Sub-pixel mapping of disturbance and tree mortality dynamics from Sentinel-2 time series around the globe

Clemens Mosig[1], Teja Kattenborn[2], David Montero[1, 3], Janusch Vanja-Jehle[2], John Brandt[4], Nathan Jacobs[5], Subash Khanal[5], Eric Xing[5], Martin Schwartz[6], Helene C. Muller-Landau[7], Mirela Beloiu[8], Aurora Bozzini[9], Yan Cheng[10], Keenan Ganz[11], Björn Grüning[12], Henrik Hartmann[13, 14, 15], Jan Hempel[1], Stéphanie Horion[10], Samuli Junttila[16, 17], Kirill Korznikov[18], Guido Kraemer[1], Milena Mönks[1], Davide Nardi[9], Paul Neumeier[1], Jonathan Schmid[2], Salim Soltani[2], Marie Therese-Schmehl[19], Josh Veitch-Michaelis[20], and Miguel D. Mahecha[1, 3, 21, 22]

[1]Institute for Earth System Science and Remote Sensing, Leipzig University, Germany [2]Chair of Sensor-based Geoinformatics (geosense), University of Freiburg, Germany [3]Helmholtz Centre for Environmental Research (UFZ), Leipzig, Germany [4]World Resources Institute, 10 G St NE #800, Washington, DC 20002, USA [5]Washington University in St. Louis, United States [6]Laboratoire des Sciences du Climat et de l'Environnement, LSCE/IPSL, CEA-CNRS-UVSQ, Université Paris-Saclay, France [7]Smithsonian Tropical Research Institute, Panama [8]Forest Resources Management, Institute of Terrestrial Ecosystems, Department of Environmental Systems Science, ETH Zurich, Switzerland [9]Department of Agronomy, Food, Natural Resources, Animals and Environment (DAFNAE), University of Padova, Italy [10]Department of Geosciences and Natural Resource Management, University of Copenhagen, Denmark [11]School of Environmental and Forest Sciences, University of Washington, Seattle, United States [12]Bioinformatics Group, Department of Computer Science, Albert-Ludwigs-University Freiburg, Germany [13]Institute for Forest Protection, Julius Kühn Institute (JKI) - Federal Research Centre for Cultivated Plants, Germany [14]Faculty of Forest sciences and Forest Ecology, Georg-August-University Göttingen, Germany [15]Department of Biogeochemical Processes, Max Planck Institute for Biogeochemistry, Germany [16]Department of Forest Sciences, University of Helsinki, Finland [17]Department of Remote Sensing and Photogrammetry, Finnish Geospatial Institute FGI, The National Land Survey of Finland, Finland [18]Department of functional ecology, Institute of Botany of the Czech Academy of Sciences, Třeboň, Czech Republic [19]Department of Hydrology and Climatology, Potsdam University, Germany [20]University of Florida, United States [21]German Centre for Integrative Biodiversity Research (iDiv), Leipzig, Germany [22]Center for Scalable Data Analytics and Artificial Intelligence (ScaDS.AI), Leipzig, Germany

This is a preprint manuscript submitted to EarthArXiv and has not been peer reviewed.

Abstract

Elevated forest disturbances and excess tree mortality are increasingly reported worldwide. Yet existing assessments are either based on patchy terrestrial observations or on large-scale satellite products, which are limited in resolution to pixel-level, binary tree loss detection. This leaves a blind spot on fine-scale disturbances where only a few trees are declining in an otherwise intact canopy. Here, we present a methodology for annually mapping sub-pixel fractional cover of standing deadwood and trees from rolling four-year windows of Sentinel-2 time series. Fractional cover is the proportion of each pixel covered by dead or live tree crowns. To obtain globally distributed sub-pixel reference labels, we leveraged the crowd-sourced archive *deadtrees.earth* of centimeter-scale drone orthophotos with two globally calibrated semantic segmentation models to derive tree and standing deadwood masks, yielding 6.2 million labeled Sentinel-2 pixels. Spatial block cross-validation shows high performance for tree cover, Pearson's r = 0.58–0.64, across biomes, and moderate performance for standing deadwood cover, Pearson's r = 0.30–0.56. Viewed as a binary classification task, the model achieves a mean precision of 91% and recall of 90% for tree cover, and a precision of 54% and recall of 82% for standing deadwood. Validation against independently obtained and ground-validated data in all major biomes reproduced known disturbance patterns and timing at unprecedented spatial detail compared to state-of-the-art products. Our method provides the critically needed missing link between fine-scale ground observations and low-resolution remote sensing products, allowing more realistic estimates of global trends in forest disturbance and tree mortality.

1 Introduction

In recent years, elevated forest disturbances and excessive tree mortality have been reported in many regions around the world (Allen et al. 2015; Hartmann et al. 2022, 2025; Senf et al. 2018; Trumbore et al. 2015). These events are the result of combined impacts due to climate change, pests, pathogens, and past forest management (International Tree Mortality Network et al. 2025). In order to better understand current and future forest dynamics, we need to advance our understanding of past mortality patterns (Migliavacca et al. 2025b). A key prerequisite for such progress is the development of spatially and temporally coherent data on tree mortality events (Hartmann et al. 2022; International Tree Mortality Network et al. 2025; Yan et al. 2024).

Field-based monitoring is the status quo for assessing global tree mortality trends, because it provides the highest accuracy at the level of individual trees. However, current monitoring networks differ strongly among and within countries in sampling design, plot size, revisit intervals (often on the order of 5–10 years), and field protocols, which limits comparability and makes it difficult to resolve year-to-year dynamics and rapid disturbance events (Hammond et al. 2022; International Tree Mortality Network et al. 2025). The access to field-inventory data is also often restricted, which limits data sharing and hampers harmonized large-scale synthesis (Hlásny et al. 2025; International Tree Mortality Network et al. 2025). As a result, the most recent global tree mortality assessment relied on synthesis efforts that compile restricted sets of field observations (Hammond et al. 2022), that are high quality but spatially severely constrained to only about 675 locations and 1303 plots globally.

Remote sensing has proven increasingly promising over the past decades for monitoring forest disturbances and tree mortality by providing observations at spatial scales that cannot be achieved with field plots (International Tree Mortality Network et al. 2025). High-resolution imagery from airplanes or drones enables direct identification of forest disturbances (Cheng et al. 2024; Schiefer et al. 2023). However, airborne and drone acquisitions are typically episodic and geographically patchy, and therefore do not yet provide temporally continuous, standardized observations at global scale. Rather than forming a globally systematic time series, airborne and drone data are currently most useful for local or country-scale tree mortality assessments and validation data for coarser satellite-based products (Cheng et al. 2024; Junttila et al. 2024; Khatri-Chhetri et al. 2024; Schiefer et al. 2024; Schwarz et al. 2024).

Satellite-based remote sensing is well suited for global disturbance monitoring because it provides

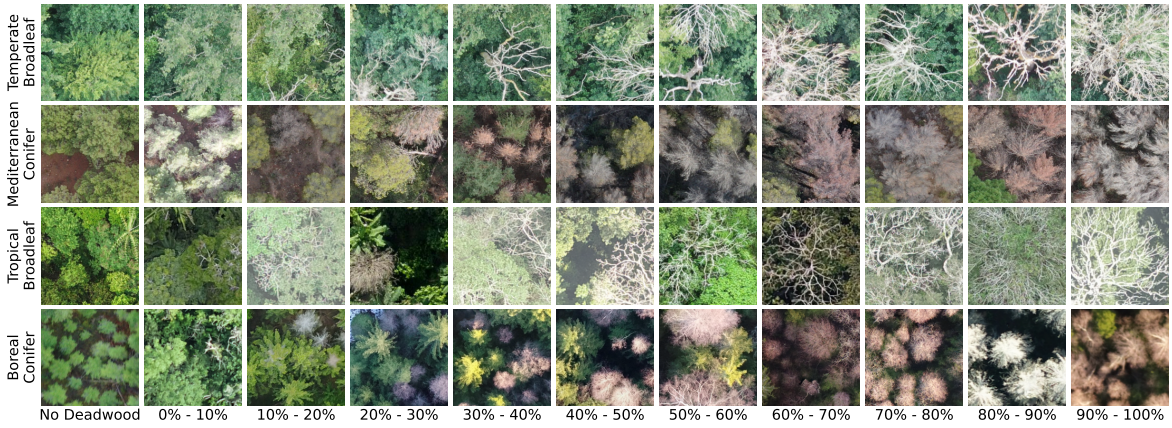

Figure 1: Drone observations of 10 $m \times 10$ m patches, the resolution that Sentinel-2 offers, with different fractions of standing deadwood cover across four major biome groups. Image data from *deadtrees.earth*.

spatially consistent multi-spectral images that enable systematic tracking of forest dynamics over time (Hlásny et al. 2021; International Tree Mortality Network et al. 2025). Repeat observations within seasons and across years provide valuable temporal context that helps to partially compensate for their coarse resolutions (Rußwurm and Körner 2020; Tseng et al. 2023). Several products monitor forest disturbances using Earth observation satellites (e.g., Sentinel-1, Sentinel-2, Landsat) at continental or global scales (M. C. Hansen et al. 2013; A. H. Pickens et al. 2025; Reiche et al. 2021; Senf and Seidl 2021; Song et al. 2024; Viana-Soto and Senf 2025; White et al. 2017). However, most existing products operate at spatial resolutions of ≥ 30 m and represent disturbance as a binary class (disturbed / intact), implicitly assuming within-pixel homogeneity. In reality, canopy condition and mortality are often spatially heterogeneous at sub-pixel scales, and diffuse mortality frequently affects only a fraction of the canopy within a pixel (Alix-Garcia and Millimet 2023; Hartmann et al. 2018; Korznikov et al. 2025; Schiefer et al. 2024; Y. Zhang et al. 2021, and see Figure 1). Consequently, scattered mortality and other small-scale disturbances can be missed, leading to underestimation of mortality and forest loss (Alix-Garcia and Millimet 2023; Cheng et al. 2024; Espírito-Santo et al. 2014; Schwartz et al. 2025).

Further, it is important to distinguish whether trees died standing and remain as standing deadwood (e.g., after drought stress) or whether canopy cover decreased because trees were physically removed or toppled (e.g., through harvesting or wind-throw). These processes reflect different drivers and often different timing, with standing mortality frequently preceding canopy loss. Separating standing deadwood from tree loss is therefore essential for robust disturbance attribution and interpre-

tation of downstream impacts (Schiefer et al. 2023; Winter et al. 2024). Current remote sensing forest disturbance products do not disentangle such dynamics, and this imposes severe constraints on studies on tree mortality and the underlying drivers and effects on forest carbon balances and limits the development of early-warning signals (International Tree Mortality Network et al. 2025; Migliavacca et al. 2025a; Yan et al. 2024).

Explicit maps of standing deadwood over several years, in combination with fractional tree cover, are required to quantify tree mortality rates. Comprehensive real-world mortality rates are critical and missing component in the initialization, calibration, and validation of individual-based forest models as well as for dynamic global vegetation models (Anders et al. 2025; Bugmann et al. 2019; Hartmann et al. 2022; Sitch et al. 2008). This need for data, and especially up-to-date tree mortality maps, is exacerbated by a changing climate and thus previously unseen conditions.

Taken together, developing temporally and spatially continuous global datasets that explicitly quantify standing deadwood alongside tree cover loss is a key frontier in tree mortality research. To fill this gap, we present a methodology to retrieve *yearly fractional tree and standing deadwood cover* at 10 m resolution from Sentinel-2 time series for the years 2017 to 2025. Specifically, we aim to *(i)* estimate yearly tree and standing deadwood cover from multi-year Sentinel-2 time series at sub-pixel resolution, *(ii)* quantify model performance and generalization across biomes, and *(iii)* systematically evaluate the impact of time-series length (look-back window) on model performance. We illustrate the performance of the developed approach through several qualitative case studies in which model output is compared to ground-validated mortality reports.

2 Methods

We developed a deep-learning model that maps Sentinel-2 reflectance time series to annual 10 m fractional tree cover and fractional standing deadwood cover, using labels derived from centimeter-scale aerial orthophotos ($<$ 10 cm). This section is organized into two main parts: data preparation (orthophoto-derived labels, phenological alignment, and Sentinel-2 preprocessing), and model training and evaluation.

Definitions: *Standing deadwood* is defined as woody biomass (twigs, branches, or stems) that has died, but has largely retained its original structure, consistent with Mosig et al. 2026. The definition of *trees* includes any standing tree, dead or alive, and explicitly excludes fallen trees. *Fractional cover*

is defined here as the proportion of a pixel that is covered by dead crowns, and enables capture of sub-pixel dynamics (Figure 1). Data preparation, training, and evaluation were performed at the level of 10 m $\times$ 10 m *Sentinel-2 pixels*.

2.1 Data preparation

The methodological development of globally generalizing deep learning models that map sub-pixel information on standing deadwood cover and tree cover presents three significant challenges: First, a data-driven approach resolving tree and mortality fractions across the diversity of global forests requires ample high-quality reference data. Second, the variability of leaf-off periods in deciduous forest types across ecosystems and biomes introduces uncertainty in distinguishing deadwood from dormancy. And third, it requires availability of globally consistent remote sensing data that are temporally continuous and with a spatial resolution that allows capturing scattered mortality, i.e., at the dimensions of typical tree crowns Fassnacht et al. 2025; Schiefer et al. 2023.

Accordingly, we assembled three core data sources that address each of those challenges. We derived georeferenced labels for standing deadwood and tree cover from drone orthophotos making use of the crowd-sourced database *deadtrees.earth* and automate the approach through deep learning segmentation models (Section 2.1.1). We determined the approximate leaf-on period for each site using remote sensing derived greenness curves (Section 2.1.2). And we acquired and preprocessed Sentinel-2 data time series in the locations where drone data is present (Section 2.1.3), Lastly, we consolidate the dataset into a high quality subset used for later training and evaluation (Section 2.1.4).

2.1.1 Orthophoto segmentation and fractional cover

We obtained masks for standing deadwood for a given orthophoto using the segmentation model by Möhring et al. 2025, which segments standing deadwood from individual branches to entire canopies. This segmentation model was trained on image resolutions ranging from 1 cm to 28 cm; however, Möhring et al. (2025) heuristically determined performance optimum at 5 cm across a subset of the *deadtrees.earth* database (F1-score of 51% to 71% across biomes). Therefore, we downsampled imagery with a resolution finer than 5 *cm* to 5 *cm* before applying the segmentation model.

Tree cover masks for each orthophoto were obtained using the segmentation model by Veitch-Michaelis et al. 2024, which segments tree canopy cover, including both dead and live trees. The segmentation model achieved a mean F1-score of 91.4% in cross-validation on a subset of the OpenAerialMap

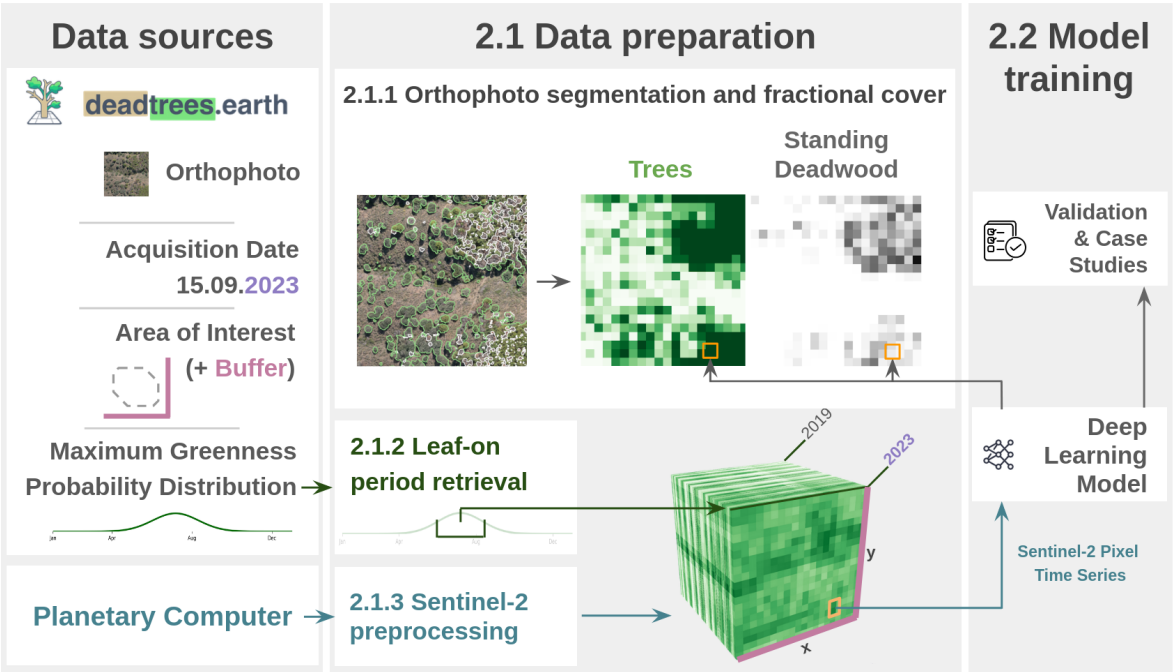

Figure 2: Data preparation flowchart for one site. Orthophotos, together with acquisition date, area of interest, and maximum greenness probability distribution, are extracted from *deadtrees.earth* (Mosig et al. 2026). Forest and standing deadwood segmentation models are applied to orthophotos, quality-checked, and then aggregated into fractional cover rasters (Möhring et al. 2025; Veitch-Michaelis et al. 2024). A season midpoint day-of-year is extracted from the maximum greenness probability distribution. The Sentinel-2 data is downloaded and preprocessed starting at the nearest season midpoint and extending 4-years into the past. Sentinel-2 cube visualization (bottom right) was created with *Lexcube* (Söchting et al. 2025).

drone database, which is partly integrated into *deadtrees.earth* (Mosig et al. 2026; OpenAerialMap Community 2024). This model was trained with a fixed spatial resolution of 10 cm in the coordinate reference system (CRS) EPSG:3095. Therefore, all orthophotos were reprojected and resampled (bilinear interpolation) to match this CRS and spatial resolution before model inference. In early tests, the threshold choice for binarization (86%) showed a trade-off between either undersegmenting individual trees in sparse forests or missing canopy gaps in dense forests. Furthermore, the preliminary tests showed that this model often failed to detect standing deadwood as trees, thus we set areas as tree cover that were already segmented as standing deadwood.

For each Sentinel-2 pixel within the area of interest, we computed *tree cover* and *standing deadwood cover* based on the centimeter-scale binary segmentation masks obtained from each orthophoto. All of the above-described analysis was restricted to the area of interest. The area of interest is the high-quality area for each orthophoto in *deadtrees.earth*, excluding regions affected by reconstruction artifacts, missing data, or poor lighting, while retaining non-forested areas (Mosig et al. 2026).

2.1.2 Leaf-on period retrieval

Leafless but live deciduous trees cannot be robustly distinguished from standing deadwood in aerial imagery (drone, airborne, or satellite). Thus, classification of standing deadwood from such imagery imagery requires first determining the leaf-on period, the time of year during which leafless trees are most likely dead. This leaf-on period varies geographically.

