forests

MDPI

Article

Toward 10 m Regional-Scale Time-Series Oil Palm Mapping in Malaysia and Indonesia (2020–2024) Using Sentinel-2 and Noise-Robust Deep Learning from Low-Resolution Historical Maps

Nuttaset Kuapanich [1], Zhiwei Zhang [1], Bohan Shi [2], Jiaying Liu [1], Jiayin Jiang [1], Jiatao Huang [1], Shenghan Tan [1] and Juepeng Zheng [1,3,*]

[1] School of Artificial Intelligence, Sun Yat-sen University, Zhuhai 519082, China; kuapanich@mail2.sysu.edu.cn (N.K.); zhangzhw65@mail2.sysu.edu.cn (Z.Z.); liujy526@mail2.sysu.edu.cn (J.L.); jiangjy79@mail2.sysu.edu.cn (J.J.); huangjt63@mail2.sysu.edu.cn (J.H.); tanshh3@mail2.sysu.edu.cn (S.T.)
[2] Khoury College of Computer Sciences, Northeastern University, Boston, MA 02115, USA; shi.boh@northeastern.edu
[3] National Supercomputing Center in Shenzhen, Shenzhen 518172, China
* Correspondence: zhengjp8@mail.sysu.edu.cn

Abstract

Accurate monitoring of oil palm plantations is important for balancing economic development with environmental conservation in Southeast Asia. However, existing plantation maps often suffer from low spatial resolution and a lack of recent temporal coverage, impeding effective surveillance of rapid land-use changes. In this study, we propose a deep learning framework to generate 10 m resolution oil palm plantation maps for Indonesia and Malaysia from 2020 to 2024, utilizing Sentinel-2 imagery without requiring new manual annotations. To address the resolution mismatch between coarse 100 m historical labels and 10 m imagery, we employ a U-Net architecture optimized with Determinant-based Mutual Information (DMI). This approach effectively mitigates the influence of label noise. We validated our method against 2058 manually verified points, achieving overall accuracies of 70.64%, 63.53%, and 60.06% for the years 2020, 2022, and 2024, respectively. The gradual decline in accuracy with time is consistent with a growing temporal mismatch between the 2016 historical reference labels and the later prediction years. At the regional scale, the mapped oil palm area suggests a peak in 2022 followed by a lower mapped extent in 2024. Land cover transition analysis further indicates exchanges with cropland and flooded vegetation, which should be interpreted together with the reported accuracy and uncertainty. Given the moderate per-year accuracies and the temporal mismatch between the 2016 supervision and the later prediction years, these results should be interpreted as regional-scale indicators rather than pixel-level change maps. The generated maps can support regional monitoring, sustainability assessment, and prioritization of areas for further validation.

Keywords: oil palm mapping; temporal trajectory analysis; deep learning; Sentinel-2; land cover changes

Check for updates

Academic Editor: Aleksandar Dj Valjarević

Received: 14 June 2026
Revised: 9 July 2026
Accepted: 9 July 2026
Published: 13 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Global demand for palm oil has risen sharply, from about 2 million tons in 1970 to nearly 80 million tons in 2020 [1]. Indonesia and Malaysia together account for over 84% of global output and hold two-thirds of the world's oil palm cultivation area, and both have rapidly expanded plantation land to meet this growing market demand. While oil palm provides farmers with higher income and supports rural economic development, this expansion has generated substantial environmental pressure. Both countries have experienced extensive deforestation, with concessions certified by the Roundtable on Sustainable Palm Oil (RSPO) in Indonesia overlapping habitats of critically endangered species and retaining only a small fraction of original forest cover [2], and Malaysia losing nearly 6 million hectares of forest between 2010 and 2018 [3]. Large-scale plantation development also accelerates carbon emissions, particularly in peat-rich regions of Indonesia where land conversion releases large amounts of stored carbon [4,5]. These intertwined economic and ecological dynamics highlight the urgent need for accurate, temporally consistent mapping of oil palm plantations to support sustainable land management and monitoring.

In recent decades, monitoring oil palm plantations has traditionally relied on extensive field-based surveys, requiring substantial labor, financial investment, and logistical effort across large and often inaccessible tropical regions. To reduce this burden, the remote sensing community has increasingly adopted machine learning and, more recently, deep learning techniques to map oil palm distribution using multispectral satellite imagery from platforms such as Sentinel-2 and Landsat [6–8]. However, the performance of these data-driven approaches critically depends on access to large, high-quality annotated datasets, and the cost of manual annotation poses a major bottleneck for scaling them to national or continental extents [9,10]; a detailed review of these mapping methods is provided in Section 2.

A promising alternative to new manual annotation is to reuse existing map products as supervisory signals. Historical oil palm products can be reused as weak supervision, but they introduce two forms of uncertainty: a spatial mismatch between coarse labels and 10 m Sentinel-2 imagery, and a temporal mismatch between historical reference labels and contemporary plantation conditions. Although noisy-label learning strategies have shown promising performance in large-scale land cover classification [11,12], most existing methods are tailored to broad land cover or land use mapping [13–15], where class semantics are relatively generic and spatial patterns are less crop-specific. In contrast, oil palm plantations exhibit distinct phenological cycles and spatial structures that are not explicitly represented in existing global or regional maps, leading to substantial label noise when these products are used as supervision. Consequently, despite substantial progress in remote sensing segmentation and noisy-label learning, accurate large-scale oil palm mapping remains challenging, and current products still suffer from limited classification accuracy. This gap highlights the urgent need for dedicated, noise-aware approaches specifically designed for oil palm plantation detection in complex tropical landscapes.

