TIANWEN FU

tianwenf - at - usc.edu | twfu.me

RESEARCH INTEREST

I am interested in integrating reasoning into generative modeling to develop video generation algorithms grounded in physical and spatial reality. I aim to leverage the physical commonsense and causal reasoning of foundation models to enforce spatiotemporal consistency and physical plausibility, including global illumination effects, in generated scenes.

EDUCATION

University of Souther California — Los Angeles, CA
Ph.D. in Computer Science — May 2029

Carnegie Mellon University — Pittsburgh, PA
M.Sc. in Computer Vision — Dec 2023
GPA: 4.29/4.33

The Chinese University of Hong Kong — Hong Kong
B.Sc. in Mathematics and Information Engineering — Jul 2022
Major GPA: 3.98/4.00 | GPA: 3.88/3.40 | Rank: 1/16

University of Pennsylvania — Philadelphia, PA
Exchange Student, Computer Engineering — May 2020
GPA: 4.00/4.00

PUBLICATIONS

* denotes equal contribution

Mind-to-Face: Neural-Driven Photorealistic Avatar Synthesis via EEG Decoding
Tianwen Fu*, Haolin Xiong*, P. Prasad, Y. Cai, W. Teng, H. Xiao, H. Chen, Y. Zhao.
In submission to the IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**) 2026.

RDD: Robust Feature Detector and Descriptor using Deformable Transformer
Gonglin Chen, **Tianwen Fu**, Haiwei Chen, Wenbin Teng, Hanyuan Xiao, Yajie Zhao.
IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**) 2025.

Unraveling the Complexities of Simplicity Bias: Mitigating and Amplifying Factors
Xuchen Gong*, **Tianwen Fu***.
Neural Information Processing Systems (**NeurIPS**) Mathematics of Modern Machine Learning Workshop 2023.

AutoLoss-Zero: Searching Loss Functions from Scratch for Generic Tasks
Hao Li*, **Tianwen Fu***, Jifeng Dai, Hongsheng Li, Gao Huang, Xizhou Zhu.
IEEE/CVF Conference on Computer Vision and Pattern Recognition (**CVPR**) 2022.

Computational Design and Optimization of Non-Circular Gears
Hao Xu*, **Tianwen Fu***, Peng Song, Mingjun Zhou, Niloy J. Mitra, Chi-Wing Fu.
Computer Graphics Forum (**CGF**), full paper (**Eurographics**) 2020.

EXPERIENCES

Vision and Graphics Lab, USC — Los Angeles, CA
Ph.D. Student | Advisor: Yajie Zhao — Aug 2024 – Present

Neural Avatar Synthesis from Biological Signals.

- Developed animatable 3D avatars with Gaussian splatting driven by textureless mesh files generated from EEG signals, submitted to CVPR 2026

- Surveyed video understanding models for feature extraction for emotion-eliciting movies

Multi-view Stereo and Novel View Synthesis.

- Conducted multi-view stereo and dataset collection for RDD, published in CVPR 2025
- Tuned diffusion models for novel view synthesis in outdoor scenes for an IARPA project

Imaging and Rendering Lab, CMU Pittsburgh, PA

Extern (Transitional Graduate Student) | Advisor: Ioannis Gkioulekas Feb 2024 — May 2024

HDR Imaging with Event Cameras.

- Established the hardware and algorithm for inferring an HDR intensity image from event cameras, including DAVIS346 and EVK4, by measuring the inter-event time with a blinking light source

Robots Perceiving and Doing Lab, CMU Pittsburgh, PA

Graduate Student (Capstone Project) | Advisor: David Held Feb 2024 — May 2024

3D Understanding with Depth Cameras for Robotic Perception.

- Implemented and evaluated algorithms and targets for extrinsics calibration of multiple Azure Kinect cameras
- Established a synchronized multi-camera system that captures videos of point clouds with human activities
- Built a MuJoCo-based digital twin of 3D scenes with depth cameras, facilitating verification and unit testing
- Proposed methods to match 2D points and segmentation masks among views of the same scene via geometry and bipartite matching, enabling the generation of 3D segmentation from 2D foundation models

Ambarella Santa Clara, CA

Algorithm Engineer Intern | Supervisor: Hua Lin Feb 2024 — May 2024

Diffusion-based Data Augmentation for Object Detection.

