

Junyu Zhang

✉ junyuz6@illinois.edu |  [jyzhang1208](https://github.com/jyzhang1208) |  [Homepage](#) |  [Google Scholar](#)



EDUCATION

University of Illinois Urbana-Champaign

Illinois, US

M.S. in Computer Science (Research Track), Siebel Scholar (Class of 2026)

Aug 2024 - May 2026

- **GPA: 4.0/4.0**
- **Research Advisor:** Prof. Huan Zhang

Huazhong University of Science and Technology

Hubei, China

*B.E. in Artificial Intelligence, **Honor Class** of Artificial Intelligence*

Sept 2020 - June 2024

- **GPA: 3.89/4; Rank: 1/29** (selected from 360 students in the school)
- **English Proficiency:** TOEFL 105 (Speaking 24); GRE 325+3.5 (Verbal 156, Quantitative 169)

RESEARCH INTERESTS

In the era of large-scale foundation models, my research goal is to incentivize the capabilities of pre-trained models by empowering them to perform *long-term reasoning* and *learn from interaction*, thereby building scalable and generalizable AI for the open world.

- **Large Language Models Reasoning:** efficient and theoretically grounded reasoning; embodied spatial reasoning for agentic decision-making; mathematical reasoning with scalable evaluation.
- **Learning from Interactions:** empirical reinforcement learning algorithms and their applications in robotics, including morphology–behavior co-optimization and robotic manipulation.

PUBLICATIONS AND PREPRINTS

1. **Junyu Zhang***, Yifan Sun*, Tianang Leng*, Jingyan Shen*, Ziyin Liu[†], Paul Pu Liang[†], Huan Zhang[†], “When Reasoning Meets Its Laws”, *Under review*
2. **Junyu Zhang***, Runpei Dong*, Han Wang, Xuying Ning, Haoran Geng, Peihao Li, Xialin He, Yutong Bai, Jitendra Malik, Saurabh Gupta, Huan Zhang, “AlphaOne: Reasoning Models Thinking Slow and Fast at Test Time”, in **EMNLP 2025 Main** [PDF] [Website]
3. Rui Yang*, Hanyang Chen*, **Junyu Zhang***, Mark Zhao*, Cheng Qian, Kangrui Wang, Qineng Wang, Teja Venkat Koripella, Marziyeh Movahedi, Manling Li, Heng Ji, Huan Zhang, Tong Zhang, “EmbodiedBench: Comprehensive Benchmarking Multi-modal Large Language Models for Vision-Driven Embodied Agents”, in **ICML 2025 Oral** [PDF] [Website] [Dataset]
4. Chengke Zou*, Xingang Guo*, Rui Yang*, **Junyu Zhang**, Bin Hu, Huan Zhang, “DynaMath: A Dynamic Visual Benchmark for Evaluating Mathematical Reasoning Robustness of Vision Language Models”, in **ICLR 2025** [PDF] [Website] [Dataset]
5. Heng Dong*, **Junyu Zhang***, Chongjie Zhang, “Leveraging Hyperbolic Embeddings for Coarse-to-Fine Robot Design”, in **ICLR 2024** [PDF] [Website]
6. Heng Dong, **Junyu Zhang**, Tonghan Wang, Chongjie Zhang, “Symmetry-Aware Robot Design with Structured Subgroups”, in **ICML 2023** [PDF] [Website]
7. Jianhao Wang*, Jin Zhang*, Haozhe Jiang, **Junyu Zhang**, Liwei Wang, Chongjie Zhang, “Offline Meta Reinforcement Learning with In-Distribution Online Adaptation”, in **ICML 2023** [PDF]

HONORS AND AWARDS

- Siebel Scholar Award (top 1%) - 2025
- Outstanding Graduate Honor - 2024
- Outstanding Undergraduate Student Award (top 1%) - 2022
- Excellent Academic Scholarship - 2021
- Science and Technology Innovation Scholarship - 2022, 2023
- The Second Prize of the World Robot Contest-BCI Brain Control Robot Contest - 2021
- Honorable Mention in Mathematical Contest in Modeling - 2022

RESEARCH EXPERIENCE

Research Intern - MIT Media Lab

Massachusetts, US

Supervisor: Prof. Paul Pu Liang

June 2025 - Present

Physics of Large Language Models Reasoning

- Developed a unified framework that formalizes intrinsic LRM reasoning via two fundamental laws.
- Introduced LoRe-Bench to evaluate whether current LRMs follow the laws, and proposed a simple yet effective fine-tuning approach that empirically validated our theory.
- Our project is submitted to ICLR 2026.

