

## ZHOU Pan

School of Computing and Information Systems  
Singapore Management University (SMU)  
80 Stamford Road  
Singapore 178902

Email: [panzhou@smu.edu.sg](mailto:panzhou@smu.edu.sg)  
Office Phone: 6808 5213



## Education

PhD, National University of Singapore, Singapore, 2020  
Master, Peking University, China, 2016  
Bachelor of Computer Science and Technology, China University of Geosciences (Wuhan), China, 2013

## Academic Appointments

Assistant Professor of Computer Science, School of Computing and Information Systems, SMU, Jan 2024 - Present

## Other Positions Held

Senior Research Scientist at SEA AI Lab, Singapore, Apr 2021 – Dec 2023  
Research Scientist at Salesforce, Singapore, Oct 2019 - Mar 2021

## RESEARCH

### Research Interests

My target is “towards an efficient and effective artificial general intelligent system (AGI)” with three research interests:-

- 1) **Network Architecture Design:** develop innovative network topology that posses high capacity and efficiency for acquiring knowledge, thereby improving the overall model capacity of AI/AGI.
- 2) **Learning Framework:** design effective learning framework / training task / loss to formulate a problem so that the network model can learn general and transferable features to handle diverse tasks, brining AGI closer.
  - a) **Self-Supervised Learning:** design effective self-supervised learning (SSL) framework that enables AI model to learn general vision knowledge and achieve human's data recognition and analysis ability, bringing AGI closer.

- b) **Generative Learning:** design generative models like diffusion models that empowers AI models with imagination and creativity akin to that of humans, bringing AGI closer.
  - c) **Meta In-Context Learning:** design new meta-learning and prompt learning methods to aid a (pretrained) model in quickly learning from a few data, improving few-shot learning ability of AGI.
- 3) **Parameter Optimizer:** design efficient optimizers to train networks efficiently, making AI and AGI model efficient.

## Publications

### Journal Articles [Refereed]

- [1]. Weihao Yu, Chenyang Si, Pan Zhou, Mi Luo, Yichen Zhou, Jiashi Feng, Shuicheng Yan, and Xinchao Wang. MetaFormer Baselines for Vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. (Accepted)
- [2]. Junbin Xiao, Pan Zhou, Angela Yao, Yicong Li, Richang Hong, Shuicheng Yan, Tat-Seng Chua. Contrastive Video Question Answering via Video Graph Transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11),13265 – 13280, 2023. <https://doi.org/10.1109/TPAMI.2023.3292266> (Published)
- [3]. Hanlin Zhang, Shuai Lin, Weiyang Liu, Pan Zhou, Jian Tang, Xiaodan Liang, Eric P. Xing. Iterative Graph Self-Distillation. *IEEE Transactions on Knowledge and Data Engineering*, 273, 414-423, 2022. <https://doi.org/10.1109/TKDE.2023.3303885> (Published)
- [4]. Lin Shuai, Liu Chen, Pan Zhou, Hu Zi-yuan, Wang Shuojia, Zhao Ruihui, Zheng Yefeng, Lin Liang, Xing Eric, Liang Xiaodan. Prototypical graph contrastive learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. <https://doi.org/10.1109/TNNLS.2022.3191086> (Advance Online)
- [5]. Fanhua Shang, Bingkun Wei, Hongying Liu, Yuanyuan Liu, Pan Zhou, and Maoguo Gong. Efficient Gradient Support Pursuit With Less Hard Thresholding for Cardinality-Constrained Learning. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12), 7806-7817, 2022. <https://doi.org/10.1109/TNNLS.2021.3087805> (Published)
- [6]. Pan Zhou, Canyi Lu, Jiashi Feng, Zhouchen Lin, Shuicheng Yan. Tensor low-rank representation for data recovery and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(5), 1718 – 1732, 2021. <https://doi.org/10.1109/TPAMI.2019.2954874> (Published)
- [7]. Pan Zhou, Xiaotong Yuan, Shuicheng Yan, Jiashi Feng. Faster first-order methods for stochastic non-convex optimization on Riemannian manifolds. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2), 459 – 472, 2021. <https://doi.org/10.1109/TPAMI.2019.2933841> (Published)
- [8]. Pan Zhou, Canyi Lu, Zhouchen Lin, Chao Zhang. Tensor factorization for low-rank tensor completion. *IEEE transactions on image processing*, 27(3), 1152-1163, 2017. <https://doi.org/10.1109/TIP.2017.2762595> (Published)
- [9]. Pan Zhou, Cong Fang, Zhouchen Lin, Chao Zhang, Edward Y. Chang. Dictionary learning with structured noise. *Neurocomputing*, 273, 414-423, 2020. <https://doi.org/10.1109/TCC.2020.2989296> (Published)
- [10]. Pan Zhou, Chao Zhang, Zhouchen Lin. Feature learning via partial differential equation with applications to face recognition. *Pattern Recognition*, 69, 14-25, 2017. <https://doi.org/10.1016/j.patcog.2017.03.034> (Published)
- [11]. Pan Zhou, Chao Zhang, Zhouchen Lin. Bilevel Model Based Discriminative Dictionary Learning for Recognition. *IEEE Transactions on Image Processing (TIP)*, 26(3), 1173-1187, 2016. <https://doi.org/10.1109/TIP.2016.2623487> (Published)
- [12]. Pan Zhou, Zhouchen Lin, Chao Zhang. Integrated Low-Rank-Based Discriminative Feature Learning for Recognition. *IEEE Transactions on Neural Networks and Learning Systems (TNNLS)*, 27(5), 1080-1093, 2016. <https://doi.org/10.1109/TNNLS.2015.2436951> (Published)

