

YIFAN YANG

College Station • yyf990925@gmail.com • (213)590-8972 • <https://rayford295.github.io/>

Education Background

Texas A&M University, Ph.D. Student, GIScience (GeoAI), GPA 3.80/4 2024-now, College Station, USA
University of Southern California, MS, Spatial Data Science, GPA 3.96/4 May 2024, Los Angeles, USA
Hainan University, BE, Software Engineering, GPA 87.4/100 June 2022, Haikou, China

Technical Skills

Applications: PyTorch, TensorFlow, Scikit-learn, HuggingFace Transformers, Stable Diffusion, Visual Studio Code, Git/GitHub, Jupyter Notebook, Anaconda, Linux/Unix, Docker, Esri ArcGIS Pro, MS SQL Server, SQL Server Management Studio, PgAdmin, PostgreSQL & PostGIS

Programming: C, C++, Java, Python, SQL, R, ArcGIS API (JavaScript, Python), JavaScript, HTML

Professional (Work) Experience

Data Engineering and AI Model Fine-tuning, Vitality AI Feb 2023 – May 2023, Remote

- **Engineered end-to-end data processing pipelines** for large-scale image datasets, automating critical preprocessing steps including normalization, augmentation, and embedding extraction to ensure training data quality.
 - **Built scalable Python workflows** to orchestrate model training and inference deployment, integrating **PyTorch** ecosystems to streamline the transition from experimentation to product prototyping.
 - **Systematized model configuration management**, establishing standard procedures for batch processing and environment setup, significantly reducing the turnaround time for iterative model experiments.
-

Project Experience

Sat2Street-DisasterGen — Satellite-to-Street View Synthesis for Post-Disaster Assessment

GitHub: <https://github.com/rayford295/Sat2Street-DisasterGen>

Developed a multimodal generative framework for satellite-to-street view synthesis under post-disaster scenarios. The system benchmarks cross-view image generation strategies, integrating diffusion models, ControlNet, and VLM-guided semantic conditioning to synthesize realistic street-view imagery from post-disaster satellite inputs.

- Designed a unified evaluation framework ensuring consistent satellite inputs, data splits, and protocols across generative methods
- Implemented multiple cross-view generation pipelines, including Pix2Pix (cGAN), SD1.5 + ControlNet, VLM-guided synthesis, and Disaster-MoE (Mixture-of-Experts)
- Built a multi-stage LoRA training pipeline for damage-severity-aware generation (Mild / Moderate / Severe)
- Evaluated outputs using pixel-level metrics (FID, SSIM, LPIPS), semantic consistency (ResNet-based classification), and VLM perceptual scoring
- Developed reproducible dataset structure, inference scripts, and evaluation modules for controlled benchmarking

GeoAgent4Disaster — Multi-Agent GeoAI Framework for Multimodal Disaster Assessment

GitHub: <https://github.com/rayford295/GeoAgent4Disaster>

Developed an autonomous multi-agent GeoAI framework for hyperlocal, interpretable, and near-real-time disaster damage assessment. The system orchestrates multiple specialized agents leveraging vision-language foundation models to perform cross-view perception, damage recognition, reasoning, and automated report generation without task-specific retraining.

- Designed a modular multi-agent architecture, including Disaster Perception, Image Restoration, Damage Recognition, and Disaster Reasoning agents
- Integrated multimodal inputs: satellite imagery, street-view imagery, textual cues, and temporal change signals
- Implemented agent-based orchestration for zero-/few-shot disaster analysis using VLMs (e.g., GPT, Gemini)
- Enabled cross-view semantic alignment and causal reasoning for damage assessment and recovery recommendation
- Built an interpretable evaluation and reporting pipeline for structured disaster intelligence outputs

DamageArbiter — CLIP-Enhanced Multimodal Framework for Hurricane Damage Assessment

GitHub: <https://github.com/rayford295/DamageArbiter>

Developed a CLIP-enhanced multimodal arbitration framework for hyperlocal hurricane damage assessment using bi-temporal street-view imagery. The system integrates vision-language models (CLIP) with deep vision backbones (e.g., ViT-B/16) to improve damage classification robustness under real-world post-disaster uncertainty.

- Designed a multimodal arbitration mechanism to reconcile conflicting visual and semantic signals
- Implemented image classification pipelines using ViT-B/16 and transformer-based architectures
- Evaluated performance on hurricane street-view datasets with precision, recall, F1, and cross-validation
- Built a reproducible codebase, including dataset structure, model training scripts, and figure generation

BiTemporal-StreetView-Damage — Hyperlocal Disaster Assessment via Bi-Temporal Street-View Imagery

GitHub: <https://github.com/rayford295/Bi-Temporal-StreetView>

Developed a bi-temporal street-view image analysis framework for hyperlocal disaster damage assessment using pre- and post-disaster imagery. The system employs dual-channel architectures with pre-trained vision backbones (Swin Transformer, ConvNeXt) to improve classification accuracy and interpretability of damage severity estimation.

- Designed dual-channel fusion architecture for comparative reasoning between pre- and post-disaster imagery
- Constructed dataset of 2,249 labeled street-view image pairs with fine-grained severity annotations
- Achieved performance improvement from 66.14% (post-only) to 77.11% (bi-temporal) classification accuracy
- Applied Grad-CAM visualization to demonstrate enhanced spatial attention using pre-disaster inputs
- Evaluated multiple fusion strategies (cross-attention, weighted fusion, Siamese difference)
- Enabled fine-grained damage mapping for climate-resilient urban planning

Industry Engagement:

The research is currently being written up and promoted by **Mosaic – Geospatial Imaging Leaders (Prague, Czech Republic)**, highlighting its relevance for 360° street-level surveying, urban mapping, and disaster intelligence applications.