Tianwen Fu

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

EDUCATION

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

Carnegie Mellon University — Pittsburgh, PA
Master of Science in Computer Vision — Dec. 2023
GPA: 4.29/4.33

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

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

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

Carnegie Mellon University — Pittsburgh, PA
Algorithm Engineer Intern | Supervisor: Ioannis Gkioulekas — Feb. 2024 – May 2024

- 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

Ambarella — Santa Clara, CA
Algorithm Engineer Intern | Supervisor: Hua Lin — 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]
- Revised interpolation methods of Shape-guided diffusion, improving the performance over small objects
- Adapted generative methods, including Shape-guided Diffusion, CoMoGAN, and Grounded Diffusion, 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

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 SAR

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 SAR

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

3D Kitchen Understanding (Capstone Project) 2023

- 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 generation of 3D segmentation from 2D foundation models

Noise Model for Continuous-wave Time-of-flight Sensors [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 pattern of different albedo and at various 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

Machine Learning, Deep Learning, and Information Retrieval

- 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
- Principle of Communication Systems

SELECTED AWARDS AND SCHOLARSHIPS

Charles Kao Top Performance Awards HK$10000 (US$1280) | 2020

- Awarded to the top final year student with academic excellence in the major program

Arthur and Louise May Scholarship HK$50000 (US$6400) | 2020

- Awarded to top students participating in outgoing exchange programs

Charles K. Kao Research Exchange Scholarship HK$50000 (US$6400) | 2020

- Awarded to at most 7 students in the Faculty of Engineering for research exchange programs
- Selection conducted by faculty based on the applicants' academic performance and research potential

SERVICE

Reviewer for NeurIPS 2023 Workshop M3L, CVPR 2024, ECCV 2024.

SKILLS

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

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

Geometry-based Computer Vision: Projective and Epipolar Geometry, Structure from Motion, Geometrical Optics, Radiometry

Deep Learning Algorithms: AutoML, Contrastive Learning, Reinforcement Learning, Denoising Diffusion Models, Object Detection, Instance Segmentation