Maximum greenness periods, as derived from remote sensing data, approximately corresponds to the leaf-on season (Körner et al. 2023). The VIIRS Global Land Surface Phenology Product (X. Zhang et al. 2020) reports a maximum greenness period based on the enhanced vegetation index (EVI) (Huete et al. 2002) at 500 m resolution between 2013 and 2022. *deadtrees.earth* derives based on this product for each site a probability for each day-of-year (DOY) to be within the reported maximum greenness period aggregating to 40 km resolution and across years, while only considering forested pixels (Mosig et al. 2026). We employed this probability distribution $\mathbf{y}$ to identify an estimated leaf-on season for each 40 km pixel, by defining any DOY with a probability above $(\max(\mathbf{y}) - \min(\mathbf{y}))/2$ to be in leaf-on season. If this yields multiple leaf-on seasons, we keep the longest. This estimated leaf-on season is later used to filter drone images (see Section 2.1.4) and its center used during model training and inference (see Section 2.2.1 and Figure S1).

2.1.3 Sentinel-2 preprocessing

For each orthophoto location, we obtained a Sentinel-2 data cube (Mahecha et al. 2020; Montero et al. 2024b), with spatial bounds defined by a 100 m outer buffer around the orthophotos's area of interest in the respective UTM zone (see Figure 2). Each data cube consists of four years of historical Sentinel-2 images starting from the nearest site-specific leaf-on season center (see Section 2.1.2) and has a consistent 7-day temporal resolution. For example, for an acquisition date in Sep 15th 2023, and a respective leaf-on season center in Jun 15th, the data cube extends from Jun 15th 2019 to 2023.

Data cleaning steps for optical time series have been shown to be an important step to improve model generalization capabilities (Rußwurm and Körner 2020). Before cleaning we corrected the Sentinel-2 L2A data for processing baseline differences by subtracting 1,000 from scenes with a baseline version greater than 4.0 (European Space Agency 2021). We then converted digital numbers to top-of-atmosphere reflectance in $[0, 1]$ using a scaling factor of 10,000. To obtain a high quality time series, we retained only clear-sky observations and removed pixels affected by clouds, cloud shadow, and snow. We classified clouds and their shadows using the CloudSEN12 segmentation model (Cesar Aybar et al. 2022), applied to individual patches ($732\ px \times 732\ px$) with all 12 Sentinel-2 bands. Lastly, we identified and masked snow at pixel level using Fmask (v3.2) (Zhu and Woodcock 2012; Zhu et al. 2015).

Reprojection followed data cleaning, with observations outside the local UTM zone reprojected using nearest-neighbor resampling. Specifically, we obtained Sentinel-2 data in a globally tiled format (Military Grid Reference System, MGRS Bauer-Marschallinger and Falkner 2023), where each tile spans 109.8 km and is stored in its local UTM zone. Due to this fixed tiling scheme across multiple coordinate reference systems, a single location can be covered by up to six separate tiles (Bauer-Marschallinger and Falkner 2023). Thus, to obtain a coherent data cube for every site, a subset of tiles was reprojected in overlapping regions.

The final preprocessing step consisted of the temporal composites generation, where we aggregated the spatially aligned Sentinel-2 observations into 7-day intervals by computing the per-pixel, per-band mean. Using fixed 7-day bins yields equally spaced timesteps, enabling more efficient temporal encoding schemes for deep-learning models (see Section 2.2.1). The number of cloud- and snow-free observations contributing to each interval varies with location and data availability; Zhou et al. (2025) report ~ 40 cloud-free Sentinel-2 observations for most regions in 2022, corresponding

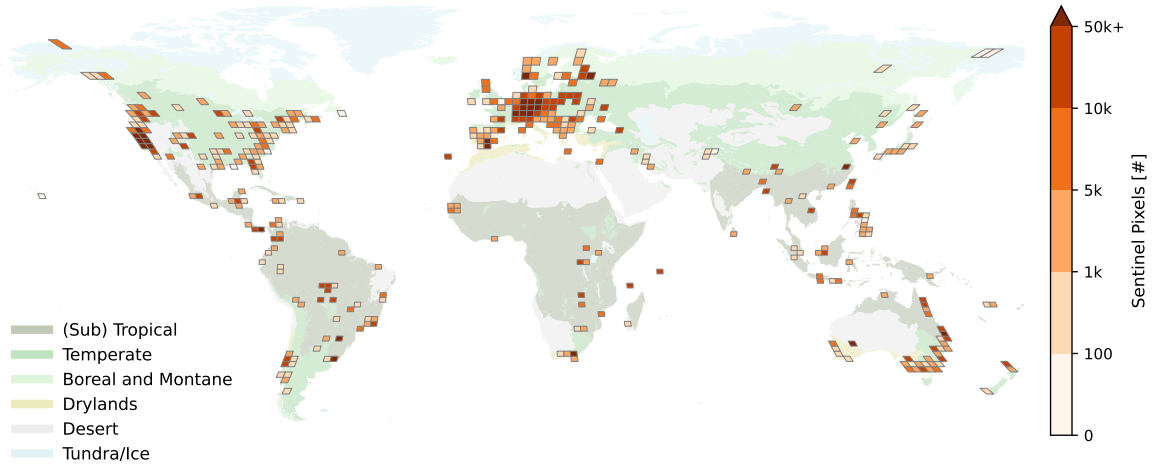

Figure 3: Global distribution of available reference data. The color indicates the number of labeled Sentinel-2 pixels within the respective 200 km bin. The spatial bins are created in EPSG:8857, an equal area projection. Different colors in the background depict different biomes (Olson et al. 2001).

to on average fewer than one observation per 7-day interval. We discarded intervals without valid observations.

The full processing pipeline, called *sentle*, is publicly available at github.com/cmosig/sentle.

2.1.4 Dataset filter and summary

We applied this workflow to all 5252 orthophotos in the *deadtrees.earth* database (Mosig et al. 2026) and then filtered for three criteria. First, we filtered for orthophotos within the temporal scope 2017-2025. Second, during a qualitative audit, we excluded orthophotos where tree or standing deadwood segmentation masks (see Section 2.1.1) were incorrect or missing for approximately more than one quarter of the area of interest. Third, we excluded orthophotos that were outside the above-described local leaf-on season (see Section 2.1.2). For orthophotos with evergreen vegetation only, seasonality does not matter for standing deadwood segmentation, and can still be used as a reliable reference. Therefore, in temperate, boreal, and dryland biomes (Olson et al. 2001), local experts selected additional orthophotos based on the visually estimated genus, acquisition date, and historical Google Earth imagery. This yielded 2641 orthophoto sites and 6.2 million derived Sentinel-2 pixels (62,000-hectare) for training and validation (Figure 3).

Biome separation is based on (Olson et al. 2001), and consists of boreal and montane, dryland, temperate and (sub) tropical forests.

2.2 Model training and evaluation

2.2.1 Model Definition

We trained a per-pixel model that maps a Sentinel-2 reflectance time series (12 bands per 10 m pixel) to tree and standing deadwood fractional cover. The reflectance time series starts at the center of the nearest leaf-on period relative the orthophoto acquisition date, and each timestep represents a 7-day composite (see Section 2.1.3). The covered timespan is set to least 4 weeks and at most 4 years, and depends on data availability; Sentinel-2 is available from mid-2015 such that for predictions in 2017 the time series will only cover 2-years. At a temporal resolution of one aggregated timestep per 7 days, this corresponds to between 4 and 208 timesteps, depending on the number of cloud-free and snow free observations.

The model independently embeds each timestep: a two-layer ReLU feed-forward network with batch normalization (Ioffe and Szegedy 2015; Nair and Hinton 2010) maps the 12-band input to 2048 units and then down to a 128-dimensional embedding. It then adds sinusoidal positional encodings tailored to the temporal sequence (indexed by 7-day intervals) to reflect the temporal order and missing timesteps (Tseng et al. 2023; Yuan et al. 2022). The encoded sequence with a prepended classification token (CLS) is passed into a transformer with two encoder layers, a single self-attention head, embeddings size 128, and a 512-wide feed-forward sub-layer. Variable length sequences, e.g. through data availability and cloud cover, are enabled by a key padding mask. The previously inserted CLS token is then extracted from the transformer output and passed into a linear projection that maps the token (size 128) to the two target fractional cover values.

2.2.2 Loss function and weighting

We employed a distributional histogram loss, HL-Gauss, where regression values are represented as a categorical distribution, which was shown to consistently outperform standard regression loss functions in related work (Farebrother et al. 2024; Imani and White 2018; Imani et al. 2024). We chose 50 bins (2% cover bin width) for the loss formulation, at which performance was shown to converge (Imani and White 2018; Imani et al. 2024). The conversion from a regression target y to a categorical distribution uses a Gaussian distribution as follows:

$$P(b_i|y_p,\sigma_p) = \frac{\Phi\left(\frac{u_i-y_p}{\sigma_p}\right) - \Phi\left(\frac{l_i-y_p}{\sigma_p}\right)}{\Phi\left(\frac{u_{max}-y_p}{\sigma_p}\right) - \Phi\left(\frac{l_{min}-y_p}{\sigma_p}\right)} \tag{1}$$

where p is the pixel identifier, y_p is the target fractional cover, b_i is the i-th bin, l_i and u_i are the lower and upper bounds of bin i, $\Phi(\cdot)$ is the standard normal Cumulative Distribution Function (CDF), l_{min} and u_{max} are the bounds of the extended support (with 10 extra bins), and σ_p is pixel-specific standard deviation. We add support outside the valid value range totaling 70 bins, as it is important to enable the model to predict exactly 0% as a pixel can contain exactly no standing deadwood or no forest, which becomes especially relevant when providing statistics over large areas. Different from Imani and White 2018, σ_p is computed dynamically from the fractional cover of neighboring pixels to represent georeferencing uncertainty (see Section 5.3):

$$\sigma_p = \text{std}(\{y_o : o \in \mathcal{N}(p)\}) \tag{2}$$

where $\mathcal{N}(p)$ is the neighborhood around pixel p containing the set of identifiers of directly adjacent pixels.

Lastly, given the naturally long-tailed distribution of fractional cover we applied the inverse frequency of cover value ranges as weight to the loss. We used coarser 20% bins, different from the above 2% bins, for increased weight stability with small sample sizes. Specifically, the loss weight for each target value y is computed as: $w(y) = \frac{1}{f(g(y))}$, where $g(y) = \lfloor y/0.2 \rfloor$ maps the fractional cover value to its corresponding 20% bin and f computes the batch-wise frequency of a 20% bin. We also apply a multiplier of 10 to the bin representing exactly zero cover as a way to balance precision and recall across different shares of fractional cover (for parameter selection see Section A.4).

The final HL-Gauss loss function is:

$$\mathcal{L}_{HL-Gauss} = \frac{1}{|B|} \sum_{p \in B} w(y_p) \cdot \text{CrossEntropy}(z_p, P(\cdot|y_p, \sigma_p)) \tag{3}$$

where B is the batch and z_p are the predicted logits.

To convert the categorical distribution back to a regression value, we apply a softmax to obtain bin probabilities and then compute the dot product of these probabilities with the bin-center vector; the resulting expectation is clipped to the range 0%–100%.

$$\hat{y} = \text{clip}\left(\sum_{i=1}^{N} \text{softmax}(z_i) \cdot c_i, 0, 1\right) \tag{4}$$

where z_i are the model logits for bin i, $c_i = \frac{l_i+u_i}{2}$ is the center of bin i, and N is the total number of bins (60 bins: 50 main bins plus 10 padding bins).

Because the model predicts a full categorical distribution over ordered bins, we can derive a distributional spread (standard deviation) in the original regression space, which can be used as an uncertainty indicator. We computed standard deviation from the categorical distribution is computed as:

$$\hat{\sigma} = \sqrt{\sum_{i=1}^{N} \text{softmax}(z_i) \cdot (c_i - \hat{y})^2} \tag{5}$$

2.2.3 Hyperparameters

We train with AdamW (Loshchilov and Hutter 2019) at a base learning rate of 5×10^{-6} and weight decay 0.1, applied only to weights while excluding biases, normalization parameters, positional encodings, and the CLS token. The schedule warms up linearly for the first 100 update steps from 1×10^{-7} to 5×10^{-6}, then follows a single-cycle cosine annealing decay over the remaining steps of the 200 training epochs. No early stopping was applied. Batch size is 512. Dropout is 0.2 in both the token embedding stage and the transformer layers, and we additionally apply timestamp dropout with probability 0.1.

2.2.4 Data augmentation

We employ several data augmentation techniques to reflect real-world uncertainties and data variability. For Sentinel-2 time series, we crop to a random temporal context length between four weeks and four years, ensuring stable inference for 2017–2018, where far less than four years of spectral history is available. To mimic cloudier conditions, 10% of timesteps are randomly dropped. Reflectance values are shifted by a random gain factor drawn from a normal distribution ($\mu = 0$, $\sigma = 0.03$), consistent with reported Sentinel-2 calibration results ((ESA) 2023, December 11). The extracted peak-season day of year is shifted by a random number of days sampled from $\mathcal{N}(0, 7\text{d})$, and target fractional cover values are jittered with $\mathcal{N}(0, 0.1)$.

2.2.5 Epoch sampling

The dataset employed here required careful balancing before model training to account for differing geographic and standing deadwood distributions. First, the crowd-sourced drone data are unevenly distributed across the globe (Figure 3), reflecting independent collection decisions by database contributors. Second, there is wide variation in the size of the area that a single orthophoto covers, and thus also in the available labeled Sentinel-2 pixels. For example, the 100 km $\times$ 100 km area with the most data represented 4% of the total data set (248,000 Sentinel-2 pixels). Additionally, the largest orthophoto (122,000 Sentinel-2 pixels) covers three orders of magnitude more area than the median (120 Sentinel-2 pixels).

Because of this, we limited the number of Sentinel-2 pixels that a single site and region can contribute to each training epoch to a fixed Sentinel-2 pixel count determined as follows: We first split the Sentinel-2 pixels into two groups: those with 0% deadwood cover and those with $> 0\%$ deadwood cover, enabling us to reduce the representation of non-deadwood pixels in the training dataset. Then, for both splits, we limited the number of Sentinel-2 pixels sampled per orthophoto to 100 (1 ha). Additionally, we reduced geographic imbalances in the data by sampling a maximum of 500 Sentinel-2 pixels (5 ha) per global grid cell of $10,000\ km^2$. Instead of applying these caps once, we decided to resample during training with these caps from the full dataset in every epoch. Across epochs, samples for $> 0\%$ standing deadwood cover remain largely the same as they are scarce, and samples for $= 0\%$ standing deadwood cover will vary across epochs. This naturally regularizes the model training, while keeping the distribution balanced and geographically representative. The outer-folds for evaluation were subsampled once with this methodology.

2.2.6 Evaluation

To avoid optimistic model performance estimates due to spatial autocorrelation, we employ spatial block cross-validation (Kattenborn et al. 2022; Ploton et al. 2020). We divide the dataset into 50 km spatial blocks in the equal-area coordinate reference system EPSG:8857 and randomly assign blocks to three outer folds. Within each outer fold, 30% of blocks are held out as a validation split during hyperparameter tuning. Results are reported across the concatenated outer-fold test sets, and variability across folds is expressed as either standard deviation or minimum and maximum.

First, as a regression task, we evaluated how well the exact fractional cover value is recovered, by

computing Pearson's correlation coefficient (r) and the root mean square error (RMSE). To reduce the influence of the strongly imbalanced label distribution, we report *bin-resampled* estimates for these metrics: we partitioned the reference cover into 0.1-wide bins over $[0, 1]$ and assigned each sample a weight inversely proportional to the number of samples in its corresponding label bin (computed per cross-validation fold), such that each label bin contributes equally within a fold. Weighted r and weighted RMSE were then computed independently for each fold and summarized as mean and standard deviation across folds. In addition, we report the mean error (ME) as the unweighted mean of $(\hat{y} - y)$ computed over all pixels within each biome, to reflect the average bias under the natural data distribution.

Second, we evaluated the predictions as a binary classification task. Viewing the problem additionally as a binary classification task is critical in the context of scattered mortality (or individual trees), where uncovering the presence/absence of a (dead) tree is much more important than getting the exact fractional cover value (see Section 5.3). Here, we converted labels and predictions into binary form, where one class represented a cover value of exactly 0% and the other class any non-zero cover value. We then computed precision (for the non-zero cover class) for each predicted fractional range of 5% and vice versa for the recall and target fractional cover. We only considered 5% bins with more than 20 observations for result stability. This means that a recall of 50% at a *target* fractional cover value of 20% corresponds to the model predicting a $> 0\%$ cover in 50% of cases where the cover is 20%. Conversely, a precision of 50% at a *predicted* cover value of 20% corresponds to the model being right in 50% cases that any (dead) tree is present ($> 0\%$). Practically, precision can be interpreted as the probability that a location with a given predicted deadwood cover indeed contains at least one (dead) tree. And recall indicates the share of pixels with a certain fractional cover that the model predicts as any standing deadwood.

2.3 Temporal context experiment

To quantify the value of increasing historical reflectance context, we evaluated model performance as a function of the input time-series *timespan*, defined as the duration between the first and last timestamp, irrespective of the number of cloud-free observations within that interval. For each trained model, we passed all samples from the cross-validation splits through the network and recorded predictions for progressively cropped timespans between 16 and 192 weeks in 16-week increments, and for the 4-week timespan as an additional minimum-context setting. To ensure comparability across

timespans, we used only samples with valid predictions for all evaluated durations and in that process aggregated to 16-week bins to maintain sufficient coverage in regions with comparatively few cloud-free observations (see Figure S7). The minimum evaluated timespan of 4 weeks corresponds to the shortest temporal context encountered during training via data augmentation (see Section 2.2.4). Performance was assessed using two complementary metrics: weighted Pearson's correlation coefficient (r) to evaluate recovery of fractional cover in the regression setting, and the F1-score computed from mean precision and recall across 5% fractional-cover bins for the binary classification task, consistent with the general evaluation protocol (see Section 2.2.6).

3 Results

This section reports cross-validation binary classification and fractional cover regression performance across biomes (Section 3.1 and Section 3.2). We further investigate how model performance responds to increasing lengths of historical Sentinel-2 time series, thereby quantifying the value of temporal context for disturbance and mortality mapping (Section 3.3).

3.1 Classification performance

For tree cover detection, we observed a mean precision of 91% and recall of 90% in the binary classification task (zero vs. nonzero) across predicted and target fractions, respectively. In general, high fractional cover predictions were more likely to contain a forest pixel (precision), and higher target cover was more likely to be detected by the model (recall) (Figure 4). Precision was highest in Mediterranean forests, and recall was highest in temperate and boreal forests. In general, both precision and recall converged towards 100% above 50% predicted and target cover values, respectively.

For standing deadwood cover, we observed a mean precision of 54% and a recall of 82% on the binary classification task across all biomes and fractional cover values (Figure 4b). For predicted fractional cover values below 40%, the reference pixel contained standing deadwood in less than 50% of cases. Precision rapidly increased with predicted fractional cover for Mediterranean and temperate forest. For boreal and tropical forests, precision remained below 50% regardless of predicted standing deadwood cover. Recall consistently increased with target deadwood cover in all forest types, indicating that higher values of target deadwood cover are more likely to be detected by the model. For Mediterranean and temperate forests, recall converges to near 100% above 70% standing deadwood

(a) Precision and recall for fractional tree cover.

