

# Zitong Yang (杨紫童)

zitong@berkeley.edu • <https://zitongyang.github.io>

## Experience

**Thinking Machines Lab** - *Member of technical staff*  
[Research] Synthetic pretraining.

Stanford, CA, U.S.  
2025.12 - present

**Apple Inc.** - *Research scientist (part-time)*  
[Apple AI/ML] [Apple Foundation Model] Large-scale synthetic data pretraining.

Cupertino, CA, U.S.  
2025.02 - 2025.12

**Apple Inc.** - *Research scientist (intern)*  
[Apple AI/ML] [Apple Foundation Model] Language model synthetic data post-training.

Cupertino, CA, U.S.  
2024.06 - 2024.09

**Google Inc.** - *Research scientist (intern)*  
[Google Research] [BigML] Retrieval augmented language model.

New York, NY, U.S.  
2022.05 - 2022.09

**The Voleon Group** - *Quantitative researcher (full-time)*  
[Fixed Income Fund] [Research] Causal market impact modeling and alpha forecasting.

Berkeley, CA, U.S.  
2021.05 - 2022.05

**Lawrence Berkeley National Laboratory** - *Undergraduate researcher (part-time)*  
[Physics Division][ATLAS] Jet clustering for massive particles with pseudo-rapidity.

Berkeley, CA, U.S.  
2016.09 - 2017.05

## Education

**Stanford University**  
*Doctor of Philosophy in Statistics*  
Advisors: Emmanuel Candès and Tatsunori Hashimoto

Stanford, CA, U.S.  
2022.09 - 2025.12

**University of California, Berkeley**  
*Master of Science in Electrical Engineering*  
Advisors: Yi Ma and Jacob Steinhardt

Berkeley, CA, U.S.  
2020.08 - 2021.05

**University of California, Berkeley**  
*Bachelor of Arts in Mathematics (Highest honors)*  
*Bachelor of Arts in Computer Science*  
*Bachelor of Arts in Physics*

Berkeley, CA, U.S.  
2016.08 - 2020.05

## Publications

### 1. Synthetic bootstrapped pretraining

**Zitong Yang**\*, Aonan Zhang\*, Hong Liu, Tatsunori Hashimoto, Emmanuel Candès, Chong Wang, Ruoming Pang  
arXiv preprint, 2025 ([Paper](#))

### 2. s1: Simple test-time scaling

Niklas Muennighoff\*, **Zitong Yang**\*, Weijia Shi\*, Xiang Lisa Li\*, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, Tatsunori Hashimoto  
EMNLP, Oral Presentation, 2025 ([Paper](#), [Code](#), [Slides](#))

### 3. Synthetic Continued Pretraining

**Zitong Yang**\*, Neil Band\*, Shuangping Li, Emmanuel Candès, Tatsunori Hashimoto  
ICLR, Oral Presentation (1.8% of submitted papers), 2025 ([Paper](#), [Code](#), [Slides](#))

4. Bellman Conformal Inference: Calibrating Prediction Intervals For Time Series  
**Zitong Yang**, Emmanuel Candès, Lihua Lei  
arXiv preprint, 2024 ([Paper](#), [Code](#))
5. ResMem: Learn what you can and memorize the rest  
**Zitong Yang**, Michal Lukasik, Vaishnavh Nagarajan, Zonglin Li, Ankit Singh Rawat, Manzil Zaheer, Aditya Krishna Menon, Sanjiv Kumar  
NeurIPS, 2023 ([Paper](#), [Slides](#))
6. Predicting Out-of-Distribution Error with the Projection Norm.  
Yaodong Yu\*, **Zitong Yang**\*, Alex Wei, Yi Ma, Jacob Steinhardt  
ICML, 2022 ([Paper](#), [Code](#), [Slides](#), [Talk](#))
7. Understanding Generalization in Adversarial Training via the Bias-Variance Decomposition  
Yaodong Yu\*, **Zitong Yang**\*, Edgar Dobriban, Jacob Steinhardt, Yi Ma  
arXiv preprint, 2021 ([Paper](#), [Code](#), [Talk](#))
8. Exact Gap between Generalization Error and Uniform Convergence in Random Feature Models  
**Zitong Yang**, Song Mei, Yu Bai  
ICML, 2021 ([Paper](#), [Talk](#))
9. Rethinking the Bias-Variance Dilemma for Generalization of Neural Networks  
**Zitong Yang**\*, Yaodong Yu\*, Chong You, Jacob Steinhardt, Yi Ma  
ICML, 2020 ([Paper](#), [Talk](#))
10. Complete Dictionary Learning via  $L^4$ -Norm Maximization over the Orthogonal Group  
Yuxiang Zhai, **Zitong Yang**, Zhengyu Liao, John Wright, Yi Ma  
JMLR, 2020 ([Paper](#))

## Awards and Honors

- Discretionary research grant from [DST Global](#) (\$50,000), 2025
- Albion Walter Hewlett Stanford Graduate Fellowship, Stanford University, 2022
- High Distinction in General Scholarship, University of California, Berkeley, 2020
- Highest Distinction in Mathematics, Department of Mathematics, University of California, Berkeley, 2020
- Percy Lionel Davis Award, Department of Mathematics, University of California, Berkeley, 2020
- Pomerantz Endowment Award, Department of Physics (\$1000), University of California, Berkeley, 2020
- Summer Undergraduate Research Fellowships (\$5000), University of California, Berkeley, 2019

## Invited Talks

### “Two frontiers in AI scaling”

- Commonwealth Bank of Australia, *AI Frontier Series*, 2025
- Apple, *Foundation Model Tech Talk*, Cupertino, CA, 2025
- Citadel, *GQS*, Chicago, IL, 2025
- BlackRock, *PAIL*, Palo Alto, CA, 2025

### “s1: 1000 thoughts elicit test-compute scaling”

- Google DeepMind, *BigML Weekly*, New York, NY, 2025
- Nvidia, *Nemo LLM Reasoning*, Santa Clara, CA, 2025
- Beijing Academy of Artificial Intelligence (BAAI), *Zhiyuan Talk*, 2025

### “Scaling language model beyond human data limit”

- Beijing Academy of Artificial Intelligence (BAAI), *Zhiyuan Talk*, 2025
- Stanford University, *Stanford Industrial Affiliates Annual Conference*, Stanford, CA, 2024
- Stanford University, *Statistics Department Seminar*, Stanford, CA, 2024

**“Memorization for generalization: practical algorithm from theoretical insights”**

- Stanford University, *Stanford Industrial Affiliates Annual Conference*, Stanford, CA, 2023
- Google Research, *BigML Weekly*, New York, NY, 2022

## Professional Service

**Conference reviewer (13):**

- International Conference of Machine Learning (ICML, 10)
- International Conference of Learning and Representation (ICLR, 3)
- IEEE International Symposium on Information Theory (ISIT, 1)

**Journal reviewer (2):**

- Journal of Machine Learning Research (JMLR, 1)
- Journal of the Royal Statistical Society. Series B (JRSS-B, 1)

**Workshop organizer:**

- Stanford-Berkeley Joint Colloquium, 2023