

Tyler Zhu

✉ tylerzhu@princeton.edu | 🏠 tylerzhu.com/research | 📧 tyleryzhu | 🌐 tyleryzhu

Education

Princeton University

PHD IN COMPUTER SCIENCE
AUG 2023 - MAY 2028 (EXP)

Advised by Olga Russakovsky.

UC Berkeley

MASTER'S OF SCIENCE IN EECS
AUG 2022 - MAY 2023

Advised by Jitendra Malik.

UC Berkeley

BACHELOR'S OF SCIENCE W/ HONORS
AUG 2018 - MAY 2022

Electrical Engineering and Computer Science

Selected Berkeley Coursework

CS 182: Deep Learning A

CS 189: Machine Learning A-

CS 194-26: Intro Comp. Vision A

CS 271: Randomized Algorithms A

CS 280: Computer Vision A

CS 294-220: Computational Learning Theory A

EECS 126: Probability Theory A

EECS 127: Optimization Models A

Math 104: Real Analysis A

Service

Transformers for Vision Workshop

CVPR 2024, Co-organizer

Broadening Research Collaborations Workshop

NeurIPS 2022, Co-organizer

ECCV 2022 Reviewer

Wrote two reviews on behalf of an Outstanding Reviewer.

Selected Awards

Berkeley Putnam Team Dec. 2019

Top 4 at Berkeley, Rank 168/3428 nationally

ICPC PacNW 7th Place Nov. 2018

5-time AIME Qualifier Mar. 2018

USAMO Qualifier Apr. 2017

Top 300 in the nation

USACO Platinum Feb. 2017

Top 300 in the nation

Publications and Preprints

Unifying Specialized Visual Encoders for Video Language Models

Jihoon Chung*, Tyler Zhu*, Max Gonzalez Saez-Diez, Juan Carlos Niebles, Honglu Zhou, Olga Russakovsky. *In submission 2024, arXiv:2501.01426*

- Made a VideoLLM, *MERV*, w/ multiple visual experts for broader coverage of visual concepts. Up to 3.8% better than Video-LLaVA w/ same data on open-ended video QA. Efficient architecture w/ detailed ablations for design.
- Qualitative analysis shows orthogonal encoder capture, i.e., motion and general understanding, and different encoders activated for specific videos.

xT: Nested Tokenization for Larger Context in Large Images

Tyler Zhu*, Ritwik Gupta*, Shufan Li*, Jitendra Malik, Trevor Darrell, Karttikeya Mangalam. *ICML 2024, arXiv:2403.01915*

- Adapt vision models from small to large images using LLM-style encoders; new acc-param frontier on large images w/ up to 8.6% classif. acc boost.
- Allows modeling of 8x larger images using our recurrent region aggregation.

PaReprop: Fast Parallelized Reversible Backpropagation

Tyler Zhu* and Karttikeya Mangalam*. *Spotlight in Transformers for Vision Workshop at CVPR 2023, arXiv:2306.09342*

- Developed a CUDA-level boost to backpropagation which reduces training memory by 8x with up to 20% speedups over prior reversible methods.
- Our technique works across vision + NLP, four model families, and all model sizes from tiny to giant. Awarded Spotlight at T4V Workshop @ CVPR 2023.

The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization

D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dordono, R. Desai, T. Zhu, S. Parajuli, M. Guo, D. Song, J. Steinhardt, J. Gilmer. *ICCV 2021.*

- Collected a new dataset, ImageNet-R, for measuring robustness to renditions.
- Our new synthetic augmentation method raised robustness on ImageNet-R from 36.1% to 42.2% (SOTA); better than methods using 1,000x labeled data.

Experience

Discrete Math Head uGSI

UC BERKELEY EECS DEPARTMENT

Berkeley, CA

Aug. 2019 - Dec. 2020

- Head TA for discrete math course; ran logistics, made weekly section notes.
- Received 50+ detailed, positive reviews at end of Sp20; rated 4.52 vs. avg 4.41.

President

MACHINE LEARNING @ BERKELEY

Berkeley, CA

Jan. 2019 - May 2021

- Led student machine learning organization, managed industry partnerships, organized research talks and projects, and initiated mentorship programs.
- Established high school ML bootcamp for educating students new to CS.

Investment Trading Intern

CITADEL SECURITIES

New York, NY

June - Aug. 2020

- Worked on the semi-systematic single-stock options desk.

Software Engineering Intern

GOOGLE

Sunnyvale, CA

May - Aug. 2019

- Created an automated tool for removing unused and deprecated products.