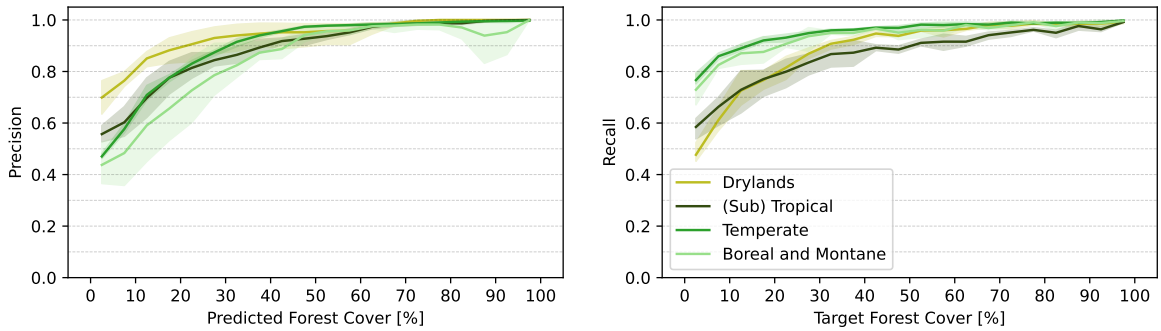

(b) Precision and recall for fractional deadwood cover.

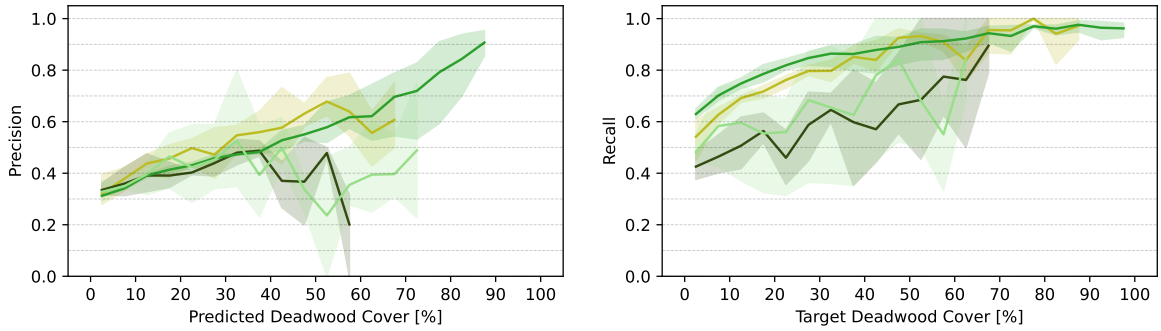

Figure 4: Mean precision and recall in the binary performance task across fractional cover bins (width = 5%). For the precision, the x-axis represents the predicted fractional cover (left). For the recall, the x-axis represents the target fractional cover (right). The displayed range in the y-dimension denotes the minimum and maximum performance across the three cross-validation sets. Only results for 5% bins with more than 20 samples are displayed.

cover, meaning dense mortality locations are almost guaranteed to be mapped by the model in these biomes.

3.2 Sub-pixel fractional cover estimation performance

The following results compare predicted and target fractional cover. For tree cover, regression performance was stable across biomes, $r = 0.58 - 0.64$ and RMSE $= 0.26 - 0.31$, with very small fold-to-fold variability ($\sigma_r = 0.01 - 0.03$ and $\sigma_{\text{RMSE}} = 0.01 - 0.02$), indicating robust generalization for fractional tree cover. Tree cover showed consistent underestimation (ME $= (-0.9) - (-0.13)$) compared the drone derived masks.

In contrast, standing deadwood cover showed a stronger biome dependence. The highest correlation was obtained in temperate forests ($r = 0.56$, RMSE $= 0.32$). RMSE fold-to-fold variation

Table 1: Mean and standard deviation (across cross-validation folds) of bin-resampled Pearson's r and RMSE, and unresampled mean error (ME), for tree and standing deadwood cover. Metrics are reported across all samples within a fold.

Biome	Metric	Tree	Standing Deadwood
(Sub) Tropical	Pearson	0.59 (0.03)	0.46 (0.05)
	RMSE	0.31 (0.01)	0.46 (0.05)
	ME	-0.11 (0.01)	0.03 (0.01)
Boreal and Montane	Pearson	0.58 (0.03)	0.30 (0.11)
	RMSE	0.28 (0.02)	0.37 (0.10)
	ME	-0.09 (0.01)	0.04 (0.02)
Drylands	Pearson	0.64 (0.02)	0.43 (0.28)
	RMSE	0.29 (0.02)	0.39 (0.07)
	ME	-0.13 (0.05)	0.04 (0.01)
Temperate	Pearson	0.63 (0.02)	0.56 (0.04)
	RMSE	0.26 (0.01)	0.32 (0.03)
	ME	-0.10 (0.00)	0.06 (0.01)

was lowest for dryland and temperate forests ($\sigma_{\text{RMSE}} = 0.03 - 0.05$) and highest for (sub) tropical and boreal forests ($\sigma_{\text{RMSE}} = 0.07 - 0.10$), consistent with data availability. Standing deadwood was consistently overestimated (ME $= 0.03 - 0.06$) compared to drone derived masks.

Overall, fold-to-fold variation was lower for tree cover than for standing deadwood cover. Performance was also consistently higher for tree cover.

3.3 Performance impact of time series length

Overall, performance increased with longer input timespans (see Figure 5). This improvement was primarily driven by substantial gains within the first year (52 weeks) of historical context. Extending the time series beyond one year further improved standing deadwood classification performance (F1-score) up to approximately two years and enhanced standing deadwood cover regression performance (Pearson's r) up to approximately three years. In contrast, tree cover regression showed only marginal gains beyond one year, while tree cover classification largely converged after 32 weeks. Variability across cross-validation folds was greater for standing deadwood cover than for forest cover, yet the overall trend of increasing performance with longer timespans remained consistent.

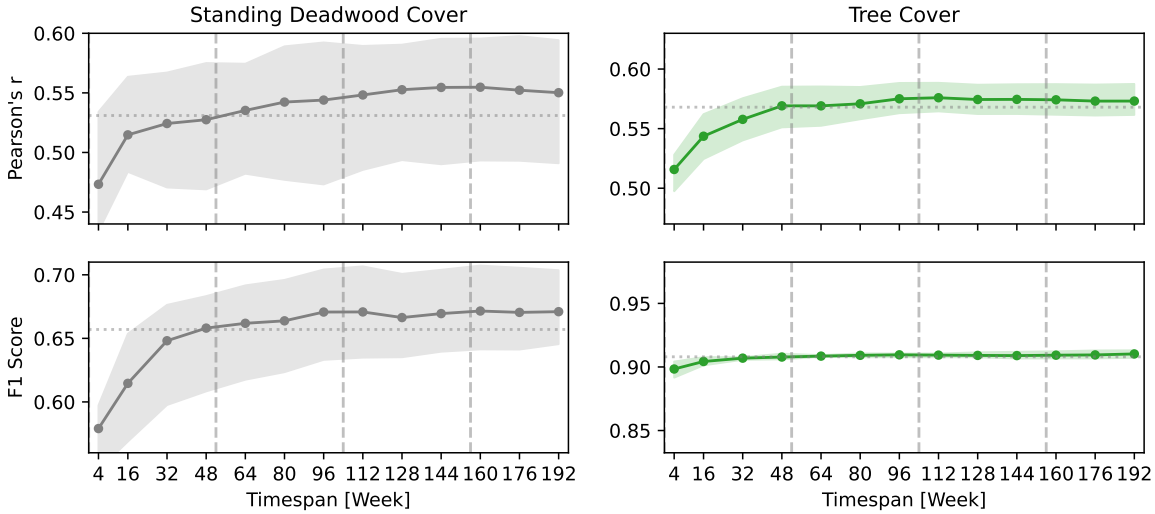

Figure 5: Mean model performance across folds and the covered timespan. F1-score measures binary classification of Sentinel-2 pixels and Pearson's *r* exact fractional cover retrieval. The displayed range in the y-dimension denotes the minimum and maximum performance across folds. The extent on the y-axis has the same height between standing deadwood cover and tree cover for the same metrics to enable comparison. Vertical dashed lines are drawn at year boundaries. A horizontal dotted line is drawn at the respective performance with one year of temporal context.

4 Case studies with ground-based validation

In this section, we compare our standing deadwood maps against independent and ground-validated tree mortality and disturbance data collected by local experts. To allow for representative evaluation, the maps were generated with custom models where all orthophoto sites within 10 km (if present) of the respective case study were excluded during training.

4.1 Case study 1: insect-driven mortality in boreal taiga (Sakhalin, Russia)

This collection of tree mortality sites is located on Sakhalin Island (Russia), in the Northwest Pacific region, a boreal taiga biome (Korznikov et al. 2022, 2025). The sites were first delineated through Pléiades-2, WorldView-2, and WorldView-3; and subsequently confirmed on the ground by local researchers. Further, the local surveys identified the affected tree species and the mortality agent (see Table 2).

The model mapped a standing deadwood cover above $\geq$ 20% for 74%, 92%, 99%, and 1% of the affected area across the sites A-D, respectively. The site that was affected by silkworms (site D) obtained fractional cover values averaging 1.4% standing deadwood and 30% for tree cover, thus not

Table 2: List of ground validated sites on Sakhalin Island (Russia). Column year denotes year of observation and map creation. Observation (Obs.) is the date where the observation was confirmed on the ground. Inference (Inf.) date is the date where the product was created on the respective location. Bark beetle refers to *Ips typographus*, silkworm refers to *Dendrolimus superans*. Species names: spruce – *Picea jezoensis*, *Picea glehnii* (only in the cite C), fir – *Abies sachalinensis*, larch – *Larix gmelinii*

.

Site	Year	Obs. date	Inf. date	Affected species	Mortality agent	Centroid
A	2020	Sep 12th	Jul 5th	Spruce	bark beetles	142.6205°N 47.4857°E
B	2021	Jun 9th	Jul 5th	Spruce	bark beetles	142.4572°N 47.4134°E
C	2021	Aug 2nd	Jul 1st	Spruce	bark beetles	141.9992°N 47.4586°E
D	2018	Jul 9th	Jul 15th	Fir, Spruce, Larch	silkworms	142.0757°N 48.5547°E

mapping the standing deadwood correctly. It is worth noting that the silkworm infestation happened before the Sentinel-2 launch between 2011 and 2014, and hence within the temporal spectral context, the trees were not alive at any point and only standing deadwood. All sites affected by bark beetles were accurately recovered by the model and mapped outbreak patches closely matched the provided polygons (see Figure S2). In comparison, Global Forest Watch (M. C. Hansen et al. 2013) was only sensitive to larger patches without a signal for medium mortality patches (see Figure S5). Note that the provided data only report tree mortality within the delineated sites; it cannot be assumed that tree mortality is absent in regions outside the polygons.

4.2 Case Study 2: fire–drought–management disturbance cascade (northern Germany)

This site in northeastern Germany was initially dominated by Scots pine and was subject to several fires and silvicultural treatments (Heinken et al. 2024; Jouy et al. 2025; Schmehl et al. 2025). Ground-validated data was made available in Schmehl et al. 2025 and is displayed atop our generated multi-year maps in Figure 6.

The first fire occurred in August of 2018, shortly after the prediction date (center of leaf-on period) of the product in 2018. Accordingly, the extent of the fire is visible in the 2019 product in the form of both tree cover loss (compared to 2018) and a dense standing deadwood patch. The majority of the north-eastern forest was cleared before June 2019 (Heinken et al. 2024). In the southwestern part thinning and clear-cut was applied to selected plots (Schmehl et al. 2025) during autumn 2019, as it is evident by reduced fractional tree cover in 2020. In the most southern part of the 2019 burn extent, the fire was predominantly on the ground resulting in no immediate visible overstory mortality (dashed line in Figure 6, surface fire). Exceptional drought and infestation contributed to the die-off

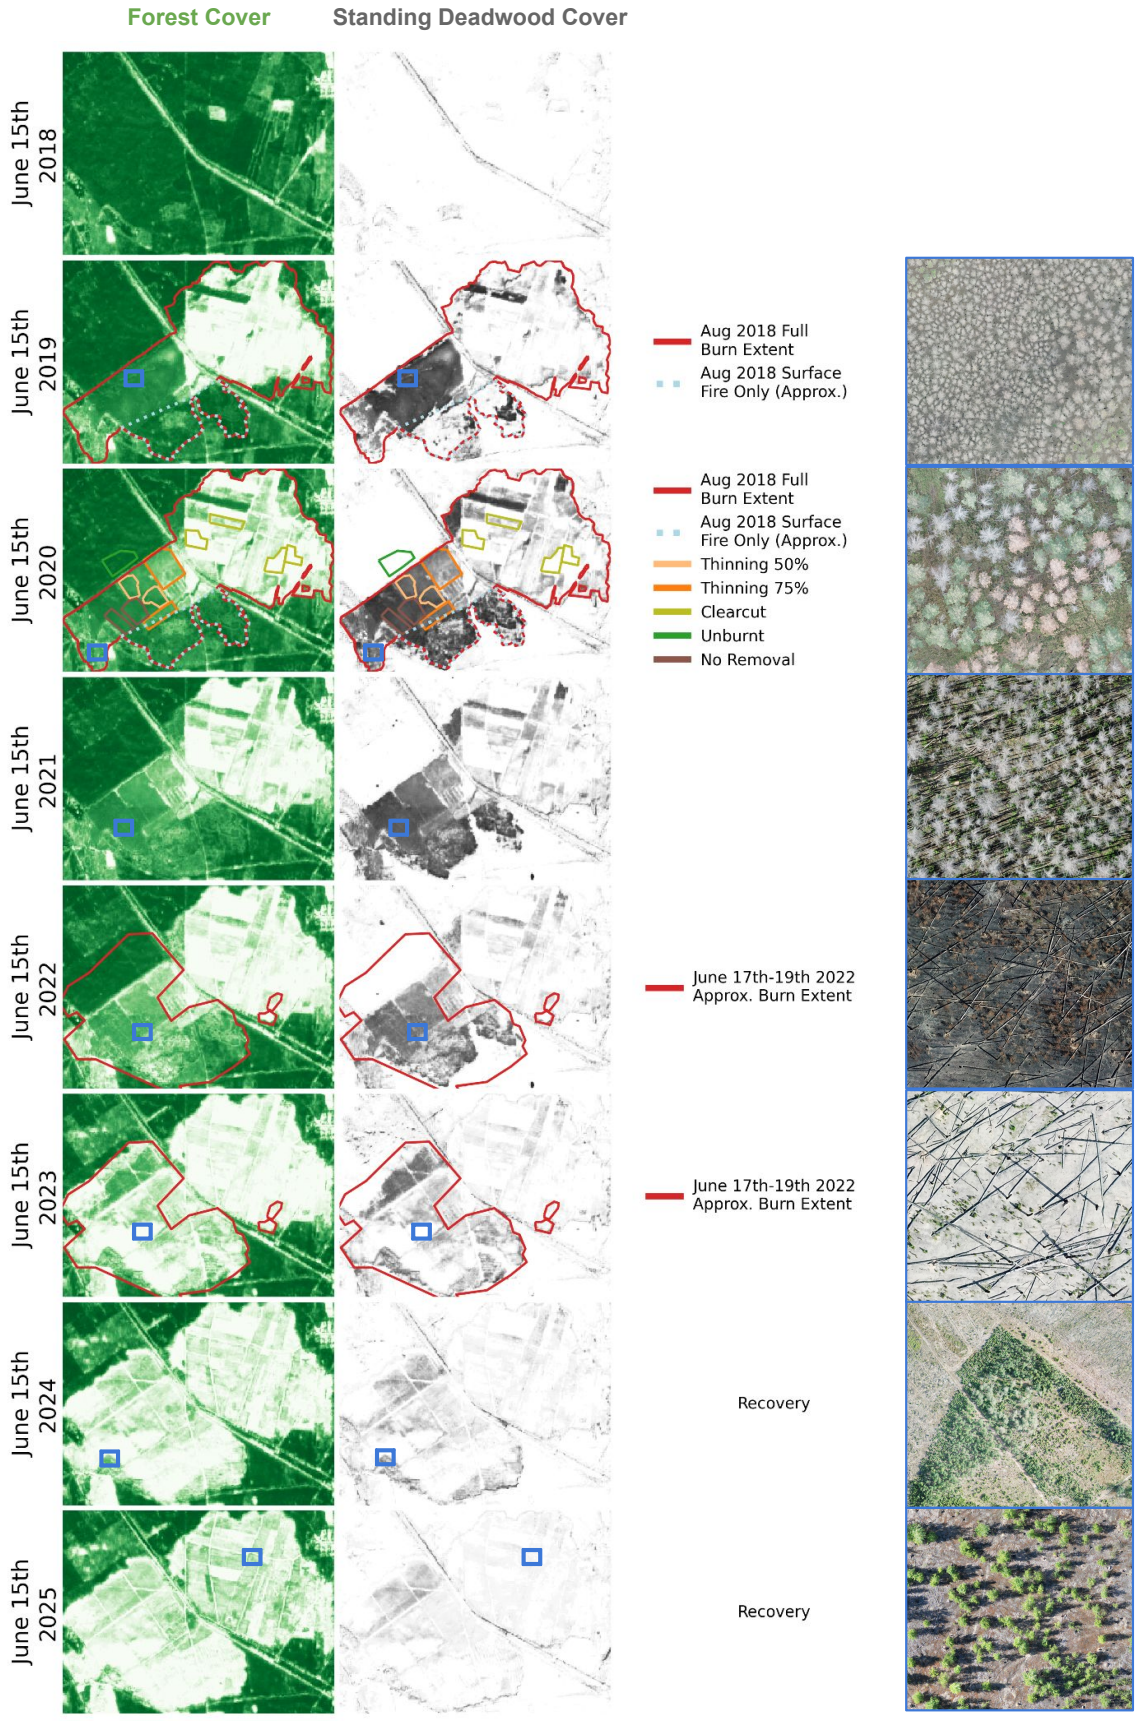

Figure 6: Fractional tree cover and standing deadwood product from 2018 to 2025 for a disturbance site in north-eastern Germany (52.04331°N, 12.92220°E). Ground-validated burn extent and monitoring plots on different silvicultural treatments were derived from Schmehl et al. 2025 (colored boxes in 2020). Rightmost column displays drone orthophotos within the respective year on multiple sites (blue).

of those trees, as is visible in the standing deadwood product in 2020. A second fire lasted two days from June 17th to 19th in 2022, just after inference date, further decreased tree cover within the burn extent for 2023. The forest product shows an increase in the northern cover for the years 2024 to 2025 suggesting forest recovery was also visible in drone orthophotos in the same years. Overall, this case study demonstrates that combined mapping of tree and standing deadwood cover enables to reveal fine-scale degradation patterns as well as recovery trajectories. In comparison, the European Forest Disturbance Atlas (Viana-Soto and Senf 2025) and Global Forest Watch (M. C. Hansen et al. 2013) only reveal that the location was disturbed but do not reflect mortality, thinning, or recovery dynamics (see Figure S4a and Figure S4b).

4.3 Case study 3: tree-level field inventory in bark beetle outbreak (Italian Alps)

Following a bark beetle outbreak, three areas in the Italian Alps were surveyed with drones (about 100 ha per area were covered) in June 2024. A ground survey of 20 plots in total was performed within the surveyed area in October 2024. In each plot, approximately 1 ha in size, the surveyors labeled each tree in the June drone imagery by recording the state of mortality and tree species. The dominant tree species on these plots were *Picea abies* (64%), *Abies alba* (20%), *Fagus sylvatica* (11%), and *Larix decidua* (2%). For the same sites we mapped standing deadwood cover and tree cover for June 5th 2024 (maximum greenness day of year).

Comparisons of the field recorded state of mortality with Sentinel-2 mapped fractional cover yielded strong overlaps (see Figure 7). Every dense tree mortality hotspot was recovered in the maps. For the plots Asiago-1, Asiago-14, and Visdende-5, the model mapped tree mortality with low standing deadwood cover ($< 50\%$) where there were only healthy trees present in the survey. For the sites Visdende-2, Ciapela-1, and Ciapela-3 the mapped cover is consistently shifted by $5 - 10$ m, which is within expectation of the registration accuracy of Sentinel-2.

