Fusing LEO and GEO observations for agricultural monitoring

[1] **Riccardo D'Ercole**, [2] David Seu, [1] Chiara Maria Cocchiara

[1] **European Space Agency Φ-lab, Frascati, Italy; [2] Co2 Angels, Cluj-Napoca, Romania**

esa Φ-lab

The goal of the work is to leverage geostationary (GEO) and low-orbiting satellite (LEO) imagery to generate higher-resolution NDVI series using historical information and physical conditioning. The model performs super-resolution and temporal prediction using diffusion models conditioned on environmental variables and past observations.

Target applications include vegetation monitoring, drought prediction, agricultural monitoring. The system aims to reconstruct fine spatial vegetation structure from coarser satellite data and contextual information.

Previous work

Mostly due to their higher spatial resolution, researchers and operators usually rely on vegetation products derived from AVHRR, SPOT/VEGETATION, VIIRS, MODIS, Landsat, Sentinel-2 but cloud contamination remains a major factor limiting vegetation dynamics monitoring from polar-orbiting satellites (Fensholt et al., 2011)[1];

Current systems do not take advantage of geostationary products (Fensholt et al. 2011)[1], which have interesting potential for mesoscale agricultural applications. Fensholt shows that in a 2008 East Africa drought (March 1-June 30) 5-days NDVI composites from AVHRR, whereas MODIS 16-days composites almost 100% in some areas. 5-days SEVIRI deliver much less cloud covered pixels (1-10% maximum)

In previous work we have demonstrated that generative models (e.g. diffusion) can be used for robust vegetation status forecasting;

Additionally, the proposed algorithm (D'Ercole et al. 2024) [2] to derive a vegetation status dataset from a geostationary satellite imagery shows a good ability to reduce the effect of clouds

The trained diffusion model is more robust to extreme drought scenarios (e.g. Horn of Africa extreme drought event in 2021-2023) compared to other spatio-temporal models (e.g. Graph Neural Network, ConvLSTM)