To the best of our knowledge, available oil palm plantation maps predominantly have spatial resolutions of 30 m and 100 m and cover the period from 2001 to 2019, and country-scale, crop-specific 10 m oil palm maps for Indonesia and Malaysia from 2020 onwards remain scarce. Existing products include annual 100 m oil palm maps for Malaysia and Indonesia from 2001 to 2016 [16], regional and multi-country maps of plantation extent and detection year [6], and high-resolution global oil palm datasets based on optical satellite imagery [8,17–19]. These datasets provide important baselines, but they do not fully resolve the monitoring need addressed here. This positions our work as a dedicated, crop-specific oil palm mapping effort operating at 10 m resolution for this region in the post-2019 period: whereas prior noisy-label remote sensing studies target generic land cover classes, we tailor

a noise-robust framework to the distinct phenological and spatial signature of oil palm. This temporal and spatial resolution gap is particularly significant, as higher-resolution mapping is crucial for accurate plantation boundary delineation and monitoring, as illustrated in Figure 1. Therefore, this study aims to develop comprehensive 10 m resolution oil palm plantation mapping products for Malaysia and Indonesia for the period 2020–2024 without requiring additional manual labeling. Rather than presenting the products as a final pixel-level change-detection solution, we frame them as regional-scale, crop-specific mapping outputs whose temporal patterns should be interpreted together with accuracy assessment and uncertainty analysis.

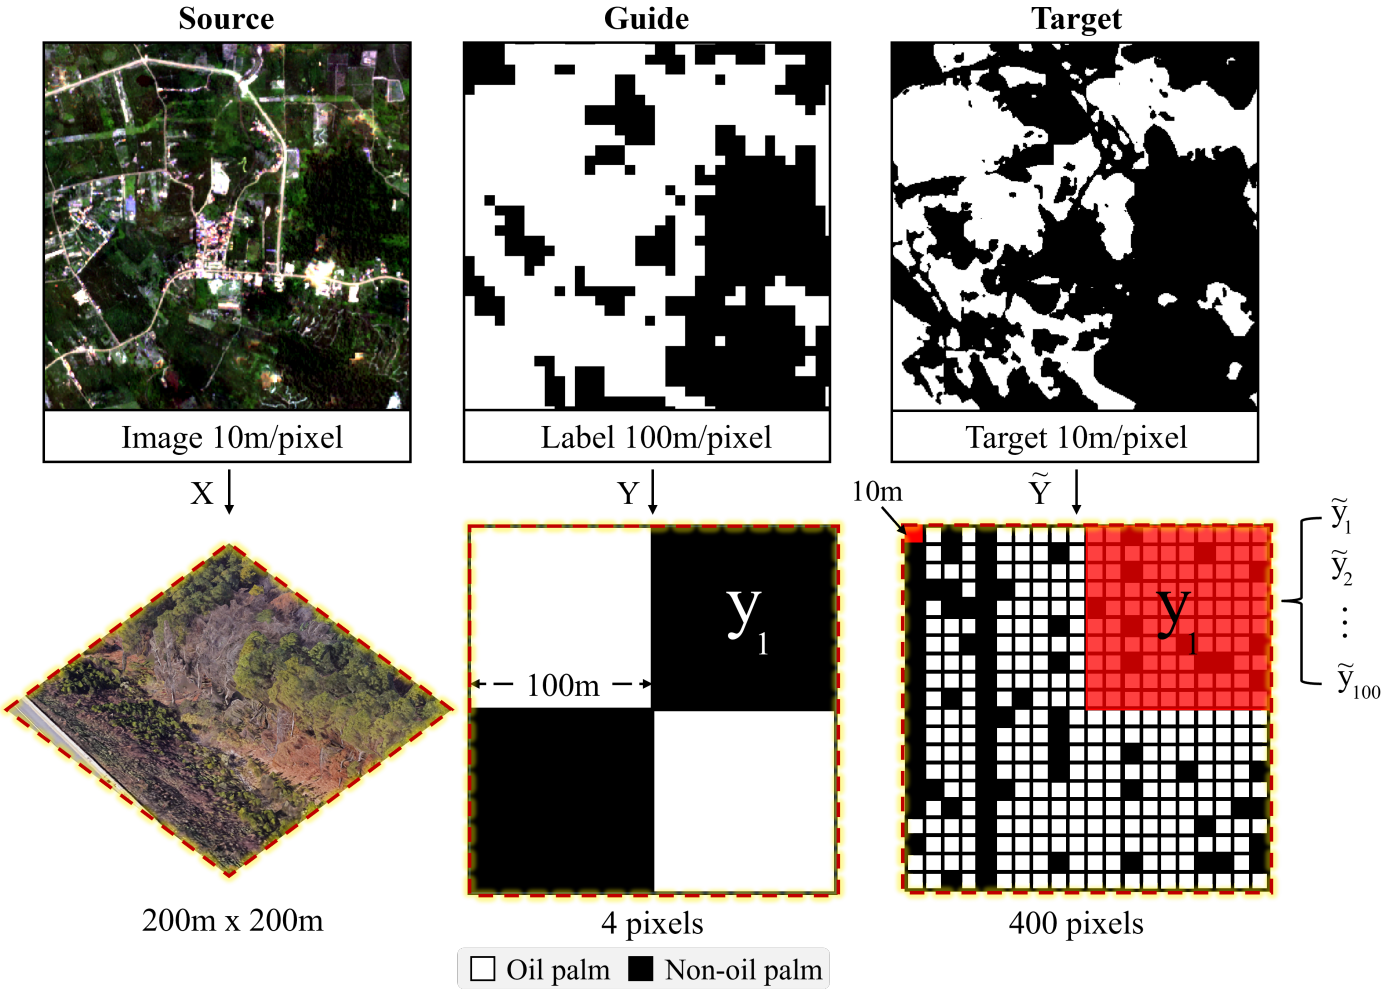


Figure 1. Precision comparison of oil palm plantation maps between 100 m and 10 m resolution.