In comparison, the European Forest Disturbance Atlas (Viana-Soto and Senf 2025) and Global Forest Watch (M. C. Hansen et al. 2013) do show coarse disturbance signals for a subset of patches, while the majority of scattered mortality is not mapped (see Figure S3a and Figure S3b).

4.4 Case Study 4: small-scale logging and forest recovery (Northern Brazil)

Tropical forests are being altered rapidly, and monitoring forest loss is essential because it is tightly linked to biodiversity decline, carbon emissions, and cascading impacts on ecosystem services and

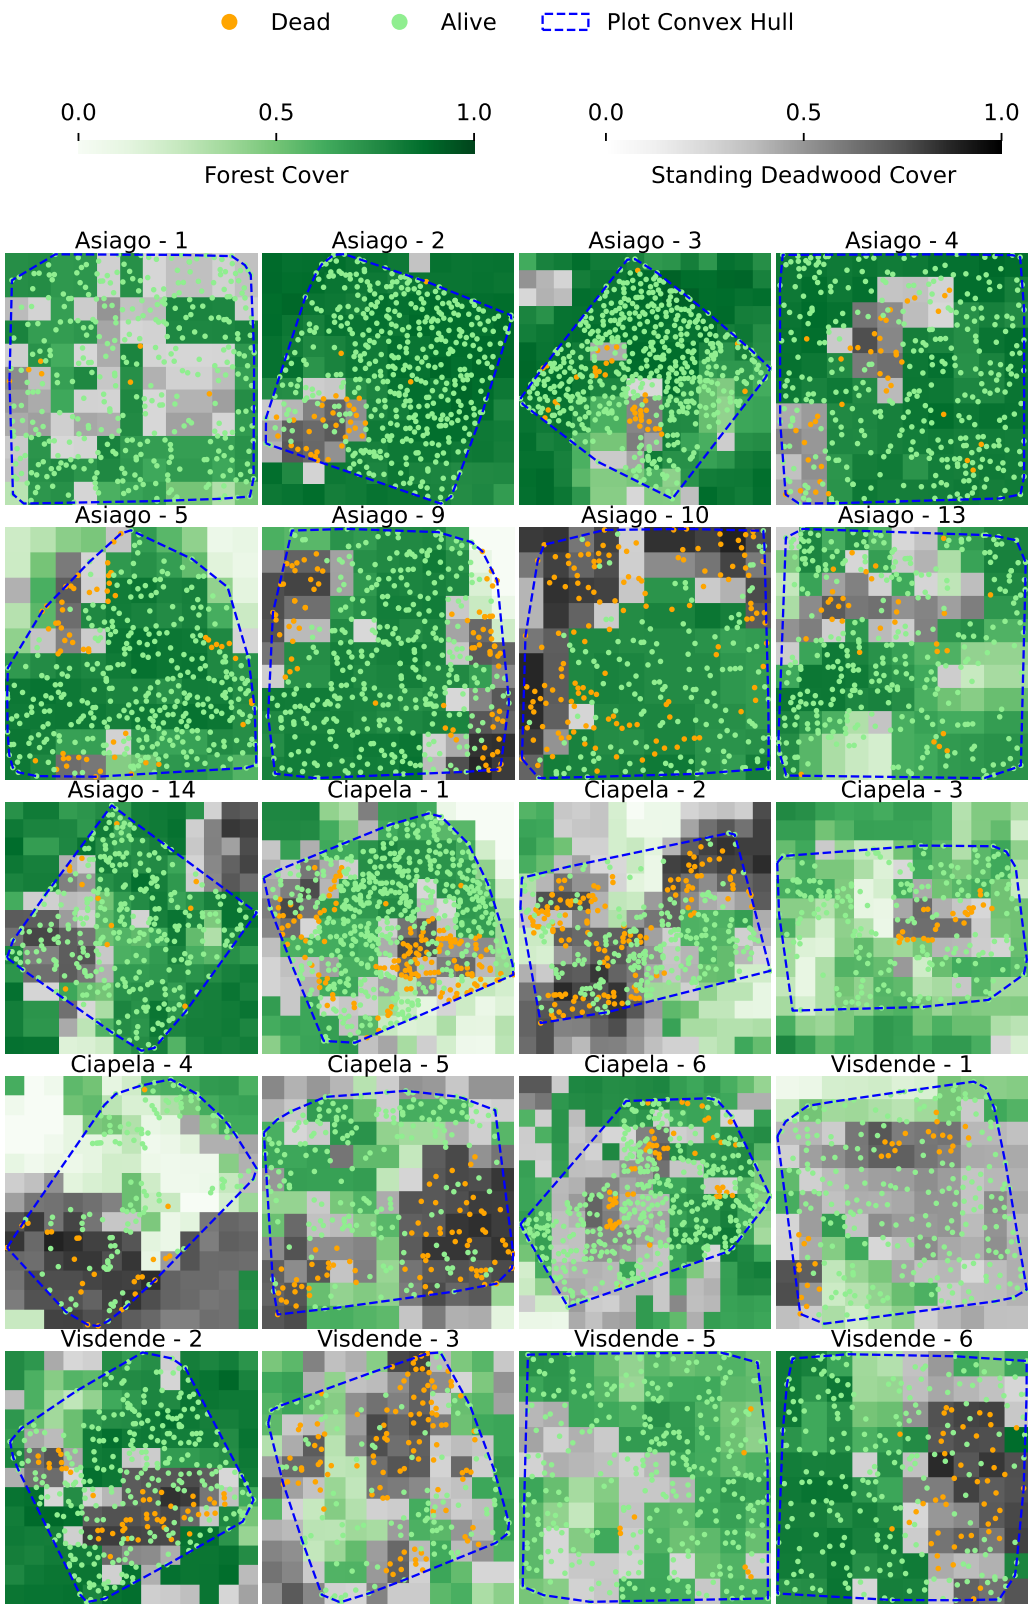

Figure 7: Comparison of mapped standing and tree cover against field-inventory for 20 plots in northern Italy in 2024. Each point corresponds to a tree, where dead trees are marked with orange and green otherwise. Mapped deadwood cover is displayed atop tree cover for pixels where the mapped deadwood cover was larger than 25%. The bounds of each plot are denoted by the convex hull (blue). Centroid locations in Table 3.

local livelihoods. To demonstrate the potential of our products even for deforestation monitoring, we generated forest cover predictions for a region in the central Amazon that showed evidence of small-scale forest loss and likely logging activities in high-resolution satellite images (see Figure 8). To detect forest cover loss, we computed the first year for every pixel where mapped tree cover dropped below 50%. In comparison to existing state of the art products (M. C. Hansen et al. 2013; Reiche et al. 2021; Vancutsem et al. 2021), our map revealed more complete and accurate representations of deforestation patches. The boundaries of detected disturbances are sharper and their shapes appear to more precisely match the patterns in the high-res reference image.

Recovery and regrowth of tropical forests after logging are highly heterogeneous, depending strongly on subsequent human activities. Hence their quantification is important for policy decisions and accounting. Given that our model maps fractional tree cover year over year, it can unveil tree regrowth in the years after a disturbance. In Figure 8, dark patches indicate older disturbances, and the example the bottom left patch (disturbance between 2019 and 2020) shows tree cover for the map in 2024 (top left), matching observations in the high-resolution image in 2023 (top right).

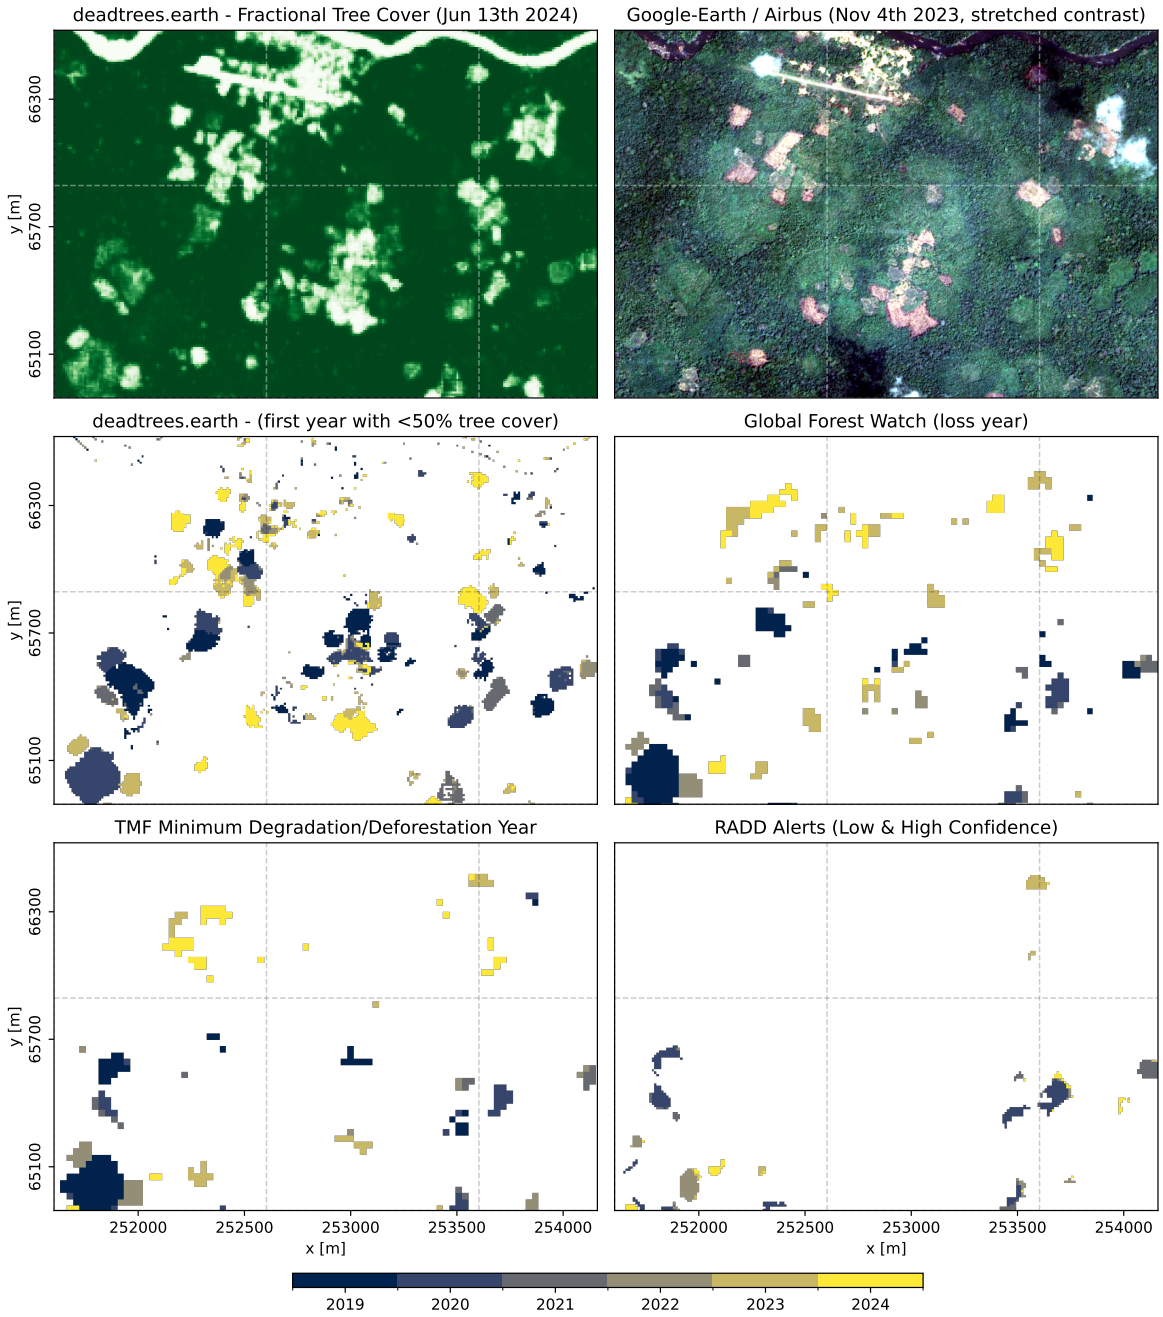

Figure 8: Comparison of disturbance timing and spatial detail across independent forest-change products for a tropical logging area. Top row: our *deadtrees.earth* Sentinel-2 fractional tree cover map for 13 June 2024 (left) and very-high-resolution optical imagery from Google Earth/Airbus (4 Nov 2023; contrast-stretched, right). Middle row: the first year in which the mapped fractional tree cover drops below 50% (derived from *deadtrees.earth*, left) and Global Forest Watch tree-cover loss year (right; (M. C. Hansen et al. 2013)). Bottom row: Tropical Moist Forests (TMF) minimum degradation/deforestation year (left; (Vancutsem et al. 2021)) and RADD alert year (low & high confidence; (Reiche et al. 2021)). For all products only losses recorded for the period 2019 to 2024 are visualized to enable fair comparison. All panels show the same spatial extent and projection (UTM zone 21N).

5 Discussion

In the following, we contextualize the scope, strengths, and limitations of the presented methodology. We first delineate the functional capabilities and inherent constraints of the methodology, discuss assumptions on phenological dynamics, forest structure, and georeferencing uncertainty. We then reflect on the length of satellite time series, evaluation limits imposed by weak supervision, and biome-specific performance patterns. Finally, we highlight research and operational applications enabled by globally consistent, annually resolved canopy-fraction products.

5.1 Scope of the resulting maps and methodological limitations

From a methodological perspective, the resulting annual maps are designed to represent canopy-level forest state and its evolution through time by jointly estimating fractional standing deadwood cover and fractional tree cover at 10 m resolution. This paired representation enables detection of top-of-canopy mortality signals (increases in standing deadwood cover) and subsequent canopy loss decline (decrease in tree cover), and, critically, allows these processes to be tracked as multi-year continuous trajectories rather than discrete classification of single-events. Because both variables are estimated as continuous sub-pixel fractions, the methodology enables to reflect gradual disturbances and dieback within a pixel, as well as the subsequent recovery of tree cover during regrowth as crowns expand and gaps fill in. In addition, fractional tree cover provides an explicit measure of canopy openness, enabling differentiation between naturally open-canopy systems and dense, closed-canopy forests, which is important for interpreting disturbance impacts across biomes. Thus, the presented method provides sub-pixel metrics that enable year-over-year tracking of mortality, disturbance and recovery dynamics.

In this study, we assume that leaf-loss and browning in the leaf-on season is a reliable indicator of tree death (see Section 2). However, defoliation does not always indicate recent or proximate tree mortality. In some cases, trees may defoliate and refoliate within the same season, as observed in the response of *Pinus radiata* to the disease red needle cast (Dick et al. 2014; Watt et al. 2024), and of Mediterranean beach to a late frost (D'Andrea et al. 2019). In other cases, pest and pathogen impacts may be confined to the upper crown, such that residual or resprouting foliage persists below the dead crown and may be partially obscured in top-down imagery, causing stressed but living trees to appear dead when viewed from above. Although single defoliation events are often not sufficient to cause

tree death, they can weaken trees, deplete carbohydrate reserves, and reduce growth rates, thereby increasing susceptibility to secondary stressors such as pests and diseases, and elevating tree mortality rates (Silvestro et al. 2025). While these phenomena complicate the detection of tree mortality, the temporal context of fractional yearly standing deadwood cover for relatively small areas aids in distinguishing temporary leaf-loss from continued leaflessness and death. Downstream analysis will need to take such phenomenon into account for accurate interpretation.

The proposed methodology maps and validates standing deadwood cover at the midpoint of the regional leaf-on season in each year. Leaf-on season is defined at the level of 40-km pixels, and thus encompasses regional and stand-level averages. However, leaf phenology including the timing of the leaf-on phase in deciduous individuals can vary substantially among species and trees within stands and regions, particularly in water-limited biomes (Park et al. 2026). This may partially explain the lower model performance in tropical or dryland ecosystems with high phenological variation within stands and/or among years (see Table 1). In these ecosystems, it is rarely immediately clear whether leaflessness corresponds to tree death, even with in-situ visual inspection. Furthermore, the regional mid-point leaf-on date may not be a useful point of reference for some trees and years, given complex and variable phenological patterns. Time series of standing deadwood cover and tree cover should nonetheless provide the best basis for mapping tree mortality. We expect that mortality predictions will be improved in the future through development of more complex models incorporating finer temporal resolution and/or more detailed or adaptive phenological characterization.

Top-down optical imagery from drones and satellites has inherent limitations in detecting tree mortality in dynamic and multi-layered forest canopies. Aerial optical imagery is dominated by the top-most vegetation layer, and tree mortality in understory layers can be completely missed if the upper layer is unaffected. Even mortality in the canopy is harder to detect in multi-layered canopies, as disturbed patches can retain a green-dominated spectral signature if one or more understory layers remain (Coppin and Bauer 1996; Stephen Frolking et al. 2009). Further, rapid vegetation regrowth and ingrowth, often accompanied by rapid decomposition or collapse of canopy deadwood, can quickly erase any signal of canopy mortality events. Hence, substantial tree mortality may be accompanied by only small and transient shifts in cover (both standing deadwood and tree cover), which can be difficult to distinguish from seasonal deciduousness or other transient defoliation. As a result, it can be assumed that the method is only sensitive to fractional tree cover when the canopy gap fraction substantially increases (see Section 4.4 and Section 3).

5.2 The value of long spectral time series

Sentinel-2 lacks the spatial resolution that is necessary to explicitly resolve exact standing deadwood or tree cover within a given pixel, but it has the advantage of high temporal resolution which provides temporal context. This temporal information adds value in two distinct ways. First, the temporal trajectories in multispectral reflectance display the spectral fingerprint of forest dynamics. For example, a disturbance event in a temperate forest is expected to reveal itself as a sudden breakpoint in what would otherwise be a stable seasonal pattern (M. C. Hansen et al. 2013; Schiefer et al. 2023). Second, high temporal resolution provides more robust reflectance information. Reflectance values can vary without any relevant change to the underlying vegetation due to sensor noise, atmospheric scattering, cirrus clouds, or other factors. Dense time series containing multiple observations within seasons provide redundancy and thus enable more robust quantification of reflectance changes relevant to the target variable (Rußwurm and Körner 2020; Tseng et al. 2023). These advantages of fine temporal resolution thus make it easier to detect and characterize subtle changes, and offset the disadvantage of coarse spatial resolution.

Our temporal context experiment clearly underlined the benefit of increased temporal resolution and extent (see Section 3.3). In the context of deep learning and vegetation mapping from Sentinel-2 time series, mostly time spans of up to two years are considered (Montero et al. 2024a, 2025; Rußwurm and Körner 2020; Schiefer et al. 2023; Tseng et al. 2023; H. K. Zhang et al. 2025). Our findings highlight that the first year of historical temporal context is the most critical, but additional years further boost the performance for standing deadwood cover estimation. We hypothesize that for standing deadwood longer timeseries are more important as they can capture the sudden or gradual death process. It is important to note that all experiments we conducted with at minimum four weeks and up to four years of temporal coverage before the disturbance event. Consequently, these results therefore show the value of temporal context pre-disturbance, but do not provide comprehensive comparison to single-timestep models or potential inclusion of observations post-disturbance.

