Physics-informed methane plume screening from EnMAP and PRISMA SWIR hyperspectral imagery

Juan Bettinelli*, Dmitry Efremenko
Atmospheric Spectroscopy Department, Earth Observation Center,
Deutsches Zentrum für Luft- und Raumfahrt e. V.
*juan.bettinelli@dlr.de

Central Message

EnMAP and PRISMA SWIR observations contain enough spectral information to detect and map strong localized methane enhancements. Operational plume screening, however, needs methods that are both physically meaningful and computationally scalable. This work compares fast statistical retrievals with BIRRA non-linear least-squares retrievals, evaluates where each method is reliable, and explores machine-learning surrogates as a route toward faster large-scene and archive-scale screening.

Motivation and Questions

Localized methane emitters are important targets for greenhouse-gas monitoring because strong sources can contribute substantially to regional emissions and may be mitigated if detected reliably. Imaging spectrometers such as EnMAP and PRISMA observe spatially resolved SWIR spectra around the 2.3 µm methane absorption region.

- How well do fast retrievals reproduce BIRRA plume morphology?
- Which method is most robust across surfaces, geometry, clouds, snow, water, and sensor artifacts?
- Can machine-learning surrogates improve high-ΔXCH_4 retrieval where statistical detectors become limited?
- Can machine learning provide a fast BIRRA-like surrogate?

Methods Compared

BIRRA non-linear least squares. Physically motivated reference retrieval that fits the measured SWIR spectrum with radiative-transfer-based methane absorption and nuisance parameters. It is interpretable, but slower than screening methods.

Linear and generalized least squares. Fit a methane target signature plus background terms to each pixel spectrum. GLS uses the background covariance to down-weight noisy or correlated spectral structures and is the robust statistical default.

Matched filter and SVD-based retrievals. Matched filters score the similarity between a spectrum and a methane target while suppressing background variability. SVD removes dominant background modes before fitting, which can stabilize results but may also remove plume signal.

Machine-learning surrogate. PyTorch MLP regressor trained on BIRRA-derived XCH_4 labels from selected pixels. Inputs are either method-specific scores with scene features or preprocessed SWIR channels 202–216; tail-weighted training emphasizes high-ΔXCH_4 plume pixels where statistical scores can saturate.

