

Dr. Zhaoxuan Wu



Position:

Postdoctoral Associate
Singapore-MIT Alliance for Research and Technology (SMART)

Contact: +65 84385708

Email: wu.zhaoxuan@u.nus.edu / zxwu@mit.edu

Office: 1 CREATE Way, #09-03 CREATE Tower, Singapore 138602

Website: <https://zhaoxuanwu.github.io>

RESEARCH INTERESTS

- Data-centric artificial intelligence (e.g., data valuation & selection, collaborative machine learning, incentives, fairness)
- Resource-efficient machine learning (e.g., Bayesian optimization)
- Large language models (e.g., inference-time techniques, prompting)
- Deep learning & applications

ACADEMIC QUALIFICATIONS

- **Doctor of Philosophy in Data Science** Singapore
National University of Singapore Aug 2020 – Oct 2024
 - CAP: 5.00/5.00
 - Thesis Title: **Data-Centric AI: Through the Lens of Data Valuation and Beyond**
 - Supervisor: Prof. Bryan Kian Hsiang Low
 - Thesis Advisory Committee: Prof. See-Kiong Ng, Prof. Vincent Yan Fu Tan, Prof. Bryan Kian Hsiang Low
- **Bachelor of Science (Honors) in Data Science & Analytics** Singapore
National University of Singapore Aug 2016 – Jun 2020
 - Minor in Computer Science
 - CAP: 4.82/5.00; Honors (Highest Distinction)
 - Thesis Title: Deep Learning for Glaucoma Diagnosis
 - Supervisor: Prof. Alexandre Hoang Thiery
 - Award: Best Academic Project in Data Science & Analytics Discipline

SCHOLARSHIPS

- Aug 2023 – Aug 2024 NUSGS Research Incentive Award
- Feb 2023 – Feb 2024 Singapore Data Science Consortium (SDSC) Dissertation Research Fellowship
- Aug 2020 – Aug 2024 President's Graduate Fellowship (Ph.D.)
- Jan 2018 – May 2018 UTown Scholarship - Tin Ka Ping Foundation Scholarship
- Nov 2011 – Nov 2015 Singapore SM1 School-based Scholarship (Secondary & Pre-U)

PUBLICATIONS

* = equal contribution / co-first authorship

- **Zhaoxuan Wu***, Zijian Zhou*, Arun Verma, Alok Prakash, Daniela Rus, and Bryan Kian Hsiang Low (2025). **TETRIS: Optimal Draft Token Selection for Batch Speculative Decoding**. *In Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL-25)*.
- Rui Qiao, **Zhaoxuan Wu**, Jingtian Wang, Pang Wei Koh, and Bryan Kian Hsiang Low (2025). **Mitigating Spurious Correlations via Group-robust Sample Reweighting**. *In Proceedings of the 13th International Conference on Learning Representations (ICLR-25)* [32.1% Acceptance Rate].
- Bingchen Wang, **Zhaoxuan Wu**, Fusheng Liu, and Bryan Kian Hsiang Low (2025). **Paid with Models: Optimal Contract Design for Collaborative Machine Learning**. *In Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI-25)* [23.4% Acceptance Rate]. **Oral Presentation**.
- **Zhaoxuan Wu***, Xiaoqiang Lin*, Zhongxiang Dai, Wenyang Hu, Yao Shu, See-Kiong Ng, Patrick Jaillet, and Bryan Kian Hsiang Low (2024). **Prompt Optimization with EASE? Efficient Ordering-aware Automated Selection of Exemplars**. *In Advances in Neural Information Processing Systems 37: 38th Annual Conference on Neural Information Processing Systems (NeurIPS-24)* [25.8% Acceptance Rate].
- Wenyang Hu, Yao Shu, Zongmin Yu, **Zhaoxuan Wu**, Xiangqiang Lin, Zhongxiang Dai, See-Kiong Ng, and Bryan Kian Hsiang Low (2024). **Localized Zeroth-Order Prompt Optimization**. *In Advances in Neural Information Processing Systems 37: 38th Annual Conference on Neural Information Processing Systems (NeurIPS-24)* [25.8% Acceptance Rate]. **Spotlight Presentation**.
- Xinyi Xu, **Zhaoxuan Wu**, Rui Qiao, Arun Verma, Yao Shu, Jingtian Wang, Xinyuan Niu, Zhenfeng He, Jiangwei Chen, Zijian Zhou, Gregory Kang Ruey Lau, Hieu Dao, Lucas Agussurja, Rachael Hwee Ling Sim, Xiaoqiang Lin, Wenyang Hu, Zhongxiang Dai, Pang Wei Koh, and Bryan Kian Hsiang Low (2024). **Position Paper: Data-Centric AI in the Age of Large Language Models**. *In Proceedings of the Conference on Empirical Methods in Natural Language Processing 2024 (EMNLP-24) Findings*.
- Xiaoqiang Lin*, **Zhaoxuan Wu***, Zhongxiang Dai, Wenyang Hu, Yao Shu, See-Kiong Ng, Patrick Jaillet, and Bryan Kian Hsiang Low (2024). **Use Your INSTINCT: INSTruction optimization for LLMs usIng Neural bandits Coupled with Transformers**. *In Proceedings of the 41st International Conference on Machine Learning (ICML-24)* [27.5% Acceptance Rate].
- Xiaoqiang Lin, Xinyi Xu, **Zhaoxuan Wu**, See-Kiong Ng, and Bryan Kian Hsiang Low (2024). **Distributionally Robust Data Valuation**. *In Proceedings of the 41st International Conference on Machine Learning (ICML-24)* [27.5% Acceptance Rate].
- **Zhaoxuan Wu**, Mohammad Mohammadi Amiri, Ramesh Raskar, and Bryan Kian Hsiang Low (2024). **Incentive-Aware Federated Learning with Training-Time Model Rewards**. *In Proceedings of the 12th International Conference on Learning Representations (ICLR-24)* [31% Acceptance Rate].
- Xinyi Xu, **Zhaoxuan Wu**, Arun Verma, Chuan Sheng Foo, and Bryan Kian Hsiang Low (2023). **FAIR: Fair Collaborative Active Learning with Individual Rationality for Scientific Discovery**. *In Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS-23)* [29.0% Acceptance Rate].