5.3 Evaluation limits from georeferencing uncertainty

Our ability to fit and evaluate a Sentinel-2-based fractional cover model from drone-derived labels is limited by georeferencing uncertainty in both datasets. That is, the training and evaluation targets are both approximate, meaning that model training is a weakly supervised task. The exact georeferencing

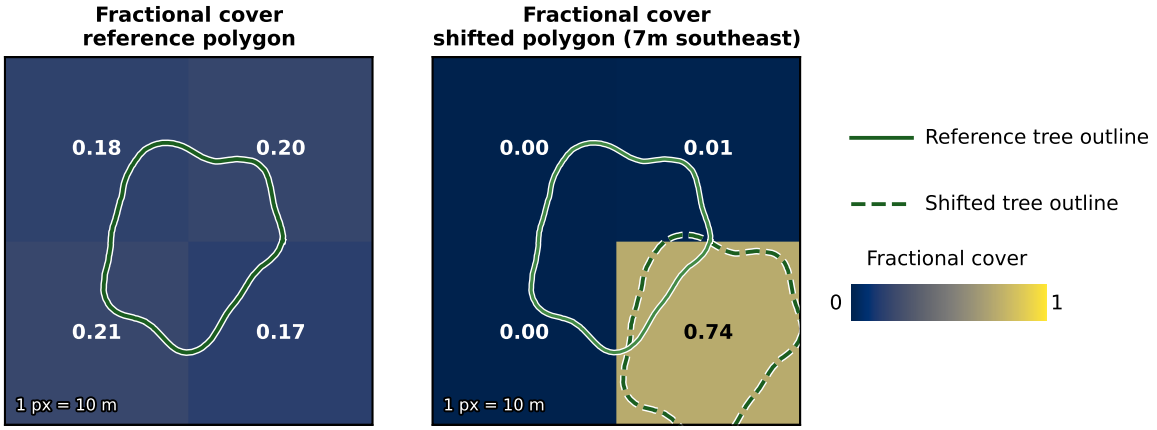

Figure 9: Illustration of how small spatial offsets propagate into large pixel-wise errors for 10 m fractional cover mapping. A single tree crown (solid outline) is converted to fractional cover in a 2 × 2 pixel window (left). Shifting the crown by 7 m southeast (dashed outline; right) changes the distribution of fractional cover across pixels, yielding an RMSE of 0.33 when comparing to the reference raster.

accuracy of both Sentinel-2 observations and crowd-sourced drone orthophotos is uncertain at the meter scale. To bound this uncertainty, *deadtrees.earth* retains only orthophotos with less than 15 m offset relative to publicly available aerial basemaps (Mosig et al. 2026). For Sentinel-2, Copernicus reports $<$ 12 m absolute geolocation error at 95.5% confidence, $<$ 5 m multitemporal registration error at 95.5% confidence, and $<$ 0.3 px multispectral registration error at 99.7% confidence (Choi et al. 2025; Copernicus n.d.). Even within these bounds, small spatial shifts can strongly affect pixel-wise fractional cover labels: a shift of only 7 m can redistribute cover across adjacent 10 m pixels and substantially change per-pixel values, which is particularly consequential for scattered standing deadwood and very sparse forests (Fig. 9). As a result, pixel-to-pixel evaluation can overestimate error and imposes an upper bound on model performance scores even with infinite model capacity and training data. Across the full dataset, for a pixel-to-pixel evaluation we obtained weighted RMSE values of 0.335 for standing deadwood cover and 0.285 for tree cover (mean across biomes) Table 1.

In contrast, spatial shifts conserve area totals, such that large-area bias estimates are considerably more stable: at biome level, our model overestimated standing deadwood cover by only 4% and underestimated tree cover by 11% Table 1, where tree cover underestimation is driven by little sensitivity to crown gaps of the drone tree segmentation model (see Section 2.1.1). This indicates that further improvements in pixel-level correlation or mean-error metrics are not necessarily the most meaningful optimization target. The issue is exacerbated for binary disturbance products, where meter-scale misregistration directly flips class labels (Alix-Garcia and Millimet 2023) and such maps remain com-

mon in operational monitoring (M. C. Hansen et al. 2013; Viana-Soto and Senf 2025). Future work should therefore consider training targets and evaluation protocols that explicitly account for georeferencing uncertainty, rather than enforcing strict per-pixel agreement; related ideas exist in learning with spatial noise (Mnih and Hinton 2012), although our preliminary adaptation to sub-pixel shifts primarily produced smoother maps without improving accuracy (not shown). Lastly, these georeferencing uncertainties are partly mitigated by scale: as the reference archive expands across regions and years, the model can learn invariant temporal–spectral patterns and become increasingly robust to geolocation offsets.

5.4 Current model capabilities, limitations, and future model improvements

Precisely mapping fractional cover values from Sentinel-2 in scattered mortality or open canopy scenarios is a challenging task because tree crown diameters are often smaller than the Sentinel-2 resolution of 10 m (Jucker et al. 2025). Further, the area of a crown, represented by the fractional cover value, is linked to the reflectance value of that pixel, that generally also contains signals of the neighboring pixels due to light scattering in every observation (Choi et al. 2025) and that is inconsistent across time due to georeferencing shifts (also see Section 5.3). Conditions across sites (functional composition, stand structure, understory, and heterogenity) introduce additional variability that is not linked to the respective cover value. Despite these challenges, we observed consistent precision and recall performance for cover values above 50%, where the model generally performed much better for tree cover than standing deadwood (see Figure 4).

Performance however rapidly decreased for low fractional cover values ($< 50\%$, also see Figure 1), especially for standing deadwood. A key difference between tree and standing deadwood cover was reference data availability. Overall, the available training data for fractional tree cover was about four times larger than for standing deadwood cover (4.92 vs. 1.4 million pixels for cover $> 0\%$) and about 44 times larger for cover values above 50% (3.1 vs. 0.07 million for cover $> 50\%$). This inequality is naturally inherited from the problem definition where standing deadwood cover is a subset of tree cover, and most sites only have low levels of tree mortality.

The model performed worse in tropical and boreal forests than in other biomes. In these forests, reliably mapping standing deadwood, especially scattered deadwood, remained challenging due to limited cloud-free observations (see Figure S7). In the tropics specifically, this is likely exacerbated by strong vertical vegetation stratification, high phenological heterogeneity, and predominantly scattered

mortality patterns (Espírito-Santo et al. 2014; Xu et al. 2026). For boreal forests, smaller crowns and generally more open canopy forests with lower forest cover likely increase the difficulty. Lastly, reference data on standing deadwood was comparably underrepresented in both biomes, showing promising room for future improvement.

In the reference data and in model predictions, standing deadwood cover above 70% was rare for non-temperate biomes, such that we were not able to compute stable precision and recall metrics for those cases (see Figure 4 and Section 2.2.6). While this partly reflects a limitation of the dataset, it also mirrors ecological constraints and natural upper bounds in fractional deadwood cover across these biomes. Dryland and boreal ecosystems commonly have open forests where parts of the forest floor are visible (i.e., a fractional tree cover of far less than 100%). Here, even if the entire forest dies off, the fractional deadwood may never be higher than 50%. Also, if sites only contain scattered mortality, the derived reference fractional cover values will only very rarely and coincidentally be high when pixel boundaries are aligned with the crown (see Figure 9).

Taken together, these patterns suggest that the main limitations can be attributed to underrepresented regions and conditions, and should resolve with a more comprehensive training data set. *deadtrees.earth* is a constantly growing drone database platform, such that every contribution in a new region (or new year) has the potential to improve the robustness and generalization capabilities of the model. Commonly, training datasets are created and harmonized only once as a preprocessing step. However, in this work, the data flow and balancing scheme during model training was specifically designed to dynamically scale with the growing archive of drone data (see Section 2.2.5). Both tree cover and standing deadwood segmentation models have been integrated into *deadtrees.earth* to be automatically applied to new contributions. Sentinel-2 data preprocessing is fully automated through *sentle* (github.com/cmosig/sentle). Auditing segmentation quality and orthophoto validity are already an integral part of the *deadtrees.earth* platform. In total, apart from manual quality checks of new contributions, the presented overall workflow enables regular and fully automatic model retraining, such that every contribution to *deadtrees.earth* naturally improves the Sentinel-2 based maps.

5.5 Applications of map products beyond quantification

Applying the proposed methodology at scale opens up a variety of new research avenues beyond tree mortality quantification that were not feasible with current state-of-the-art products or restricted to small, well monitored sites.

Tree mortality and recovery trajectories are a key uncertainty in carbon-cycle assessments because mortality interrupts CO_2 uptake, reducing gross primary production, while simultaneously enhancing heterotrophic respiration. Together, these processes can flip ecosystems from carbon sinks to carbon sources (Schulze 2006; Wijas et al. 2024). However, for this our data product needs to be analyzed in tandem with carbon flux data or model products. Other downstream applications could aim to compute large-scale tree mortality rates or identify the onset of a tree mortality event and the temporal development of subsequent tree cover loss and ecosystem recovery via succession. Spatio-temporal disturbance patterns that differentiate between loss and standing mortality will enable us to better attribute disturbances to specific causes such as fire, harvest, or drought-induced mortality. Annual trajectories further help estimating quantifying recovery rates as well as persistence and decay of standing deadwood.

Precise tree mortality and disturbance maps of the current year can function as an early-warning signal for larger insect outbreaks and tree loss in remote regions in later years (Korznikov et al. 2025). They can also guide more detailed on-the-ground or aerial surveying efforts such as the Insect and Disease Survey (U.S. Forest Service 2025) or forest inventories. Further they can inform potential sanitation efforts to reduce risk of spreading wildfire as standing deadwood is fuel for fire.

These remote sensing derived disturbances maps will augment traditional ground-based inventory instead of replacing them (The International Tree Mortality Network et al. 2024). Section 4 confirms that known forest disturbance dynamics can be recovered accurately but, critically, with globally consistent metrics that can be intercompared between regions. While deep learning applied to remote sensing data does not yield immaculate results, its derived maps have a quantifiable bias and error, a property that is near impossible to manage across independent and internationally spread teams with varying goals. The fully open nature of the reference data, methodology, future derived maps, and the input Sentinel-2 data enables accountability and full reproducibility by any institution. This contrasts with non-public forest inventory data or pay-to-access higher-resolution satellite archives such as Maxar, SPOT, Planet, or Pléiades.

Beyond disturbance related applications, annual fractional tree cover maps can contribute to research on treelines, forest fragmentation, trees outside forests, and tree cover distributions (Garbarino et al. 2023; Liu et al. 2023). The model is particularly suited for such applications because it does not rely on external forest masks as inference boundaries, as is common in related approaches (Lange et al. 2024; Schiefer et al. 2023; Schiller et al. 2024). Year-over-year tree cover maps can further help

to track canopy closure and thus growth in plantations, thereby supporting monitoring of restoration efforts. Both standing deadwood and tree cover layers can also serve as priors or training targets for weakly supervised segmentation in higher-resolution imagery (Schmitt et al. 2020), and as reference products for linking to coarser remote-sensing systems that offer longer temporal archives (e.g., Landsat) or higher revisit frequency (e.g., MODIS).

Comparing to existing state-of-the-art continental-scale products, our model delivered qualitatively superior performance in the case studies (see Section 4.4 and Section A.3). The proposed methodology enables mapping at finer spatial resolution while explicitly representing sub-pixel dynamics through fractional cover estimates, thereby capturing gradual change and partial disturbances that are typically missed by binary disturbance masks. Moreover, by producing annual maps of both standing deadwood cover and tree cover, the approach provides complementary information on canopy-visible mortality versus subsequent canopy loss or removal (e.g., see Figure S3b and Section A.3), which is essential for interpreting disturbance agents and trajectories. The resulting maps open the door to more spatially granular analyses of forest degradation, including fine-scale selective logging and subsequent recovery in the tropics (see Section 4.4). The presented comparisons are however only qualitative and limited to a small set of sites. Future work should therefore focus on exhaustive quantitative benchmarking of disturbance products and validation against centimeter-scale images and ground-validated data, to more comprehensively characterize performance patterns across disturbance types, severities, and biomes.

6 Conclusion

We introduce a globally scalable approach to map yearly fractional tree cover and fractional standing deadwood cover at 10 m resolution (2018–2025) from four-year Sentinel-2 multispectral pixel time series. We trained and evaluated the satellite-based models with centimeter-scale drone orthophotos from *deadtrees.earth*, which were used to create a globally distributed standing deadwood cover and tree cover dataset using semantic segmentation models. This represents the first globally consistent mapping approach that simultaneously tracks fractional tree cover and standing deadwood cover across biomes, thereby enabling new perspectives on disturbance trajectories and post-disturbance ecosystem dynamics.

Across biomes, the satellite-based models showed consistently strong performance for tree cover

and moderate performance for standing deadwood, with clear biome dependence and the expected sensitivity to reference data availability. We show that the length of satellite time series has a strong imprint on the model performance, highlighting the value of continuous Earth observation products. Multiple case studies indicate that the maps recover known, fine disturbance patterns and timing with far higher sensitivity to fine-scale dynamics compared to existing state-of-the-art disturbance products.

Finally, the presented pipeline is designed to improve continuously: *deadtrees.earth* is growing, automated segmentation in drone orthophotos is integrated into the platform, and the training workflow is explicitly built to scale with new contributions, enabling routine retraining as the global reference archive expands. Together with the open processing stack, this provides a reproducible path towards progressively reliable, globally consistent fine-scale disturbance metrics for ecology, carbon cycle studies, risk assessment, and targeted field campaigns.

Code and Data Availability

The Sentinel-1 and -2 processing pipeline (see Section 2.1.3) is publicly available in the form of a python package: github.com/cmosig/sentle. The version used in this paper is 2025.10.1. The code to reproduce the results and the model will be made public after acceptance of the manuscript. The preprocessed training dataset including Sentinel-2 data and references rasters will be made available after the acceptance of the manuscript. Previews of maps generated with the introduced methodology will be made available and continuously updated on deadtrees.earth/deadtrees.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process.

During the preparation of this work the author(s) used ChatGPT for language editing. During analysis, the authors employed GitHub Copilot and OpenAI Codex for programming assistance. After using these tools, the author(s) reviewed and edited the content and code as needed and take(s) full responsibility for the content of the published article and code.

Acknowledgements

C.M., T.K., M.D.M., J.V.J. thank the European Space Agency for funding the “DeepFeatures” project via the AI4SCIENCE activity and the “FORTRACK” project via the ESA CLIMATE SPACE: Climate-Biodiversity studies.. C.M. and M.D.M. acknowledge financial support from the German Aerospace Centre (DLR) on behalf of the Federal Ministry for Economic Affairs and Climate Action (BMWK) for project ML4Earth (FKZ 50EE2201B). T.K. and J.S. acknowledge funding from the German Research Foundation via the “PANOPS” (grant no. 504978936) and EXC 3127 Future Forests. T.K., J.V.J. and S.S. acknolwedge support by the Ministry for Food, Rural Areas, and Consumer Protection, Baden-Württemberg, Germany (grant no. 52-8670.00) and the German Research Foundation via the NFDI4Earth pilot project GeoLabel (DFG grant no. 460036893). K.K. acknowledges the financial support from the Czech Science Foundation (25-15727S). D.N. was partially supported by the European Union’s Horizon Europe Research and Innovation Programme under grant agreement 101134200 “Forest surveillance with artificial intelligence and digital technologies – FORSAID”. MB acknowledges support from SNSF-Sinergia grant number CRSII5_216562 and SNSF grant number IZCOZ0_213355. S.H. and Y.C. acknowledged funding from Villum Fonden (DRYTIP project, grant agreement no. 37465). We thank Ian McGregor and Evan Gora of the Cary Institute of Ecosystem Studies for providing helpful insights and perspectives about mortality and tree loss detectability with this approach on Barro Colorado Island.

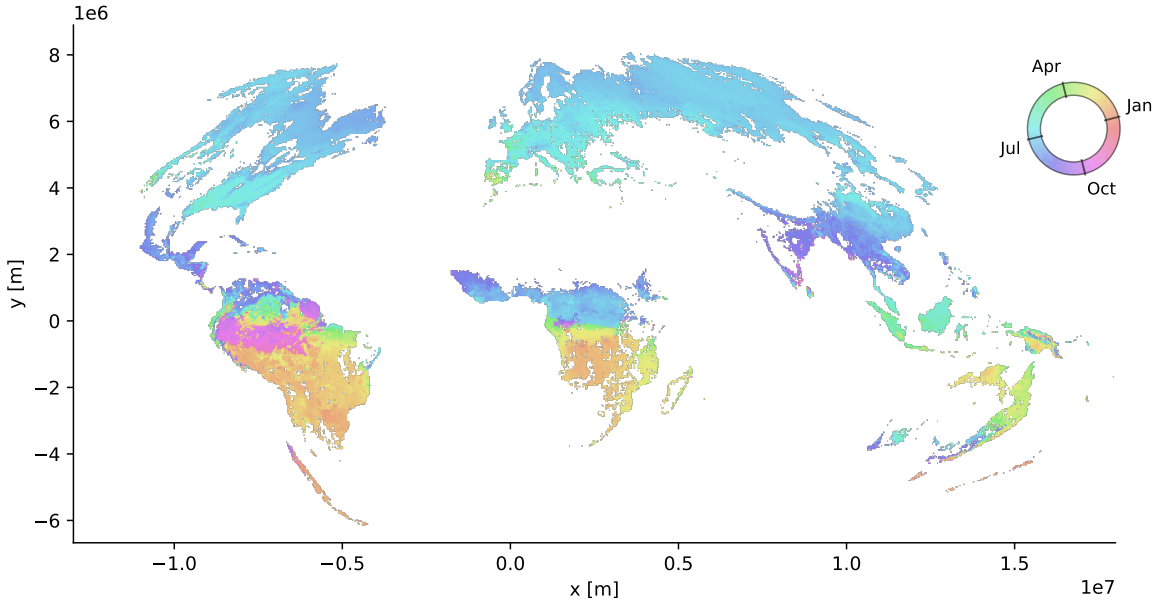

Figure S1: Computed training and inference day-of-year (see Section 2.1.2) based on phenology indicator on *deadtrees.earth* (Mosig et al. 2026). Projection is MODIS Sinusoidal (ESRI:54008).

A Appendix

A.1 Center of leaf-off period global map

A.2 Qualitative results on Sakhalin island

For methodology, quantitative results, and site descriptions see Section 4.1 and Table 2.

A.3 Comparison against Global Forest Watch (GFW) and European Forest Disturbance Atlas (EFDA) for case studies 2 and 3.

A.4 Ablation Precision-Recall Tradeoff

Applying extra weight to samples where fractional cover is exactly zero acts as a mechanism to balance precision and recall (see Supplementary Figure S6). For deadwood, increased weight on the zero class reduces recall, but has limited impact on the prediction precision. For forest cover, increases weight increases the precision by similar amounts by which the recall is decreased.

