Tianren Zhang

Website: trzhang0116.github.io

✉: zhangtr22@mails.tsinghua.edu.cn

EDUCATION

Tsinghua University   Beijing, China

Ph.D. in Automation (Combined Master's and Ph.D. Program)   2019 – 2025

- Thesis: "Representation Learning: An Out-of-Distribution Generalization Perspective"
- Advisor: Feng Chen

Tsinghua University   Beijing, China

B.S. in Automation   2015 – 2019

- Thesis: "Curiosity-Driven Algorithms in Reinforcement Learning"

RESEARCH INTERESTS

I am interested in two intertwined areas of machine learning: **representation learning** and **generalization**. I believe that a good representation is all we need for human-level generalization and is key to achieving general intelligence.

My recent research primarily focuses on the **theoretical foundations** of *learning world representations* and *out-of-distribution generalization*, as well as translating these theoretical principles into efficient algorithms to solve real-world problems. I also work on **reinforcement learning**.

SELECTED RESEARCH (* indicates equal contributions, † indicates equal correspondance)

1. **Tianren Zhang**, Guanyu Chen, and Feng Chen. When do Neural Networks Learn World Models? *ICLR 2025 Workshop on World Models*, 2025.

   TL;DR: We show that in a multi-task setting, prediction models with a low-degree bias can provably identify latent data-generating variables under mild assumptions.

2. **Tianren Zhang**, Yizhou Jiang, and Feng Chen. A Method of Supervised Learning from Conflicting Data with Hidden Contexts. *arXiv preprint*, 2025.

   TL;DR: A formulation and a theoretically grounded method for the problem of open-ended training on data with hidden contexts.

3. **Tianren Zhang**, Guanyu Chen, Chujie Zhao, and Feng Chen. Gradient Canonicalization: Can You Trust Your Gradients in Fine-Tuning? *In submission to ICML 2025.*

   TL;DR: We identify that the selectivity of neurons can induce spurious gradient updates in fine-tuning foundation models and provide an effective fix.

4. Yizhou Jiang, **Tianren Zhang**, Yihan Li, Yuqian Liu, Haichuan Gao, Ying Fang[†], and Feng Chen[†]. Adaptive Fission: Post-Training Encoding for Low-Latency SNNs. *In submission to ICML 2025.*

5. Guanyu Chen*, Peiyang Wang*, Yizhou Jiang, ..., **Tianren Zhang**[†], and Feng Chen[†]. Exploring the Hidden Reasoning Process of Large Language Models by Misleading Them. *arXiv preprint*, 2025.

TL;DR: We show that LLMs can generalize false mathematical reasoning rules to real-world problems, implying the existence of an "abstract-then-reason" process in LLMs.

6. **Tianren Zhang***, Chujie Zhao*, Yizhou Jiang, and Feng Chen. Feature Contamination: Neural Networks Learn Uncorrelated Features and Fail to Generalize. In *International Conference on Machine Learning (ICML)*, 2024.
TL;DR: We identify that neural networks can learn task-irrelevant features due to an implicit bias of SGD, resulting in a failure to generalize under distribution shifts.

7. Chujie Zhao*, **Tianren Zhang***, Guanyu Chen, Yizhou Jiang, and Feng Chen. M^{3}PL: Identifying and Exploiting View Bias of Prompt Learning. In *Transactions on Machine Learning Research (TMLR)*, 2024.
TL;DR: We identify a view bias in prompt learning of foundation models, i.e., it may extract only a partial subset of useful features while ignoring others, and provide an effective fix.

8. Yizhou Jiang*, Kunlin Hu*, **Tianren Zhang**, Haichuan Gao, Yuqian Liu, Ying Fang[†], and Feng Chen[†]. Spatio-Temporal Approximation: A Training-Free SNN Conversion for Transformers. In *International Conference on Learning Representations (ICLR)*, 2024.
TL;DR: We propose the first training-free method for converting transformers to purely event-driven spiking neural networks.

9. Haichuan Gao, **Tianren Zhang**, Zhile Yang, Yuqing Guo, Jinsheng Ren, Shangqi Guo[†], and Feng Chen[†]. Fast Counterfactual Inference for History-Based Reinforcement Learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2023.
TL;DR: We propose a tree-based counterfactual inference method for learning history representations in reinforcement learning.

