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* 2027
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: Topology-Aware Modeling of Terrain and Volumetric Data: From Discrete to Neural Implicit Representations.*

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 2026 – August 2026
Sunnyvale, CA *Mentors & Manager: Peimeng Guan, Dae Yeol Lee & Guan-Ming Su*

Greenfield research internship. Proposed, scoped, and led a research project on native asset authoring for volumetric video in the TVGS format.

- **Native Gaussian Asset Insertion & Authoring for Volumetric Video**:
  - Proposed inserting authored text and icon elements into dynamic volumetric video as native Time-Varying Gaussian Splatting (TVGS) assets, so that inserted content follows the same rendering path as the host scene rather than being pasted on as a 2D overlay.

- Implemented single-pass depth-sorted compositing that produces correct mutual occlusion between inserted assets and the host scene, verified in behind-glass occlusion scenarios.
- Brought physically based shading to inserted assets with a Cook–Torrance material model and image-based environment lighting, validated against Blender references.
- Automated the conversion of font glyphs into a reusable TVGS asset library with per-asset acceptance checks, and studied video-codec compression of the resulting assets.
- Developed a Viser-based interactive scene editor featuring a keyframe timeline and a text-track schema for timed, multi-instance annotation reuse; batched neural decoding brought over 10× speed-up and made paragraph-scale text editing interactive.

◇ **Feed-Forward Amortized Asset Authoring**:

- Identified per-asset optimization as the scaling bottleneck of the authoring workflow, and formulated an amortized alternative where a shared encoder–decoder predicts renderable TVGS assets directly from canonical anchor point clouds, leaving only a compact anchor group as the per-asset payload.
- Designed the representation around deterministic on-surface anchor sampling and a compact single-plane latent, reaching normal map prediction accuracy under 2 degrees on held-out assets.
- Engineered the training stack from scratch, including a streaming dataset pipeline, batched rasterization, and multi-GPU training, then scaled it to a corpus of over 27,000 physically based assets with leak-free splits that generalize to fonts never seen in training.
- Devised a compressed asset-group delivery format with roughly 28× storage reduction at negligible quality loss, plus a test-time refinement stage that further improved every asset in the test library.

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

◇ **Geometry-Guided Camera Motion Understanding in VideoLLMs** (accepted at CVPR 2026 Workshop PVUW):

- Developed a geometry-guided framework to augment Video Large Language Models (VideoLLMs) with cinematographic camera motion understanding.
- 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 VideoLLMs.

◇ **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 VideoLLMs, 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 VideoLLM-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

Journal Article(s)

[1] **Feng, H.**, Song, Y., & De Floriani, L. (2026). **A Parallel Scale-Space Method for Critical Features Tracking on Triangulated Irregular Networks.** *ACM Transactions on Spatial Algorithms and Systems (TSAS)*. https://doi.org/10.1145/3812549. (🔗, 🌐)

Conference Articles

[1] **Feng, H.**, Aldana, D., Novello, T., & De Floriani, L. (2026). **SASNet: Spatially-Adaptive Sinusoidal Neural Networks.** In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2026)*. (🔗)

[2] **Feng, H.**, Musunuri, S. H., & Su, G.-M. (2026). **Geometry-Guided Camera Motion Understanding in VideoLLMs.** In *CVPR 2026 Workshop on Pixel-level Video Understanding in the Wild (PVUW)*. (🔗)

[3] **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, 🔗, 📄)**

[4] **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*, 🔗)**

[5] 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) & Other(s)

[1] **Feng, H.**, Xu, X., & De Floriani, L. (2026). **Rethinking Amortized Neural Representations for High-Resolution Terrain Elevation Data.** *arXiv preprint arXiv:2606.00404*. (🔗)

[2] **Feng, H.**, Xu, X., & De Floriani, L. (2026). **ImplicitTerrainV2: Wavelet-Guided Spatially Adaptive Neural Terrain Representation.** *arXiv preprint arXiv:2605.22556*. (🔗)

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

[4] **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, 2026)
- IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR 2026)
- IEEE/CVF International Conference on Computer Vision (ICCV 2026)
- Pacific Graphics (2026)

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

- Jacob K. Goldhaber Travel Grant (CVPR 2026), University of Maryland Graduate School 2026
- CS Department Travel Award (CVPR 2026), University of Maryland 2026
- Chair's Graduate Fellowship 2021 - 2023

July 23, 2026

- 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.

July 23, 2026