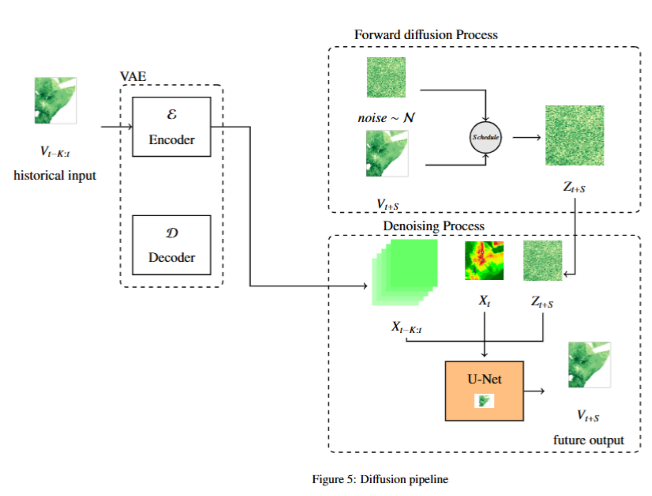

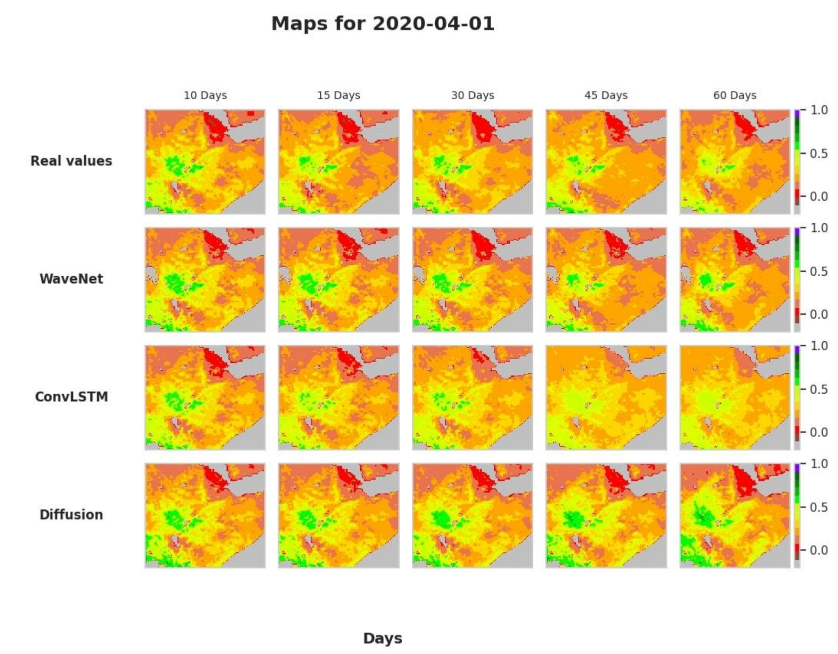

Data collection & pre-processing

Geostationary observations were obtained from the Meteosat Third Generation (MTG) Flexible Combined Imager (FCI), whereas the MOD09GQ product was used as a target.

To reduce cloud contamination and capture peak surface reflectance conditions, five acquisitions per day were selected over the 08:00–13:00 UTC window.

A daily maximum compositing (MVC) strategy was applied, whereby the maximum reflectance value across the selected time window was retained for each pixel. This approach increases the likelihood of selecting cloud-free observations and enhances vegetation signal quality.

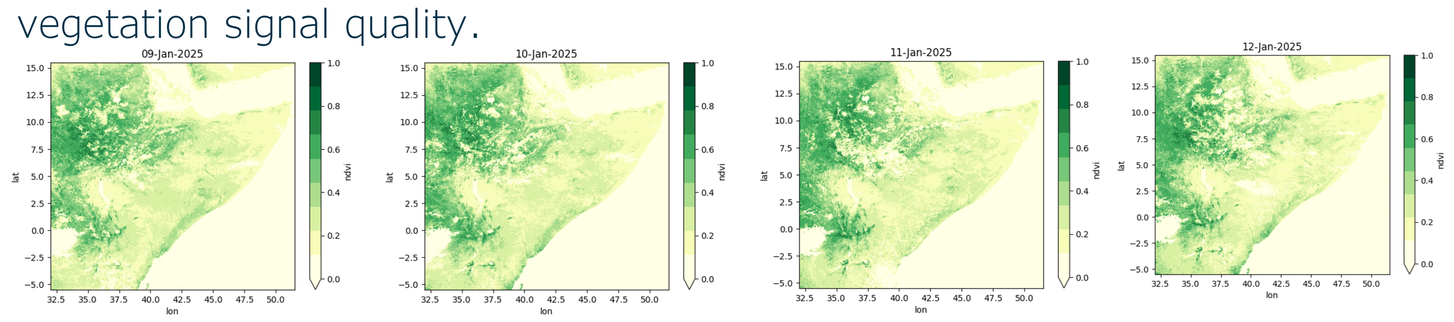

MODIS monthly vegetation products were used to provide a coarse but temporally stable seasonal prior, complementing the high-temporal-resolution MTG observations. All datasets were harmonized onto the MajorTOM grid, to ensure spatial consistency across sensors.

Additional static and contextual layers were incorporated, including land cover classification, latitude and longitude coordinates, and day-of-year (DOY), enabling the model to account for geographic and seasonal variability in vegetation patterns.

Cloud screening

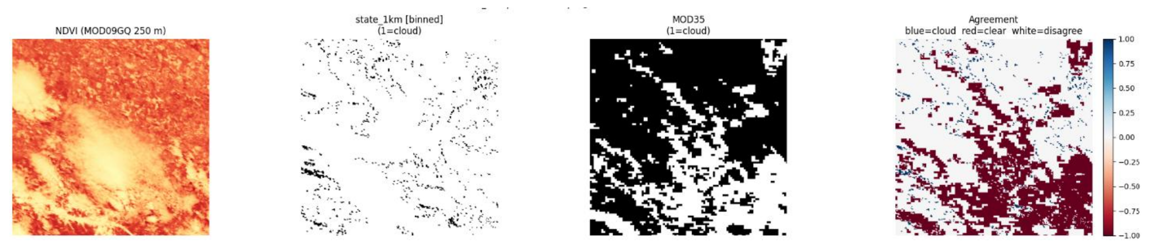

MOD09QA and MOD35 are the two standard MODIS cloud masking products, but both exhibit well-documented limitations, including sensitivity to optically thin features, land cover biases, and processing artifacts [3].

