

Education

- 2018–2020 **PhD**, Machine Learning Department (MLD), *Carnegie Mellon University (CMU)*, USA
Dissertation: Towards Data-Efficient Machine Learning
Thesis committee: Eduard Hovy (advisor), Tom Mitchell, Ruslan Salakhutdinov, and Quoc Le
- 2016–2018 **Master's**, Language Technologies Institute (LTI), *Carnegie Mellon University (CMU)*, USA
Advisor: Eduard Hovy
Funded by PwC and DARPA projects
- 2012–2016 **B.Eng.**, ACM Honored Class (Computer Science), Zhiyuan College, *Shanghai Jiao Tong University (SJTU)*, China
Zhiyuan Excellence Award (Highest Honors)
Advisors: Kai Yu (X-LANCE), Yong Yu

Employment

- 2024–present **Assistant Professor**, *National University of Singapore (NUS)*, Singapore
Department of Computer Science, School of Computing (SoC)
- 2022–2024 **Research Assistant Professor**, *National University of Singapore (NUS)*, Singapore
Department of Computer Science, SoC
Funded by Sea, NExT, and SoC
- 2020–2022 **Co-founder**, *Startup explorations*, USA & China
- 2018–2020 **Student Researcher**, *Google Brain (now Google DeepMind)*, Mountain View & Pittsburgh, USA
Managers: Quoc Le, Minh-Thang Luong
- Fall 2015 **Research Intern**, *Microsoft Research Asia, Machine Learning*, Beijing, China
Managers: Tie-Yan Liu, Bin Gao, Di He

Notable Projects

- 2025 **Diffusion Language Models (DLMs): Super Data Learners and Quokka**, *NUS*
- Explored alternative paradigms to autoregressive LLMs. Showed that DLMs have the potential to alleviate the **data-scarcity problem** for pretraining, and established the **scaling law for DLMs** (Quokka).
 - Sparked wide interest on X and in the AI community in general.
- 2018–2020 **Semi-supervised Learning: Noisy Student and UDA**, *Google Brain*
- Noisy Student is one of the first pieces of evidence that **supervision from AI could work**. It is a general method that achieved state-of-the-art results on ImageNet, speech recognition and object detection, and is also deployed on Google Search. A similar technique contributed to the superior performance of AlphaFold 2 (2024 Nobel Prize in Chemistry).
 - UDA uses semi-supervised learning to achieve the same performance with 1,000x less data.
 - The results are the outcome of more than 6 months of exploration on up to **8,192 Google TPU v3 chips**, carried out as a solo project involving substantial engineering work on infrastructure and resource management.

2017–2018 **Evaluation from Exams: RACE and CLOTH, Carnegie Mellon University**

- RACE and CLOTH are among the first NLP evaluations whose data was collected from exams, and among the first datasets to emphasize general **human-level reasoning**.
- RACE was used to evaluate the **original GPT, Llama, Megatron-LM**, etc.

Selected Awards, Grants, and Honors

2026	Razer–NUS Joint AI Research Lab (Co-Lead PI, with Shuicheng Yan)	<i>Razer Inc. & NUS</i>
2026	AI4S Challenge Grant (Investigator)	<i>National Research Foundation, Singapore</i>
2025	MOE AcRF Tier 1 Grant (Lead PI)	<i>Singapore Ministry of Education</i>
2025	AIVP Grant (Co-Lead PI, with Tomas Mikolov; relinquished)	<i>MDDI, Singapore</i>
2025	Best Paper Honorable Mention (for SynthRL)	<i>ICML DataWorld Workshop</i>
2024	NUS Start-Up Grant (Lead PI)	<i>NUS</i>
2016	Best Presentation Honorable Mention, Student Research Symposium	<i>CMU LTI</i>
2016	Zhiyuan Excellence Award	<i>SJTU</i>
2016	Chun-Tsung Scholar (endowed by Nobel laureate Tsung-Dao Lee)	<i>SJTU</i>
2016	Outstanding Graduate Award	<i>SJTU</i>
2016	Star of Tomorrow Internship	<i>Microsoft Research Asia</i>
2014, 2015	Academic Excellence Scholarship	<i>SJTU</i>
2013	KoGuan Scholarship	<i>SJTU</i>
2013	Merit Student	<i>SJTU</i>
2011	Silver Medal, National Olympiad in Informatics (NOI)	<i>China</i>
2011	Silver Medal, IOI Chinese Team Selection Contest (CTSC)	<i>China</i>
2010	First Prize, National Olympiad in Informatics in Provinces (NOIP)	<i>Hunan, China</i>

