Haoan Feng

Maryland, US hfengac@terpmail.umd.edu homepage Profile ORCID

ABOUT ME

I am a dedicated researcher with a deep passion for exploring the interdisciplinary fields of computer vision and geospatial data analysis. My research interests broadly span *neural representations of geospatial data, neural rendering, topological analysis, generative models, and AI for Science topics*. Believing that collaboration across disciplines is key to driving innovation, my experience working with diverse teams has strengthened my commitment to interdisciplinary research. I am constantly expanding my knowledge in these areas, driven by curiosity and a desire to contribute to advancements in the field.

EDUCATION

Doctor of Philosophy in Computer Science 2021 - *expect* 2026
University of Maryland, College Park, United States

- *GPA*: 4.0/4.0
- *Related coursework*: Geographical Information Systems and Spatial Databases, Advanced Techniques in Visual Learning and Recognition.
- *Thesis: (provisional) Neural Representations of Geospatial Data: Analysis, Generation, and Beyond*.

Master of Philosophy in Computer Science and Engineering 2018 - 2020
Hong Kong University of Science and Technology, Hong Kong, China

- *GPA*: 3.9/4.3
- *Thesis: Linear structure vectorization in large-scale landscape point cloud*.

Bachelor of Engineering in Computer Science Engineering and Electronic and Computer Engineering 2014 - 2018
Hong Kong University of Science and Technology, Hong Kong, China

- *GPA*: 3.9/4.3
- *Related coursework*: Discrete Math Tools, Advanced Computer Graphics, Data Visualization, Probability and Random Processes in Engineering, Signals and Systems.

RESEARCH PROJECTS

PhD Research Intern at Dolby ATG Laboratories May 2025 – August 2025
Sunnyvale, CA *Mentor & Manager: Harsha Musunuri Sri & Guan-ming Su*

- **Camera-Aware Video Understanding Framework**:
  - Designed a modular, training-free pipeline that augments Vision-Language Models (VLMs) with cinematographic awareness, focusing on camera motion, shot composition, and spatial geometry.

- Constructed a benchmark using synthetic 3D video clips rendered in Unreal Engine 5 with ground-truth camera metadata, enabling multi-label classification of canonical camera motions.
- Probed pretrained vision encoders with a Q-former to evaluate whether camera motion cues are preserved, demonstrating systemic neglect of spatial perception in current VLMs.

◇ **Prompt-Based Integration of Camera Perception**:

- Developed a plug-and-play strategy that injects camera motion and depth-aware spatial descriptors as textual prefixes to VLMs, enabling richer filmmaker-centric video understanding without retraining.
- Enabled downstream tasks such as camera-aware VQA, descriptive video services (DVS), and scene-level plagiarism detection through structured cinematographic conditioning.
- Validated that explicit camera prompts improve temporal reasoning and stylistic sensitivity in VLM-generated outputs, outperforming baselines on benchmarks like CameraBench and CineTechBench.

Analytical Neural Representations of Geospatial Data 2021 - Present
University of Maryland, College Park *Advisor: Prof. Leila De Floriani*

◇ **Survey on Neural Representations of Geospatial Data**:

- Conducting a comprehensive survey using the PyTorch framework on *neural representations* for geospatial data storage, rendering, and analysis.
- Evaluating implicit and explicit representations to enhance flexibility, scalability, parallel computation, and support for physical simulation in geospatial contexts.

◇ **Implicit Neural Representation for Terrain Surface Modeling**: Published a practical continuous surface model for terrain data using implicit neural representations.

- Achieved accurate terrain surface reconstruction with 25% storage and 4 times training speed with a progressive training strategy.
- Conducted topological analysis, and topographical feature extraction on the implicit surface function represented by the neural network.
- Evaluated and visualized experimental results using OpenCV and Matplotlib and collaborated with teammates on the W&B platform for project and experiment tracking.

◇ **Topological Feature Tracking on Triangulated Irregular Networks (TINs) Using a Scale-Space Approach**: Designed and implemented an adaptive scale-space algorithm to track topologically critical features on 2D manifolds discretized as TINs.

- Adapted previous regular grid-based scale-space algorithm to work with TINs for more accurate critical feature tracking results and less overall computational resources.
- Accelerated the triangular mesh smoothing process by approximately 100 times using C++ data structure extracting adjacency graph and custom GPU kernels via PyTorch.
- Implemented a local geometry adaptive sampling method to construct efficient TINs from point cloud for topological analysis.

Researcher and Developer at Vision and Graphics Laboratory 2017 - 2020
Hong Kong University of Science and Technology *Advisor: Prof. Long Quan*

- ◇ **Large-Scale Point Cloud Processing, Information Extraction, Semantic Segmentation**: Developed a comprehensive pipeline for processing large-scale noisy point clouds, enhancing PointNet's ability to segment thin and neglected structures. Created algorithms and a GUI tool for feature extraction, clustering, and recovering lost linear structures.
  - Enhanced PointNet with handcrafted features to better identify and segment thin structures in point clouds.
  - Designed a point cloud processing pipeline that statistically removes noisy data, analyzes local geometric features and adaptively segments the cloud for downstream tasks like pointwise classification, surface detection, and structure preservation.
  - Implemented a robust PCA algorithm (Fast-MCD) for fast data clustering and feature extraction as a pretext task.
  - Created algorithms and GUI tools for recovering lost linear structures (e.g., high-voltage powerlines), achieving modeling accuracy comparable to laser-scanning techniques.
- ◇ **3D Web Application for Large-scale Landscape Reconstruction**: Implemented a 3D web application, which loads 3D reconstructions of large-scale landscapes efficiently, and provides simulation of the Earth for user interaction and engineering measurement.
  - Built an efficient 3D data loading pipeline utilizing level-of-detail (LOD) to reduce data streaming by 90%, achieving 60fps on standard devices by optimizing data processing with web workers.
  - Integrated Apple ARKit for AR mode, designing user interactions for large-scale scenarios with precise transformation models.
  - Implemented high-accuracy landscape measurements for point-to-point distances and area calculations, offloading computational tasks to GPU via WebGL for a low-latency experience.
  - Enabled real-time interaction with complex 3D scenes containing billions of triangles using hidden-frame rendering and masking algorithms.

