Yang Fei

Address: Hong Kong University of Science and Technology, Hong Kong SAR, China
Email: yfeiac@connect.ust.hk
Homepage: sunfly04.github.io

Research Interests

- My research focuses on **motion-centric generative video models** as interactive neural simulators. I work on **motion-preserving video tokenizers** and on grounding diffusion-based generators in **structural priors** to achieve temporally consistent, physics-aware video synthesis for agents that need to act and plan in complex environments.

Education

- **Hong Kong University of Science and Technology** — Hong Kong SAR
  B.Sc. in Computer Science and Mathematics — Sept. 2022 – June. 2026 (Expected)
  Overall GPA: 4.064/4.3 (**Top 2 in CS & MATH**, Class of 2026) Major GPA: 4.151/4.3

- **University of Washington** — Seattle, WA
  College of Engineering, Exchange Student — Jan. 2025 – Mar. 2025
  GPA: 3.97/4

Selected Coursework: Advanced Deep Learning Architectures (Grad Level, A), Computer Graphics (UW, 3.9/4.0), Honors Design and Analysis of Algorithms (A+), Honors OOP and Data Structures (A+), Applied Statistics (A+), Statistical Computing (A+), Multivariable Calculus (A+), Linear Algebra (A+).

Publications

* Denotes equal contribution

[1] **Structure From Tracking: Distilling Structure-Preserving Motion for Video Generation** **[paper] [page]**
Yang Fei, George Stoica, Jingyuan Liu, Qifeng Chen, Ranjay Krishna, Xiaojuan Wang, Benlin Liu
Under Review at *CVPR 2026*

[2] **VideoVAE+: Large Motion Video Autoencoding with Cross-modal Video VAE** **[paper] [page] [code]**
Yazhou Xing*, **Yang Fei***, Yingqing He*, Jingye Chen, Jiaxin Xie, Xiaowei Chi, Qifeng Chen
In *ICCV 2025*.

Research Experiences

RAIVN Lab, University of Washington — Feb. 2025 – Present
Advisor: Ranjay Krishna, Assistant Professor at UW

- **Structure-Preserving Distillation for Video Generation** (Lead Author)
  - Spearheaded the development of **SAM2VideoX**, a novel framework that distills structure-preserving motion priors from tracking models (SAM2) into Video Diffusion Transformers. This work is currently under review at **CVPR 2026** [1].
  - Proposed a Bidirectional Feature Fusion mechanism to bridge the architectural gap between causal trackers and bidirectional generators, and designed a Local Gram Flow loss to enforce relational motion consistency.
  - Curated a large-scale motion dataset (∼10k clips) focusing on articulated objects (e.g., animals, humans) to facilitate robust representation learning.
  - Reduced FVD by 22.5% compared to baselines, and secured a 71.4% win rate in human preference evaluations.

Visual Intelligence Lab, HKUST — Jan. 2024 – Present
Advisor: Qifeng Chen, Associate Professor at HKUST

- **Large Motion Video Autoencoding with Cross-modal Video VAE** (Co-first Author)
  - Developed a robust Video Variational Autoencoder (VAE) that achieves high-fidelity reconstruction, particularly in challenging large-motion scenarios. This work was accepted to **ICCV 2025** [2].
  - Implemented the first cross-modal integration strategy incorporating textual semantics into Video VAEs, enhancing alignment with text-video paired data.
  - Conducted extensive ablation studies proving the method's superiority over state-of-the-art baselines.

Projects

- **VideoTuna: Video Generation Toolkit (500+ Stars)** [code]
  Advisor: Qifeng Chen July 2024 – Jan. 2025
  - Engineered a comprehensive toolkit for fine-tuning video generation models with support for various VAE architectures.
  - Maintained the open-source codebase, optimizing data processing pipelines for large-scale video datasets.
- **AIGC Showcase Platform** [page]
  Advisor: Qifeng Chen June 2023 – Aug. 2023
  - Built a web platform demonstrating SOTA video generation capabilities and conducted hyperparameter tuning experiments to optimize generation quality.

Teaching

- **COMP 1023 - Python Programming** Fall 2024
  Undergraduate Teaching Assistant HKUST
  - Developed two laboratory exercises and one Programming Assignment (PA) from the ground up.
  - Prepared self-assessment tests for future cohorts to evaluate students' understanding of the material.
- **COMP 2012 - Object-Oriented Programming and Data Structures** Spring 2024
  Undergraduate Teaching Assistant HKUST
  - Facilitated weekly lab sessions and provided one-on-one technical support to help students debug code and understand algorithmic concepts.
- **COMP 2011 - Programming with C++** Fall 2023
  Undergraduate Teaching Assistant HKUST
  - Conducted lab tutorials on C++ fundamentals and resolved technical inquiries regarding pointers and memory management.

Honors & Awards

- **HKSAR Government Scholarship (Top 1%)** 2023 – Present
- **Hong Kong, China - Asia-Pacific Scholarship** 2024/25
- **Tse Cheuk Ng Tai Scholarship** Fall 2025
- **HKSAR Government Scholarship Fund - Reaching Out Award** Spring 2025
- **Best UROP Mini-Conference Paper Award** Fall 2024
- **Dean's List** All Semesters (2022 – Present)
- **First Prize**, National Olympiad in Informatics in Provinces (NOIP) 2020

Skills

- **Languages**: Python, C/C++, Java, LaTeX
- **Tools**: PyTorch, Git, HTML/CSS
- **English**: TOEFL 102 (Speaking 25)