Invited Talks

Title: *Scaling Beyond Autoregression*

Jan 2026	Seminar @ Adobe Research, hosted by Trung H. Bui	<i>Online</i>
Dec 2025	Seminar @ OpenAI, hosted by Zhili Feng	<i>San Francisco, USA</i>
Dec 2025	Seminar @ Nvidia, hosted by Fitsum Reda	<i>Santa Clara, USA</i>
Dec 2025	Seminar @ Google DeepMind, hosted by Thang Luong	<i>Mountain View, USA</i>
Dec 2025	Monetization AI Speaker Series @ Meta, hosted by Fei Tian	<i>Menlo Park, USA</i>

Title: *Research on Mechanisms of AI Scaling and Capabilities*

May 2025	SJTU ACM Class Alumni Symposium, hosted by Yong Yu	<i>Shanghai, China</i>
Apr 2025	SJTU X-LANCE Alumni Symposium, hosted by Kai Yu	<i>Shanghai, China</i>

Title: *Towards Data-Efficient Machine Learning*

Feb 2023	Seminar @ NUS, hosted by See-Kiong Ng	<i>Singapore</i>
Mar 2021	Seminar @ Google Brain, hosted by Thang Luong and Quoc Le	<i>Online</i>
Mar 2021	Seminar @ Nvidia, hosted by Mohammad Shoeybi	<i>Online</i>

Mar 2021	Seminar @ Microsoft, hosted by Ahmed Hassan Awadallah	Online
Mar 2021	Seminar @ Adobe, hosted by Trung H. Bui	Online
Mar 2021	Seminar @ Amazon	Online
Mar 2021	Seminar @ Waymo, hosted by Charles Qi	Online
Mar 2021	Seminar @ NIO, hosted by Shaoqing Ren	Online
Mar 2021	Seminar @ Sea	Online
Feb 2021	Seminar @ ByteDance, hosted by Zhijie Chen	Online
Feb 2021	Seminar @ Snap, hosted by Sergey Tulyakov	Online
Feb 2021	Seminar @ Pony.ai	Online
Feb 2021	Seminar @ WeRide	Online
Feb 2021	Seminar @ Nuro, hosted by Nicholas Carlevaris-Bianco	Online
Feb 2021	Seminar @ Kuaishou, hosted by Yan Li	Online

Other Talks

Aug 2026	Panel discussion on AI for research @ Annual Academic Retreat of NUS SoC	Singapore
Apr 2026	Research sharing talk @ NUS–Meta MoU Signing Event	Singapore
Apr 2023	LLM, ChatGPT and Beyond @ Chat on GPT (NUS public outreach event)	Singapore

Publications

Published under names: *Qizhe Xie, Michael Qizhe Shieh, Michael Shieh, and Michael Qizhe Xie.*

Journal articles

1. Filip Ilievski, Eduard Hovy, Piek Vossen, Stefan Schlobach, **Qizhe Xie**. The role of knowledge in determining identity of long-tail entities. *Journal of Web Semantics*, 2020
2. Kai Yu, Lu Chen, Kai Sun, **Qizhe Xie**, Su Zhu. Evolvable dialogue state tracking for statistical dialogue management. *Frontiers of Computer Science*, 2016
3. Kai Sun, **Qizhe Xie** and Kai Yu. Recurrent Polynomial Network for Dialogue State Tracking. *Dialogue & Discourse*, 2016