Therefore, the contributions of this paper mainly include the following three points:

1. We propose a new noise learning method for oil palm plantation mapping that can be completed without any manual labels. Specifically, we employ a U-Net architecture optimized with Determinant-based Mutual Information (DMI).
2. We develop a set of maps for 2020–2024 and manually confirmed 2058 validation points, achieving an Overall Accuracy (OA) of 70.64% in 2020, 63.53% in 2022 and 60.06% in 2024, showing significant improvement compared to other methods
3. We conduct a comprehensive spatiotemporal analysis of oil palm dynamics, revealing a peak in plantation area in 2022 followed by a contraction in 2024. Furthermore, we quantified land cover transitions, identifying that while agricultural crop rotation remains the dominant land cover exchange, there is a distinct and increasing

trend of oil palm encroachment into flooded vegetation (wetland) areas, particularly in Kalimantan.

The remainder of this paper is organized as follows: We review oil palm mapping products, mapping methods, and noisy learning from low-resolution historical labels in Section 2. We describe the study area in Section 3 and introduce the data and methods in Section 4. Section 5 presents the results, including classification accuracy, mapped oil palm areas, regional patterns, and land-cover transitions. We then discuss the comparison with official statistics, the potential drivers of the observed trends, and the reliability and uncertainty of the results in Section 6, and finally summarize this paper in Section 7.

2. Related Work

2.1. Oil Palm Mapping Products

Existing map products provide valuable references but differ in spatial resolution, temporal coverage, class definition, and intended use. The AOPD product provides annual oil palm maps for Malaysia and Indonesia from 2001 to 2016 at a 100 m resolution [16], making it useful as historical supervision but too coarse for 10 m boundary delineation. Multi-country products for Indonesia, Malaysia, and Thailand provide broader regional context [6], while global high-resolution oil palm datasets have mapped closed-canopy plantations and planting years using optical satellite imagery [8,17]. General 10 m land-cover time series, such as the Esri Sentinel-2 product [20,21], are helpful for analyzing surrounding land-cover transitions, but their broad class taxonomy does not explicitly identify oil palm plantations. These datasets provide important baselines; however, none of them offers updated, crop-specific 10 m oil palm maps for Malaysia and Indonesia in the post-2019 period, which is the gap addressed in this study.

2.2. Oil Palm Mapping Methods

The previous work related to oil palm mapping can be divided into three parts: traditional image processing methods, traditional machine learning approaches, and deep learning methods. Traditional image processing methods primarily rely on hand-crafted features and geometric rules, for example template matching [22], which correlates a crown template to find trees, and morphological analysis combined with watershed segmentation [23] to delineate tree crowns. Traditional machine learning approaches have been extensively employed for land-cover classification using multi-spectral low-resolution imagery. These methods encompass random forest [24], XGBoost [25], and decision tree algorithms [26]. However, these approaches are highly sensitive to data quality, as label noise can significantly degrade classification performance and reduce model efficiency [27]. With the advancement of data-driven semantic segmentation techniques, deep learning methodologies have gained considerable prominence [9]. Convolutional neural networks and semantic segmentation architectures such as U-Net, DeepLab, and fully convolutional networks have consistently outperformed traditional machine learning models in delineating plantation boundaries, identifying palm crowns, and capturing fine-scale plantation structures; for instance, Descals et al. [8] demonstrated that DeepLabv3+ achieves significantly higher accuracy than Random Forest classifiers when detecting oil palm plantations using high-resolution imagery. These approaches have also demonstrated superior performance across various remote sensing applications, including cloud shadow identification [28], gas quantification [29], and plant detection [30]. Nevertheless, deep learning models typically require extensive training datasets to achieve optimal performance: Rodríguez et al. [9] manually annotated approximately 3.4×10^6 oil palm instances, while Putra et al. [10] created nearly 100,000 manually labeled bounding boxes across 507 images, with both requiring substantial human effort, domain expertise, and computational resources. This

annotation burden significantly limits the development of more advanced deep learning models for oil palm mapping, especially in regions where ground truth data are scarce or inconsistently available.