We combine them into a cloud confidence score through spatial dilation, morphological closing, and a blue band veto. Dilation buffers cloud edges to absorb misregistration between the two products, closing fills interior disagreement holes, and anomalously high MODIS Band 3 reflectance relative to a clear-sky climatology flags clouds where the two masks disagree. The final soft score weights MOD35 confidence, the MOD09QA flag, and the blue band anomaly

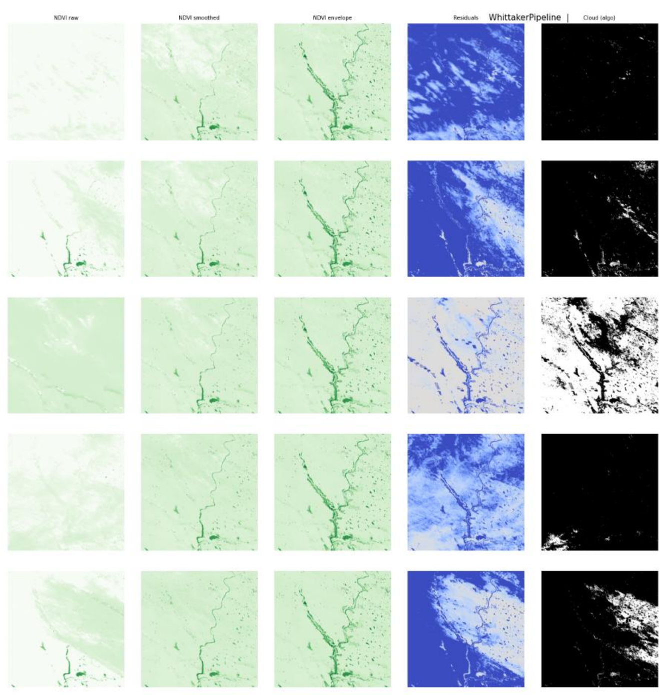

We fit an asymmetric Whittaker smoother [4] to the daily NDVI time series per pixel, using MOD35 confident-clear flags as initial weights and iteratively down-weighting negative residuals. A pixel is flagged as cloudy if its residual falls more than k standard deviations below the local noise floor

The soft cloud confidence score enters the diffusion pipeline in two roles: as a pixel-wise loss weight, and as an explicit conditioning channel alongside MTG FCI and ERA5. Tiles exceeding 30% cloud fraction are discarded, while those below this threshold are retained with soft weighting, reframing residual cloud contamination as structured uncertainty rather than a data quality failure.