Data and Retrieval Workflow

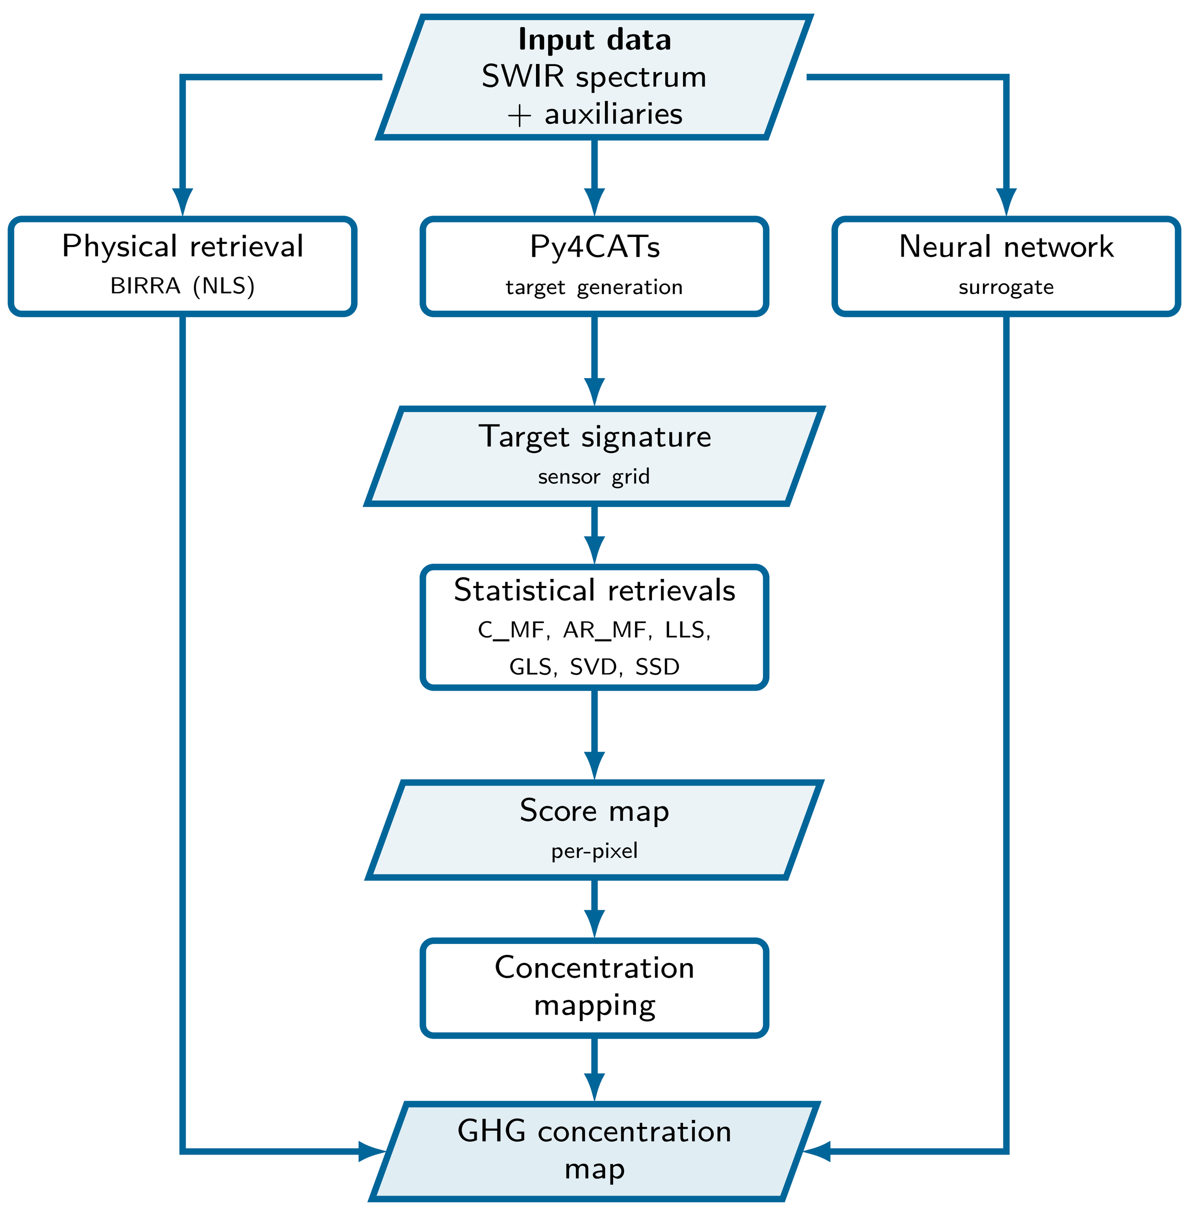

Evaluation Strategy

- **Pixelwise correlation** tests amplitude consistency with BIRRA.
- **IoU and Dice** evaluate plume-mask overlap.
- **FSS** measures scale-dependent plume morphology agreement.