Conference Proceedings

- [1]. Zhongzhan Huang, Pan Zhou, Shuicheng Yan, Liang Lin. Towards More Stable Training of Diffusion Model via Scaling Network Long Skip Connection. Neural Information Processing Systems (NeurIPS), 2023
- [2]. Shanghua Gao, Pan Zhou, Ming-Ming Cheng, Shuicheng Yan. Masked Diffusion Transformer is a Strong Image Synthesizer. International Conference on Computer Vision (ICCV), 2023
- [3]. Ming Li, Xiangyu Xu, Hehe Fan, Pan Zhou, Jun Liu, Jia-Wei Liu, Jiahe Li, Jussi Keppo, Mike Zheng Shou, Shuicheng Yan. STPrivacy: Spatio-Temporal Privacy-Preserving Action Recognition. International Conference on Computer Vision (ICCV), 2023
- [4]. Alex Jinpeng Wang, Pan Zhou, Mike Zheng Shou, Shuicheng Yan. Position-guided Text Prompt for Vision-Language Pre-training. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2023
- [5]. Pan Zhou, Xingyu Xie, Shuicheng Yan. Win: Weight-Decay-Integrated Nesterov Acceleration for Adaptive Gradient Algorithms. International Conference on Learning Representations (ICLR), 2023 (oral)
- [6]. Jiachun Pan\*, Pan Zhou\*, Shuicheng Yan. Towards Understanding Why Mask Reconstruction Pretraining Helps in Downstream Tasks. International Conference on Learning Representations (ICLR), 2023 (\* equal contribution)
- [7]. Bowen Dong, Pan Zhou, Shuicheng Yan, Wangmeng Zuo. LPT: Long-tailed Prompt Tuning for Image Classification. International Conference on Learning Representations (ICLR), 2023
- [8]. Chenyang Si\*, Weihao Yu\*, Pan Zhou, Yichen Zhou, Xinchao Wang, Shuicheng Yan. Inception Transformer. Neural Information Processing Systems (NeurIPS), 2022 (oral) (\*equal contribution)
- [9]. Yuxuan Liang, Pan Zhou, Roger Zimmermann, Shuicheng Yan. DualFormer: Local-Global Stratified Transformer for Efficient Video Recognition. European Conference on Computer Vision (ECCV), 2022
- [10]. Junbin Xiao, Pan Zhou, Tat-Seng Chua, Shuicheng Yan. Video Graph Transformer for Video Question Answering. European Conference on Computer Vision (ECCV), 2022
- [11]. Bowen Dong, Pan Zhou, Shuicheng Yan, Wangmeng Zuo. Self-Promoted Supervision for Few-Shot Transformer. European Conference on Computer Vision (ECCV), 2022
- [12]. Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, Shuicheng Yan. MetaFormer is Actually What You Need for Vision. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2022 (oral)
- [13]. Pan Zhou, Caiming Xiong, Xiaotong Yuan, Steven Hoi. A Theory-Driven Self-Labeling Refinement Method for Contrastive Representation Learning. Neural Information Processing Systems (NeurIPS), 2021 (spotlight)
- [14]. Pan Zhou, Hanshu Yan, Xiaotong Yuan, Jiashi Feng, Shuicheng Yan. Towards Understanding Why Lookahead Generalizes Better Than SGD and Beyond. Neural Information Processing Systems (NeurIPS), 2021
- [15]. Pan Zhou, Yingtian Zou, Xiaotong Yuan, Jiashi Feng, Caiming Xiong, and Steven Hoi. Task Similarity Aware Meta Learning: Theory-inspired Improvement on MAML. International Conference on Uncertainty in Artificial Intelligence (UAI), 2021
- [16]. Guolin Zheng, Yubei Xiao, Ke Gong, Pan Zhou, Xiaodan Liang, and Liang Lin. Wav-BERT: Cooperative Acoustic and Linguistic Representation Learning for Low-Resource Speech Recognition. Conference on Empirical Methods in Natural Language Processing (EMNLP), 2021 (Findings)
- [17]. Yu Bai, Minshuo Chen, Pan Zhou, Tuo Zhao, Jason D. Lee, Sham Kakade, Huan Wang, Caiming Xiong. How Important is the Train-Validation Split in Meta-Learning?. International Conference on Machine Learning (ICML), 2021
- [18]. Junnan Li, Pan Zhou, Caiming Xiong, Richard Socher, and Steven Hoi. Prototypical Contrastive Learning of Unsupervised Representations. International Conference on Learning Representations (ICLR), 2021
- [19]. Shuai Lin, Pan Zhou, Xiaodan Liang, Jianheng Tang, Ruihui Zhao, Ziliang Chen and Liang Lin. Graph-Evolving Meta-Learning for Low-Resource Medical Dialogue Generation. Association for the Advancement of Artificial Intelligence (AAAI), 2021
- [20]. Yubei Xiao, Ke Gong, Pan Zhou, Guolin Zheng, Xiaodan Liang and Liang Lin. Adversarial Meta Sampling for Multilingual Low-Resource Speech Recognition. Association for the Advancement of Artificial Intelligence (AAAI), 2021
- [21]. Pan Zhou, Caiming Xiong, Richard Socher, and Steven Hoi. Theory-Inspired Path-Regularized Differential Network Architecture Search. Neural Information Processing Systems (NeurIPS), 2020 (oral)
- [22]. Pan Zhou, Jiashi Feng, Chao Ma, Caiming Xiong, Steven Hoi, and Weinan E. Towards Theoretically Understanding Why SGD Generalizes Better Than ADAM in Deep Learning. Neural Information Processing Systems (NeurIPS), 2020
- [23]. Yue Wu, Pan Zhou, Andrew Gordon Wilson, Eric Xing, and Zhiting Hu. Improving GAN Training with Probability Ratio Clipping and Sample Reweighting. Neural Information Processing Systems (NeurIPS), 2020