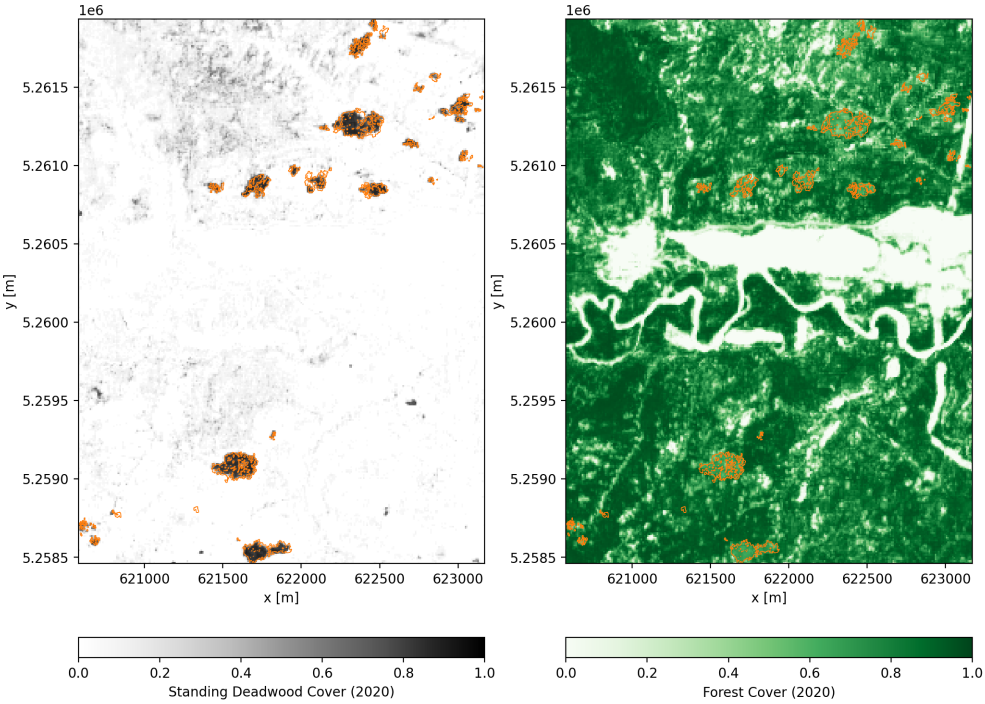

(a) Site A.

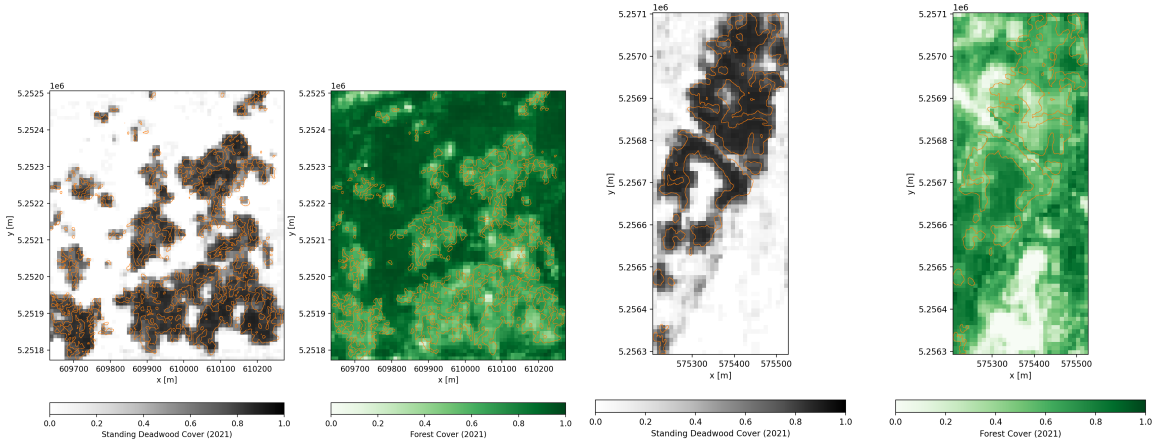

(b) Site B.

(c) Site C.

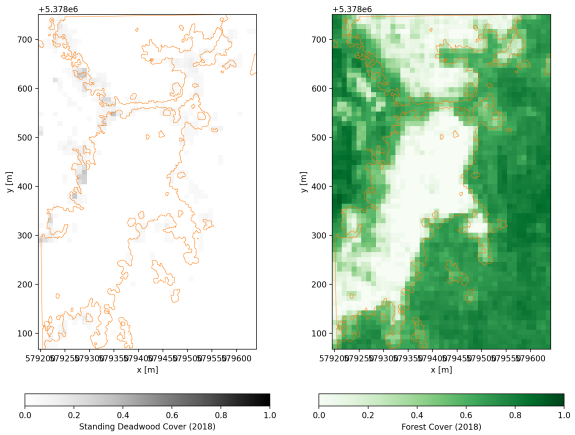

(d) Site D.

Figure S2: Fractional tree and standing deadwood cover for the sites on Sakhalin Island (Russia). Coordinates in UTM zone 54N. Ground-validated disturbance sites are delineated in orange.

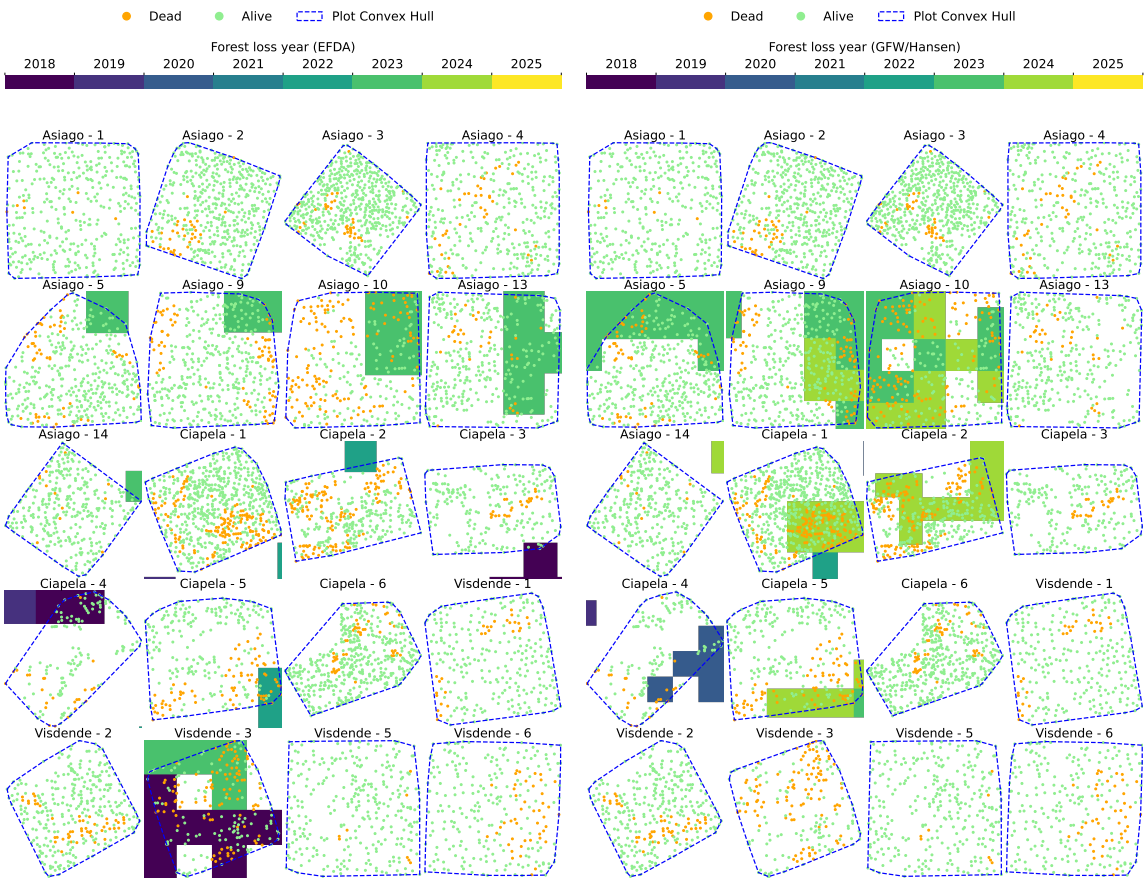

(a) EFDA (Viana-Soto and Senf 2025).

(b) GFW (M. C. Hansen et al. 2013).

Figure S3: Case Study 3: Identical visualization to Supplementary Figure 7 but disturbance map was obtained from Global Forest Watch (GFW) (M. C. Hansen et al. 2013) and European Forest Disturbance Atlas (EFDA) (Viana-Soto and Senf 2025).

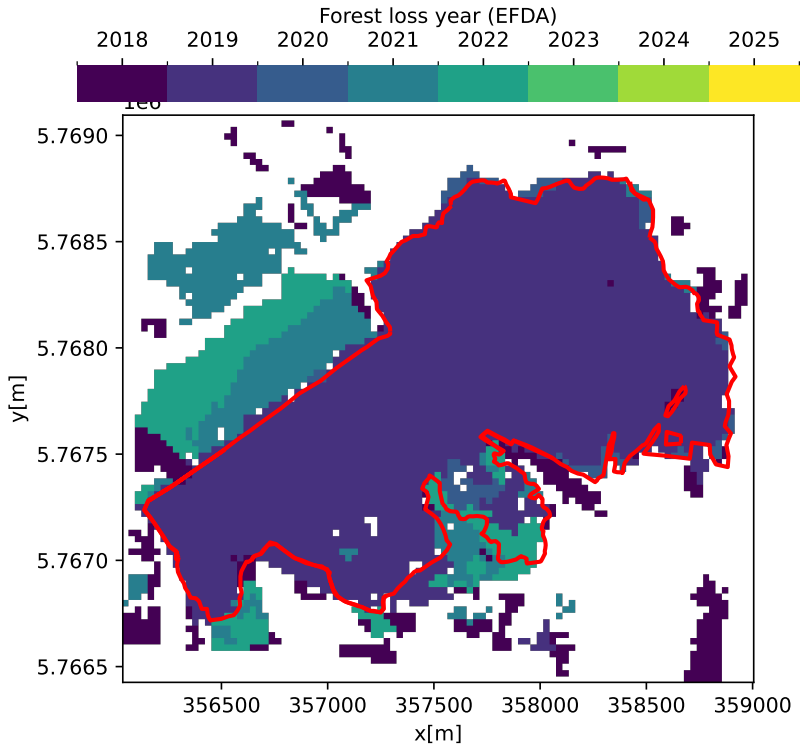

(a) EFDA (Viana-Soto and Senf 2025).

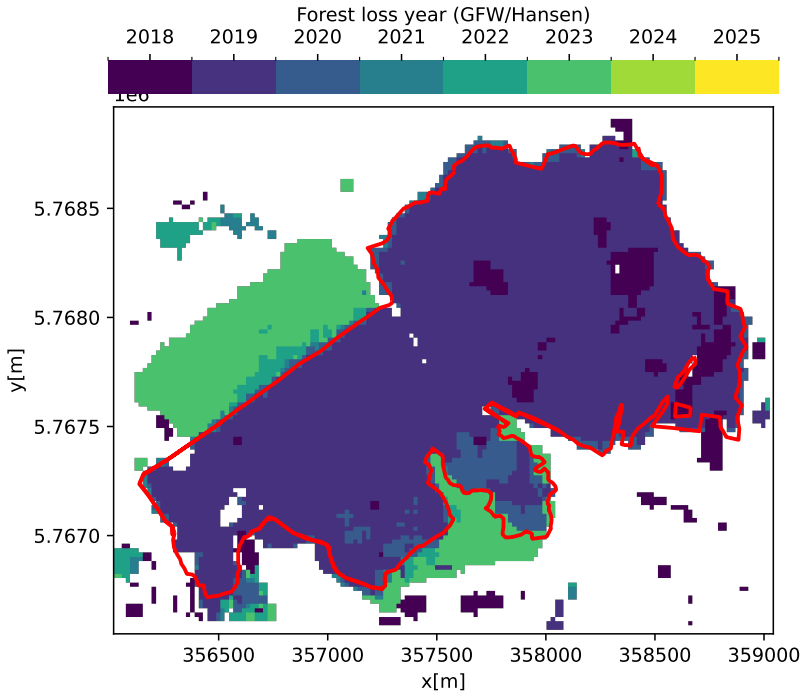

(b) GFW (M. C. Hansen et al. 2013).

Figure S4: Case Study 2: Aggregated visualization for site in Section 4.2, with disturbance map from obtained from Global Forest Watch (GFW) (M. C. Hansen et al. 2013) and European Forest Disturbance Atlas (EFDA) (Viana-Soto and Senf 2025). Red line shows extent of fire for 2018 (see Section 4.2).

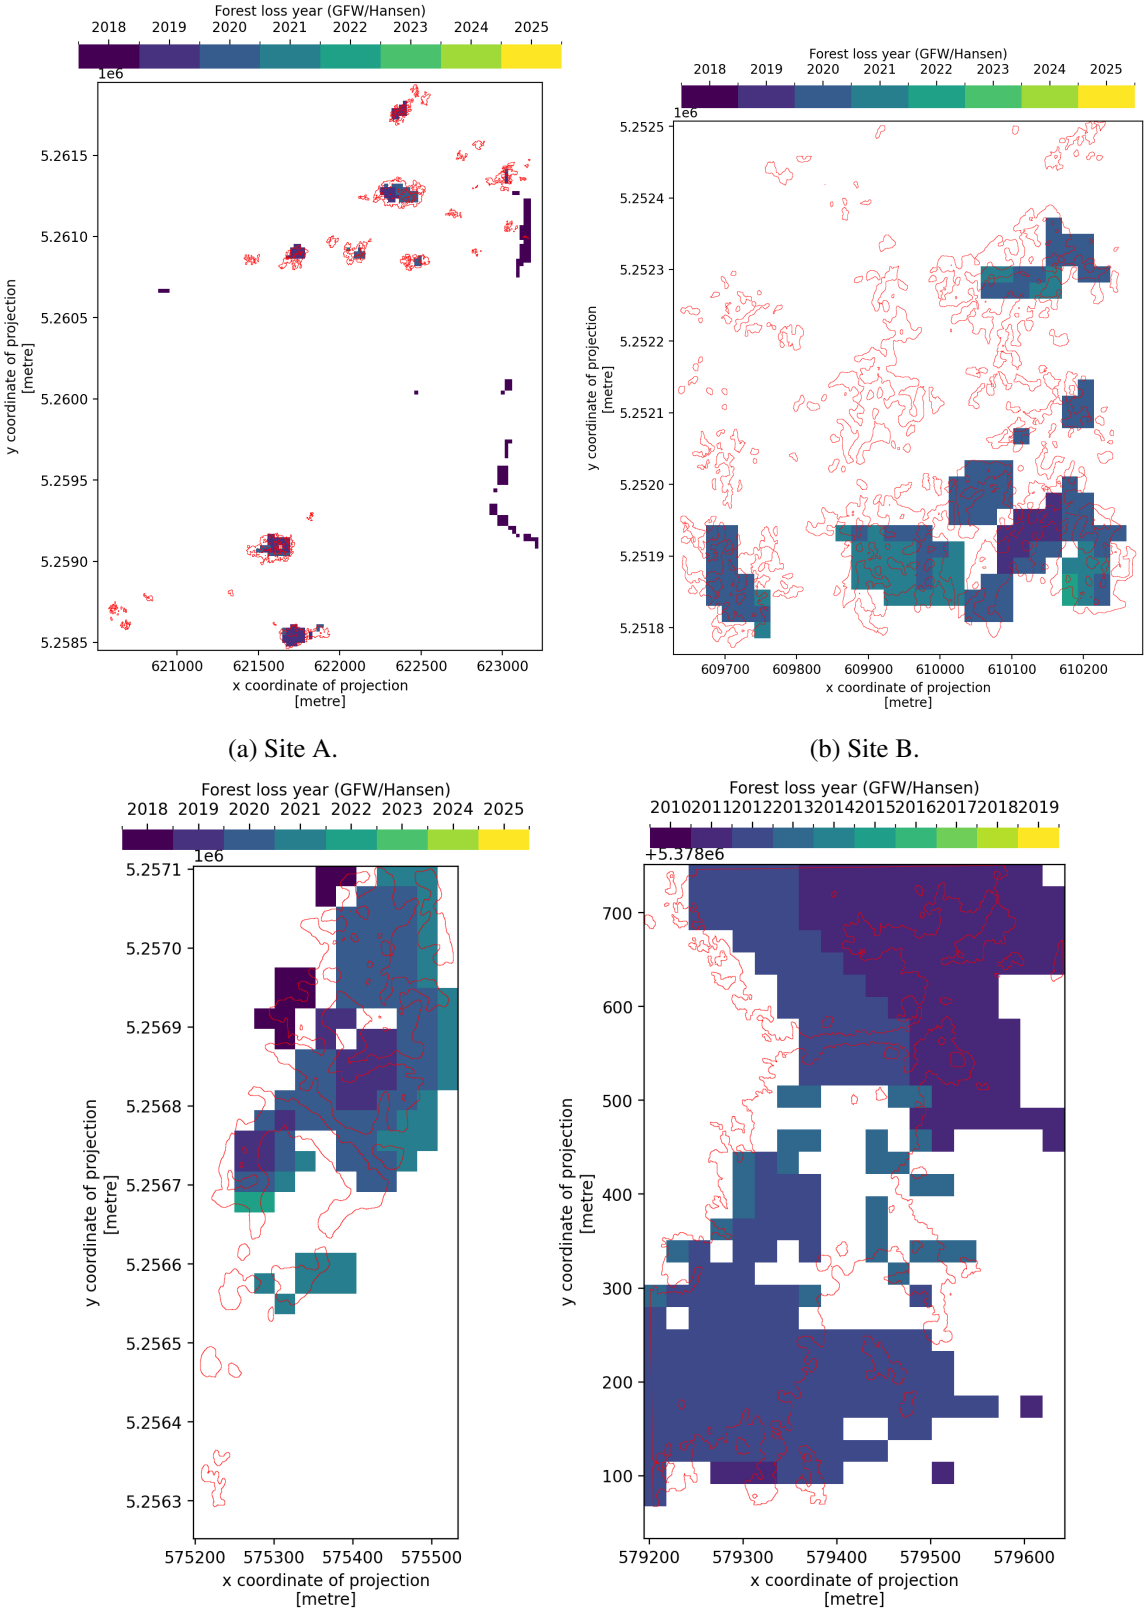

(a) Site A.

(b) Site B.

(c) Site C.

(d) Site D.

Figure S5: Case Study 3: Identical visualization to Supplementary Figure S2 but disturbance map was obtained from GlobalForestWatch (M. C. Hansen et al. 2013).

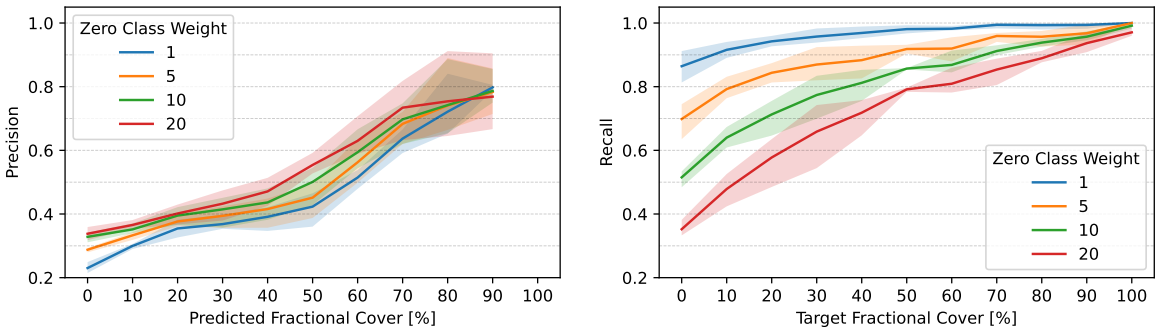

(a) Fractional deadwood cover.