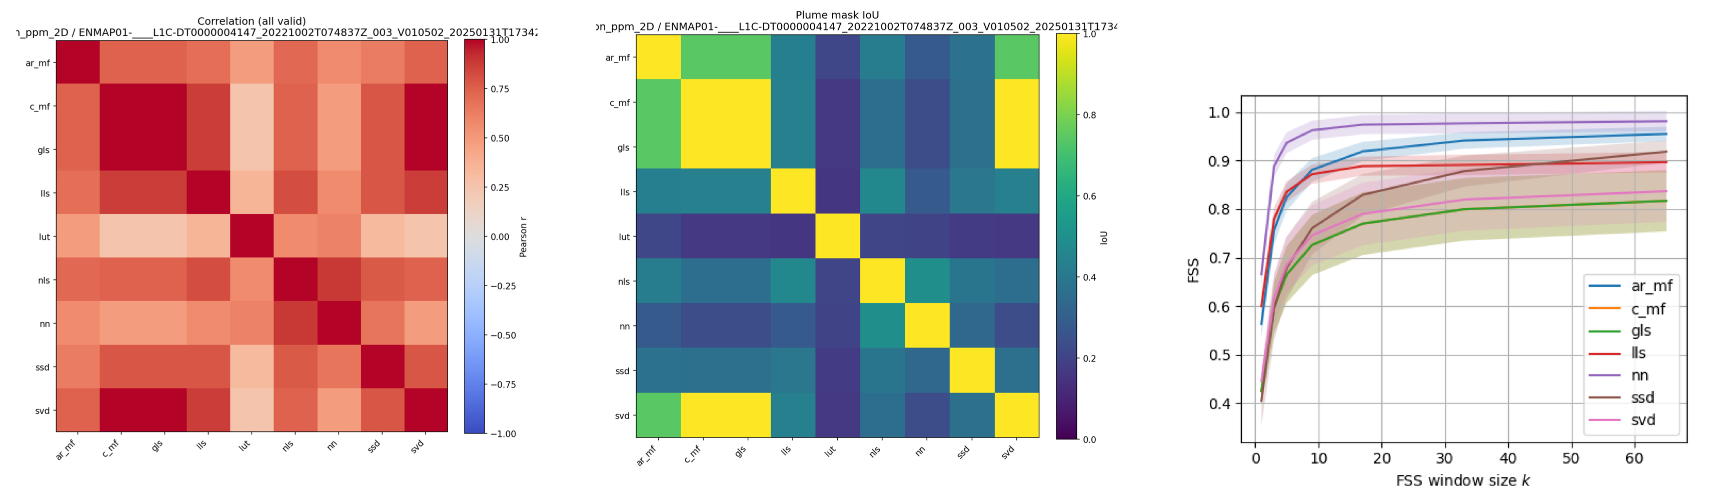

Correlation IoU FSS

Main Findings

- **Fast retrievals find the main plume core.** They are well suited for rapid scene-wide screening and hotspot selection.
- **Morphology differs.** Matched-filter variants are sensitive but noisy; LLS/SSD are more compact; SVD can reduce background structure but may fragment plumes.
- **Quantification is limited.** Linear correction reduces bias, but high-ΔXCH_4 pixels retain scatter because statistical scores cannot fully represent non-linear radiative transfer.
- **The NN gives the best dense approximation to BIRRA.** Tail-weighted BIRRA training improves plume morphology while keeping inference fast.

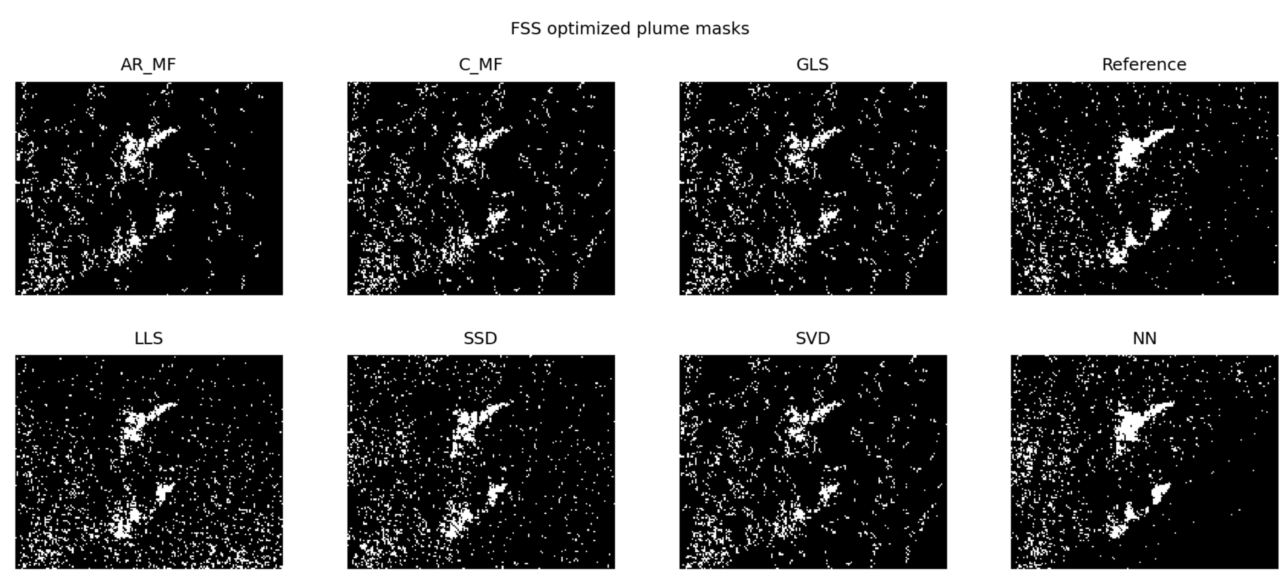

Take-Home and Next Steps

- Statistical retrievals for full-scene plume screening.
- BIRRA for validation, calibration, and ML labels.
- Work in progress: expand training capacity with physics-based Gaussian plume augmentation; use U-Net denoising for retrieval-independent background suppression; train U-Net plume/emitter detection for fast scene scanning and retrieval targeting.