Modeling Framework

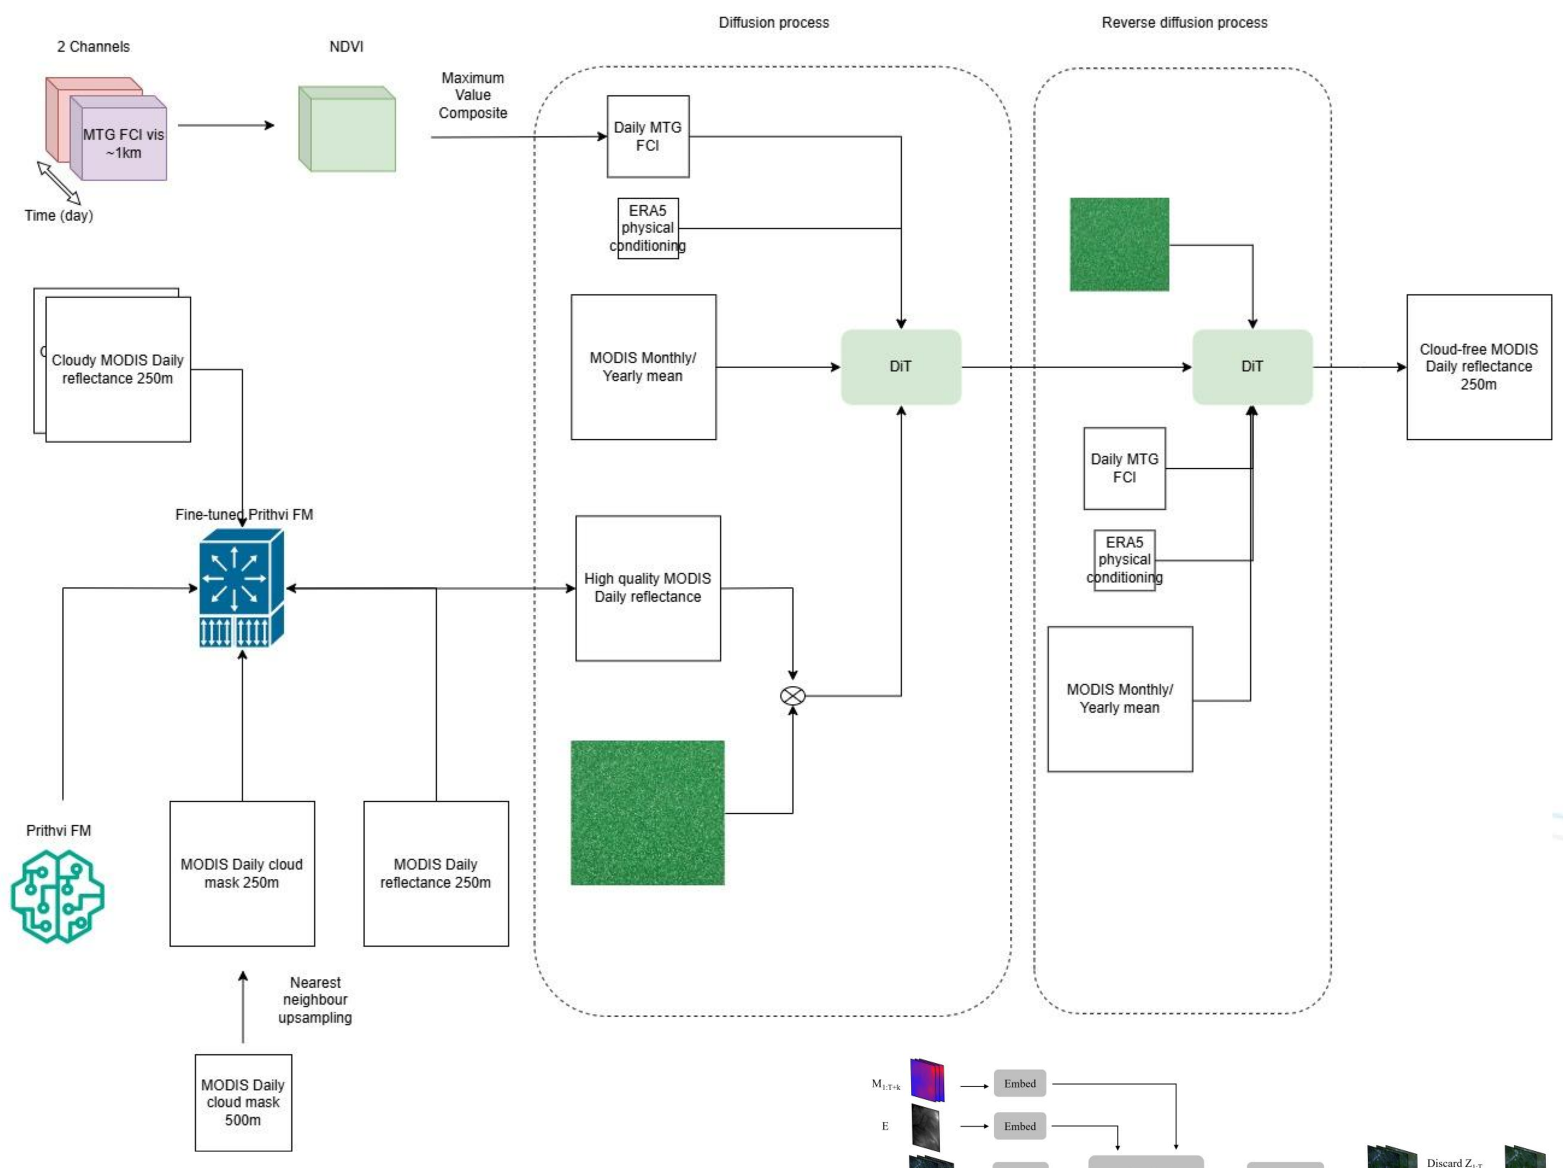

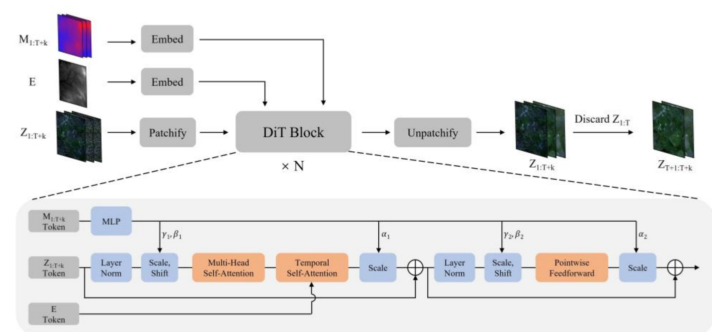

VegDiff DiT Architecture [5]

Future steps

- Finalize the data pipeline and train the diffusion model on MODGQ Band 1 and Band 2 with MTG FCI and ERA5 as conditioning
- Validate the output with local parcel data from Romania and Brazil
- Implement a second super-resolution stage using Sentinel-2 imagery to push reconstruction from 250m to 10m resolution
- Extend the reconstruction targets to additional MODIS bands beyond Red and NIR, enabling a broader set of vegetation and surface indices
- Benchmark the full pipeline against established gap-filling baselines and evaluate temporal and spatial consistency of the reconstructed time series
- Explore extension to additional biophysical variables beyond NDVI such as LAI, FAPAR

Open-source code

Eumetsearch
A wrapper around EUMETSAT eumdac, to easily download, reproject and export to zarr format MTG FCI imagery.

Modis-majortom
Package for downloading, processing, and archiving MODIS satellite data with cloud masking capabilities and MajorTom grid integration.