PUBLICATIONS

Conference Articles

[1] **Feng, H.**, Song, Y., & De Floriani, L. (2024). **Critical Features Tracking on Triangulated Irregular Networks by a Scale-Space Method.** In *The 32nd ACM International Conference on Advances in Geographic Information Systems (SIGSPATIAL '24)*, October 29-November 1, 2024, Atlanta, GA, USA. ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3678717.3691218. **(Best paper runner-up, Oral Presentation, 🔗, 📄)**

[2] **Feng, H.**, Xu, X., & De Floriani, L. (2024). **ImplicitTerrain: a Continuous Surface Model for Terrain Data Analysis.** In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 899-909). **(Oral Presentation @ *1st Implicit Neural Representation for Vision Workshop*, 🔗)**

[3] Zhen, M., Li, S., Zhou, L., Shang, J., **Feng, H.**, Fang, T., & Quan, L. (2020). **Learning Discriminative Feature with CRF for Unsupervised Video Object Segmentation.** In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVII 16* (pp. 445-462). Springer International Publishing.

Preprint(s)

[1] **Feng, H.**, Novello, T., Aldana, D., & De Floriani, L. **SASNet: Spatially-Adaptive Sinusoidal Neural Networks**. **(🔗, 📄)**

[2] Aldana, D., **Feng, H.**, Novello, T., & De Floriani, L. **Structured Pruning in Implicit Neural Representations**. *(Under review)*.

[3] **Feng, H.** (2020). **Linear Structure Vectorization in Large-Scale Landscape Point Cloud.** *(MPhil dissertation, 🔗)*.

PRESENTATIONS

1. Critical Features Tracking on Triangulated Irregular Networks by a Scale-Space Method., *ACM SIGSPATIAL 2024*, Atlanta, United States (October 30th, 2024).
2. ImplicitTerrain: a Continuous Surface Model for Terrain Data Analysis, *CVPR 2024 Workshop on Implicit Neural Representation for Vision*, Seattle, United States (June 18, 2024).

SKILLS

Programming

- ◇ Daily programming languages: Python and C++.
- ◇ Parallel computation toolkits: CUDA and OpenMP.
- ◇ Experiment and data visualization toolkits: OpenCV, Matplotlib, W&B, D3.js, and Tableau.

Miscellaneous

- ◇ Professional engineering software: QGIS, MATLAB, MeshLab, and Paraview.
- ◇ Development primarily in Unix (CentOS) environments, managing computational tasks with SLURM.
- ◇ Database (MySQL, MongoDB) and web development (NodeJS, PHP, ReactJS).

Languages

- ◇ Mandarin (native), English (fluent), Cantonese & Japanese (amateur)

ACADEMIC SERVICE

Conference reviewer:

- ◇ International Conference on Pattern Recognition (ICPR 2024)
- ◇ Conference on Neural Information Processing Systems (NeurIPS 2025)

EXPERIENCE

Guest Lectures

- **Machine Learning Algorithms for Point Clouds** in the course *CMSC401: Algorithms for Geospatial Computing*, Spring 2024, 2025

Teaching Assistant

University of Maryland, College Park, MD, US 2021 - 2025

- CMSC414: Computer and Network Security, Fall 2022 & Fall 2024, 2025
- CMSC401: Algorithms for Geospatial Computing, Spring 2023, 2024, 2025
- CMSC416: Introduction to Parallel Computing, Fall 2023
- CMSC454: Algorithms for Data Science, Spring 2022
- CMSC427: Computer Graphics, Fall 2021

Hong Kong University of Science and Technology 2018 - 2020

- Introduction to Computer Science in Python and Multimedia Computing
- Object-Oriented Programming and Data Structures

Moodle Software Developer and Technical Support 2020

Hong Kong University of Science and Technology, Hong Kong, China

- Implemented LaTeX file compilation, encoding, and distribution features in the Moodle course management system as part of a copyright protection scheme.
- Developed web app user interfaces and system plugins for diverse applications using PHP and MySQL.

CERTIFICATIONS & AWARDS

- Chair's Graduate Fellowship 2021 - 2023
- Postgraduate Studentship 2018 - 2020
- Simatelex Charitable Foundation Scholarship 2015 - 2018
- University's Scholarship Scheme for Continuing Undergraduate Students (HKUST top 2% CGA Award) 2015 - 2018
- Champion Team in HackUST (Healthcare Theme), Hong Kong 2017
- First Runner-up in VAST Challenge 2016 (Global Data Visualization Competition) 2016
- First Prize in Chinese Western Mathematical Olympiad 2012

EXTRACURRICULAR ACTIVITIES

- Volunteer teaching and cultural experience program at Bali, Indonesia 2017
- General Secretary of the Model United Nations Club, HKUSTSU 2015 - 2016

REFERENCES

References available upon request.

August 18, 2025