2.3. Noisy Learning from Low-Resolution Historical Training Data

To address the limited availability of high-resolution (HR) annotations, an increasing number of studies have explored the use of abundant low-resolution (LR) land cover products as noisy supervisory signals. Techniques such as pseudo-label refinement, consistency-based learning, and k-nearest-neighbor-based noise correction have shown promising performance in large-scale land cover classification tasks when training with imperfect labels [11,12], and more advanced frameworks further incorporate label smoothing, uncertainty estimation, or teacher–student models to enhance robustness against mislabeled pixels [13,14]. Early work focused on alleviating the resolution gap by integrating multi-scale representations to smooth predictions [31], or by casting the problem as label super-resolution through joint feature–label distribution alignment [32,33]. However, directly training HR models with LR labels inevitably introduces substantial noise, motivating the development of noise-robust learning strategies. To this end, ref. [34] proposed a confident area selection mechanism to identify reliable regions within LR maps, while ref. [35] introduced a weakly supervised framework that transfers high-confidence samples from LR products to generate pseudo-labels for HR training. These methods explicitly reduce the influence of mislabeled pixels and enable effective learning under imperfect supervision. More recently, such noise-aware frameworks have been scaled to large-area and meter-level land cover mapping. Ref. [36] employed self-supervised consistency losses to refine noisy constraints for 1 m resolution mapping, and ref. [37] further extended this paradigm through parallel feature extraction without requiring temporally aligned ground truth. Despite these advances, existing approaches predominantly target general land cover classes and do not explicitly account for the temporal dynamics and plantation-specific spatial patterns of oil palm systems, where the use of historical LR labels can introduce even more severe noise. In particular, when a historical product such as the 2016 AOPD is used to supervise 2020–2024 imagery, temporal bias can arise from new plantations, replanting cycles, immature stands, abandonment, or land conversion, so this temporal mismatch must be treated as a core source of uncertainty rather than a peripheral limitation.

3. Study Area

The study area covers Malaysia and Indonesia, the two dominant oil palm-producing countries in Southeast Asia. The region spans a humid tropical environment with frequent cloud cover, heterogeneous land-use mosaics, and extensive coastal and lowland areas where oil palm, forest, peatland, smallholder agriculture, and other perennial crops often occur in close proximity. These conditions make the region both economically important and technically challenging for optical remote-sensing-based plantation mapping.

In Malaysia, oil palm plantations are distributed across Peninsular Malaysia and Malaysian Borneo, with Sabah and Sarawak forming major production zones and Johor, Pahang, and Perak representing important plantation regions on the peninsula. In Indonesia, plantations are concentrated mainly in Sumatra and Kalimantan, including large production areas such as Riau and several Kalimantan provinces, while Java and the eastern islands generally contain smaller and more fragmented plantation areas. These regional names and abbreviations are used consistently in the maps, figures, tables, and text throughout this paper.

Geographically, the study area extends across maritime Southeast Asia between approximately 95°–120° E and 6° S–7° N, encompassing three principal landmasses where

oil palm is concentrated: the Malay Peninsula, the island of Sumatra, and the island of Borneo, which comprises the Malaysian states of Sabah and Sarawak together with the Indonesian region of Kalimantan (Figure 2). Administrative subdivisions (Malaysian states and Indonesian provinces) are outlined in Figure 2 and provide the spatial units used for the regional area statistics reported in Section 6. Although the two countries share a similar humid tropical climate, they differ in plantation structure: Malaysia is characterized by a comparatively mature, estate-dominated sector, whereas Indonesia combines large private and state estates with a substantial and expanding smallholder component, which contributes to the more fragmented plantation patterns observed in parts of Sumatra and Kalimantan.