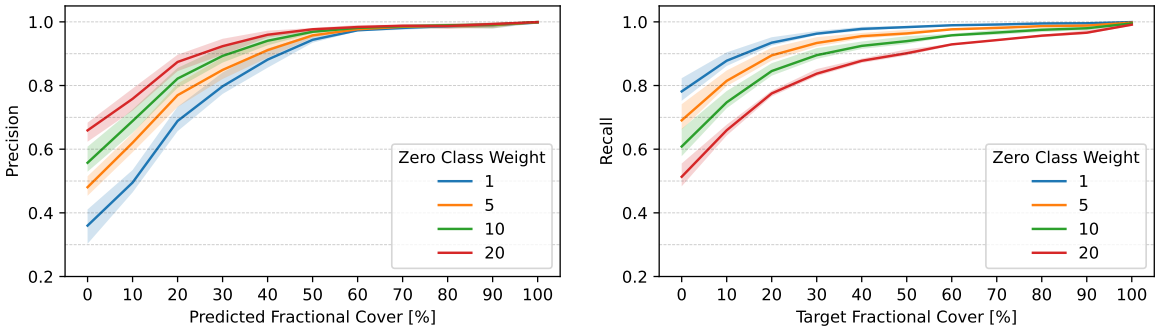

(b) Fractional forest cover.

Figure S6: Precision and recall for across predicted and target fractional cover as a function of applying extra weight in the loss function to cases with exactly 0% fractional deadwood cover. The displayed range on the y-axis is the respective minimum and maximum across all folds.

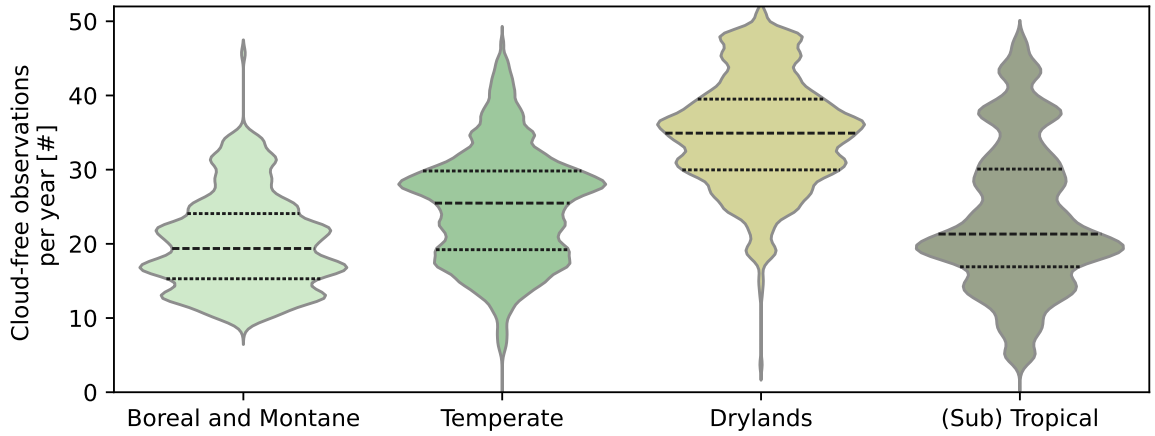

Figure S7: Violinplot of cloud-free and snow free observations per year across biome groups, where plot width depicts kernel density estimate. Dashed lines corresponds to median (thick dashed) and 25th and 75th quantile (thin dashed).

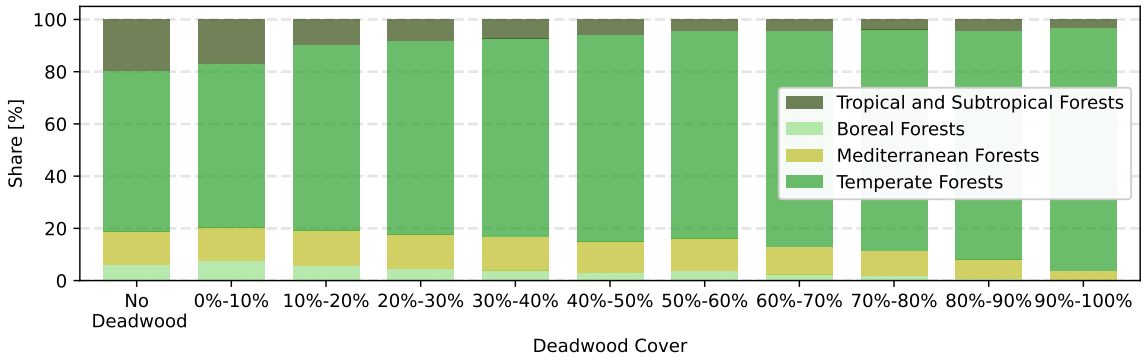

Figure S8: Distribution across biomes for each 10% share group of standing deadwood in a training/validation epoch after spatial sampling caps have been applied. Averaged across folds.

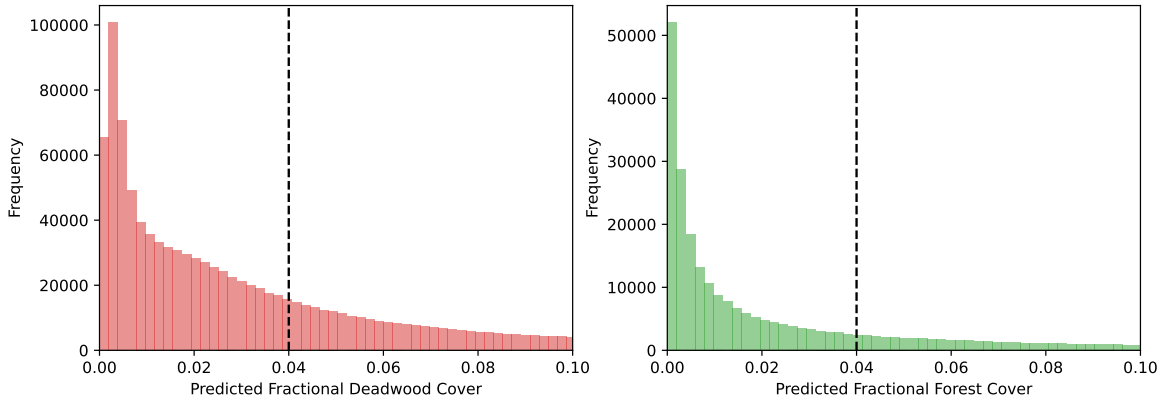

Figure S9: Frequency of each predicted fractional cover value for samples for where the target cover is 0%. Vertical line shows postprocessing cutoff below which predictions are clipped to zero during inference. Displayed distribution is across all values in the cross-validation sets.

A.5 Number of available timesteps per biome per year

A.6 Epoch-level biome/fractional cover distribution

The implemented sampling scheme (see Section 2.2.5) yields a distribution that is largely dominated by Temperate Forests (Supplementary Figure S8), even more for high deadwood shares.

A.7 Location Details, Case Study 3

A.8 Predicted fractional cover distribution and clipping

Provides basis for clipping fractional cover values below 4% to 0% to be able to compute binary classification metrics and avoid overestimation in later spatial aggregates.

Table 3: Centroid locations of each plot in forest inventory comparison in Section 4.3

Plot Name	Number	Centroid
Asiago	1	45.9384°N 11.6252°E
Asiago	2	45.9377°N 11.6223°E
Asiago	3	45.9336°N 11.5993°E
Asiago	4	45.9375°N 11.6158°E
Asiago	5	45.9389°N 11.6145°E
Asiago	9	45.9362°N 11.6110°E
Asiago	10	45.9367°N 11.6048°E
Asiago	13	45.9333°N 11.6041°E
Asiago	14	45.9342°N 11.6171°E
Ciapela	1	46.4204°N 11.9337°E
Ciapela	2	46.4222°N 11.9261°E
Ciapela	3	46.4235°N 11.9286°E
Ciapela	4	46.4217°N 11.9310°E
Ciapela	5	46.4233°N 11.9255°E
Ciapela	6	46.4189°N 11.9361°E
Visdende	1	46.6195°N 12.6587°E
Visdende	2	46.6140°N 12.6528°E
Visdende	3	46.6133°N 12.6559°E
Visdende	5	46.6212°N 12.6590°E
Visdende	6	46.6180°N 12.6581°E

A.9 Example of multi-year bark beetle outbreak tracking

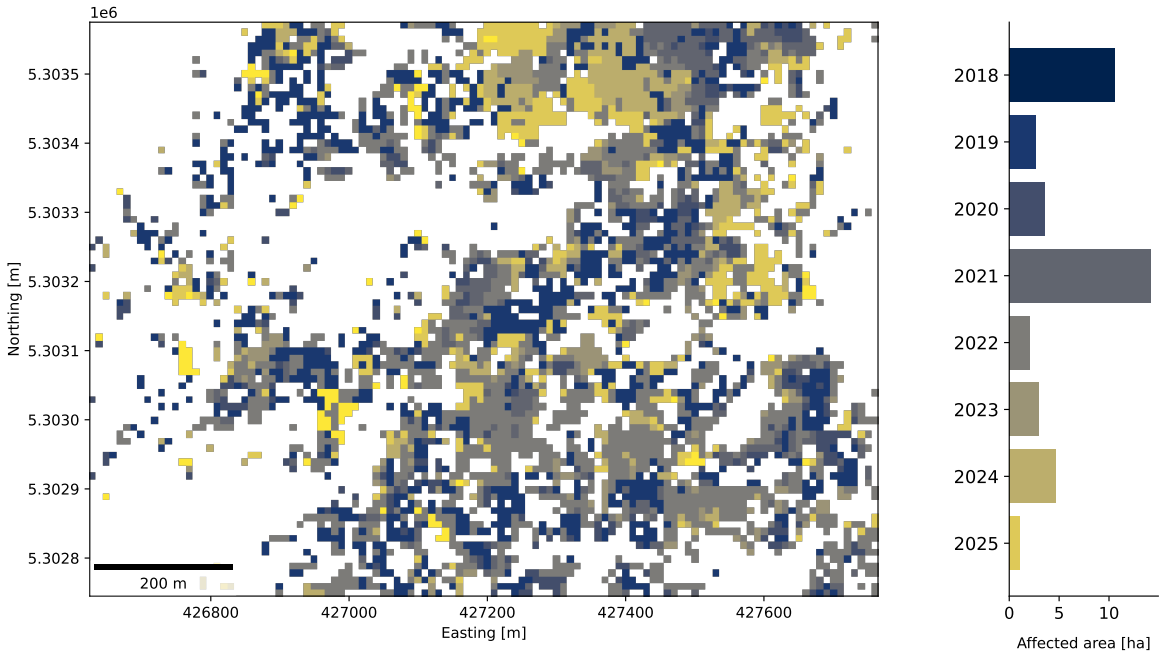

Figure S10: Bark beetle induced mortality spreading from 2018 to 2025 on the Feldberg, Germany (47°52′47.3″N 8°01′01.7″E). Coordinates are in UTM Zone 32N. The color indicates the earliest year where the respective pixel reported above 50% fractional deadwood cover. The barplot on the right displays the number of pixels that were detect first within a specific year.

References

Alix-Garcia, J., & Millimet, D. L. (2023). Remotely incorrect? accounting for nonclassical measurement error in satellite data on deforestation. *Journal of the Association of Environmental and Resource Economists*, *10*(5), 1335–1367.

Allen, C. D., Breshears, D. D., & McDowell, N. G. (2015). On underestimation of global vulnerability to tree mortality and forest die-off from hotter drought in the anthropocene. *Ecosphere*, *6*(8), 1–55.

Anders, T., Hetzer, J., Knapp, N., Forrest, M., Langan, L., Tölle, M. H., Wellbrock, N., & Hickler, T. (2025). Modelling past and future impacts of droughts on tree mortality and carbon storage in norway spruce stands in germany. *Ecological Modelling*, *501*, 110987.

Aybar, C. [Cesar], Ysuhuaylas, L., Loja, J., Gonzales, K., Herrera, F., Bautista, L., Yali, R., Flores, A., Diaz, L., Cuenca, N., et al. (2022). Cloudsen12, a global dataset for semantic understanding of cloud and cloud shadow in sentinel-2. *Scientific data*, *9*(1), 782.

Bauer-Marschallinger, B., & Falkner, K. (2023). Wasting petabytes: A survey of the sentinel-2 utm tiling grid and its spatial overhead. *ISPRS Journal of Photogrammetry and Remote Sensing*, *202*, 682–690.

Bugmann, H., Seidl, R., Hartig, F., Bohn, F., Bruna, J., Cailleret, M., François, L., Heinke, J., Henrot, A.-J., Hickler, T., et al. (2019). Tree mortality submodels drive simulated long-term forest dynamics: Assessing 15 models from the stand to global scale. *Ecosphere*, *10*(2), e02616.

Cheng, Y., Oehmcke, S., Brandt, M., Rosenthal, L., Das, A., Vrieling, A., Saatchi, S., Wagner, F., Mugabowindekwe, M., Verbruggen, W., et al. (2024). Scattered tree death contributes to substantial forest loss in california. *Nature communications*, *15*(1), 641.

Choi, W., Ryu, Y., Kong, J., Jeong, S., & Lee, K. (2025). Evaluation of spatial and temporal variability in sentinel-2 surface reflectance on a rice paddy landscape. *Agricultural and Forest Meteorology*, *363*, 110401.

Copernicus. (n.d.). Level-2a dataset specification. https://sentinels.copernicus.eu/sentinel-data-access/sentinel-products/sentinel-2-data-products/collection-0-level-2a

Coppin, P. R., & Bauer, M. E. (1996). Digital change detection in forest ecosystems with remote sensing imagery. *Remote sensing reviews*, *13*(3-4), 207–234.

D'Andrea, E., Rezaie, N., Battistelli, A., Gavrichkova, O., Kuhlmann, I., Matteucci, G., Moscatello, S., Proietti, S., Scartazza, A., Trumbore, S., et al. (2019). Winter's bite: Beech trees survive complete defoliation due to spring late-frost damage by mobilizing old c reserves. *New Phytologist*, *224*(2), 625–631.

Dick, M. A., Williams, N. M., Bader, M. K.-F., Gardner, J. F., & Bulman, L. S. (2014). Pathogenicity of phytophthora pluvialis to pinus radiata and its relation with red needle cast disease in new zealand. *New Zealand Journal of Forestry Science*, *44*(1), 6.

(ESA), S. (2023, December 11). *Data quality report – sentinel-2 l1c msi, december 2023* (tech. rep. No. Ref. OMPC.CS.DQR.01.11-2023, Issue 94.0). Copernicus Space Component Sentinel Optical Imaging Mission Performance Cluster Service.

Espírito-Santo, F. D., Gloor, M., Keller, M., Malhi, Y., Saatchi, S., Nelson, B., Junior, R. C. O., Pereira, C., Lloyd, J., Frolking, S. [Steve], et al. (2014). Size and frequency of natural forest disturbances and the amazon forest carbon balance. *Nature communications*, *5*(1), 1–6.

European Space Agency. (2021). Copernicus Sentinel-2: Major Products Upgrade Upcoming [Accessed: 2026-01-20].

Farebrother, J., Orbay, J., Vuong, Q., Taïga, A. A., Chebotar, Y., Xiao, T., Irpan, A., Levine, S., Castro, P. S., Faust, A., et al. (2024). Stop regressing: Training value functions via classification for scalable deep rl. *arXiv preprint arXiv:2403.03950*.

Fassnacht, F. E., Mager, C., Waser, L. T., Kanjir, U., Schäfer, J., Buhvald, A. P., Shafeian, E., Schiefer, F., Stančič, L., Immitzer, M., et al. (2025). Forest practitioners' requirements for remote sensing-based canopy height, wood-volume, tree species, and disturbance products. *Forestry: An International Journal of Forest Research*, *98*(2), 233–252.

Frolking, S. [Stephen], Palace, M. W., Clark, D., Chambers, J. Q., Shugart, H., & Hurtt, G. C. (2009). Forest disturbance and recovery: A general review in the context of spaceborne remote sensing of impacts on aboveground biomass and canopy structure. *Journal of Geophysical Research: Biogeosciences*, *114*(G2).

Garbarino, M., Morresi, D., Anselmetto, N., & Weisberg, P. J. (2023). Treeline remote sensing: From tracking treeline shifts to multi-dimensional monitoring of ecotonal change. *Remote Sensing in Ecology and Conservation*, *9*(6), 729–742.

Hammond, W. M., Williams, A. P., Abatzoglou, J. T., Adams, H. D., Klein, T., López, R., Sáenz-Romero, C., Hartmann, H., Breshears, D. D., & Allen, C. D. (2022). Global field observations of tree die-off reveal hotter-drought fingerprint for earth's forests. *Nature Communications*, *13*(1), 1761.

Hansen, M. C. [M. C.], Potapov, P. V., Moore, R., Hancher, M., Turubanova, S. A., Tyukavina, A., Thau, D., Stehman, S. V., Goetz, S. J., Loveland, T. R., Kommareddy, A., Egorov, A., Chini, L., Justice, C. O., & Townshend, J. R. G. (2013). High-resolution global maps of 21st-century forest cover change. *Science*, *342*(6160), 850–853. https://doi.org/10.1126/science.1244693

Hartmann, H., Schuldt, B., Sanders, T. G., Macinnis-Ng, C., Boehmer, H. J., Allen, C. D., Bolte, A., Crowther, T. W., Hansen, M. C. [Matthew C], Medlyn, B. E., et al. (2018). Monitoring global tree mortality patterns and trends. report from the vw symposium 'crossing scales and disciplines to identify global trends of tree mortality as indicators of forest health'. *New Phytologist*, *217*(3), 984–987.

Hartmann, H., Bastos, A., Das, A. J., Esquivel-Muelbert, A., Hammond, W. M., Martínez-Vilalta, J., McDowell, N. G., Powers, J. S., Pugh, T. A., Ruthrof, K. X., et al. (2022). Climate change risks to global forest health: Emergence of unexpected events of elevated tree mortality worldwide. *Annual Review of Plant Biology*, *73*, 673–702.

Hartmann, H., Battisti, A., Brockerhoff, E. G., Bełka, M., Hurling, R., Jactel, H., Oliva, J., Rousselet, J., Terhonen, E., Ylioja, T., et al. (2025). European forests are under increasing pressure from global change-driven invasions and accelerating epidemics by insects and diseases. *Journal für Kulturpflanzen*, *77*(02), 6–24.

Heinken, T., Blumröder, J. S., Adhikari, Y., Balthasar, C., Binder, A., Birkhofer, K., Bischoff, A., Brosinsky, A., Bubner, B., Clerc, D., et al. (2024). Pyrophob-a post-fire ecosystem research project to inform management for resilient forest development. *Tuexenia*, *44*.

Hlásny, T., König, L., Krokene, P., Lindner, M., Montagné-Huck, C., Müller, J., Qin, H., Raffa, K. F., Schelhaas, M.-J., Svoboda, M., et al. (2021). Bark beetle outbreaks in europe: State of knowledge and ways forward for management. *Current Forestry Reports*, *7*(3), 138–165.

Hlásny, T., Perunová, M., Modlinger, R., Blake, M., Brazaitis, G., Csóka, G., de Groot, M., Duduman, M.-L., Faccoli, M., Georgieva, M., et al. (2025). Perspectives: State of national forest damage survey programmes in europe and ways toward improved harmonization and data sharing. *Forest Ecology and Management*, *597*, 123111.

Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the modis vegetation indices. *Remote sensing of environment*, *83*(1-2), 195–213.

Imani, E., Luedemann, K., Scholnick-Hughes, S., Elelimy, E., & White, M. (2024). Investigating the histogram loss in regression. *arXiv preprint arXiv:2402.13425*.

Imani, E., & White, M. (2018). Improving regression performance with distributional losses. *International conference on machine learning*, 2157–2166.

