Xiaofeng Zhang

xiaofeng.zhang@mila.quebec | Personal Website | LinkedIn
Montréal, Québec, Canada

Research Interests

My research focuses on generative modeling and representation learning for images and videos, with an emphasis on visual quality, text-image consistency and physical plausibility. I am also interested in world models, action-conditioned models and their applications. My work spans diffusion and flow-based models, visual tokenization for autoregressive generation, and benchmarking and evaluation.

Selected Publications

C=Conference, W=Workshop, S=Submitted

A full list of publications is available on my Google Scholar profile.

[S.1] **Xiaofeng Zhang**, Samuel Lavoie, Nicolas Beltran-Velez, et al. (2026). **FeaTok: Feature-Loss Tokenization for Image Generation**. Submitted to *ICLR 2027*.

[C.1] Nicolas Beltran-Velez, Felix Friedrich, **Xiaofeng Zhang**, et al. (2026). **The Reward Was in Your Data All Along: Correcting Flow Matching with Discriminator-Guided RL**. In *NeurIPS 2026*, Sydney, Australia.

[C.2] Jianhao Yuan, **Xiaofeng Zhang**, Felix Friedrich, et al. (2026). **Inference-time Physics Alignment of Video Generative Models with Latent World Models** (Highlight). In *CVPR 2026*, Denver, USA.

[C.3] **Xiaofeng Zhang**, Aaron Courville, Michal Drozdzal, et al. (2026). **The Intricate Dance of Prompt Complexity, Quality, Diversity, and Consistency in T2I Models**. In *ICLR 2026*, Rio de Janeiro, Brazil.

[C.4] Nicola Dall'Asen, **Xiaofeng Zhang**, Reyhane Askari Hemmat, et al. (2025). **Increasing the Utility of Synthetic Images through Chamfer Guidance**. In *NeurIPS 2025*, San Diego, USA.

[C.5] Zeyu Huang, Yikang Shen, **Xiaofeng Zhang**, et al. (2023). **Transformer-patcher: One mistake worth one neuron**. In *ICLR 2023*, Kigali, Rwanda.

[C.6] **Xiaofeng Zhang**, Yikang Shen, Zeyu Huang, et al. (2022). **Mixture of attention heads: Selecting attention heads per token**. In *EMNLP 2022*, pp. 4150–4162. ACL., UAE.

[W.1] Tom Stanic, Marco Jiralerspong, **Xiaofeng Zhang**, et al. (2026). **Beyond Reward Maximization: Evaluating the Diversity of Trajectories in Reinforcement Learning with Temporal Vendi Score**. In *ICLR 2026 Lifelong Agent Workshop*, Rio de Janeiro, Brazil.

Education

Mila (Québec AI Institute), Université de Montréal — Sept. 2023 – Present
PhD candidate in Computer Science — Montréal, Canada

- Supervisor: Aaron Courville
- Research interests: visual and multimodal generation, representation learning, diffusion and flow models, and world models
- GPA: 4.30/4.30

École Centrale de Pékin, Beihang University — Sept. 2020 – June 2023
Master's in Electronics & Information (China) & Diplôme d'ingénieur (France) — Beijing, China

- Supervisors: Zhang Xiong and Wenge Rong
- Thesis: Transformer architecture optimization for pre-trained language models

École Centrale de Pékin, Beihang University — Sept. 2016 – June 2020
Bachelor's in Information and Computing Science — Beijing, China

- Selected coursework: Mathematical Analysis, Algebra, Geometry, Probability and Statistics, Differential Equations, Theoretical Mechanics, Structural Mechanics, Electronics, Electromagnetism, Optics, and Automatic Control
- GPA: 3.87/4.00; average grade: 92.64/100; class rank: 1/87

Research Experience

FAIR, Meta Sept. 2024 – Sept. 2026
Visiting Researcher, AI Mentorship Program; supervised by Adriana Romero-Soriano Montréal, Canada

- Developed FeaTok, a discrete visual tokenizer built on pretrained vision encoders, with feature supervision and global-token conditioning for autoregressive image generation. The method shows consistent improvement in image generation with discrete tokens across four pretrained vision encoders.
- Contributed to WMReward on both the method and experiments, which uses JEPA world models to guide video generation at inference time, improving physical plausibility without fine-tuning the generator. WMReward is used by NVIDIA's Cosmos3 and Black Forest Labs' FLUX 3 in Physics-IQ evaluations.
- Contributed to the methodology of Discriminator-Guided RL, which uses discriminator rewards in pretrained feature spaces to post-train flow-based image generators, improving generation across SiT, JiT, REPA, and RAE.
- Designed a systematic evaluation for synthetic visual data from text-to-image models across levels of prompt complexity, comparing quality, diversity, and image–prompt consistency with real-world data.

WeChat, Tencent Jan. 2022 – June 2022
Research Intern; supervised by Yikang Shen Beijing, China

- Proposed mixture of attention heads, a sparse attention architecture that matched the performance of dense models with similar parameter counts while using substantially less computation.
- Contributed to Transformer-Patcher, a model editing method that continually updates knowledge in masked and autoregressive language models while preserving knowledge unrelated to the edits.

Honors and Awards

Doctoral Research Scholarship, Nature and Technology - C$ 58 334 April 2026
Fonds de recherche Québec, QC, CA

EDI in Research Scholarship - C$ 24 000 April 2025
Mila (Québec AI Institute), QC, CA

Tuition Exemption Scholarship - C$ 21 593.35 per year Sept. 2023
Université de Montréal, QC, CA

National Scholarship (Graduate) - C$ 1 600 Oct. 2021
Ministry of Education, CN

National Scholarship (Undergraduate) - C$ 1 000 Oct. 2017 and Oct. 2018
Ministry of Education, CN

Academic Service

Reviewer: ICLR 2027, ICML 2026 (Silver Reviewer), CVPR 2026, ICLR 2025, ICML 2025, NeurIPS 2025, EMNLP 2025, MAIS 2024

Organizer: Mila World Models Reading Group and First World Modelling Workshop

Booth support: Meta at ICLR 2026 (Rio de Janeiro, Brazil), Mila at NeurIPS 2024 (Vancouver, Canada), and NeurIPS 2025 (San Diego, USA)

Additional Information

Programming and typesetting: Python, Java, C, LaTeX

ML frameworks and tools: PyTorch, Transformers, Diffusers, vLLM, NumPy, Weights & Biases, and Submitit

Languages: Mandarin (native), English (professional), French (professional)