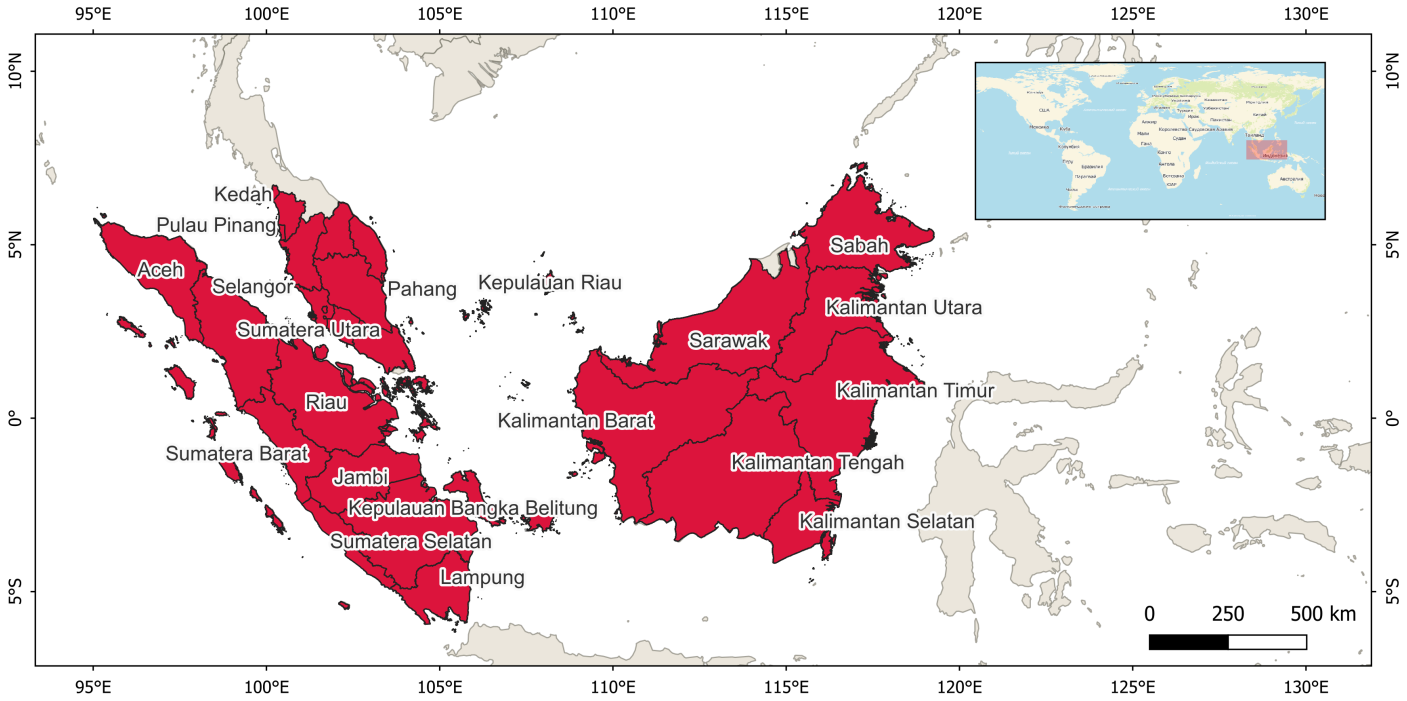


Figure 2. Study area covering Malaysia and Indonesia in maritime Southeast Asia. The two countries are shown in red, with internal lines denoting administrative subdivisions (Malaysian states and Indonesian provinces). The inset locates the study area within a global context, and the scale bar and graticule (latitude/longitude) provide geographic reference. Oil palm plantations are concentrated across the Malay Peninsula, Sumatra, and Borneo (Sabah, Sarawak, and Kalimantan).

4. Data and Methods

4.1. Overview

This study presents a comprehensive deep learning-based approach for mapping oil palm plantations across Indonesia and Malaysia using satellite imagery. Figure 3 shows the overall flowchart of our proposed method, including four major steps, i.e., data preprocessing, neural network architecture, Determinant-based Mutual Information (DMI) loss, and post-processing for final results. We summarized these four major steps as follows:

1. **Data preprocessing**. We used the Annual Oil Palm Area Dataset (AOPD) [16] at 100 m resolution as noisy supervision, together with 10 m Sentinel-2 surface reflectance imagery obtained from Google Earth Engine. The imagery was cropped into 512 × 512 patches and split into training and validation sets.
2. **Neural network architecture**. We evaluated U-Net [38] together with five state-of-the-art semantic segmentation models—MMSFormer [39], DeepLabv3+ [40], UNetFormer [41], HRNet [42], and PyramidMamba [43]—under a unified training protocol with consistent hyperparameters, including z-score normalization and SGD optimization.

3. **DMI loss for noisy supervision**. To address the resolution mismatch between AOPD labels and Sentinel-2 imagery, we adopted the determinant-based mutual information (DMI) loss [44], which provides robustness to noisy labels by maximizing mutual information between predictions and annotations without requiring explicit noise modeling.
4. **Post-processing and mapping**. The trained models were applied to Sentinel-2 imagery from 2020 to 2024. Bayesian inference was used to fuse model outputs with AOPD priors, and patch-level predictions were mosaicked to produce wall-to-wall oil palm plantation maps for Indonesia and Malaysia.

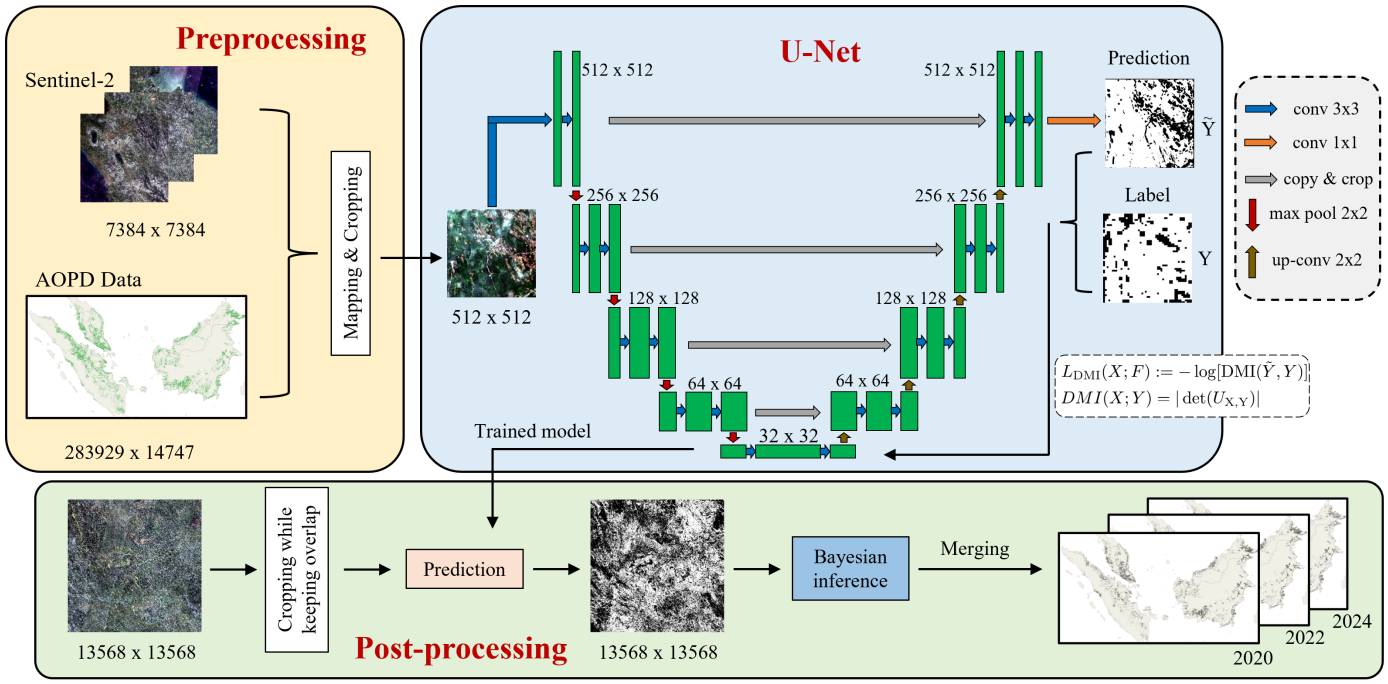


Figure 3. Process for developing a high-resolution oil palm plantation map.