- [24]. Pan Zhou and Xiaotong Yuan. Hybrid Stochastic-Deterministic Minibatch Proximal Gradient: Less-Than-Single-Pass Optimization with Nearly Optimal Generalization. International Conference on Machine Learning (ICML), 2020.
- [25]. Pan Zhou, Xiaotong Yuan, Huan Xu, Shuicheng Yan, Jiashi Feng. Efficient Meta Learning via Minibatch Proximal Update. Neural Information Processing Systems (NeurIPS), 2019 (spotlight)
- [26]. Hu Zhang, Pan Zhou, Yi Yang, Jiashi Feng. Generalized Majorization-Minimization for Non-Convex Optimization. International Joint Conference on Artificial Intelligence (IJCAI), 2019.
- [27]. Pan Zhou, Xiaotong Yuan, Jiashi Feng. Faster First-Order Methods for Stochastic Non-Convex Optimization on Riemannian Manifolds. International Conference on Artificial Intelligence and Statistics (AISTATS), 2019
- [28]. Pan Zhou, Xiaotong Yuan, Jiashi Feng. Efficient Stochastic Gradient Hard Thresholding. Neural Information Processing Systems (NeurIPS), 2018
- [29]. Pan Zhou, Xiaotong Yuan, Jiashi Feng. New Insight into Hybrid Stochastic Gradient Descent: Beyond With-Replacement Sampling and Convexity. Neural Information Processing Systems (NeurIPS), 2018
- [30]. Pan Zhou, Jiashi Feng. Understanding Generalization and Optimization Performance of Deep CNNs. International Conference on Machine Learning (ICML), 2018
- [31]. Pan Zhou, Yunqing Hou, Jiashi Feng. Deep Adversarial Subspace Clustering. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018 (Published)
- [32]. Pan Zhou, Jiashi Feng. Empirical Risk Landscape Analysis for Understanding Deep Neural Networks. International Conference on Learning Representations (ICLR), 2018
- [33]. Jianshu Li, Pan Zhou, Yunpeng Chen, Jian Zhao, Sujoy Roy, Yan Shuicheng, Jiashi Feng, and Terence Sim. Task Relation Networks. IEEE Winter Conference on Applications of Computer Vision (WACV), 2019
- [34]. Pan Zhou, Jiashi Feng. Outlier-Robust Tensor PCA. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.