Class-Conditioned Consensus Loss for Multi-Expert Annotation of Low Surface Brightness structures

Renaud Vancoellié[1], Adeline Paiement[1], and Pierre-Alain Duc[2]

[1]Laboratoire d'Informatique et des Systèmes, Université de Toulon, France
[2]Observatoire d'astrophysique, Université de Strasbourg, France

April 10, 2026

Abstract

Tidal structures produced by gravitational stripping during galaxy interactions are among the most direct observational tracers of hierarchical mass assembly. Their detection is intrinsically difficult: surface brightnesses typically fall below $\mu_r \sim 26\text{–}30\,\mathrm{mag\,arcsec^{-2}}$, morphologies are heterogeneous and multi-scale. Hand annotation is still the most common way datasets of such structures are created, but at the data volumes of forthcoming surveys such as Euclid or LSST, manual expert inspection is entirely intractable, motivating the need for robust automated pipelines.

We propose a class-conditioned consensus loss that treats annotation disagreement. The expert annotations are aggregated into a per-pixel consensus score $y \in [0, 1]$, encoding the fraction of annotators in agreement. A piecewise loss function built on Focal Loss then (i) amplifies the gradient contribution of high-consensus pixels (unanimous or near-unanimous agreement), (ii) applies standard supervision on majority-agreement regions, and (iii) leaves entirely unsupervised boundary-uncertain pixels. The decision thresholds are optimized on a held-out validation set, allowing the loss to adapt to the different annotation difficulty of each class.

The consensus loss is embedded in a modified ResNet-50 backbone whose first convolutional layer is replaced by a dual-channel, three-scale stem: the same image is processed simultaneously at its original scale and three other arcsinh scalings learned by the network during training. This provides sensitivity to both compact high-surface-brightness features and large-scale diffuse emission within a single forward pass. Images are drawn from the MATLAS survey. To address class imbalance and annotation scarcity, the training set is balanced through class-aware oversampling.