- Zhaoxuan Wu, Yao Shu, and Bryan Kian Hsiang Low (2022). **DAVINZ: Data Valuation using Deep Neural Networks at Initialization**. In *Proceedings of the 39th International Conference on Machine Learning (ICML-22)* [21.9% Acceptance Rate].
- Yao Shu, Zhongxiang Dai, Zhaoxuan Wu, and Bryan Kian Hsiang Low (2022). **Unifying and Boosting Gradient-Based Training-Free Neural Architecture Search**. In *Advances in Neural Information Processing Systems 35: 36th Annual Conference on Neural Information Processing Systems (NeurIPS-22)* [25.6% Acceptance Rate].
- Xinyi Xu*, Zhaoxuan Wu*, Chuan Sheng Foo, and Bryan Kian Hsiang Low (2021). **Validation Free and Replication Robust Volume-based Data Valuation**. In *Advances in Neural Information Processing Systems 34: 35th Annual Conference on Neural Information Processing Systems (NeurIPS-21)* [25.7% Acceptance Rate].
- Quoc Phong Nguyen*, Zhaoxuan Wu*, Bryan Kian Hsiang Low, and Patrick Jaillet (2021). **Trusted-Maximizers Entropy Search for Efficient Bayesian Optimization**. In *Proceedings of the 37th Conference on Uncertainty in Artificial Intelligence (UAI-21)* [26.5% Acceptance Rate].

BOOK CHAPTERS

- Zhaoxuan Wu, Xinyi Xu, Rachael Hwee Ling Sim, Yao Shu, Xiaoqiang Lin, Lucas Agussurja, Zhongxiang Dai, See-Kiong Ng, Chuan-Sheng Foo, Patrick Jaillet, Trong Nghia Hoang, and Bryan Kian Hsiang Low (2024). **Data Valuation in Federated Learning**. In *L. M. Nguyen, T. N. Hoang, P.-Y. Chen, editors, Federated Learning: Theory and Practice, chapter 15, pages 281-296, Academic Press*.
- Xiaoqiang Lin, Xinyi Xu, Zhaoxuan Wu, Rachael Hwee Ling Sim, See-Kiong Ng, Chuan-Sheng Foo, Patrick Jaillet, Trong Nghia Hoang, and Bryan Kian Hsiang Low (2024). **Fairness in Federated Learning**. In *L. M. Nguyen, T. N. Hoang, P.-Y. Chen, editors, Federated Learning: Theory and Practice, chapter 8, pages 143-160, Academic Press*.
- Rachael Hwee Ling Sim, Sebastian Shenghong Tay, Xinyi Xu, Yehong Zhang, Zhaoxuan Wu, Xiaoqiang Lin, See-Kiong Ng, Chuan-Sheng Foo, Patrick Jaillet, Trong Nghia Hoang, and Bryan Kian Hsiang Low (2024). **Incentives in Federated Learning**. In *L. M. Nguyen, T. N. Hoang, P.-Y. Chen, editors, Federated Learning: Theory and Practice, chapter 16, pages 299-309, Academic Press*.