Conference proceedings

4. Zijian Wu, Xiangyan Liu, Xinyuan Zhang, Lingjun Chen, Fanqing Meng, Lingxiao Du, Yiran Zhao, Fanshi Zhang, Yaoqi Ye, Jiawei Wang, Zirui Wang, Jinjie Ni, Yufan Yang, Arvin Xu, **Michael Qizhe Shieh**. MCPMark: A Benchmark for Stress-Testing Realistic and Comprehensive MCP Use. *ICLR*, 2026
5. Guanzheng Chen, **Michael Qizhe Shieh**, Lidong Bing. LongRLVR: Long-Context Reinforcement Learning Requires Verifiable Context Rewards. *ICLR*, 2026
6. Ziyao Guo, Kaipeng Zhang, **Michael Qizhe Shieh**. Improving Autoregressive Image Generation through Coarse-to-Fine Token Prediction. *CVPR Findings*, 2026
7. Zijian Wu*, Jinjie Ni*, Xiangyan Liu*, Zichen Liu, Hang Yan, **Michael Qizhe Shieh**. SynthRL: Scaling Visual Reasoning with Verifiable Data Synthesis. *ACL Findings*, 2026. **Best Paper Honorable Mention at the ICML 2025 DataWorld Workshop.**

8. Zichen Liu, Anya Sims, Keyu Duan, Changyu Chen, Simon Yu, Xiangxin Zhou, Haotian Xu, Shaopan Xiong, Bo Liu, Chenmien Tan, Weixun Wang, Hao Zhu, Weiyan Shi, Diyi Yang, **Michael Qizhe Shieh**, Yee Whye Teh, Wee Sun Lee, Min Lin. GEM: A Gym for Generalist LLMs. *ICLR*, 2026
9. Jiawei Gu, Yunzhuo Hao, Huichen Will Wang, Linjie Li, **Michael Qizhe Shieh**, Yejin Choi, Ranjay Krishna, Yu Cheng. ThinkMorph: Emergent Properties in Multimodal Interleaved Chain-of-Thought Reasoning. *ICLR*, 2026
10. Guanzheng Chen*, Qilong Feng*, Jinjie Ni, Xin Li, **Michael Qizhe Shieh**. Long-Context Inference with Retrieval-Augmented Speculative Decoding. *ICML*, 2025 (**Spotlight**)
11. Jinjie Ni, Yifan Song, Deepanway Ghosal, Bo Li, David Junhao Zhang, Xiang Yue, Fuzhao Xue, Zian Zheng, Kaichen Zhang, Mahir Shah, Kabir Jain, Yang You, **Michael Shieh**. MixEval-X: Any-to-Any Evaluations from Real-World Data Mixtures. *ICLR*, 2025 (**Spotlight**)
12. Xiangyan Liu*, Jinjie Ni*, Zijian Wu*, Chao Du, Longxu Dou, Haonan Wang, Tianyu Pang, **Michael Qizhe Shieh**. NoisyRollout: Reinforcing Visual Reasoning with Data Augmentation. *NeurIPS*, 2025
13. Yuxin Chen, Yiran Zhao, Yang Zhang, An Zhang, Kenji Kawaguchi, Shafiq Joty, Junnan Li, Tat-Seng Chua, **Michael Qizhe Shieh**, Wenxuan Zhang. The Emergence of Abstract Thought in Large Language Models Beyond Any Language. *NeurIPS*, 2025
14. Keyu Duan*, Yiran Zhao*, Zhili Feng, Jinjie Ni, Tianyu Pang, Qian Liu, Tianle Cai, Longxu Dou, Kenji Kawaguchi, Anirudh Goyal, J Zico Kolter, **Michael Qizhe Shieh**. Unnatural Languages Are Not Bugs but Features for LLMs. *ICML*, 2025
15. Guanzheng Chen, Xin Li, **Michael Qizhe Shieh**, Lidong Bing. LongPO: Long Context Self-Evolution of Large Language Models through Short-to-Long Preference Optimization. *ICLR*, 2025
16. Yiran Zhao, Wenxuan Zhang, Yuxi Xie, Anirudh Goyal, Kawaguchi Kenji, **Michael Qizhe Shieh**. Understanding and enhancing safety mechanisms of LLMs via safety-specific neuron. *ICLR*, 2025
17. Keyu Duan, Zichen Liu, Xin Mao, Tianyu Pang, Changyu Chen, Qiguang Chen, **Michael Qizhe Shieh**, Longxu Dou. Efficient Process Reward Model Training via Active Learning. *COLM*, 2025
18. Leon Lin, Hannah Brown, Kenji Kawaguchi, **Michael Shieh**. Single character perturbations break llm alignment. *AAAI*, 2025
19. Xiangyan Liu, Bo Lan, Zhiyuan Hu, Yang Liu, Zhicheng Zhang, Fei Wang, **Michael Shieh**, Wenmeng Zhou. Codexgraph: Bridging large language models and code repositories via code graph databases. *NAACL*, 2025
20. Yiran Zhao, Wenyue Zheng, Tianle Cai, Xuan Long Do, Kenji Kawaguchi, Anirudh Goyal, **Michael Qizhe Shieh**. Accelerating Greedy Coordinate Gradient and General Prompt Optimization via Probe Sampling. *NeurIPS*, 2024
21. Xuan Long Do*, Yiran Zhao*, Hannah Brown*, Yuxi Xie, James Xu Zhao, Nancy F Chen, Kenji Kawaguchi, **Michael Qizhe Xie**, Junxian He. Prompt optimization via adversarial in-context learning. *ACL*, 2024
22. Hongfu Liu, Yuxi Xie, Ye Wang, **Michael Shieh**. Advancing Adversarial Suffix Transfer Learning on Aligned Large Language Models. *EMNLP*, 2024
23. Esther Gan, Yiran Zhao, Liying Cheng, Yancan Mao, Anirudh Goyal, Kenji Kawaguchi, Min-Yen Kan, **Michael Shieh**. Reasoning Robustness of LLMs to Adversarial Typographical Errors. *EMNLP*, 2024
24. Yuxi Xie, Kenji Kawaguchi, Yiran Zhao, James Xu Zhao, Min-Yen Kan, Junxian He, **Michael Qizhe Xie**. Self-Evaluation Guided Beam Search for Reasoning. *NeurIPS*, 2023

