ZHOU Pan

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

Email: panzhou@smu.edu.sg
Office Phone: 68085213

Education

PhD, National University of Singapore, Singapore, 2020

Master of Natural Science, Peking University, China, 2016

Bachelor of Engineering, China University of Geosciences, China, 2013

Academic Appointments

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

RESEARCH

Research Interests

My research aims to build ”efficient and effective artificial intelligent systems” so that machines can cognize, understand, and interact with the environment. Currently, I mainly focus on three research topics across machine learning, computer vision, and optimization.

1) Learning Framework like Self-Supervised (multi-modal) Learning and Generative Models: design an effective learning framework/training task/loss to formulate a problem so that the AI models can learn desired knowledge to handle general/specific tasks.

2) Network Architecture Design: develop innovative network topology that posses high capacity and efficiency for acquiring knowledge, thereby improving the overall model capacity of AI.

3) Parameter Optimizer: design efficient optimizers to train AI models efficiently.

Publications

Journal Articles [Refereed]

Gamba: Marry Gaussian splatting with Mamba for single-view 3D reconstruction, by SHEN, Qiuhong; WU, Zike; YI, Xuanyu; ZHOU, Pan; ZHANG, Hanwang; YAN, Shuicheng; WANG, Xinchao. (2025). *IEEE Transactions on Pattern Analysis and Machine Intelligence,* 1-14.

https://doi.org/10.1109/TPAMI.2025.3569596 (Advance Online)

LoCo: Low-bit communication adaptor for large-scale model training, by XIE, Xingyu; LIN, Zhijie; TOH, Kim-chuan; ZHOU, Pan. (2025). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 47* (16), 4285-4298. https://doi.org/10.1109/TPAMI.2025.3544764 (Published)

A causality-aware paradigm for evaluating creativity of multimodal large language models, by HUANG, Zhongzhan; ZHONG, Shanshan; ZHOU, Pan; GAO, Shanghua; ZITNIK, Marink; LIN, Liang. (2025). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 47* (5), 3830-3846. https://doi.org/10.1109/TPAMI.2025.3539433 (Published)

Adan: Adaptive Nesterov Momentum Algorithm for faster optimizing deep models, by XIE, Xingyu; ZHOU, Pan; LI, Huan; LIN, Zhouchen; YAN, Shuicheng. (2024). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 46* (12), 1-34. https://doi.org/10.1109/TPAMI.2024.3423382 (Advance Online)

Enhancing visual grounding in vision-language pre-training with position-guided text prompts, by WANG, Alex Jinpeng; ZHOU, Pan; SHOU, Mike Zheng; YAN, Shuicheng. (2024). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 46* (5), 3406-3421. https://doi.org/10.1109/TPAMI.2023.3343736 (Published)

Towards understanding convergence and generalization of AdamW, by ZHOU, Pan; XIE, Xingyu; LIN, Zhouchen; YAN, Shuicheng . (2024). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 46* (9), 1-8. https://doi.org/10.1109/TPAMI.2024.3382294 (Advance Online)

Win: Weight-decay-integrated Nesterov acceleration for faster network training, by ZHOU, Pan; XIE, Xingyu; LIN, Zhouchen; TOH, Kim-Chuan; YAN, Shuicheng. (2024). *Journal of Machine Learning Research, 25* 1-74. https://www.jmlr.org/papers/v25/23-1073.html (Published)

Iterative graph self-distillation, by ZHANG, Hanlin; LIN, Shuai; LIU, Weiyang; ZHOU, Pan; TANG, Jian; LIANG, Xiaodan; XING, Eric . (2024). *IEEE Transactions on Knowledge and Data Engineering, 36* (3), 1161-1169. https://doi.org/10.1109/TKDE.2023.3303885 (Published)

Instant3D: Instant Text-to-3D Generation, by LI, Ming; ZHOU, Pan; LIU, Jia-Wei; KEPPO, Jussi; LIN, Min; YAN, Shuicheng; XU, Xiangyu. (2024). *International Journal of Computer Vision, 132* (10), 1-23. https://doi.org/10.1007/s11263-024-02097-5 (Advance Online)

MetaFormer baselines for vision, by YU, Weihao; SI, Chenyang; ZHOU Pan; LUO, Mi; ZHOU, Yichen; FENG, Jiashi; YAN, Shuicheng; WANG, Xinchao. (2024). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 46* (2), 896-912. https://doi.org/10.1109/TPAMI.2023.3329173 (Published)