4.2. Data Preprocessing

We obtained oil palm plantation labels at 100 m resolution for Indonesia and Malaysia from the Annual Oil Palm Area Dataset (AOPD) [16], using data from 2016, which represents the most recent year available in this dataset. The AOPD was created using a combination of geometric pre-processing (radiometric and terrain correction and speckle filtering of the radar mosaics) and machine learning classification. For the periods 2007–2010 and 2015–2016, AOPD trained a random forest classifier with four classes (oil palm, other vegetation, water, and other) using PALSAR and PALSAR-2 backscatter parameters. For input imagery, we utilized Sentinel-2 surface reflectance (Level-2A) data from the Copernicus program at 10 m resolution, accessed through Google Earth Engine (GEE). The imagery covers areas in Indonesia and Malaysia for 2020, which serves as the starting year for our dataset updates, and the same procedure was repeated for the 2022 and 2024 prediction years. Because the study region is characterized by persistent cloud cover, we did not rely on individual scenes but instead built an annual cloud-free median composite for each target year; no seasonal or single-scene imagery was used. The compositing window for each target year spans the full calendar year (1 January–31 December). Specifically, for each year we collected all Level-2A scenes acquired within this window over the study area; given the five-day joint revisit cycle of the Sentinel-2A/B constellation, this corresponds to roughly 70 potential acquisitions per orbital tile per year. Clouds and cirrus were masked using the scene classification layer (SCL) and the QA60 cloud-mask band together with

the cloud-probability layer, and cloud shadows were removed using the cloud-shadow class of the SCL band. The per-pixel median of the remaining clear observations was then computed. All available Level-2A scenes within each annual window were used, rather than a fixed subset of selected scenes. Compositing over a full year of repeat passes ensures that pixels affected by clouds, haze, or shadow in some acquisitions are filled by clear observations from other dates, so that missing observations at any single date do not leave data gaps in the final composite. We used 10 spectral bands: Blue, Green, Red, Red Edge 1–4, Near Infrared, and Short Wave Infrared 1–2. Among these, the Blue, Green, Red, and Near Infrared bands are natively acquired at 10 m, whereas the four Red Edge and two Short Wave Infrared bands are natively acquired at 20 m; the latter were resampled to 10 m using nearest-neighbor interpolation so that all bands share a common 10 m grid.

To create a consistent dataset from the 100 m resolution oil palm labels and 10 m resolution satellite imagery, we resampled the oil palm labels from 100 m to 10 m resolution using nearest-neighbor interpolation, which preserves the original discrete label values without introducing intermediate classes. Next, we cropped both the labels and imagery into corresponding 512 × 512 pixel patches based on their geographic coordinates; at 10 m resolution each patch covers a ground area of 5.12 km × 5.12 km. This tiling yielded a total of 191,236 patches (161,668 non-oil-palm and 29,568 oil-palm patches) prior to class balancing.

To address significant class imbalance—identifiable in the initial dataset of 161,668 non-oil palm images versus only 29,568 oil palm images—we employed a random undersampling strategy. We retained all positive samples (oil palm) and randomly selected an equivalent number of negative samples (non-oil palm) to construct a balanced dataset of 59,136 images. Given that the input labels are coarse (100 m) while the target predictions are fine-grained (10 m), we frame this as a weakly supervised learning task; this resolution discrepancy serves to mitigate spatial autocorrelation and label leakage. The final balanced dataset was partitioned into training (70%) and validation (30%) subsets. Prior to model ingestion, z-score normalization (detailed in Section 5) was applied to all imagery. While we experimented with several deep learning architectures, U-Net provided the best accuracy, as discussed in Section 5.3.

4.3. Recall of U-Net

The U-Net architecture was originally proposed in 2015 [38] for biomedical image segmentation. The original architecture consists of 18 convolutional layers with 3 × 3 kernels, each followed by ReLU activation functions. The network employs a symmetric encoder–decoder structure with four 2 × 2 max pooling operations for downsampling in the encoder path and four 2 × 2 upsampling layers in the decoder path. The final output layer uses a 1 × 1 convolution for classification. A key innovation of U-Net is the use of skip connections that concatenate feature maps from the encoder to the corresponding decoder layers, enabling the network to preserve fine-grained spatial information crucial for accurate localization.

Although originally designed for biomedical image segmentation, U-Net has proven highly effective for various segmentation tasks, including remote sensing applications as utilized in our study. Our implementation follows the original architecture with a modification, batch normalization [45] is inserted between each convolutional layer and ReLU activation function to improve training stability. For training, we employed Stochastic Gradient Descent (SGD) as the optimizer with a learning rate of 0.001, which was the optimal hyperparameter configuration for this experiment.