25. Hai Ye, **Qizhe Xie**, Hwee Tou Ng. Multi-Source Test-Time Adaptation as Dueling Bandits for Extractive Question Answering. *ACL*, 2023
26. James Xu Zhao, Yuxi Xie, Kenji Kawaguchi, Junxian He, **Michael Qizhe Xie**. Automatic Model Selection with Large Language Models for Reasoning. *EMNLP findings*, 2023
27. Hieu Pham, Zihang Dai, **Qizhe Xie**, Minh-Thang Luong and Quoc V. Le. Meta Pseudo Labels. *CVPR*, 2021
28. **Qizhe Xie**, Eduard Hovy, Minh-Thang Luong and Quoc V. Le. Self-training with Noisy Student Improves ImageNet Classification. *CVPR*, 2020
29. **Qizhe Xie**, Zihang Dai, Eduard Hovy, Minh-Thang Luong and Quoc V. Le. Unsupervised Data Augmentation for Consistency Training. *NeurIPS*, 2020
30. Xiang Kong, **Qizhe Xie**, Zihang Dai and Eduard Hovy. Fast and Simple Mixture of Softmaxes with BPE and Hybrid-LightRNN for Language Generation. *AAAI*, 2019
31. **Qizhe Xie***, Guokun Lai*, Zihang Dai and Eduard Hovy. Large-scale Cloze Test Dataset Created by Teachers. *EMNLP*, 2018 (Oral) (*: equal contribution)
32. Zihang Dai*, **Qizhe Xie*** and Eduard Hovy. From Credit Assignment to Entropy Regularization: Two New Algorithms for Neural Sequence Prediction. *ACL*, 2018 (*: equal contribution)
33. **Qizhe Xie**, Zihang Dai, Yulun Du, Eduard Hovy and Graham Neubig. Controllable Invariance through Adversarial Feature Learning. *NeurIPS*, 2017
34. Guokun Lai*, **Qizhe Xie***, Hanxiao Liu, Yiming Yang and Eduard Hovy. RACE: Large-scale ReAding Comprehension Dataset From Examinations. *EMNLP*, 2017 (*: equal contribution)
35. **Qizhe Xie**, Xuezhe Ma, Zihang Dai and Eduard Hovy. An Interpretable Knowledge Transfer Model for Knowledge Base Completion. *ACL*, 2017 (Oral)
36. **Qizhe Xie**, Kai Sun, Su Zhu, Lu Chen and Kai Yu. Recurrent Polynomial Network for Dialogue State Tracking with Mismatched Semantic Parsers. *SigDial*, 2015 (Oral)