Contrastive video question answering via video graph transformer, by XIAO, Junbin Xiao; ZHOU, Pan; YAO, Angela; LI, Yicong; HONG, Richang; YAN, Shuicheng; CHUA, Tat-Seng. (2023). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 45* (11), 13265-13280. https://doi.org/10.1109/TPAMI.2023.3292266 (Published)

Prototypical graph contrastive learning, by LIN, Shuai; LIU, Chen; ZHOU, Pan; HU, Zi-Yuan; WANG, Shuojia; ZHAO, Ruihui; ZHENG, Yefeng; LIN, Liang; XING, Eric; LIANG, Xiaodan. (2024). *IEEE Transactions on Neural Networks and Learning Systems, 35* (2), 2747-2758. https://doi.org/10.1109/TNNLS.2022.3191086 (Published)

Efficient gradient support pursuit with less hard thresholding for cardinality-constrained learning, by SHANG, Fanhua; WEI, Bingkun; LIU, Hongying; LIU, Yuanyuan; ZHOU, Pan; GONG, Maoguo. (2022). *IEEE Transactions on Neural Networks and Learning Systems, 33* (12), 7806-7817. https://doi.org/10.1109/TNNLS.2021.3087805 (Published)

A hybrid stochastic-deterministic minibatch proximal gradient method for efficient optimization and generalization, by ZHOU, Pan; YUAN, Xiao-Tong; LIN Zhouchen; HOI, Steven C. H.. (2021). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 44* (10), 5933-5946. https://doi.org/10.1109/TPAMI.2021.3087328 (Published)

Tensor low-rank representation for data recovery and clustering, by ZHOU, Pan; LU, Canyi; FENG, Jiashi; LIN, Zhouchen; YAN, Shuicheng. (2021). *IEEE Transactions on Pattern Analysis and Machine Intelligence, 43* (5), 1718-1732. https://doi.org/10.1109/TPAMI.2019.2954874 (Published)

Faster first-order methods for stochastic non-convex optimization on Riemannian manifolds, by ZHOU, Pan; YUAN, Xiao-Tong; YAN, Shuicheng; FENG, Jiashi. (2019). *IEEE Transactions on Pattern Analysis and*

Machine Intelligence, 43 (2), 459-472. https://doi.org/10.1109/TPAMI.2019.2933841 (Published)

Dictionary learning with structured noise, by ZHOU, Pan; FANG, Cong; LIN, Zhouchen; ZHANG, Chao; CHANG, Y. Edward . (2018). *Neurocomputing, 273* 414-423. https://doi.org/10.1016/J.NEUCOM.2017.07.041 (Published)

Tensor factorization for low-rank tensor completion, by ZHOU, Pan; LU, Canyi; LIN, Zhouchen; ZHANG, Chao. (2017). *IEEE Transactions on Image Processing, 27* (3), 1152-1163. https://doi.org/10.1109/TIP.2017.2762595 (Published)

Feature learning via partial differential equation with applications to face recognition, by FANG, Cong; ZHAO, Zhenyu; ZHOU, Pan; LIN, Zhouchen. (2017). *Pattern Recognition, 69* 14-25. https://doi.org/10.1016/J.PATCOG.2017.03.034 (Published)

Bilevel model-based discriminative dictionary learning for recognition, by ZHOU, Pan; ZHANG, Chao; LIN Zhouchen . (2016). *IEEE Transactions on Image Processing, 26* (3), 1173-1187. https://doi.org/10.1109/TIP.2016.2623487 (Published)

Integrated low-rank-based discriminative feature learning for recognition, by ZHOU, Pan; LIN, Zhouchen; ZHANG, Chao. (2015). *IEEE Transactions on Neural Networks and Learning Systems, 27* (5), 1080-1093. https://doi.org/10.1109/TNNLS.2015.2436951 (Published)

Book Chapters

Tensor principal component analysis, by ZHOU, Pan; LU, Canyi; LIN, Zhouchen. (2022). In LIU, Yipeng (Ed.), *Tensors for data processing: Theory, methods, and applications* (pp. 153-213) Elsevier. https://doi.org/10.1016/B978-0-12-824447-0.00012-1 (Published)

Conference Proceedings

SoPo: Text-to-motion generation using semi-online preference optimization, by TAN, Xiaofeng; WANG, Hongsong; GENG, Xin; ZHOU, Pan . (2025.0). *Proceedings of the 39th Conference on Neural Information Processing Systems (NeurIPS 2025), San Diego, December 2-7,* (pp. 1-33) San Diego, USA: https://openreview.net/forum?id=IjDlvzBSVp (Published)