4.4. Dmi Loss

Due to the resolution mismatch between the coarse 100 m oil palm training data and both the satellite imagery and target outputs (10 m), we treat this low-resolution training data as noisy supervision rather than as exact ground truth. To address this challenge, we employ the DMI (Determinant-based Mutual Information) loss function [44] for learning with noisy labels without requiring additional information, as shown in Equation (1):

$$L_{\mathrm{DMI}}(X; F) := -\log[\mathrm{DMI}(\tilde{Y}, Y)], \tag{1}$$

where X represents the input image, F denotes the U-Net model, $\tilde{Y}$ is the model's prediction, and Y is the ground truth. DMI (Determinant-based Mutual Information) [46], which employs a DMI-Mechanism to evaluate the agreement between model predictions and labels. This mechanism encodes the relationship between predictions and labels into a 2×2 matrix $U_{\mathrm{X,Y}}$, as defined in Equations (2) and (3). Each entry in $U_{\mathrm{X,Y}}$ represents the joint occurrence probabilities of specific predicted and actual values. The determinant of this matrix quantifies the dependency structure: a positive determinant indicates positive correlation, while a negative determinant reflects negative correlation.

$$DMI(X; Y) = |\det(U_{\mathrm{X,Y}})|, \tag{2}$$

$$U_{\mathrm{X,Y}} = Pr[X = x, Y = y], \tag{3}$$

where $Pr[X = x, Y = y]$ represents the matrix form of the joint probability distribution of X and Y.

We selected DMI over alternative noise-robust objectives because of its information-theoretic properties. DMI is provably robust to instance-independent label noise and, unlike loss-correction approaches such as Forward/Backward correction [47] or Generalized Cross-Entropy (GCE) [48], it does not require estimating a noise transition matrix or tuning a noise-rate hyperparameter—both of which are difficult to specify when the true noise level is unknown, as in our resolution-mismatch setting. This property makes DMI particularly suitable for weakly supervised mapping from coarse historical products.

4.5. Post-Processing for Final Results

We predicted oil palm plantation maps for Indonesia and Malaysia from 2020 to 2024 at 10 m resolution using our trained model. Raw satellite imagery was downloaded from GEE for the years 2020, 2022, and 2024. To accommodate the model's input requirements, we partitioned these large images into smaller 512×512 pixel tiles with a 20 pixels overlap to ensure seamless predictions at tile boundaries.

Since the AOPD dataset [16] was created using remote sensing methods with different preprocessing approaches—specifically, the AOPD model was trained on multiclass classification while ours uses binary classification—we leveraged the AOPD product to enhance our map accuracy through Bayesian inference. This approach combines the strengths of both methods, as shown in Equations (4) and 5, where the UNet prediction serves as the prior, AOPD provides the evidence, and the posterior yields a more accurate oil palm plantation map:

$$p(\theta \mid \{y_{\mathrm{i}}\}) = \frac{p(\{y_{\mathrm{i}}\} \mid \theta) p(\theta)}{p(\{y_{\mathrm{i}}\})}, \tag{4}$$

$$p(\{y_i\} \mid \theta) = \prod_i p(y_i \mid \theta). \tag{5}$$

where θ represents the probability of $y_i = 1$ (indicating oil palm presence). In this formulation, $p(\theta|y_i)$ represents the posterior probability, $p(y_i)$ represents the evidence, $p(y_i|\theta)$ signifies the likelihood, and $p(\theta)$ denotes the prior probability.

Concretely, the fusion is applied independently at each 10 m pixel. The prior $p(\theta)$ is the per-pixel oil palm probability produced by the sigmoid output of the trained U-Net for the target year, i.e., $p(\theta) = \hat{p}_i$ for the oil palm hypothesis (θ: $y_i = 1$) and $1 - \hat{p}_i$ for the non-oil-palm hypothesis. The evidence is provided by the AOPD product, which is a discrete oil palm/non-oil palm label $a_i \in \{0, 1\}$; the same 2016 AOPD layer (the most recent year available) is used as evidence for all three prediction years (2020, 2022, and 2024), after being resampled to the 10 m prediction grid. Because AOPD is a hard binary label rather than a calibrated probability, we convert it into a per-pixel likelihood using a single agreement parameter α (the assumed reliability of the AOPD label), as shown in Equation (6):

$$p(a_i \mid \theta) = \begin{cases} \alpha, & \text{if } a_i = \theta, \\ 1 - \alpha, & \text{if } a_i \neq \theta, \end{cases} \tag{6}$$

so that the likelihood reinforces the U-Net probability where the two sources agree and attenuates it where they disagree. We adopt the standard non-informative default $\alpha = 0.5$, which corresponds to a symmetric, unbiased treatment of the AOPD evidence rather than a value tuned on the validation set; this choice deliberately keeps the fused posterior driven primarily by the current-year U-Net prior and uses AOPD only as a weak agreement term. The unnormalized posterior for the oil palm hypothesis is $\tilde{p}(\theta = 1 \mid a_i) = p(a_i \mid \theta = 1)\, \hat{p}_i$, and it is normalized over the two mutually exclusive hypotheses:

$$p(\theta = 1 \mid a_i) = \frac{p(a_i \mid \theta = 1)\, \hat{p}_i}{p(a_i \mid \theta = 1)\, \hat{p}_i + p(a_i \mid \theta = 0)\, (1 - \hat{p}_i)}, \tag{7}$$

so that the denominator is exactly the evidence $p(\{y_i\})$ in Equation (4). A pixel is finally labeled as oil palm when this normalized posterior probability exceeds a decision threshold of 0.5; otherwise it is labeled as non-oil palm. The 0.5 threshold is the conventional default for a two-class Bayesian decision (i.e., choosing the maximum-a-posteriori class under equal misclassification costs) and was not selected empirically; we did not tune it against the validation points, so as to avoid optimizing the reported accuracies on the same data used for evaluation. Because the AOPD evidence is fixed at 2016 while the U-Net prior is year-specific, the fused posterior remains driven primarily by the current-year imagery, and the 2016 prior mainly suppresses isolated false positives in areas that were never oil palm; we note that using a fixed 2016 evidence layer can also limit the detection of genuinely new plantations established after 2016, and this design choice is one contributor to the temporal-bias behavior discussed in Section 6. Finally, all 13,568 × 13,568 pixel images were merged to create comprehensive oil palm plantation maps for Indonesia and Malaysia, resulting in a final map size of 278,789 × 148,117 pixels.

