

2017

Host: Dr. Yang Wang

CONFERENCE PUBLICATIONS

(* denotes equal contribution or alphabetical ordering)

1. **Liu Ziyin**. Symmetry leads to structured constraint of learning. *ICML*, 2024
2. **Liu Ziyin** and Zihao Wang. Sparsity by redundancy: Solving l_1 with sgd. *ICML*, 2023
3. James B. Simon, Maksis Knutins, **Liu Ziyin**, Daniel Geisz, Abraham J. Fetterman, and Joshua Albrecht. On the stepwise nature of self-supervised learning. *ICML*, 2023
4. **Liu Ziyin**, Ekdeep Singh Lubana, Masahito Ueda, and Hidenori Tanaka. What shapes the loss landscape of self-supervised learning? *International Conference on Learning Representations*, 2022
5. Zihao Wang* and **Liu Ziyin***. Posterior collapse of a linear latent variable model. In *NeurIPS*, 2022 ([Oral](#))
6. **Liu Ziyin**, Botao Li, and Xiangming Meng. Exact solutions of a deep linear network. In *NeurIPS*, 2022
7. **Liu Ziyin**, Kentaro Minami, Kentaro Imajo, et al. Theoretically motivated data augmentation and regularization for portfolio construction. *3rd ACM International Conference on AI in Finance*, 2022
8. Takashi Mori, **Liu Ziyin**, Kangqiao Liu, and Masahito Ueda. Power-law escape rate of sgd. In *ICML*, 2022
9. **Liu Ziyin**, Botao Li, James B Simon, and Masahito Ueda. SGD can converge to local maxima. In *International Conference on Learning Representations*, 2022 ([Spotlight](#))
10. **Liu Ziyin**, Kangqiao Liu, Takashi Mori, and Masahito Ueda. Strength of minibatch noise in sgd. In *International Conference on Learning Representations*, 2022 ([Spotlight](#))
11. Zhang Zhiyi* and **Liu Ziyin***. On the distributional properties of adaptive gradients. In *UAI 2021*, 2021
12. Kangqiao Liu*, **Liu Ziyin***, and Masahito Ueda. Noise and fluctuation of finite learning rate stochastic gradient descent. In *ICML 2021*, 2021
13. Paul Pu Liang, Peter Wu, **Liu Ziyin**, Louis-Philippe Morency, and Ruslan Salakhutdinov. Cross-modal generalization: Learning in low resource modalities via meta-alignment. *ACM Multimedia*, 2021. ([Oral](#))
14. **Liu Ziyin**, Tilman Hartwig, and Masahito Ueda. Neural networks fail to learn periodic functions and how to fix it. *Advances in Neural Information Processing Systems*, 33, 2020
15. **Liu Ziyin**, Zhikang Wang, Paul Pu Liang, Russ R Salakhutdinov, Louis-Philippe Morency, and Masahito Ueda. Deep gamblers: Learning to abstain with portfolio theory. In *NeurIPS 2019*, pages 10622–10632, 2019

16. Paul Pu Liang, **Liu Ziyin**, AmirAli Bagher Zadeh, and Louis-Philippe Morency. Multimodal language analysis with recurrent multistage fusion. In *EMNLP 2018*, pages 150–161, 2018. ([Oral](#))

WORKSHOP PAPERS

1. Yizhou Xu and **Liu Ziyin**. Three mechanisms of feature learning in an analytically solvable model, 2024
2. **Liu Ziyin**, Ekdeep Singh Lubana, Masahito Ueda, and Hidenori Tanaka. What shapes the loss landscape of self-supervised learning? *NeurIPS 2022 Workshop on Symmetry and Geometry in Neural Representations*, 2022
3. **Liu Ziyin**, Ekdeep Singh Lubana, Masahito Ueda, and Hidenori Tanaka. Loss landscape of self-supervised learning. *NeurIPS 2022 Workshop: Self-Supervised Learning - Theory and Practice*, 2022
4. Paul Pu Liang, Peter Wu, **Liu Ziyin**, Louis-Philippe Morency, and Ruslan Salakhutdinov. Learning in low-resource modalities via cross-modal generalization. *Meta-Learning Workshop of NeurIPS 2020*, 2020
5. Paul Pu Liang, Terrance Liu, **Liu Ziyin**, Ruslan Salakhutdinov, and Louis-Philippe Morency. Think locally, act globally: Federated learning with local and global representations. *Federated Learning Workshop of NeurIPS 2019*, 2019 ([Distinguished Paper Award](#))
6. **Ziyin Liu**, Yao-Hung Hubert Tsai, Makoto Yamada, and Ruslan Salakhutdinov. Semi-supervised pairing via basis-sharing wasserstein matching auto-encoder. *Bayesian Deep Learning Workshop of NeurIPS 2018*
7. **Ziyin Liu**, Junxiang Chen, Paul Pu Liang, and Masahito Ueda. Relational attention networks via fully-connected conditional random fields. *Bayesian Deep Learning Workshop of NeurIPS 2018*