10. **Tianren Zhang***[†], Shangqi Guo*[†], Tian Tan, Xiaolin Hu, and Feng Chen[†]. Adjacency Constraint for Efficient Hierarchical Reinforcement Learning. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2022.

11. **Tianren Zhang**, Yizhou Jiang, Xin Su, Shangqi Guo, and Feng Chen. Subjective Learning for Conflicting Data. In *ICLR 2022 Workshop on Agent Learning in Open-Endedness*, 2022.
TL;DR: An initial attempt of formulating and addressing the problem of data conflicts in open-ended learning.

12. Chongkai Gao, Haichuan Gao, Shangqi Guo, **Tianren Zhang**, and Feng Chen. CRIL: Continual Robot Imitation Learning via Generative and Prediction Model. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
TL;DR: A continual imitation learning method for robot learning based on generation and prediction models.

13. **Tianren Zhang***, Shangqi Guo*, Tian Tan, Xiaolin Hu[†], and Feng Chen[†]. Generating Adjacency-Constrained Subgoals for Hierarchical Reinforcement Learning. In *Conference on Neural Information Processing Systems (NeurIPS)* (**spotlight**), 2020.
TL;DR: We show that a state representation based on state adjacency can significantly improve the sample efficiency of hierarchical reinforcement learning.

14. Guangyi Chen, **Tianren Zhang**, Jiwen Lu, and Jie Zhou. Deep Meta Metric Learning. In *International Conference on Computer Vision (ICCV)*, 2019.

Skill Set

- **Machine Learning:** Deep Learning, Reinforcement Learning, Fine-Tuning Foundation Models, Neural Network Quantization, Brain-Inspired Computing & Spiking Neural Networks

- **Programming:** Python, C/C++, MATLAB, PyTorch
- **Theory:** Statistical Learning Theory, Deep Learning Theory (Optimization & Generalization & Representation Learning), Reinforcement Learning Theory

EXPERIENCE

Qianjue Robots   Beijing, China
CTO   June 2024 – Present

- Start-Up Company on Brain-Inspired Embodied Intelligence
- Leading the technical team in training, fine-tuning, and deployment of the "Qianjue-Thalamus" and "Qianjue-Brain" perception & decision models on neuromorphic chips.
- The project has secured **tens of millions** of RMB in financing.

Cambricon   Beijing, China
Research Intern   July 2018 – September 2018

- Research Area: Quantization of Deep Neural Networks
- Implemented 2-bit weight quantization for deep neural networks such as VGG and AlexNet, along with 8-bit quantization for activations, errors, and gradients.
- Successfully trained quantized networks on MNIST, CIFAR-10, and ImageNet datasets.

Tsinghua University   Beijing, China
Research Intern at i-Vision Group, advised by Jiwen Lu   March 2018 – July 2018

- Research Area: Metric Learning
- Introduced a meta-learning formulation into the metric learning framework, with a novel objective function to better handle the relationships between positive and negative samples.
- Improved the robustness of the learned metric space, achieving state-of-the-art performance in person re-identification and vehicle re-identification tasks.

TEACHING

- **Teaching Assistant** for "Matrix Analysis and Applications", Tsinghua University   2020 – 2024
  *Co-authored the handout for the course; now co-writing a **textbook** on matrix analysis.*

ACADEMIC SERVICES

Conference Reviewer

- International Conference on Machine Learning (ICML)   2025
- Conference on Neural Information Processing Systems (NeurIPS)   2024
- International Conference on Representation Learning (ICLR)   2025
- International Conference on Artificial Intelligence and Statistics (AISTATS)   2025
- ACM Multimedia (ACM MM)   2024
- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)   2024 – 2025

    – International Conference on Computer Vision (ICCV)   2025

Journal Reviewer

    – IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)

    – Artificial Intelligence Review

SCHOLARSHIPS AND AWARDS

- Tsinghua University Comprehensive Excellence Scholarship   2024
- Tsinghua University Comprehensive Excellence Scholarship   2021 – 2022
- Tsinghua University Outstanding Teaching Assistant Award (**top 4** in the department)   2023
- Tsinghua University Academic Excellence Scholarship   2017
- Second Prize for Student Laboratory Construction Contribution at Tsinghua University   2017
- Third Place in Tsinghua University Ma Yuehan Cup Go Competition   2017