

Figure S1: Study region, training sample distribution, and landscape heterogeneity in Tanzania and Rwanda. (a) Location of Tanzania and Rwanda within Africa, marking the national extent used for training. (b) Distribution of training samples derived from the 100 m above-ground biomass map by Xu et al. (2021) and PlanetScope images at 3 m resolution. Each red dot represents the centroid of a 100 m \times 100 m tile associated with a scalar above-ground biomass label, enabling high-resolution learning without dense per-pixel supervision. The sample captures diverse ecological gradients across Tanzania and Rwanda, supporting generalization across climate zones, land use types, and vegetation structures. (c) Land cover classification from the ESA CCI Land Cover product (ESACCI-LC-L4-LC10-map-20m-P1Y-2016-v1.0), illustrating the structural heterogeneity of both forest and non-forest ecosystems. Map data: © OpenStreetMap contributors 2024, Distributed under the Open Data Commons Open Database License (ODbL) v1.0; land cover data: © ESA Climate Change Initiative, Land Cover project 2017.




Figure S2: Spatial interpretability analysis of the trained weakly supervised framework for two randomly selected study sites. Column (a) shows composites from Google Earth serving as land cover reference. (b) displays the corresponding PlanetScope imagery, which serves as model input. Column (c) presents the predicted above-ground biomass density maps, where the yellow-to-blue colour gradient represents high-to-low above-ground biomass, with spatial patterns consistent with canopy structural variation noticeable in the reference imagery. Column (d) shows vanilla gradient saliency maps, a first-order gradient-based attribution method that identifies the input pixels to which the predicted above-ground biomass is most sensitive, thereby providing a spatially explicit diagnostic of the network’s learned feature importance. Across both sites, the highest attribution values localise over tree crowns, canopy gaps, and woody vegetation boundaries, ecologically meaningful structures associated with above-ground biomass variability, rather than over radiometrically uniform surfaces. This spatially consistent attribution pattern provides evidence that the weak supervision framework successfully guided the architecture toward learning meaningful structures from PlanetScope imagery, without requiring high-resolution training labels. The upper row shows a forested site with localised canopy disturbance and the lower row a landscape characterised by a sharp forest–non-forest boundary. Map data: © Google Earth 2024; images: © CNES/Airbus, Maxar Technologies, Airbus.