

Self-supervised background priors and the inductive erasure of the rare class: lessons from a low-surface-brightness detection pipeline (or: the model is trying to erase what I'm looking for).

Monday, 15 June 2026 10:00 (60)

We present an in-progress open-world detection pipeline for low-surface-brightness sources in multi-band astronomical imaging (DES DR1, 5 bands, 222 tiles). The architecture composes three stages: a Masked Image Modeling (MIM) self-supervised background prior; a multi-scale matched-filter detector on the MIM residual, with PSF-convolved Sersic kernels per band and combined-band SNR; and a hybrid classification head built from a foundation-model embedding, a Mamba sequence block, and engineered shape and color statistics. The pipeline reaches 95–99% recall on existing LSB catalogs and produces a working pool of novel candidates. Several ML-interesting design tensions emerged along the way — trade-offs between background-model expressivity and rare-class preservation, multi-channel feature integration, and calibration under partial ground truth — which we present in the hope of initiating discussion.

Presenter(s) : NAWROCKI, Krzysztof (NCBJ)