GRIFFIN: Effective token alignment for faster speculative decoding, by HU, Shijing; LI, Jingyang; XIE, Xingyu; LU, Zhihui; TOH, Kim-Chuan; ZHOU, Pan. (2025.0). *Proceedings of the 39th Conference on Neural Information Processing Systems (NeurIPS 2025), San Diego, December 2-7,* (pp. 1-26) San Diego, USA: https://openreview.net/forum?id=JwnAItQF9v (Accepted)

Probabilistic prototype calibration of vision-language models for generalized few-shot semantic segmentation, by LIU, Jie; SHEN, Jiayi; ZHOU, Pan; SONKE, Jan-Jakob; GAVVES, Stratis. (2025.0). *Proceedings of the 2025 International Conference on Computer Vision, ICCV, Honolulu, HawaiI, October 19-23,* (pp. 1-16) Honolulu, HI, USA: https://doi.org/10.48550/arXiv.2506.22979 (Accepted)

Memory-efficient 4-bit preconditioned stochastic optimization, by LI, Jingyang; DING, Kuangyu; TOH, Kim-Chuan; ZHOU, Pan. (2025.0). *Proceedings of the 2025 International Conference on Computer Vision, ICCV, Honolulu, Hawaii, October 19-23,* Honolulu, HI, USA: (Accepted)

Zeroth-order fine-tuning of LLMs in random subspaces, by YU, Ziming; ZHOU, Pan; WANG, Sike; LI, Jia; TIAN, Mi; HUANG, Hua. (2025.0). *Proceedings of the 2025 International Conference on Computer Vision, ICCV, Honolulu, HawaiI, October 19-23,* (pp. 1-27) Honolulu, HI, USA: https://doi.org/10.48550/arXiv.2410.08989 (Accepted)

HPS: Hard preference sampling for human preference alignment, by ZOU, Xiandong; LIN, Wanyu; LI, Yuchen; ZHOU, Pan. (2025.0). *Proceedings of the 42nd International Conference on Machine Learning, ICML 2025, Vancouver, Canada, July 13-19* , (pp. 1-24) Vancouver, Canada: https://openreview.net/forum?id=hLvWwRZkok (Published)

Probabilistic interactive 3D segmentation with hierarchical neural processes, by LIU, Jie; ZHOU, Pan; XIAO, Zehao; SHEN, Jiayi; YIN, Wenzhe; SONKE, Jan-Jakob; GAVVES Efstratios. (2025.0). *Proceedings of the Forty-second International Conference on Machine Learning, ICML 2025, Vancouver, Canada, July 13* , (pp. 1-20) Vancouver, Canada: https://openreview.net/forum?id=6qNbVtKGY2 (Published)

Collaborative tree search for enhancing embodied multi-agent collaboration, by ZU, Lizheng; LIN, Lin; FU,

Song; ZHAO, Na; ZHOU, Pan. (2025.0). *Proceedings of the 2025 IEEE Conference on Computer Vision and Pattern Recognition, Nashville, TN, June 11-15,* (pp. 29513-29522) Nashville, USA: https://openaccess.thecvf.com/content/CVPR2025/papers/Zu_Collaborative_Tree_Search_for_Enhancing_Embodied_Multi-Agent_Collaboration_CVPR_2025_paper.pdf (Published)

Towards understanding why FixMatch generalizes better than supervised learning, by LI, Jingyang; PAN, Jiachun; TAN, Vincent; TOH, Kim-chuan; ZHOU, Pan. (2025.0). *Proceedings of the Thirteenth International Conference on Learning Representations, Singapore, 2025 April 24-28,* (pp. 1-38) Singapore: https://openreview.net/forum?id=25kAzqzTrz (Published)

CaPo: Cooperative plan optimization for efficient embodied multi-agent cooperation, by LIU, Jie; ZHOU, Pan; DU, Yingjun; TAN, Ah-Hwee; SNOEK, Cees; SONKE, Jan-Jakob; GAVVES, Efstratios. (2025.0). *Proceedings of the Thirteenth International Conference on Learning Representations, Singapore, 2025 April 24-28,* (pp. 1-25) Singapore: https://openreview.net/forum?id=KRv9NubipP (Published)