PROFESSIONAL SERVICE

- Conference reviewer/PC member for
 - AAAI Conference on Artificial Intelligence (AAAI), 2024, 2025
 - International Conference on Artificial Intelligence and Statistics (AISTATS), 2024
 - International Joint Conference on Artificial Intelligence (IJCAI), 2024
 - International Conference on Learning Representations (ICLR), 2023, 2024, 2025
 - International Conference on Autonomous Agents and Multiagent Systems (AAMAS), 2023, 2024, 2025
 - International Conference on Machine Learning (ICML), 2022, 2023, 2024, 2025
 - Conference on Neural Information Processing Systems (NeurIPS), 2022, 2023, 2024
 - Asian Conference on Machine Learning (ACML), 2022, 2023
- Received **Top Reviewer** for NeurIPS 2023
- Received **Notable Reviewer** for ICLR 2025

INVITED TALKS

- “Use Your INSTINCT: INSTRUCTION optimization usIng Neural bandits Coupled with Transformers”. Deep Learning and Optimization Seminar (jointly organized by Westlake University, CityU, Peking University), virtual, Oct 24, 2023.

TEACHING EXPERIENCE

- CS3244 (Machine Learning), NUS
 - Teaching Assistant for 1 tutorial class
- CS3244 (Machine Learning), NUS
 - Teaching Assistant for 2 tutorial classes
- DSA2102 (Essential Data Analytics Tools: Numerical Computation), NUS
 - Teaching Assistant

Spring 2022

Spring 2021

Fall 2020

HONORS AND AWARDS

- **NUS School of Computing Research Achievement Award 2024**
- **Lijen Industrial Development Medal AY2019/20**
 - Being the Honors year student with the *best academic exercise/project* in the Data Science and Analytics discipline in the Faculty of Science, NUS
 - In my Honors project, I designed a multi-task U-Net architecture for learning three tasks on Optical Coherence Tomography (OCT) images simultaneously
 - Achieved an overall average test accuracy of 91.4% across tasks and further developed an algorithm to reconstruct a more realistic predicted eye structure
- **Faculty of Science Dean’s List Recipient for Semester 2 AY2019/20, Semester 1 AY2018/19 and Semester 2 AY2017/18**
 - Awarded to students in the *top 5 percent* of the total undergraduate Science cohort
- **NUS Science Diamond Jubilee Student Award 2019**
 - A testimony of excellent academic track records both in NUS and the Student Exchange Program to Northwestern University, IL, USA
- Gold Award in Nanyang Research Program 2014
 - Awarded for the Electrical & Electronic Engineering project on Nanowires Silicon/PEDOT:PSS Hybrid Solar Cells after months of experiments, written report, and oral presentation
- High Distinction in National Economics & Financial Management Competition 2015
- Young Engineers & Scientist (YES) Academic Award Physics 2013
 - Awarded by the Defence Science & Technology Agency of Singapore
- Silver Award in Singapore Junior Physics Olympiad 2012

EMPLOYMENT HISTORY

- **NUS AI Innovation & Commercialization Center** Suzhou, China
Research Intern *May 2019 – Aug 2019*
 - **Supervisors:** Prof. Teck Khim Ng and Prof. Yin Xu
 - **AutoML:** Contributed to the development of ***Rafiki***, an open-source distributed system that offers automated Machine Learning (AutoML) model training, tuning and deployment services
 - **ASR:** Enriched Rafiki's base of supported tasks to Automated Speech Recognition (ASR) and integrated a ready-to-use DeepSpeech model into the Rafiki framework
 - **Impact:** Enable users with minimal background knowledge in AI to train, tune and deploy an ASR application with a Word Error Rate of less than 10%
- **Insignia Ventures Partners** Singapore
Full-Stack Developer Intern *Jan 2018 – Jul 2018*
 - **Supervisors:** Dr. Yinglan Tan and Mr. Ridy Lie
 - **Web Development:** Designed and developed features in the company's web application under the engineering team, including KPIs, web scraping, securities and third-party application integration, thus improving the user-friendliness of the application and the efficiency of the investment process
- **Pteris Global Limited** Singapore
Software Developer Intern *Mar 2016 – May 2016*
 - **VBA:** Designed and developed VBA programs to generate templates for project costing estimate, manpower costing estimate and procurement list, resulting in a much more reliable automated costing calculation free of human error, and at the same time, increased the productivity by reducing labor hours