Chia-Hsiang Kao

ck696@cornell.edu | Google Scholar | Website

Ithaca, New York (14850)

Vision and Interest

My research vision is to build systems that infer and learn the hidden dynamics of real-world interactions. Specifically, I am interested in **Vision Reasoning**, **Visual–Symbolic Learning (image/video-to-code), Video Editing**, and **AI Agent**.

I am currently a **Cornell Foundational AI PhD Fellow (2025)**, a Meta Student Researcher, and was a Applied Scientist Intern at **Amazon**.

Education

Cornell University, *Ithaca, NY*   *2023 - Present*

- Ph.D. Student in Computer Science
- Advisor: Professor Bharath Hariharan

National Yang-Ming Chiao-Tung University, *Taipei, Taiwan*   *2015 - 2022*

- Doctor of Medicine

Work Experience

Applied Scientist Intern, Amazon; Sunnyvale, CA   *May 2025 -Aug. 2025*

- I developed VLMs that fully reconstructs the real-world rigid-body dynamics in MuJoCo simulator.
- We leveraged language as a structural visual representation and enhanced performance with event reasoning.

Student Researcher (part-time), Meta; remote   *Dec 2025 -May. 2026*

- I will develop video-to-video editing system for controlling and changing the real-world physics and motion in the video.

Publications

[U.1] **Inferring Physics from Videos for Rigid Object Motion Capture and Editing**
Chia-Hsiang Kao, Chien-Yi Wang, Cong Phuoc Huynh, Bharath Hariharan
Submitted to CVPR 2026

[S.2] **Towards LLM Agents for Earth Observation.**
Chia-Hsiang Kao, Wenting Zhao, Shreelekha Revankar, Samuel Speas, Snehal Bhagat, Rajeev Datta, Cheng Perng Phoo, Utkarsh Mall, Carl Vondrick, Kavita Bala, Bharath Hariharan
ICML 2025 workshop

[C.6] **Counter-Current Learning: A Biologically Plausible Dual Network Approach for Deep Learning**
Chia-Hsiang Kao, Bharath Hariharan
NeurIPS 2024

[C.5] **AllClear: A Comprehensive Dataset and Benchmark for Cloud Removal in Satellite Imagery**
Chia-Hsiang Kao*, Hangyu Zhou*, Cheng Perng Phoo, Utkarsh Mall, Bharath Hariharan, Kavita Bala
NeurIPS Datasets and Benchmarks Track 2024

[C.4] **Caduceus: Bi-Directional Equivariant Long-Range DNA Sequence Modeling**
Yair Schiff, Chia-Hsiang Kao, Aaron Gokalsan, Tri Dao, Albert Gu, Volodymyr Kuleshov
ICML 2024

[C.3] **Advancing DNA Language Models: The Genomics Long-Range Benchmark**

Chia-Hsiang Kao*, Evan Trop*, McKinley Polen*, Yair Schiff*, Bernardo P. de Almeida, Aaron Gokaslan, Thomas Pierrot, Volodymyr Kuleshov

AAAI (workshop) 2023

[S.1] **FedBug: A Bottom-Up Gradual Unfreezing Framework for Federated Learning**

Chia-Hsiang Kao, Yu-Chiang Frank Wang

Preprint

[C.2] **MAML Is a Noisy Contrastive Learner in the Classification**

Chia-Hsiang Kao, Wei-Chen Chiu, Pin-Yu Chen

ICLR 2022

[C.1] **Demystifying T1-MRI to FDG18-PET Image Translation via Representational Similarity**

Chia-Hsiang Kao, Yong-Sheng Chen, Li-Fen Chen, Wei-Chen Chiu

MICCAI 2021

Honors & Awards

- **Foundational AI PhD Fellowship**, *Cornell*   *2025*
- Student Travel Award, *MICCAI*   *2021*
- Undergraduate Research Fellowship, *National Science and Technology Council, Taiwan*   *2018, 2020*
- Summer Research Fellowship, *National Health Research Institutes*   *2018*

Services

Conference Reviewer: AAAI[25], AISTATS[25], AutoML[22], ICML[25], ICLR[25], NeurIPS[21,24,25], ICLR [25, 26]
Journal Reviewer: CVIU (2022), Comput. Electr. Eng. (2024), IEEE TETCI (2024)