

Zihan Wang

✉ zihan.wang@mila.quebec  [beikewzh.github.io](https://github.com/beikewzh)  [LinkedIn](#)  [GitHub](#)  [Google Scholar](#)
📍 Montreal, Canada

EDUCATION

Ph.D. in Computer Science, **McGill University & Mila** 2025 – 2029
GPA: 4.0/4.0. Supervisors: Doina Precup, Xiao-Wen Chang
Committee Members: Joelle Pineau, Pablo Samuel Castro

M.Sc. in Computer Engineering, **McGill University & Mila** 2023 – 2025
GPA: 4.0/4.0. Thesis: *Zero-Shot Anomaly Detection with Dual-Branch Prompt Learning*.
Supervisors: Narges Armanfard, Samira Ebrahimi Kahou

B.Sc. in Computer Engineering, **University of Alberta** 2018 – 2023
GPA: 3.7/4.0. Minor in Mathematics. Capstone: *Guess a Sketch* 

EXPERIENCE

Visiting researcher, **ServiceNow AI Research**, Montreal, Canada Oct 2026 – Jan 2027
Working with the Core AI team on recursive self-improving agents for enterprise benchmarks, enabling open-source language models to iteratively improve through interaction and experience.

Scientist in Residence, **Mila / Linarite AI**, Montreal, Canada Sep 2025 – Apr 2026
Built SIC-Agents, a clinician-facing platform that lets doctors practice pediatric serious illness communication. Two main contributions: (1) the first expert-verified benchmark suite for SIC training, validated with educators and pediatric clinicians; and (2) a self-improving agentic framework that generates a clinician-editable skill document to guide simulator behavior and outperforms static expert prompting.

Researcher Intern, **Huawei Noah's Ark Lab**, Montreal, Canada Jul 2024 – Sep 2024
Researched LLM post-training, studying supervised fine-tuning and preference-based alignment to improve reasoning and instruction-following.

PUBLICATIONS

* *Equal contribution.*

Published and Accepted Papers

Zihan Wang, and the Koala Science team. *Koala Science: A Platform for AI Reviewers in the Wild*. *AI-Native Academia Workshop, NeurIPS 2026*.

Zihan Wang, Anita Slominska, Rennie Bimman, Elizabeth Di Flumeri, Amanda Mayappo-Neeposh, Conall Francoeur, Tamara Ellen Carver, Xiao-Wen Chang, Doina Precup, Esin Darici Haritaoglu, Ismail Haritaoglu, Akshatha Arodi, Naomi Goloff. *SIC-Agents: Benchmarking and Building an Adaptive Simulator for Pediatric Serious Illness Communication Training*. *The 2026 Conference on Empirical Methods in Natural Language Processing (EMNLP 2026)* (**Oral**).

Diganta Misra, Nizar Islah, Victor May, Brice Rauby, **Zihan Wang**, Justine Gehring, Muawiz Sajjad Chaudhary, Antonio Orvieto, Eilif B. Muller, Irina Rish, Samira Ebrahimi Kahou, Massimo Caccia. *GitChameleon 2.0: Evaluating AI Code Generation Against Python Library Version Incompatibilities*. *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (ACL 2026)* (**Main**).

Wesley Chung*, **Zihan Wang***, Xiao-Wen Chang, David Meger, Doina Precup. *Perceived Regret: Evaluating Agents in Any World*. *Continual RL Workshop, Reinforcement Learning Conference (RLC 2026)*.

Mohammadreza Sadeghi, Sareh Soleimani, **Zihan Wang**, Narges Armanfard. *Unsupervised Continual Clustering via Forward-Backward Knowledge Distillation*. *European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD 2026)*(**Oral**).

Shuyuan Zhang, **Zihan Wang**, Xiao-Wen Chang, Doina Precup. *From Connectivity to Rewards: Dense Reward Learning with Directed State Graphs*. *Transactions on Machine Learning Research (TMLR) 2026*.

Zihan Wang, Samira Ebrahimi Kahou, Narges Armanfard. *Zero-Shot Anomaly Detection with Dual-Branch*

Prompt Learning. *Proceedings of the British Machine Vision Conference (BMVC 2025)* (Oral).

Shuyuan Zhang, **Zihan Wang**, Xiao-Wen Chang, Doina Precup. [Learning Subgoal Representations from State Graphs in Goal-Conditioned Hierarchical Reinforcement Learning](#). *Transactions on Machine Learning Research (TMLR)* 2025.

Kunhan Wu*, **Zihan Wang***, Jingyi Zhao, Haodong Xu, Tianming Hao, Wenzhi Lin. [Instrumental Variables Matter: Towards Causal Inference Using Deep Learning](#). *Journal of Physics: Conference Series* 2023.

Zihan Wang, Ruimin Chen, Mengxuan Liu, Guanfang Dong, Anup Basu. [SPGNet: Spatial Projection Guided 3D Human Pose Estimation in Low Dimensional Space](#). *International Conference on Smart Multimedia* 2022.

Weijie Sun, Sunil Vasu Kalmady, Nariman Sepehrvand, Luan Manh Chu, **Zihan Wang**, Amir Salimi, Abram Hindle, Russell Greiner, Padma Kaul. [Improving ECG-based COVID-19 Diagnosis and Mortality Predictions Using Pre-pandemic Medical Records at Population-Scale](#). *NeurIPS 2022 Workshop on Learning from Time Series for Health*.

TALKS

An Agent Can Only Be Evaluated Where It Has Been, short talk, [Reinforcement Learning and Philosophy Workshop \(RLPH\)](#), Montreal Aug 2026

AWARDS AND ACHIEVEMENTS

Flight PS752 Commemorative Scholarship ($\times 2$), Government of Canada; \$58,000	2025 – 2027
FRQNT Doctoral Research Scholarship , Fonds de recherche du Québec; \$100,000	2025 – 2029
McGill GREAT Award ($\times 2$), McGill University; \$1,100	2025, 2026
FRQNT Master’s Research Scholarship , Fonds de recherche du Québec; \$26,667	2024 – 2025
Graduate Excellence Fellowship , McGill University; \$24,000	2023 – 2026
Dean’s Research Award ($\times 2$), University of Alberta; \$1,000	2021, 2022
University of Alberta Academic Scholarship ($\times 3$), University of Alberta; \$24,000	2019 – 2023
President’s International Distinction Scholarship , University of Alberta; \$120,000	2018 – 2022
Graduated with Distinction , University of Alberta	2023

SKILLS

Programming	Python, C/C++, Java, Rust, SQL, JavaScript, MATLAB
Deep Learning	PyTorch, JAX, OpenCV
RL and LLMs	OpenAI Gym, HuggingFace Transformers/TRL, vLLM, DeepSpeed
Languages	English (fluent), Mandarin (native)

ACADEMIC SERVICES

Reviewer	Transactions on Machine Learning Research (TMLR, since 2024) Conference on Empirical Methods in Natural Language Processing (EMNLP 2026) Reinforcement Learning Conference (RLC 2026) IEEE/CVF Winter Conference on Applications of Computer Vision (WACV 2025)
Workshop Organizer	Finding the Frame , Reinforcement Learning Conference (RLC 2026) 2nd Workshop on Automated Spatial and Temporal Anomaly Detection (ASTAD) , WACV 2025
Teaching Assistant	COMP 550 Natural Language Processing COMP 579 Reinforcement Learning COMP 424 Artificial Intelligence
Tea Talk Organizer	Mila – Quebec AI Institute (since 2026)