MVGamba : Unify 3D content generation as state space sequence modeling, by YI, Xuanyu; WU, Zike; SHEN, Qiuhong; XU, Qingshan; ZHOU, Pan; LIM, Joo-Hwee; YAN, Shuicheng; WANG, Xinchao; ZHANG, Hanwang. (2024.0). *Proceedings of the 37th Conference on Neural Information Processing Systems (NeurIPS 2024): Vancouver Canada, December 10-15,* Canada: NeurIPS. (Accepted)

Unsupervised modality adaptation with text-to-Image diffusion models for semantic segmentation, by XIA, Ruihao; LIANG, Yu; JIANG, Peng-Tao; ZHANG, Hao; LI, Bo; TANG, Yang; ZHOU, Pan. (2024.0). *Proceedings of 38th Annual Conference on Neural Information Processing Systems (NeurIPS 2024) : Vancouver, Canada, December 10-15,* Vancouver, Canada: NeurIPS. (Accepted)

LOVA3 : Learning to visual question answering, asking and assessment, by ZHAO, Henry Hengyuan; ZHOU, Pan; GAO, Difei; SHOU, BAI, Zechen; SHOU, Mike Zheng. (2024.0). *Proceedings of 38th Annual Conference on Neural Information Processing Systems (NeurIPS 2024) : Vancouver, Canada, December 10-15,* Vancouver, Canada: NeurIPS. (Accepted)

4-bit shampoo for memory-efficient network training, by WANG, Sike; ZHOU, Pan; LI, Jia; HUANG, Hua. (2024.0). *Proceedings of 38th Annual Conference on Neural Information Processing Systems (NeurIPS 2024) : Vancouver, Canada, December 10-15,* Vancouver, Canada: NeurIPS. (Accepted)

Efficient cascaded multiscale adaptive network for image restoration, by ZHOU, Yichen; ZHOU, Pan; NG, Teck Khim. (2024.0). *Proceedings of the 18th European Conference on Computer Vision, Milan, Italy, 2024 September 29 - October 4,* Berlin : Springer. (Accepted)

Genixer : Empowering multimodal Large Language Models as a powerful data generator, by ZHAO, Henry Hengyuan; ZHOU, Pan; SHOU, Mike Zheng. (2024.0). *18th European Conference on Computer Vision (ECCV 2024) : Milan, Italy, September 29 - October 4,* Milan, Italy: European Conference on Computer Vision. (Accepted)

Efficient cascaded multiscale adaptive network for image restoration, by ZHOU, Yichen; ZHOU, Pan; NG, Teck Khim. (2024.0). *18th European Conference on Computer Vision (ECCV 2024) : Sep 29- Oct 4, Milan Italy,* Milan, Italy: European Conference on Computer Vision. (Accepted)

Few-shot learner parameterization by diffusion time-steps, by YUE, Zhongqi; ZHOU, Pan; HONG, Richang; ZHANG, Hanwang; SUN Qianru. (2024.0). *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition Conference (CVPR), Seattle, 2024 June 17-21,* (pp. 23263-23272) Seattle WA, USA: CVPR. https://openaccess.thecvf.com/content/CVPR2024/papers/Yue_Few-shot_Learner_Parameterization_by_Diffusion_Time-steps_CVPR_2024_paper.pdf (Published)

Let’s think outside the box: Exploring leap-of-thought in large language models with multimodal humor generation, by ZHONG, Shanshan; HUANG, Zhongzhan; GAO, Shanghua; WEN, Wushao; LIN, Liang; ZITNIK, Marinka; ZHOU, Pan. (2024.0). *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition Conference (CVPR), Seattle, 2024 June 17-21,* (pp. 13246-13257) Seattle WA, USA: CVPR. https://openaccess.thecvf.com/content/CVPR2024/papers/Zhong_Lets_Think_Outside_the_Box_Exploring_Leap-of-Thought_in_Large_Language_CVPR_2024_paper.pdf (Advance Online)

Consistent3D: Towards consistent high-fidelity text-to-3D generation with deterministic sampling prior, by WU, Zike; ZHOU, Pan; YI, Xuanyu; YUAN, Xiaoding; ZHANG, Hanwang. (2024.0). *Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Seattle, June 17-21,* (pp. 1-16) Los Alamitos, CA: IEEE.

https://openaccess.thecvf.com/content/CVPR2024/papers/Wu_Consistent3D_Towards_Consistent_High-Fidelity_Text-to-3D_Generation_with_Deterministic_Sampling_Prior_CVPR_2024_paper.pdf (Advance Online)