Workshops

37. Yuxi Xie, Anirudh Goyal, Wenye Zheng, Min-Yen Kan, Timothy P Lillicrap, Kenji Kawaguchi, **Michael Shieh**. Monte Carlo Tree Search Boosts Reasoning via Iterative Preference Learning. *NeurIPS Workshop*, 2024
38. Kaixin Li*, Qisheng Hu*, Xu Zhao, Yuxi Xie, Tiedong Liu, Hui Chen, **Qizhe Xie**, Junxian He. InstructCoder: Empowering Language Models for Code Editing. *ACL Workshop*, 2024

Preprints & Technical Reports

39. Fanqing Meng, Lingxiao Du, Zijian Wu, Guanzheng Chen, Xiangyan Liu, Jiaqi Liao, Chonghe Jiang, Zhenglin Wan, Jiawei Gu, Pengfei Zhou, Rui Huang, Ziqi Zhao, Shengyuan Ding, Ailing Yu, Bo Peng, Bowei Xia, Hao Sun, Haotian Liang, Ji Xie, Jiajun Chen, Jiajun Song, Liu Yang, Ming Xu, Qionglin Qiu, Runhao Fu, Shengfang Zhai, Shijian Wang, Tengfei Ma, Tianyi Wu, Weiyang Jin, Yan Wang, Yang Dai, Yao Lai, Youwei Shu, Yue Liu, Yunzhuo Hao, Yuwei Niu, Jinkai Huang, Jiayuan Zhuo, Zhennan Shen, Linyu Wu, Hannah Yao, Charles Chen, Cihang Xie, Yuyin Zhou, Jiaheng Zhang, Zeyu Zheng, Mengkang Hu, **Michael Qizhe Shieh**. ClawMark: A Living-World Benchmark for Multi-Turn, Multi-Day, Multimodal Coworker Agents. *Preprint*, 2026

40. Yaoqi Ye, Yiran Zhao, Keyu Duan, Zeyu Zheng, Kenji Kawaguchi, Cihang Xie, **Michael Qizhe Shieh**. In-Context Reinforcement Learning for Tool Use in Large Language Models. *Preprint*, 2026
41. Zijun Wang, Haoqin Tu, Letian Zhang, Hardy Chen, Juncheng Wu, Xiangyan Liu, Zhenlong Yuan, Tianyu Pang, **Michael Qizhe Shieh**, Fengze Liu, Zeyu Zheng, Huaxiu Yao, Yuyin Zhou, Cihang Xie. Your Agent, Their Asset: A Real-World Safety Analysis of OpenClaw. *Preprint*, 2026
42. Jinjie Ni, Qian Liu, Longxu Dou, Chao Du, Zili Wang, Hang Yan, Tianyu Pang, **Michael Qizhe Shieh**. Diffusion Language Models are Super Data Learners. *Preprint*, 2025
43. Jinjie Ni, Qian Liu, Chao Du, Longxu Dou, Hang Yan, Zili Wang, Tianyu Pang, **Michael Qizhe Shieh**. Training Optimal Large Diffusion Language Models. *Preprint*, 2025
44. Hannah Brown, Leon Lin, Kenji Kawaguchi, **Michael Shieh**. Self-Evaluation as a Defense Against Adversarial Attacks on LLMs. *Preprint*, 2024
45. Haonan Wang, James Zou, Michael Mozer, Linjun Zhang, Anirudh Goyal, Alex Lamb, Zhun Deng, **Michael Qizhe Xie**, Hannah Brown, Kenji Kawaguchi. Can AI Be as Creative as Humans?. *Preprint*, 2024
46. Keyu Duan, Qian Liu, Tat-Seng Chua, Shuicheng Yan, Wei Tsang Ooi, **Qizhe Xie**, Junxian He. Simteg: A frustratingly simple approach improves textual graph learning. *Preprint*, 2023
47. Filip Ilievski, Eduard Hovy, **Qizhe Xie**, Piek Vossen. The profiling machine: active generalization over knowledge. *Preprint*, 2018

