

YIHUA DU

◇ ydu967@connect.hkust-gz.edu.cn ◇ <https://hit-perfect.github.io/>

EDUCATION

Hong Kong University of Science and Technology, Guangzhou

Sep. 2025 - Present

M.Phil. student in Artificial Intelligence (Red Bird MPhil Program), advised by Prof. Ying-Cong CHEN.

Harbin Institute of Technology

Sep. 2021 - Jul. 2025

B.S. in Computer Science, Rank: 16/117, GPA: 3.787/4, advised by Prof. Shaohui Liu.

EXPERIENCES

LIGHTSPEED STUDIOS, Tencent

Jan. 2026 - Present

Qingyun Program Intern

Research Intern working on World Model, Video Generation Acceleration, advised by Shengju Qian and Zeyu Hu

Hong Kong University of Science and Technology, Guangzhou

Mar. 2025 - Aug. 2025

Research Assistant working on Controllable Video Generation, advised by Luozhou Wang and Guibao Shen

Meituan

Mar. 2024 - Aug. 2024

Research Intern working on Face Anti-Spoofing, advised by Lei Deng and Jiaxiang You

PROJECTS

StreamAccel: Automated Inference Acceleration for Streaming Causal Video Models

April. 2026 - Present

Core Contributor

- **Automated Streaming Inference:** Developed a staged, agent-guided workflow and reusable technique knowledge library to select, combine, and validate optimizations for causal video models. On Solaris and SolarWM-Wan5B, reduced steady-state DiT latency by 38.33% and 32.21%, respectively, and VAE decoding latency by up to 24.84%.

StereoPilot: Learning Unified and Efficient Stereo Conversion via Generative Priors

Mar. 2025 - Dec. 2025

Core Contributor

- **Efficient Stereo Video Generation:** Proposed StereoPilot, a feed-forward diffusion framework that directly synthesizes stereo views without explicit depth estimation, resolving ambiguity in reflective scenes. We constructed UniStereo, the first large-scale benchmark unifying parallel and converged formats, and designed a learnable domain switcher with cycle-consistency loss to adapt to diverse stereo geometries. This approach achieves state-of-the-art visual fidelity and significantly boosts efficiency, reducing inference latency to just 11 seconds.

VideoMemory: Toward Consistent Video Generation via Memory Integration

Sep. 2025 - Dec. 2025

Core Contributor

- **Multi-shot Consistent Video Generation:** Proposed VideoMemory, an entity-centric framework for multi-shot consistent video generation. It reframes long-range consistency as entity asset management, introducing a read-write Dynamic Memory Bank and a multi-agent closed loop (retrieve → generate on miss → write back) to enable cross-shot entity reuse and reduce identity drift under story-driven control. Built a 54-case multi-shot consistency benchmark spanning characters, props, and backgrounds; experiments show improved consistency and overall visual quality.

A Mechanistic View on Video Generation as World Models: State and Dynamics

Sep. 2025 - Feb. 2026

Main Contributor

- **Video World Model Survey:** Analyzed explicit and implicit state representations and causal dynamics in video generation models, then synthesized evaluation criteria and research directions for interactive world simulation.

- **Face Anti-Spoofing Defense:** Developed a tamper- and recompression-aware liveness pipeline for face verification, combining forensic cues and multimodal features; achieved a 99.99% interception rate on in-house data. Built a data-collection and labeling pipeline with 500K+ samples aligned to production conditions, and added an interpretable image-quality module for lightweight deployment.

TECHNICAL REPORT

[1] **SolarWM: Open Data and Scalable Training for Long-Horizon Video World Models** *arXiv 2026*
Junchao Huang, Guian Fang, Shengju Qian, Xianghao Kong, Zhuoran Zhao, Wei Huang, **Yihua Du**, Zixin Zhang, Justin Cui, Yuchao Gu, Yukang Chen, Xinting Hu, Tianyu He, Shaoshuai Shi, Zhuotao Tian, Xin Wang, Mike Zheng Shou, Li Jiang.

PUBLICATIONS

* denotes equal contribution.

[1] **StreamAccel: Automated Inference Acceleration for Streaming Causal Video Models** *Submitted ICLR2027*
Yihua Du, YaoYang Liu, Junchao Huang, Shengju Qian, Luozhou Wang, ZhiFei Chen, Guibao Shen, Dongyu Yan, Xianghao Kong, Zeyu HU, Xin WANG, Ying-Cong Chen.

[2] **StereoPilot: Learning Unified and Efficient Stereo Conversion via Generative Prior** *arXiv 2025*
Guibao Shen*, **Yihua Du***, Wenhang Ge*, Jing He, Chirui Chang, Donghao Zhou, Zhen Yang, Luozhou Wang, Xin Tao, Ying-Cong Chen.

[3] **VideoMemory: Toward Consistent Video Generation via Memory Integration** *arXiv 2026*
Jinsong Zhou*, **Yihua Du***, Xinli Xu*, Luozhou Wang, Zijie Zhuang, Yehang Zhang, Shuaibo Li, Xiaojun Hu, Bolan Su, Ying-Cong Chen.

[4] **A Mechanistic View on Video Generation as World Models: State and Dynamics** *arXiv 2026*
Luozhou Wang*, Zhifei Chen*, **Yihua Du**, Dongyu Yan, Wenhang Ge, Guibao Shen, Xinli Xu, Leyi Wu, Man Chen, Tianshuo Xu, Peiran Ren, Xin Tao, Pengfei Wan, Ying-Cong Chen.

[5] **ReWorld: An Interactive World Model with Long-Horizon Memory** *arXiv 2026*
Zhifei Chen*, Luozhou Wang*, Guibao Shen, Dongyu Yan, Shuai Yang, Tianshuo Xu, **Yihua Du**, Wei Wang, Tianyi Gui, Lianghua Huang, Yingcong Chen.

TECHNICAL STRENGTHS

Programming Languages	Python, Java, C++, Matlab
Deep Learning Framework	Pytorch, Tensorflow

HONORS & AWARDS

MPhil Scholarship, HKUST (Guangzhou)	RMB 240,000 over two years
National Prize, National College Students' Innovation and Entrepreneurship Program	Sep. 2024
National Second Prize, The 6th Global Campus Artificial Intelligence Algorithm Elite Competition	Nov. 2024
Outstanding Innovative Integrated Design Award — Harbin Institute of Technology (HIT) Undergraduate Thesis (Final-Year Project)	May. 2025
Top 20 (Northeast China Regional Division) — The 17th National College Students' Software Innovation Competition	Mar. 2024
Harbin Institute of Technology Outstanding Student Award	2022, 2023
Renmin Scholarship (Harbin Institute of Technology)	2022