Friendly sharpness-aware minimization, by LI, Tao; ZHOU, Pan; HE, Zhengbao; CHENG, Xinwen; HUANG, Xiaolin. (2024.0). *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition Conference (CVPR), Seattle, 2024 June 17-21,* (pp. 5631-5640) Seattle WA, USA: CVPR. https://openaccess.thecvf.com/content/CVPR2024/papers/Li_Friendly_Sharpness-Aware_Minimization_CVPR_2024_paper.pdf (Published)

Diffusion time-step curriculum for one image to 3D generation, by YI, Xuanyu; WU, Zike; XU, Qingshan; ZHOU, Pan; LIM, Joo Hwee; ZHANG, Hanwang. (2024.0). *Proceedings of the IEEE/CVF International Conference on Computer Vision and Pattern Recognition Conference (CVPR), Seattle, 2024 June 17-21,* (pp. 9948-9958) Seattle WA, USA: CVPR. https://openaccess.thecvf.com/content/CVPR2024/papers/Yi_Diffusion_Time-step_Curriculum_for_One_Image_to_3D_Generation_CVPR_2024_paper.pdf (Published)

InceptionNeXt: When Inception meets ConvNeXt, by YU, Weihao; ZHOU, Pan; YAN, Shuicheng; WANG, Xinchao . (2024.0). *2024 IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR): Seattle, June 17-21: Proceedings,* (pp. 1-12) Piscataway, NJ: IEEE. https://doi.org/10.1109/CVPR52733.2024.00542 (Published)

ScaleLong: Towards more stable training of diffusion model via scaling network long skip connection, by HUANG, Zhongzhan; ZHOU, Pan; YAN, Shuicheng; LIN, Liang. (2023.0). *Proceedings of the 37th Conference on Neural Information Processing, New Orleans, United States, December 12-14,* (pp. 1-26) New Orleans: NeurIPS . https://openreview.net/forum?id=0N73P8pH2l (Published)

EditAnything: Empowering unparalleled flexibility in image editing and generation, by GAO, Shanghua; LIN, Zhijie; XIE, Xingyu; ZHOU, Pan; CHENG, Ming-Ming; YAN, Shuicheng. (2023.0). *MM '23: Proceedings of the 31st ACM International Conference on Multimedia, Ottowa, Canada, October 29 - November 3,* (pp. 9414-9416) New York: ACM. https://doi.org/10.1145/3581783.3612680 (Published)

STPrivacy: Spatio-temporal privacy-preserving action recognition, by LI, Ming; XU, Xiangyu; FAN, Hehe; ZHOU, Pan; LIU, Jun; LIU, Jia-Wei; LI, Jiahe; KEPPO, Jussi; SHOU, Mike Zheng; YAN, Shuicheng. (2023.0). *Proceedings of the 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, 2023 October 1-6,* (pp. 5106-5115) Piscataway, NJ: IEEE. https://doi.org/10.1109/ICCV51070.2023.00471 (Published)

Masked diffusion transformer is a strong image synthesizer, by GAO, Shanghua; ZHOU, Pan; CHENG, Ming-Ming; YAN, Shuicheng. (2023.0). *Proceedings of the 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Paris, France, October 1-6,* (pp. 23164-23173) Piscataway, NJ: IEEE. https://doi.org/10.1109/ICCV51070.2023.02117 (Published)

Position-guided text prompt for vision-language pre-training, by WANG, Alex Jinpeng; ZHOU, Pan; SHOU, Mike Zheng; YAN Shuicheng. (2023.0). *Proceedings of the 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Vancouver, June 17-24 ,* (pp. 23242-23251) Piscataway, NJ: IEEE. https://doi.org/10.1109/CVPR52729.2023.02226 (Published)

LPT: Long-tailed prompt tuning for image classification, by DONG, Bowen; ZHOU, Pan; YAN, Shuicheng; ZUO, Wangmeng. (2023.0). *Proceedings of The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5,* (pp. 1-20) Kigali, Rwanda: ICLR. https://openreview.net/forum?id=8pOVAeo8ie (Published)

Towards understanding why mask reconstruction pretraining helps in downstream tasks, by PAN, Jiachun; ZHOU, Pan; YAN, Shuicheng. (2023.0). *Proceedings of the 11th International Conference on Learning Representations ICLR 2023: Kigali, Rwanda, May 1-5,* (pp. 1-48) Kigali, Rwanda: ICLR. https://openreview.net/forum?id=PaEUQiY40Dk (Published)

Win: Weight-decay-integrated nesterov acceleration for adaptive gradient algorithms, by ZHOU, Pan; XIE, Xingyu; YAN, Shuicheng . (2023.0). *Proceedings of the 11th International Conference on Learning Representations, Kigali, Rwanda, 2023 May 1-5*