* *equal contribution*

Patents

1. Hieu Hy Pham, Zihang Dai, **Qizhe Xie**, Quoc V Le. Meta pseudo-labels. US Patent App. 17/551,065
2. Thang Minh Luong, Quoc V Le, **Qizhe Xie**. Self-training technique for generating neural network models. US Patent App. 17/018,555
3. Thang Minh Luong, Quoc V Le, **Qizhe Xie**, Zihang Dai. Training machine learning models using unsupervised data augmentation. US Patent App. 17/606,190
4. Kai Yu, **Qizhe Xie**, Xueyang Wu, Wenbo Li, Yunqi Guo. A Cloud-based Speech Content Extraction Method and Device. CN105957531A
5. Kai Yu, **Qizhe Xie**, Xueyang Wu, Wenbo Li, Yunqi Guo. A Speech Abstract Generation Method and Device. CN105930471A
6. Xianran Wang, Weida Zhou, Kai Yu, **Qizhe Xie**, Da Zheng, Wenbo Li. An Interactive Input-based Data Search Method and Device. CN104699784A

Teaching

25/26 Sem 2	CS6216 Advanced Topics in Machine Learning (NUS)	<i>Instructor</i>
24/25 Sem 2	CS5260 Neural Networks and Deep Learning II (NUS)	<i>Co-instructor</i>
24/25 Sem 1	CS6280 Topics in CS: Deep Learning with Language Applications (NUS)	<i>Instructor</i>

5 of 20 course projects grew into publications: two at ICLR 2025, and one each at ICCV 2025, EMNLP 2025, and IEEE TMM

23/24 Sem 2	CS6216 Advanced Topics in Machine Learning (NUS)	<i>Instructor</i>
	2 of ~11 course projects grew into publications: one at EMNLP 2024 and one at ICML 2025	
23/24 Sem 1	CS6101 Exploration of Computer Science Research (NUS)	<i>Guest Lecturer</i>
	Week 1 lecture “Intro to LLM”; course led by Min-Yen Kan	
Spring 2019	10-315 Machine Learning (CMU)	<i>Teaching Assistant</i>
Fall 2018	10-715 Advanced Introduction to Machine Learning (CMU)	<i>Teaching Assistant</i>
Fall 2017	11-747 Neural Networks for NLP (CMU)	<i>Teaching Assistant</i>
Summer 2014	Principle and Practice of Algorithms, ACM Honored Class (SJTU)	<i>Head Teaching Assistant</i>
Spring 2014	Data Structure, ACM Honored Class (SJTU)	<i>Head Teaching Assistant</i>
Fall 2013	Data Structure, SEIEE (SJTU)	<i>Teaching Assistant</i>

Advising

Postdocs	Jinjie Ni → Google DeepMind
Current PhDs	Guanzheng Chen Lingxiao Du Keyu Duan Esther Gan (co-advised) Ziyao Guo Xiangyan Liu Fanqing Meng David Afonso Valente Zijian Wu Yaoqi Ye
Past PhDs	Hannah Brown (co-advised)
Master’s	Qisheng Hu → PhD @ NTU James Xu Zhao → PhD @ NUS Liangtao (Leon) Lin → PhD @ NTU
Undergrads	Qilong Feng → PhD @ CMU Lingjun Chen → PhD @ CUHK Zhennan Shen → PhD @ HKUST
Informal advising	Yiran Zhao → Salesforce Yuxi Xie → Research visit @ UCSB Kaixin Li → Tencent Hunyuan Hongfu Liu → Amazon

Service

Co-organizer	NeurIPS 2026 workshop on “Diffusion Language Models: Foundations, Efficiency, and Reasoning”
Area Chair	ICLR (2024, 2025), NeurIPS (2024, 2025), ICML (2025, 2026)
PC Member	ACL, EMNLP, CVPR, ICCV, ECCV, JMLR (since 2018)

NUS PhD admission committee (since 2024)
Committee Final year project committee (since 2024)
Student conduct committee (25/26 Sem 2)