- Accelerated Shape-guided Diffusion by parallelizing input prompts and refactoring suboptimal tensor operations, reducing the inference time by 35% (Pull Request)
- Revised interpolation methods of Shape-guided Diffusion, improving the performance over small objects
- Adapted generative methods, including Diffusion-based and GAN-based, to the augmentation of existing instance segmentation datasets with generated samples, with a 1.0 increase in the class-specific mask mAP

SenseTime Beijing, China

Research Intern | Supervisor: Jifeng Dai Sep 2020 — Dec 2021

Generic AutoML Algorithms.

- Proposed an efficient generic framework for loss function search with random knowledge-free initialization but performance matching the expert-designed counterparts in 192 GPU hours, published in CVPR2022
- Investigated the inconsistency between the BLEU metric and conventional cross-entropy-based losses, with a detailed report on the challenges of a possible BLEU loss
- Explored extensions of AutoLoss-Zero to automatic generation and evaluation of ViT attention patterns to achieve better trade-offs of computational resources and performance

Unsupervised Foundation Models.

- Incorporated unimodal unsupervised learning techniques into pretraining large-scale visual-linguistic models

Google Glass Applications.

- Developed a facial and speech recognition application for Google Glass to memorize and recognize names

University of Cambridge Cambridge, UK

Summer Research Intern | Supervisor: Ramji Venkataramanan Jun 2020 -– Aug 2020

- Developed an efficient implementation of the belief propagation algorithm that detects defective nodes with LPDC codes, suitable for both noiseless and noisy cases

The Chinese University of Hong Kong Hong Kong

Summer Research Intern | Supervisor: Wing Cheong Lau Jun 2019 – Aug 2019

- Automated the pipeline for analyzing the vulnerabilities of mobile payment applications, with an SQLite dataset storing the metadata and results

- Parallelized the downloading script with multiprocessing and error handling techniques, reaching a throughput of 34000 packages per day for 45 consecutive CPU days

The Chinese University of Hong Kong Hong Kong
Undergraduate Research Intern | Supervisor: Chi-wing Fu Jan 2019 – Aug 2019

- Produced the algorithm for designing and optimizing non-circular gears that resemble user input and rotate seamlessly, published in Eurographics 2020
- Programmed and evaluated an optimization-based approach in C++ for a project that aims to find non-periodic tiling schemes of arbitrary 2D shapes with blocks given by user input

PROJECTS

Noise Model for Continuous-wave Time-of-flight Sensors Report | Code | Spring 2023

- Formulated a noise model of the inaccuracy in depth measurements of continuous-wave time-of-flight sensors in low-albedo regions with Fourier analysis
- Conducted experiments on a Pico Flexx sensor with printed patterns of different albedo and at various depths

Non-Gaussian Transition Distributions in Discrete Diffusion Report | Code | Spring 2023

- Investigated effects of non-Gaussian transition distributions in the noising step of discrete diffusion models with a proposed synthetic dataset, with detailed reports on peakiness, absorbing state, and noise schedules

SELECTED COURSEWORK

Computer Vision, Computational Imaging, and Optics

- Physics-based Methods in Vision
- Geometry-based Methods in Vision
- Computer Vision
- Computational Photography
- Optics
- Optical Materials, Instruments and Devices

Machine Learning, Deep Learning, and Information Retrieval

- Advanced Natural Language Processing
- Probabilistic Graphical Models
- Introduction to Robot Learning
- Text Mining Models and Application

Foundational Computer Science and Signal Processing

- Advanced Data Structure
- Operating Systems Design and Implementation
- Compilers and Interpreters
- Information Theory
- Signals and Systems
- Principle of Communication Systems

SELECTED AWARDS AND SCHOLARSHIPS

Graduate School Fellowship for Incoming Students USC, 2024
Charles Kao Top Performance Awards CUHK, 2022
Arthur and Louise May Scholarship CUHK, 2020
Charles K. Kao Research Exchange Scholarship CUHK, 2020

SERVICE

Reviewer: NeurIPS 2023 Workshop M3L, CVPR 2024, ECCV 2024, CVPR 2025, AAAI 2025
Teaching Assistant: CSCI 578: Software Architecture (Fall 2025)

SKILLS

Programming Languages: Python, C/C++, Bash, C#, Rust
Frameworks and Toolkits: PyTorch, Jax/Flax, CUDA, Numpy, OpenCV, Slurm, Git, CUDA
Deep Learning: AutoML, Contrastive Learning, Reinforcement Learning, Diffusion Models
Computer Vision: Projective and Epipolar Geometry, Structure from Motion, Optics, Radiometry