, (pp. 1-28) USA: ICLR. https://openreview.net/pdf?id=dNK2bw4y0R (Published)

Inception transformer, by SI, Chenyang; YU, Weihao; ZHOU, Pan; ZHOU, Yichen; WANG, Xinchao; YAN, Shuicheng. (2022.0). *Proceedings of the 36th Conference on Neural Information Processing Systems (NeurIPS 2022) Track on Datasets and Benchmarks, Virtual Conference, 2022 November 28,* (pp. 1-15)

New Orleans: NeurIPS . https://proceedings.neurips.cc/paper_files/paper/2022/hash/94e85561a342de88b559b72c9b29f638-Abstract-Conference.html (Published)

Self-promoted supervision for few-shot transformer, by DONG, Bowen; ZHOU, Pan; YAN, Shuicheng; ZUO, Wangmeng. (2022.0). *Proceedings of the 17th European Conference (ECCV 2022), Tel Aviv, Israel, October 23-27,* (pp. 329-347) Cham: Springer. https://doi.org/10.1007/978-3-031-20044-1_19 (Published)

Video graph transformer for video question answering, by XIAO, Junbin; ZHOU, Pan; CHUA, Tat-Seng; YAN, Shuicheng. (2022.0). *Proceedings of the 17th European Conference (ECCV 2022), Tel Aviv, Israel, October 23-27,* (pp. 39-58) Cham: Springer. https://doi.org/10.1007/978-3-031-20059-5_3 (Published)

DualFormer: Local-global stratified transformer for efficient video recognition, by LIANG, Yuxuan; ZHOU, Pan; ZIMMERMANN, Roger; YAN, Shuicheng. (2022.0). *Proceedings of the 17th European Conference (ECCV 2022), Tel Aviv, Israel, October 23-27,* (pp. 577-595) Cham: Springer. https://doi.org/10.1007/978-3-031-19830-4_33 (Published)

MetaFormer is actually what you need for vision, by YU, Weihao; LUO, Mi; ZHOU, Pan; SI, Chenyang; ZHOU, Yichen; WANG, Xinchao; FENG, Jiashi; YAN, Shuicheng. (2022.0). *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, June 18-24,* (pp. 10819-10829) Piscataway, NJ: IEEE. https://doi.org/10.1109/CVPR52688.2022.01055 (Published)

Towards understanding why Lookahead generalizes better than SGD and beyond, by ZHOU, Pan; YAN, Hanshu; YUAN, Xiaotong; FENG, Jiashi; YAN, Shuicheng . (2021.0). *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021), Sydney, Australia, December 6-14,* (pp. 1-15) Virtual Conference: NeurIPS . https://proceedings.neurips.cc/paper/2021/hash/e53a0a2978c28872a4505bdb51db06dc-Abstract.html (Published)

A theory-driven self-labeling refinement method for contrastive representation learning, by ZHOU, Pan; XIONG, Caiming; YUAN, Xiao-Tong, HOI, Steven . (2021.0). *Proceedings of the 35th Conference on Neural Information Processing Systems (NeurIPS 2021), Virtual Conference, December 6-14,* (pp. 1-15) Virtual Conference: NeurIPS . (Published)

Wav-BERT: Cooperative acoustic and linguistic representation learning for low-resource speech recognition, by ZHENG, Guolin; XIAO, Yubei; GONG, Ke; ZHOU, Pan; LIANG, Xiaodan; LIN, Liang . (2021.0). *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Virtual Conference, November 7-11,* (pp. 2765-2777) Punta Cana: ACL. https://doi.org/10.18653/V1/2021.FINDINGS-EMNLP.236 (Published)

Task similarity aware meta learning: Theory-inspired improvement on MAML, by ZHOU, Pan; ZPU, Yingtian; YUAN, XiaoTong; FENG, Jiashi; XIONG, Caiming; HOI, Steven. (2021.0). *Proceeding of The Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, Virtual Conference, 2021 July 27-29,* (pp. 1-11) Virtual Conference: Proceedings of Machine Learning Research. https://proceedings.mlr.press/v161/zhou21a.html (Published)

How important is the train-validation split in meta-learning?, by BAI, Yu; CHEN, Minshuo; ZHOU, Pan; ZHAO, Tuo; LEE, D. Jason; KAKADE, Sham; WANG, Huan; XIONG, Caiming. (2021.0). *Proceedings of the 38th International Conference on Machine Learning, Virtual Conference, 2021 July 18-24,* (pp. 1-11) Virtual Conference: https://proceedings.mlr.press/v139/bai21a.html. (Published)

