PU YANG

yang_pu@pku.edu.cn | yangpuPKU | yangpupku.github.io | +86 15801209763

About me

I am a PhD Student at the School of Mathematical Sciences, Peking University, where I work with Prof. Bin Dong. I am generally interested in inverse problem, reinforcement learning and large language model. My recent research interests are exploring the use of synthetic data in large language models.

Education

Peking University — Beijing, China
Ph.D. major in Computational Mathematics — *Sep. 2021 – Present*

Peking University — Beijing, China
Undergraduate major in Mathematics — *Sep. 2017 – May. 2021*

Publication

L2SR: Learning to Sample and Reconstruct for Accelerated MRI via Reinforcement Learning
with Bin Dong

- Published in Inverse Problem, 2024
- We propose an alternating training framework for jointly learning a good pair of samplers and reconstructors via deep reinforcement learning for accelerated MRI.

A Tale of Tails: Model Collapse as a Change of Scaling Laws
with Elvis Dohmatob, Yunzhen Feng, Francois Charton, and Julia Kempe

- Published in International Conference on Machine Learning (ICML 2024)
- We theoretically and empirically explore how the inclusion of synthetic data in training affects neural scaling laws in large language models, revealing potential risks of model collapse.

Beyond Model Collapse: Scaling Up with Synthesized Data Requires Reinforcement
with Yunzhen Feng, Elvis Dohmatob, Francois Charton, and Julia Kempe

- Accept by International Conference on Learning Representations (ICLR 2025)
- We theoretically and empirically explore how feedback-augmented synthesized data can mitigate model collapse in fine-tuning large language models.

MoColl: Agent-Based Specific and General Model Collaboration for Image Captioning
with Bin Dong

- Under review in IEEE Transactions on Image Processing (TIP)
- We proposes a novel agent-enhanced model collaboration framework, designed to effectively integrate domain-specific and general knowledge.

Spend Wisely: Maximizing Post-Training Gains in Iterative Synthetic Data Boostrapping
with Yunzhen Feng, Ciyuan Chen, Yuhang Wu, and Zhuoyuan Li

- Under review in International Conference on Machine Learning (ICML 2025)
- We theoretically and empirically analysis how should the total budget on generation and training be allocated across iterations to maximize final performance, and demonstrate that an exponential growth policy exhibits significant advantages.

Work experience

AI Algorithm Internship 2021 Spring
In ByteDance *Beijing, China*

- I am primarily responsible for designing algorithms for audio fingerprinting, which involves extracting fingerprint features, searching through a speech library, and parallel acceleration, among other tasks. My team utilizes contrastive learning and propose a new feature extraction model that is more suitable for audio data, resulting in faster retrieval speeds and higher accuracy.

Skills

Languages: Python, C/C++, MATLAB
Human Languages: Chinese, English
Development Tools: Pytorch, Linux (the administrator of our group)
Other Tools: Latex, Markdown
Familiar Methods: computational mathematics, common machine learning methods, deep learning theory, reinforcement learning, large language model

Awards

Scholarships
in the undergraduate and doctoral program

- College Scholarship, School-level
- Outstanding Research Award, University-level

Modeling competition
in the undergraduate program

- MCM, Honorable Award

Olympiad in Informatics (OI)
in high school

- NOIP, First Prize

Mathematical Competition
in high school

- Peking University Summer Camp, Excellent camper (Admission to PKU)
- China Mathematics Competition, First Prize
- China Mathematics Olympiad, Silver Medal