

Tianwen Fu

tianwenf@cs.cmu.edu ♦ (412) 520-6178 ♦ twfu.me

EDUCATION

Carnegie Mellon University

Master of Science in Computer Vision

GPA: 4.26/4.33

Pittsburgh, PA

Dec. 2023

The Chinese University of Hong Kong

Bachelor of Science in Mathematics and Information Engineering

Stream: Engineering Leadership, Innovation, Technology and Entrepreneurship

Major GPA: 3.98/4.00 | GPA: 3.88/4.00 | Rank: 1/16

Hong Kong SAR

Jul. 2022

University of Pennsylvania

Exchange Student

GPA: 4.00/4.00

Philadelphia, PA

May 2020

PUBLICATIONS

Unraveling the Complexities of Simplicity Bias: Mitigating and Amplifying Factors

Xuchen Gong*, Tianwen Fu* (**equal contribution*)

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. (**equal contribution*)

International 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. (**equal contribution*)

Computer Graphics Forum (CGF), full paper (Eurographics) 2020.

EXPERIENCE

Ambarella

Algorithm Engineer Intern | Supervisor: Hua Lin

Santa Clara, CA

May 2023 – Aug. 2023

- Accelerated Shape-guided Diffusion by parallelizing input prompts and refactoring suboptimal tensor operations, reducing the inference time by 35% [[GitHub Pull Request](#)]
- Corrected interpolation methods of Shape-guided diffusion, improving the performance of small objects
- Assessed the performances of various generative methods in augmenting existing object detection datasets with generated samples, increasing the class-specific mask mAP by 1.0

SenseTime

Research Intern | Supervisor: Jifeng Dai

Beijing, China

Sep. 2020 – Dec. 2021

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

Summer Research Intern | Supervisor: Ramji Venkataramanan

Cambridge, UK
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

Summer Research Intern | Supervisor: Wing Cheong Lau

Hong Kong SAR
Jun. 2019 – Aug. 2019

- Automated the pipeline for analyzing the vulnerabilities of mobile payment applications, with a 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

Undergraduate Research Intern | Supervisor: Chi-wing Fu

Hong Kong SAR
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

3D Kitchen Understanding (Capstone Project)

2023 (In Progress)

- Implemented and evaluated algorithms and targets for extrinsics calibration of multiple 3D cameras
- Established a synchronized multi-camera system that captures videos of point clouds with human activities
- Facilitated unit testing and method verification with a digital twin in a Mujoco environment
- Proposed a method to generate 3D segmentation from unsupervised 2D foundation models and scene geometry

Noise Model for Continuous-wave Time-of-flight Sensors

[\[Code\]](#) | Spring 2023

- Formulated a noise model of the phase measurement of continuous-wave time-of-flight depth sensors in low-albedo regions, with experiments conducted on a Pico Flexx sensor with different albedo and depths

Non-Gaussian Transition Distributions in Discrete Diffusion

[\[Code\]](#) | 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 and Computational Imaging

- Physics-based Methods in Vision
- Geometry-based Methods in Vision*
- Computer Vision (Graduate-level)
- Computational Photography*

General Machine Learning and Deep Learning

- Introduction to Machine Learning (PhD)
- Probabilistic Graphical Models
- Introduction to Robot Learning
- Text Mining Models and Application

Foundational Computer Science and Signal Processing

- Advanced Data Structures
- Operating Systems Design and Implementation
- Compilers and Interpreters
- Information Theory
- Signals and Systems
- Principal of Communication Systems

“*” denotes courses in progress. Introduction to Robot Learning was formerly named as “Statistical Techniques in Robotics”.

SKILLS

Programming Languages: Python, C/C++, Bash Scripting, C#

Frameworks and Toolkits: PyTorch, Jax/Flax, CUDA, Numpy, OpenCV, Slurm, Git

Algorithms and Paradigms: Object Detection, Instance Segmentation, AutoML, Contrastive Learning, Denoising Diffusion Reinforcement Learning, Parallel Programming (CUDA, Multiprocessing), Projective Geometry, Structure-from-Motion