Prototypical contrastive learning of unsupervised representations, by LI, Junnan; ZHOU, Pan; XIONG, Caiming; HOI, Steven C. H.. (2021.0). *Proceedings of the Ninth International Conference on Learning Representations: ICLR 2021, Vienna, Austria, May 4-8,* (pp. 1-16) Virtual Conference: ICLR. https://openreview.net/forum?id=KmykpuSrjcq (Published)

Graph-evolving meta-learning for low-resource medical dialogue generation, by LIN, Shuai; ZHOU, Pan; LIANG, Xiaodan; TANG, Jianheng; ZHAO, Ruihui; CHEN, Ziliang; LIN, Liang. (2021.0). *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021), Virtual Conference, February 2-9,* (pp. 13362-13370) Palo Alto, CA: AAAI Press. https://doi.org/10.1609/aaai.v35i15.17577 (Published)

Adversarial meta sampling for multilingual low-resource speech recognition, by XIAO, Yubei; GONG, Ke; ZHOU, Pan; ZHENG, Guolin; LIANG, Xiaodan; LIN, Liang . (2021.0). *Proceedings of The Thirty-Fifth AAAI Conference on Artificial Intelligence, Virtual Conference, 2021 February 2-9,* (pp. 13362-13370) Virtual Conference: AAAI. https://cdn.aaai.org/ojs/17577/17577-13-21071-1-2-20210518.pdf (Published)

Improving GAN training with probability ratio clipping and sample reweighting, by WU, Yue; ZHOU, Pan; GORDON, Andrew Wilson; XING, Eric; HU, Zhiting. (2020.0). *Proceedings of the 34th Conference on Neural Information Processing Systems, Virtual Conference, 2020 December 6-12,* (pp. 1-12) Virtual Conference: NeurIPS . https://proceedings.neurips.cc/paper_files/paper/2020/hash/3eb46aa5d93b7a5939616af91addfa88-Abstract.html (Published)

Towards theoretically understanding why SGD generalizes better than ADAM in deep learning, by ZHOU, Pan; FENG, Jiashi; MA, Chao; XIONG, Caiming; HOI, Steven; E, Weinan . (2020.0). *Proceedings of the 34th Conference on Neural Information Processing Systems, NeurIPS 2020, Vancouver, Canada, December 6-12,* (pp. 1-12) Virtual Conference: NeurIPS . https://proceedings.neurips.cc/paper_files/paper/2020/hash/f3f27a324736617f20abbf2ffd806f6d-Abstract.html (Published)

Theory-inspired path-regularized differential network architecture search, by ZHOU, Pan; XIONG, Caiming; SOCHER, Richard; HOI, Steven C. H.. (2020.0). *Proceedings of the 34th Conference on Neural Information Processing Systems, Virtual Conference, 2020 December 6-12,* (pp. 1-12) Virtual Conference: NeurIPS . https://proceedings.neurips.cc/paper_files/paper/2020/hash/5e1b18c4c6a6d31695acbae3fd70ecc6-Abstract.html (Published)

Hybrid stochastic-deterministic minibatch proximal gradient: Less-than-single-pass optimization with nearly optimal generalization, by ZHOU, Pan; YUAN, Xiaotong. (2020.0). *Proceedings of the 37th International Conference on Machine Learning (ICML 2020), Virtual Conference, July 13-18,* (pp. 1-10) Virtual Conference: NeurIPS . https://proceedings.mlr.press/v119/zhou20g.html (Published)

Efficient meta learning via minibatch proximal update, by ZHOU, Pan; YUAN, Xiao-Tong; XU, Huan; YAN, Shuicheng; FENG, Jiashi. (2019.0). *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada, December 8-14,* (pp. 1-11) Vancouver, Canada: NeurIPS. https://proceedings.neurips.cc/paper_files/paper/2019/hash/8c235f89a8143a28a1d6067e959dd858-Abstract.html (Published)

Generalized majorization-minimization for non-convex optimization, by ZHANG, Hu; ZHOU, Pan; YANG, Yi; FENG, Jiashi. (2019.0). *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, Macao, China, 2019 August 10-16,* (pp. 4257-4263) Macao, China: IJCAI. https://doi.org/10.24963/IJCAI.2019/591 (Published)

