Generative AI for Geospatial Challenge

Project Lead

From Data to Insight: Predicting Tomorrow's Landscapes

Clark University's Center for Geospatial Analytics (CGA), in collaboration with Washington University in St. Louis, is developing a generative AI model to forecast future land-use and land cover (LULC) changes in the U.S., starting with data from the National Land Cover Database (NLCD) and expanding to integrate satellite imagery, elevation, temperature, and population trends.

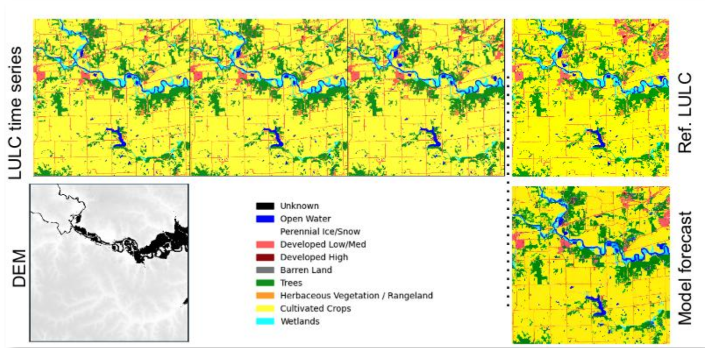

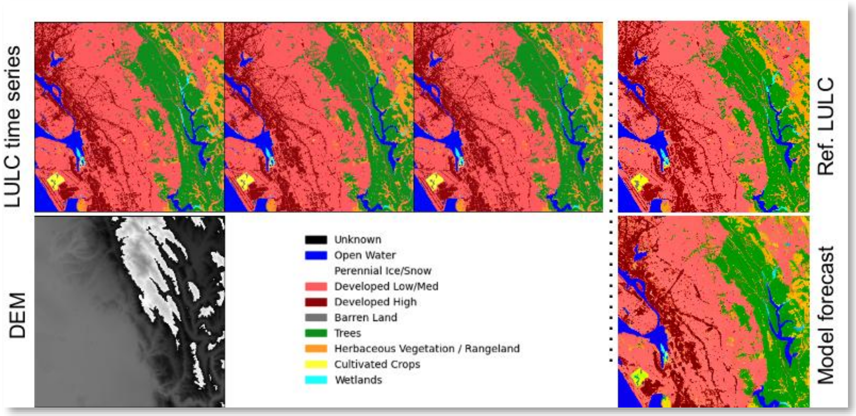

Applications

- **Urban Planning:** Anticipates expansion trends to inform infrastructure development
- **Environmental Management**: Identifies potential deforestation or habitat loss hotspots.
- **Policy & Conservation:** Supports sustainable land-use strategies and climate adaptation efforts.

Transforming LULC Prediction

Leverages diffusion-based generative AI to model landscape evolution based on environmental and societal trends.

Informing Climate Resilience

Supports proactive land-use decisions across urban planning, conservation, and policy domains.

Key Features

- **Diffusion-Based AI Modeling:** Adapts next-gen AI to produce future landscape maps grounded in real-world data.
- **Context-Aware Forecasting:** Moves beyond static, pixel-based models by training AI to recognize landscape trends.
- **Scalable, Cross-Sector Solutions:** Designed for broad application across conservation, urban development, and policy planning.

Team:
Clark University, Center for Geospatial Analytics (Hamed Alemohammad, Sam Khallaghi, Denys Godwin); **Washington University in St. Louis** (Nathan Jacobs, Srikumar Sastry, Dan Cher)

Dataset Preparation

The Clark University team constructed a time-series dataset based on the NLCD product for the years 2001, 2006, 2011, 2016, and 2021, covering the contiguous United States. The dataset is organized as 512×512 image chips, with the first four time points used for fine-tuning a diffusion model and the final time point reserved for evaluation. To enrich this dataset, the team integrated corresponding Landsat scenes, Copernicus GLO-30 digital elevation data, and LandScan population estimates. Both the dataset and the code for generating it will be publicly released with the completion of the project

Model Development and Experiments

ControlNet architecture was adapted to handle time-series inputs and ran initial experiments. The findings showed that generating RGB outputs with a diffusion model was suboptimal for categorical land cover class prediction, even after discretization. To address this, they:

- Refined model hyperparameters,
- Revised the land cover taxonomy to merge original land cover to 9 classes and a defined a tailored color palette for NLCD data,
- Developed improved conditioning prompts incorporating domain-specific knowledge.

The team also upgraded their text prompt encoder from a standard CLIP model to LongCLIP, which expanded their token capacity from 77 to 234, allowing for richer and more descriptive conditioning.

Interface Development

A prototype interface has been developed to visualize the input time series and predicted future land cover maps. It includes interactive sliders for drivers such as population, temperature, and precipitation. However, these sliders are not yet functional, pending the integration of driver-based conditioning into the model.

Next Steps

Moving forward, the team plans to:

- Improve the geometric accuracy of the diffusion model's predictions
- Integrate the driver variables into the conditioning pipeline, and,
- Finalize the interface by enabling full interactivity with the driver inputs.