JOURNAL PUBLICATIONS

(* denotes equal contribution or alphabetical ordering)

1. **Liu Ziyin** and Masahito Ueda. Zeroth, first, and second-order phase transitions in deep neural networks. *Physical Review Research*, 2023.
2. **Liu Ziyin**, Botao Li, and Xiangming Meng. Exact solutions of a deep linear network. *Journal of Statistical Mechanics: Theory and Experiment*, 2023 (**by invitation**).
3. **Liu Ziyin** and Masahito Ueda. Universal thermodynamic uncertainty relation in non-equilibrium dynamics. *Physical Review Research*, 2022
4. **Liu Ziyin**, Katsuya Ito, Kentaro Imajo, and Kentaro Minami. Power laws and symmetries in a minimal model of financial market economy. In *Physical Review Research*, 2022

PREPRINTS

1. **Liu Ziyin**, Mingze Wang, and Lei Wu. The implicit bias of gradient noise: A symmetry perspective, 2024. Preprint.
2. Yizhou Xu* and **Liu Ziyin***. When does feature learning happen? perspective from an analytically solvable model, 2024. Preprint.
3. **Liu Ziyin**, Hongchao Li, and Masahito Ueda. Law of balance and stationary distribution of stochastic gradient descent. *arXiv preprint arXiv:2308.06671*, 2023. Preprint.

4. **Liu Ziyin**, Botao Li, Tomer Galanti, and Masahito Ueda. The probabilistic stability of stochastic gradient descent. *arXiv preprint arXiv:2303.13093*, 2023. Preprint.

SEMINARS AND INVITED TALKS

1. *Parameter Symmetry and Formation of Latent Representations*
MIT IAIFI Summer Workshop Talk
2. *Symmetry and Formation of Hidden Representations*
Gatsby Computational Neuroscience Seminar
University College London, June, 2024
3. *How does loss function symmetry affect the learning of neural networks?*
Institute for Physics of Intelligence Seminar
The University of Tokyo, November, 2023
4. *What shapes the loss landscape of Self-Supervised Learning?*
Okinawa Institute of Science and Technology (OIST), April, 2023
5. *Collapse and Phase Transition in Deep Learning*
Seminar of Statistical Physics and Statistical Sciences
The Institute of Statistical Mathematics, Japan, March 07, 2023
6. *What shapes the loss landscape of Self-Supervised Learning?*
NTT Japan, Musashino R&D Center, February, 2023
7. *The Probabilistic Stability and Low-Rank Bias of SGD*
Mathematical Machine Learning Seminar
Max Planck Institute & UCLA, January, 2023
8. *Collapse and Phase Transition in Deep Learning.*
Redwood Center for Theoretical Neuroscience, University of California, Berkeley, December, 2022
9. *Collapse and Phase Transition in Deep Learning.*
Tomaso Poggio Lab, MIT, November, 2022
10. *Collapse and Phase Transition in Deep Learning*
NeuroTheory Seminar
Harvard University, October 11th, 2022
11. *Does the Refractory Period Help Learning? A Spiking Neural Network Perspective.*
Harvard-LMU Young Scientists' Forum, 2022
12. *Stochastic Gradient Descent with Multiplicative Noise.*
RIKEN AIP, 2021
13. *Careful Deep Learning: Learning to Abstain by Training on A Simple Loss Function.*
Institute for Physics of Intelligence Seminar
The University of Tokyo, 2019
14. *A Full Potential Approach to the Solution of Core States.*
American Physical Society March Conference, 2018

SCHOLARSHIP / GRANT / AWARDS

NeurIPS 2022 Scholar Award & Travel Grant ($\sim \$2000$)	2022
Japan Society for the Promotion of Science Fellowship & Grant ($\sim \$10000$ per year)	2022-2023

PROFESSIONAL SERVICE

Conference Referee:

International Conference on Machine Learning (ICML)
International Joint Conference on Artificial Intelligence (IJCAI)
Conference on Computer Vision and Pattern Recognition (CVPR)
International Conference on Computer Vision (ICCV)
International Conference on Artificial Intelligence and Statistics (AISTATS)
The Conference on Uncertainty in Artificial Intelligence (UAI)
Neural Information Processing Systems (NeurIPS)
International Conference on Learning Representations (ICLR)
Special Interest Group on Knowledge Discovery and Data Mining (KDD)
International Conference on Data Mining (ICDM)
and workshop affiliated thereto.

Journal Referee:

Transactions on Machine Learning Research (TMLR)
Journal of Machine Learning Research (JMLR)
IEEE Transactions on Signal Processing (IEEE-TSP)
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)
KDD (Special Interest Group on Knowledge Discovery and Data Mining)
IEEE Transactions on Neural Networks and Learning Systems (IEEE-TNNLS)
IEEE Signal Processing Letters

Organizer:

Institute for Physics of Intelligence Seminars (The University of Tokyo)
Advanced Statistics Study Group (The University of Tokyo)

HOBBY

Play Go (ancient Chinese/Japanese board game), Draw, Philosophy, Poems, Anime