[1] Fensholt, R., Anyamba, A., Huber, S., Proud, S. R., Tucker, C. J., Small, J., ... & Shisanya, C. (2011). Analysing the advantages of high temporal resolution geostationary MSG SEVIRI data compared to Polar Operational Environmental Satellite data for land surface monitoring in Africa. International Journal of Applied Earth Observation and Geoinformation, 13(5), 721-729.
[2] D'Ercole, R., Casella, D., Panegrossi, G., & Sanò, P. (2024). A high temporal resolution NDVI time series to monitor drought events in the Horn of Africa. International Journal of Applied Earth Observation and Geoinformation, 135, 104264.
[3] Zhang, S., Ma, Y., Chen, F., Shang, E., Yao, W., Qiu, Y., & Liu, J. (2021). Global land high-resolution cloud climatology based on an improved MOD09 cloud mask. Remote Sensing, 13(19), 3997.
[4] Eilers, P. H., Pesendorfer, V., & Bonifacio, R. (2017, June). Automatic smoothing of remote sensing data. In 2017 9th International Workshop on the Analysis of Multitemporal Remote Sensing Images (MultiTemp) (pp. 1-3). IEEE.
[5] Zhao, S., Chen, H., Zhang, X., Xiao, P., & Bai, L. (2025). VegeDiff: Latent diffusion model for geospatial vegetation forecasting. IEEE Transactions on Geoscience and Remote Sensing.