Faster first-order methods for stochastic non-convex optimization on Riemannian manifolds, by ZHOU, Pan; YUAN, Xiao-Tong; FENG, Jiashi. (2019.0). *Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics, Naha, Okinawa, Japan, 2019 April 16-18* , (pp. 1-20) Naha, Okinawa, Japan: Proceedings of Machine Learning Research. https://proceedings.mlr.press/v89/zhou19a.html (Published)

Task relation networks, by LI, Jianshu; ZHOU, Pan; CHEN, Yunpeng; ZHAO, Jian; ROY, Sujoy; YAN, Shuicheng; FENG, Jiashi; SIM, Terence . (2019.0). *Proceedings of the 2019 IEEE Winter Conference on Applications of Computer Vision, Waikoloa, HI, USA, January 7-11,* (pp. 1-9) Piscataway, NJ: IEEE. https://doi.org/10.1109/WACV.2019.00104 (Published)

New insight into hybrid stochastic gradient descent: Beyond with-replacement sampling and convexity, by ZHOU, Pan; YUAN, Xiao-Tong; FENG, Jiashi . (2018.0). *Proceedings of the 32nd Annual Conference on Advances in Neural Information Processing Systems, Montréal, Canada, 2018 December 2-8,* (pp. 1-10) Montréal, CANADA: NeurIPS . https://papers.nips.cc/paper_files/paper/2018/hash/67e103b0761e60683e83c559be18d40c-Abstract.html (Published)

Efficient stochastic gradient hard thresholding, by ZHOU, Pan; YUAN, Xiao-Tong; FENG, Jiashi . (2018.0). *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada, December 2-8,* (pp. 1-10) Montréal, Canada: NeurIPS . https://papers.nips.cc/paper_files/paper/2018/hash/ec5aa0b7846082a2415f0902f0da88f2-Abstract.html (Published)

Understanding generalization and optimization performance of deep CNNs, by ZHOU, Pan; FENG, Jiashi. (2018.0). *Proceedings of the 35th International Conference on Machine Learning, Stockholm Sweden, 2018 July 10-15,* (pp. 1-38) Stockholm, Sweden: Proceedings of Machine Learning Research. https://publons.com/wos-op/publon/52135107/ (Published)

Deep adversarial subspace clustering, by ZHOU, Pan; HOU, Yunqing; FENG, Jiashi. (2018.0). *Proceedings of the 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, Salt Lake City, USA, June 18-23 ,* (pp. 1596-1604) Piscataway, NJ: IEEE. https://doi.org/10.1109/CVPR.2018.00172 (Published)

Empirical risk landscape analysis for understanding deep neural networks, by ZHOU, Pan Zhou; FENG, Jiashi. (2018.0). *Proceedings of the 6th International Conference on Learning Representations, ICLR 2018, Vancouver, Canada, April 30 - May 3,* (pp. 1-60) Vancouver, Canada: ICLR. https://openreview.net/forum?id=B1QgVti6Z (Published)

Outlier-robust tensor PCA, by ZHOU, Pan; FENG, Jiashi . (2017.0). *Proceedings of 2017 IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, USA, July 21-26,* (pp. 1-9) Piscataway, NJ: IEEE. https://doi.org/10.1109/CVPR.2017.419 (Published)

Other Outputs and Contributions

Others

Neural network based scene text recognition [US Patent 2022/0237403 A1], by ZHOU, Pan; TANG, Peng; XU, Ran; HOI, Steven Chu Hong. (2024). In (1-18). https://patents.google.com/patent/US20220237403A1 (Published)

Systems and methods for contrastive learning with self-labeling refinement [US Patent US20220269946A1], by ZHOU, Pan; XIONG, Caiming; HOI, Chu Hong. (2022). In (1-17). (Published)

Research Grants

Singapore Management University

Towards Theoretically-Sound and Practical-Efficient Deep Learning Optimizer, SMU Internal Grant, Ministry of Education (MOE) Tier 1 , PI (Project Level): ZHOU Pan, 2024, S$120,000

Synthesis and Resilience: Generative Models for Generalizable 3D World Understanding, SMU Internal Grant, Ministry of Education (MOE) Tier 1 , PI (Project Level): ZHOU Pan, 2024, S$150,000

TEACHING

Teaching Areas

Computer Vision

Courses Taught

Singapore Management University

Undergraduate Programmes :

- Generative AI for Vision
- IS/SMT/C&L Project Experience (Applications)

Postgraduate Research Programmes :

- Empirical Research Project 1

Empirical Research Project 2

Empirical Research Project 3