Research Assistant - Assured and Trustworthy AI Research Lab

Illinois, US

Supervisor: Prof. Huan Zhang

Aug 2024 - Present

Dynamic Benchmark for Mathematical Reasoning in Vision-Language Models

- Investigated the mathematical reasoning robustness and revealed limitations of SOTA VLMs.
- Proposed a dynamic benchmark capable of generating a large number of question variants and conducted an extensive evaluation of both closed-source and open-source VLMs.
- Our work is accepted by ICLR 2025.

Vision-Driven Embodied Agent Benchmark of Multi-modal Large Language Models

- Developed a standardized, multifaceted evaluation platform for automatically quantifying the performance of task planning in MLLM-based embodied agents.
- Created capability-oriented task datasets from high-level rearrangement to low-level manipulation and performed extensive experimental evaluations to further understand MLLM-based planning.
- Our work is accepted by ICML 2025.

Test-Time Scaling of Large Reasoning Models

- Introduced a universal framework for modulating reasoning progress in LRMs at test time.
- Unified and generalized existing monotonic scaling methods by enabling flexible and dense slow-to-fast reasoning modulation, demonstrating superior reasoning capability and efficiency.
- Our work is accepted by EMNLP 2025 Main.

Visual Preference Optimization (in progress)

- Proposed a learning framework that applies explicit, dense supervision to LVLMs by jointly optimizing both visual prior and language posterior representations.
- Distilled knowledge from contextualized to non-contextualized vision language models to enhance their alignment and performance.

Research Intern - MIT-IBM Watson AI Lab

Massachusetts, US (remote)

Supervisor: Prof. Chuang Gan

April 2023 - Mar 2024

Sequential Decision Making for Robotic Manipulation

- Proposed a novel framework that enabled efficient policy generalization in the offline multi-task and imitation learning settings.
- Incorporated mixture of experts layers into the transformer model that effectively harnesses the commonalities and discriminations of multimodal data.
- Evaluated our method on the RL Bench benchmark that demonstrated great generalization ability.

Research Intern - IIIS, Tsinghua University

Beijing, China

Supervisor: Prof. Chongjie Zhang

July 2022 - Jan 2024

Offline Meta Reinforcement Learning

- Revealed theoretical insights for offline meta-RL with online adaptation.
- Generated in-distribution context using a given uncertainty quantification and performed effective task belief inference to address new tasks.
- Our work is accepted by ICML 2023.

Robot Design via Reinforcement Learning

- Designed robots with various functionalities in simulated environments by using symmetry to exploit the structure of the robot design space with symmetry.
- Proposed a novel plug-and-play transformation module to map any robot design into a learned symmetry space and provided theoretical analysis to verify its rationality.
- Our work is accepted by ICML 2023.

Multi-cellular Soft Robot Design

- Inspired from real multi-cellular organisms and developed a novel algorithm to co-design soft robots in behavior and morphology.
- Introduced coarse-to-fine robot design strategy and conducted a comprehensive analysis of its benefits in the evolution of intelligent collectives
- Our project is accepted by ICLR 2024.

SPEAKING ENGAGEMENTS

AlphaOne: Reasoning Models Thinking Slow and Fast at Test Time

- Invited talk at NVIDIA AI Reasoning Team - *Aug 2025*

Unlocking Reasoning in Foundation Models

- Invited talk at Multisensory Intelligence Group, MIT Media Lab - *Aug 2025*

ACADEMIC SERVICE

Conference Reviewer

- International Conference on Learning Representations (ICLR) 2025, 2026
- Conference on Neural Information Processing Systems (NeurIPS) 2025
- ACL Rolling Review (ACL ARR) 2025

TEACHING EXPERIENCE

Teaching Assistant

- ECE598-Advanced Topics in Machine Learning and Formal Methods, University of Illinois Urbana-Champaign - *Fall 2024*.
- CS441-Applied Machine Learning, University of Illinois Urbana-Champaign - *Spring 2025*.
- CS107-Data Science Discovery, University of Illinois Urbana-Champaign - *Fall 2025*.

SKILLS SUMMARY

- **Programming Languages** Python, C/C++, Matlab, SQL, Bash
- **Languages** Chinese, English
- **Frameworks** PyTorch, TensorFlow, Keras, OpenCV, Scikit, etc.
- **Tools** PyCharm, VS Code, Markdown, Jupyter Notebook, Mobaxterm, Kubernetes, Git