International Tree Mortality Network, Senf, C., Esquivel-Muelbert, A., Pugh, T. A., Anderegg, W. R., Anderson-Teixeira, K. J., Arellano, G., Beloiu Schwenke, M., Bentz, B. J., Boehmer, H. J., et al. (2025). Towards a global understanding of tree mortality. *New Phytologist*.

Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proceedings of the 32nd International Conference on Machine Learning (ICML)*.

Jouy, F., Schüle, M., Adhikari, Y., Binder, A., Clerc, D., Gerwin, W., Heinken, T., Raab, T., Repmann, F., Rönnefarth, S., et al. (2025). Factors impacting the variability of post-fire forest regeneration in central european pine plantations. *Restoration Ecology*, *33*(4), e70017.

Jucker, T., Fischer, F. J., Chave, J., Coomes, D. A., Caspersen, J., Ali, A., Panzou, G. J. L., Feldpausch, T. R., Falster, D., Usoltsev, V. A., Jackson, T. D., Adu-Bredu, S., Alves, L. F., Aminpour, M., Angoboy, I. B., Anten, N. P. R., Antin, C., Askari, Y., Muñoz, R., ... Zimmermann, N. E. (2025). The global spectrum of tree crown architecture. *Nature communications*, *16*(1), 4876.

Junttila, S., Blomqvist, M., Laukkanen, V., Heinaro, E., Polvivaara, A., O'Sullivan, H., Yrttimaa, T., Vastaranta, M., & Peltola, H. (2024). Significant increase in forest canopy mortality in boreal forests in southeast finland. *Forest Ecology and Management*, *565*, 122020. https://doi.org/10.1016/j.foreco.2024.122020

Kattenborn, T., Schiefer, F., Frey, J., Feilhauer, H., Mahecha, M. D., & Dormann, C. F. (2022). Spatially autocorrelated training and validation samples inflate performance assessment of convolutional neural networks. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, *5*, 100018.

Khatri-Chhetri, P., van Wagtendonk, L., Hendryx, S. M., & Kane, V. R. (2024). Enhancing individual tree mortality mapping: The impact of models, data modalities, and classification taxonomy. *Remote Sensing of Environment*, *300*, 113914.

Körner, C., Möhl, P., & Hiltbrunner, E. (2023). Four ways to define the growing season. *Ecology Letters*, *26*(8), 1277–1292.

Korznikov, K., Kislov, D., Doležal, J., & Altman, J. (2025). Hidden forest loss: Challenges in detecting wind-and insect-driven forest disturbances with global forest change landsat-based products in mixed southern boreal forests. *Remote Sensing Applications: Society and Environment*, 101798.

Korznikov, K., Kislov, D., Doležal, J., Petrenko, T., & Altman, J. (2022). Tropical cyclones moving into boreal forests: Relationships between disturbance areas and environmental drivers. *Science of the Total Environment*, *844*, 156931.

Lange, M., Preidl, S., Reichmuth, A., Heurich, M., & Doktor, D. (2024). A continuous tree species-specific reflectance anomaly index reveals declining forest condition between 2016 and 2022 in germany. *Remote Sensing of Environment*, *312*, 114323.

Liu, S., Brandt, M., Nord-Larsen, T., Chave, J., Reiner, F., Lang, N., Tong, X., Ciais, P., Igel, C., Pascual, A., et al. (2023). The overlooked contribution of trees outside forests to tree cover and woody biomass across europe. *Science Advances*, *9*(37), eadh4097.

Loshchilov, I., & Hutter, F. (2019). Decoupled weight decay regularization. *International Conference on Learning Representations (ICLR)*.

Mahecha, M. D., Gans, F., Brandt, G., Christiansen, R., Cornell, S. E., Fomferra, N., Kraemer, G., Peters, J., Bodesheim, P., Camps-Valls, G., et al. (2020). Earth system data cubes unravel global multivariate dynamics. *Earth System Dynamics*, *11*(1), 201–234.

Migliavacca, M., Grassi, G., Bastos, A., Cecchirini, G., Ciais, P., Janssens-Maenhout, G., Lugato, E., Mahecha, M. D., Novick, K. A., Peñuelas, J., Pilli, R., Reichstein, M., Avitabile, V., Beck, P. S. A., Barredo, J. L., Forzieri, G., Herold, M., Korosuo, A., Mansuy, N., … Cescatti, A. (2025a). Securing the forest carbon sink for the European Union's climate ambition [Accepted: 28 March 2025; Received: 4 September 2024]. *Nature*. https://doi.org/10.1038/s41586-025-08967-3

Migliavacca, M., Grassi, G., Bastos, A., Ceccherini, G., Ciais, P., Janssens-Maenhout, G., Lugato, E., Mahecha, M. D., Novick, K. A., Peñuelas, J., et al. (2025b). Securing the forest carbon sink for the european union's climate ambition. *Nature*, *643*(8074), 1203–1213.

Mnih, V., & Hinton, G. E. (2012). Learning to label aerial images from noisy data. *Proceedings of the 29th International conference on machine learning (ICML-12)*, 567–574.

Möhring, J., Kattenborn, T., Mahecha, M. D., Cheng, Y., Beloiu Schwenke, M., Cloutier, M., Denter, M., Frey, J., Gassilloud, M., Göritz, A., Hempel, J., Horion, S., Jucker, T., Junttila, S., Khatri-Chhetri, P., Korznikov, K., Kruse, S., Laliberté, E., Maroschek, M., … Mosig, C. (2025). Global, multi-scale standing deadwood segmentation in centimeter-scale aerial images. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, *18*, 100104. https://doi.org/https://doi.org/10.1016/j.ophoto.2025.100104

Montero, D., Mahecha, M. D., Martinuzzi, F., Aybar, C., Klosterhalfen, A., Knohl, A., Anaya, J., Mosig, C., & Wieneke, S. (2025). Transformers vs. recurrent models for estimating forest gross primary production. https://doi.org/10.48550/ARXIV.2511.11880

Montero, D., Mahecha, M. D., Martinuzzi, F., Aybar, C., Klosterhalfen, A., Knohl, A., Koebsch, F., Anaya, J., & Wieneke, S. (2024a). Recurrent neural networks for modelling gross primary production. *IGARSS 2024 - 2024 IEEE International Geoscience and Remote Sensing Symposium*, 4214–4217. https://doi.org/10.1109/igarss53475.2024.10640715

Montero, D., Kraemer, G., Anghelea, A., Aybar, C. [César], Brandt, G., Camps-Valls, G., Cremer, F., Flik, I., Gans, F., Habershon, S., et al. (2024b). Earth system data cubes: Avenues for advancing earth system research. *Environmental Data Science*, *3*, e27.

Mosig, C., Vajna-Jehle, J., Mahecha, M. D., Cheng, Y., Hartmann, H., Montero, D., Junttila, S., Horion, S., Schwenke, M. B., Koontz, M. J., Maulud, K. N. A., Adu-Bredu, S., Al-Halbouni, D., Ali, M., Allen, M., Altman, J., Amorós, L., Angiolini, C., Astrup, R., … Kattenborn, T. (2026). Deadtrees.earth - an open-access and interactive database for centimeter-scale aerial imagery to uncover global tree mortality dynamics. *Remote Sensing of Environment*, *332*, 115027. https://doi.org/https://doi.org/10.1016/j.rse.2025.115027

Nair, V., & Hinton, G. E. (2010). Rectified linear units improve restricted boltzmann machines. *Proceedings of the 27th International Conference on Machine Learning (ICML)*.

Olson, D. M., Dinerstein, E., Wikramanayake, E. D., Burgess, N. D., Powell, G. V. N., Underwood, E. C., D'amico, J. A., Itoua, I., Strand, H. E., Morrison, J. C., Loucks, C. J., Allnutt, T. F., Ricketts, T. H., Kura, Y., Lamoreux, J. F., Wettengel, W. W., Hedao, P., & Kassem, K. R. (2001). Terrestrial Ecoregions of the World: A New Map of Life on Earth: A new global map of terrestrial ecoregions provides an innovative tool for conserving biodiversity. *BioScience*, *51*(11), 933–938. https://doi.org/10.1641/0006-3568(2001)051[0933:TEOTWA]2.0.CO;2

OpenAerialMap Community. (2024). Openaerialmap [Accessed: 2026-01-20].

Park, J. Y., Lichstein, J. W., Muller-Landau, H. C., Wright, S. J., & Bohlman, S. A. (2026). Interspecific variation in leaf phenology and its relationship with plant traits in a seasonal tropical forest. *Functional Ecology*, *40*(2), 433–450.

Pickens, A. H., Hansen, M. C. [Matthew C], Song, Z., Poulson, A., Komarova, A., Baggett, A., Kerr, T., Mikus, A., Ortiz Dominguez, C., Tyukavina, A., et al. (2025). Rapid monitoring of global land change. *Nature Communications*, *16*(1), 8948.

Ploton, P., Mortier, F., Réjou-Méchain, M., Barbier, N., Picard, N., Rossi, V., Dormann, C., Cornu, G., Viennois, G., Bayol, N., et al. (2020). Spatial validation reveals poor predictive performance of large-scale ecological mapping models. *Nature communications*, *11*(1), 4540.

Reiche, J., Mullissa, A., Slagter, B., Gou, Y., Tsendbazar, N.-E., Odongo-Braun, C., Vollrath, A., Weisse, M. J., Stolle, F., Pickens, A., et al. (2021). Forest disturbance alerts for the congo basin using sentinel-1. *Environmental Research Letters*, *16*(2), 024005.

Rußwurm, M., & Körner, M. (2020). Self-attention for raw optical satellite time series classification. *ISPRS journal of photogrammetry and remote sensing*, *169*, 421–435.

Schiefer, F., Schmidtlein, S., Frick, A., Frey, J., Klinke, R., Zielewska-Büttner, K., Junttila, S., Uhl, A., & Kattenborn, T. (2023). Uav-based reference data for the prediction of fractional cover of standing deadwood from sentinel time series. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, *8*, 100034.

Schiefer, F., Schmidtlein, S., Hartmann, H. H., Schnabel, F., & Kattenborn, T. (2024). Large-scale remote sensing reveals that tree mortality in germany appears to be greater than previously expected. *Under Review*.

Schiller, C., Költzow, J., Schwarz, S., Schiefer, F., & Fassnacht, F. E. (2024). Forest disturbance detection in central europe using transformers and sentinel-2 time series. *Remote Sensing of Environment*, *315*, 114475.

Schmehl, M.-T., Adhikari, Y., Balthasar, C., Binder, A., Clerc, D., Dobkowitz, S., Gerwin, W., Günther, K., Hartong, H., Heinken, T., et al. (2025). Monitoring abiotic and biotic parameters of forest regrowth under different management regimes on former wildfire sites in northeastern germany–data from the pyrophob project. *Earth System Science Data*, *17*(11), 6295–6313.

Schmitt, M., Prexl, J., Ebel, P., Liebel, L., & Zhu, X. X. (2020). Weakly supervised semantic segmentation of satellite images for land cover mapping–challenges and opportunities. *arXiv preprint arXiv:2002.08254*.

Schulze, E.-D. (2006). Biological control of the terrestrial carbon sink. *Biogeosciences*, *3*(2), 147–166. https://doi.org/10.5194/bg-3-147-2006

Schwartz, M., Fogel, F., Besic, N., Robert, D., Geist, L., Renaud, J.-P., Monnet, J.-M., Mosig, C., Vega, C., d'Aspremont, A., et al. (2025). Formspot: A decade of tree-level, country-scale forest monitoring. *arXiv preprint arXiv:2512.17021*.

Schwarz, S., Werner, C., Fassnacht, F. E., & Ruehr, N. K. (2024). Forest canopy mortality during the 2018-2020 summer drought years in central europe: The application of a deep learning approach on aerial images across luxembourg. *Forestry: An International Journal of Forest Research*, *97*(3), 376–387.

Senf, C., Pflugmacher, D., Zhiqiang, Y., Sebald, J., Knorn, J., Neumann, M., Hostert, P., & Seidl, R. (2018). Canopy mortality has doubled in europe's temperate forests over the last three decades. *Nature Communications*, *9*(1), 4978.

Senf, C., & Seidl, R. (2021). Mapping the forest disturbance regimes of europe. *Nature Sustainability*, *4*(1), 63–70.

Silvestro, R., Rossi, S., Berthiaume, R., Bauce, É., Rademacher, T., & Deslauriers, A. (2025). The impact of forest tent caterpillar defoliation on maple fitness and syrup production. *Forest Ecology and Management*, *595*, 123011.

Sitch, S., Huntingford, C., Gedney, N., Levy, P., Lomas, M., Piao, S., Betts, R., Ciais, P., Cox, P., Friedlingstein, P., et al. (2008). Evaluation of the terrestrial carbon cycle, future plant geography and climate-carbon cycle feedbacks using five dynamic global vegetation models (dgvms). *Global change biology*, *14*(9), 2015–2039.

Söchting, M., Scheuermann, G., Montero, D., & Mahecha, M. D. (2025). Interactive earth system data cube visualization in jupyter notebooks. *Big Earth Data*, *9*(2), 321–335.

Song, Z., Hansen, M., Pickens, A. H., Poulson, A. J., Baggett, A., Komarova, A., Lima, A., Mikus, A., Ortiz Dominguez, C., & Kerr, T. (2024). Dist-alert: Near-real time monitoring system of land cover change. *AGU Fall Meeting Abstracts*, *2024*, GC41B–02.

The International Tree Mortality Network, Senf, C., Esquivel-Muelbert, A., Pugh, T. A. M., Anderegg, W. R. L., Anderson-Teixeira, K. J., Arellano, G., Beloiu Schwenke, M., Bentz, B. J., Boehmer, H. J., Bond-Lamberty, B., Bordin, K., Boson De Castro-Faria, A., Brearley, F. Q., Bussotti, F., Cailleret, M., Camarero, J. J., Chirici, G., Costa, F. R., … Zuleta, D. (2024). Towards a global understanding of tree mortality. *Under Review.*

Trumbore, S., Brando, P., & Hartmann, H. (2015). Forest health and global change. *Science*, *349*(6250), 814–818.

Tseng, G., Cartuyvels, R., Zvonkov, I., Purohit, M., Rolnick, D., & Kerner, H. (2023). Lightweight, pre-trained transformers for remote sensing timeseries. *arXiv preprint arXiv:2304.14065*.

U.S. Forest Service. (2025). Insect and disease survey (ids) — forest health protection [Accessed via Ïnsect & Disease Explorerät USDA LCMS; data publicly available under Creative Commons Attribution license]. https://apps.fs.usda.gov/lcms-viewer/IDS.html

Vancutsem, C., Achard, F., Pekel, J.-F., Vieilledent, G., Carboni, S., Simonetti, D., Gallego, J., Aragao, L. E., & Nasi, R. (2021). Long-term (1990–2019) monitoring of forest cover changes in the humid tropics. *Science advances*, *7*(10), eabe1603.

Veitch-Michaelis, J., Cottam, A., Schweizer, D., Broadbent, E., Dao, D., Zhang, C., Almeyda Zambrano, A., & Max, S. (2024). Oam-tcd: A globally diverse dataset of high-resolution tree

cover maps. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, & C. Zhang (Eds.), *Advances in neural information processing systems* (pp. 49749–49767, Vol. 37). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2024/file/58efdd77196fa8159062afa0408245da-Paper-Datasets_and_Benchmarks_Track.pdf

Viana-Soto, A., & Senf, C. (2025). The european forest disturbance atlas: A forest disturbance monitoring system using the landsat archive. *Earth System Science Data*, *17*(6), 2373–2404.

Watt, M. S., Holdaway, A., Watt, P., Pearse, G. D., Palmer, M. E., Steer, B. S. C., Camarretta, N., McLay, E., & Fraser, S. (2024). Early prediction of regional red needle cast outbreaks using climatic data trends and satellite-derived observations. *Remote Sensing*, *16*(8). https://doi.org/10.3390/rs16081401

White, J. C., Wulder, M. A., Hermosilla, T., Coops, N. C., & Hobart, G. W. (2017). A nationwide annual characterization of 25 years of forest disturbance and recovery for canada using landsat time series. *Remote Sensing of Environment*, *194*, 303–321.

Wijas, B. J., Allison, S. D., Austin, A. T., Cornwell, W. K., Cornelissen, J. H. C., Eggleton, P., Fraver, S., Ooi, M. K., Powell, J. R., Woodall, C. W., et al. (2024). The role of deadwood in the carbon cycle: Implications for models, forest management, and future climates. *Annual Review of Ecology, Evolution, and Systematics*, *55*.

Winter, C., Mueller, S., Kattenborn, T., Stahl, K., Szillat, K., Weiler, M., & Schnabel, F. (2024). Forest dieback in drinking water protection areas–a hidden threat to water quality. *bioRxiv*, 2024–08.

Xu, Y., Ciais, P., Santoro, M., Bourgoin, C., Ritter, F., Pellissier-Tanon, A., Feng, Y., Zhou, C., He, G., Heinrich, V., Besnard, S., Robinson, N., Cook-Patton, S. C., Chave, J., Aragao, L. E. O. C., Ometto, J. P., Bowring, S. P. K., Fayad, I., Zhu, L., … Li, W. (2026). Small persistent humid forest clearings drive tropical forest biomass losses. *Nature*, *649*(8096), 375–380. https://doi.org/10.1038/s41586-025-09870-7

Yan, Y., Piao, S., Hammond, W. M., Chen, A., Hong, S., Xu, H., Munson, S. M., Myneni, R. B., & Allen, C. D. (2024). Climate-induced tree-mortality pulses are obscured by broad-scale and long-term greening. *Nature Ecology & Evolution*, *8*(5), 912–923.

Yuan, Y., Lin, L., Liu, Q., Hang, R., & Zhou, Z.-G. (2022). Sits-former: A pre-trained spatio-spectral-temporal representation model for sentinel-2 time series classification. *International Journal of Applied Earth Observation and Geoinformation*, *106*, 102651.

Zhang, H. K., Camps-Valls, G., Liang, S., Tuia, D., Pelletier, C., & Zhu, Z. (2025). Preface: Advancing deep learning for remote sensing time series data analysis.

Zhang, X., Friedl, M., & Henebry, G. (2020). VIIRS/NPP Land Cover Dynamics Yearly L3 Global 500m SIN Grid V001 [[Data set]]. https://doi.org/10.5067/VIIRS/VNP22Q2.001

Zhang, Y., Ling, F., Wang, X., Foody, G. M., Boyd, D. S., Li, X., Du, Y., & Atkinson, P. M. (2021). Tracking small-scale tropical forest disturbances: Fusing the landsat and sentinel-2 data record. *Remote Sensing of Environment*, *261*, 112470.

Zhou, Q., Neigh, C. S., Ju, J., Dabney, P., Cook, B., Zhu, Z., Crawford, C. J., Gascon, F., Strobl, P., & Sridhar, M. (2025). Towards seamless global 30-meter terrestrial monitoring: Evaluating 2022 cloud free coverage of harmonized landsat and sentinel-2 (hls) v2. 0. *IEEE Geoscience and Remote Sensing Letters*.

Zhu, Z., Wang, S., & Woodcock, C. E. (2015). Improvement and expansion of the fmask algorithm: Cloud, cloud shadow, and snow detection for landsats 4–7, 8, and sentinel 2 images. *Remote sensing of Environment*, *159*, 269–277.

Zhu, Z., & Woodcock, C. E. (2012). Object-based cloud and cloud shadow detection in landsat imagery. *Remote sensing of environment*, *118*, 83–94.