5. Results

5.1. Experimental Setup

Our experiments were conducted using the PyTorch open-source deep learning library on a single NVIDIA GeForce RTX 4090 GPU. We implemented a U-Net architecture [38] designed in Section 4.3 to process input images of 512 × 512 pixels with 10 channels. The model was trained from scratch using a Stochastic Gradient Descent (SGD) optimizer [49] with a momentum of 0.9. We set the learning rate to 0.001 and trained the model for 5 epochs. We utilized the DMI loss [44] as our loss function (detailed in Section 4.4). All

input images underwent z-score normalization [50] (Equation (8)) to standardize intensity distributions and enhance model convergence.

Several of these design choices were deliberate given the weakly supervised setting. We selected U-Net as the primary architecture because its symmetric encoder–decoder with skip connections preserves fine spatial detail at plantation boundaries, and our comparative experiments (Table 1) confirmed that heavier transformer- and Mamba-based models tended to over-smooth boundaries under noisy supervision rather than improve over it. We trained for only 5 epochs because, under the coarse 100 m supervision, validation accuracy saturated early and longer training caused the model to overfit the blocky label artifacts rather than the true 10 m morphology; 5 epochs offered the best trade-off between convergence and avoiding memorization of label noise. To keep the comparison controlled, all architectures in Table 1 were trained under the identical protocol (same optimizer, learning rate, number of epochs, class-balancing, and DMI loss), so that observed differences are attributable to the architectures themselves rather than to per-model tuning; we verified that the heavier transformer- and Mamba-based models had likewise reached their validation plateau within this budget, so the shared schedule did not systematically truncate their training. We regard this uniform protocol as a deliberate design choice for a fair like-for-like comparison under noisy supervision, while acknowledging that architecture-specific hyperparameter optimization could further improve individual baselines and is left for future work. The random undersampling strategy described in Section 4 was used to counter the severe positive/negative class imbalance and to prevent the model from collapsing toward the majority non-oil-palm class.

$$\tilde{s}_{ki} = \frac{s_{ki} - \mu_i}{\sigma_i} \tag{8}$$

where s_{ki} represents the intensity value of pixel k in image i, $\tilde{s}_{ki}$ is the corresponding normalized value, and μ_i and σ_i are the mean and standard deviation of all pixel intensities in image i, respectively.

Table 1. Accuracy comparison between AOPD, U-Net, MMSFormer, DeepLabv3+, UNetFormer, HRNet and PyramidMamba in 2020, 2022 and 2024 year.

Method	2020				2022				2024			
	F1	UA	PA	OA	F1	UA	PA	OA	F1	UA	PA	OA
AOPD	0.49	34.10	86.94	35.93	0.54	39.22	88.72	40.67	0.55	40.23	88.84	41.52
U-Net	**0.69**	**55.03**	93.13	**70.64**	0.62	**53.04**	74.54	**63.53**	0.58	**50.84**	67.26	**60.06**
MMSFormer	0.50	37.54	74.05	47.23	0.54	42.76	74.70	50.03	0.55	43.41	74.55	49.91
DeepLabv3+	0.48	40.19	59.45	54.35	0.50	44.30	56.25	54.35	0.53	47.05	59.38	56.11
UNetFormer	0.63	48.20	91.92	62.19	**0.64**	50.22	87.20	60.43	**0.62**	50.14	80.65	59.33
HRNet	0.62	46.35	95.02	59.33	0.63	49.02	87.65	58.72	0.61	48.57	83.63	57.14
PyramidMamba	0.54	37.87	**95.53**	42.98	0.59	42.64	**95.43**	46.99	0.60	43.60	**95.24**	47.72

The evaluation metrics employed in this study include F1 score, User's Accuracy (UA), Producer's Accuracy (PA), and Overall Accuracy (OA). To facilitate the calculation of these metrics, we denote the number of correctly identified positive pixels as True Positives (TP), incorrectly identified positive pixels as False Positives (FP), correctly identified negative pixels as True Negatives (TN), and incorrectly identified negative pixels as False Negatives (FN).

Precision (equivalent to User's Accuracy) indicates the proportion of positive predictions that were actually correct, while Recall (equivalent to Producer's Accuracy) measures

the proportion of actual positive cases that were correctly identified. These metrics are defined in Equation (9) and (10), respectively:

$$\text{Precision} = \frac{TP}{TP + FP} \tag{9}$$

$$\text{Recall} = \frac{TP}{TP + FN} \tag{10}$$

The F1 score evaluates the classification model's performance by calculating the harmonic mean of precision and recall, as shown in Equation (11). This metric assesses the balance between correctly identifying positive cases and minimizing false detections.

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \tag{11}$$

User's Accuracy and Producer's Accuracy correspond conceptually to precision and recall but are expressed as percentages, as demonstrated in Equations (12) and (13). Overall Accuracy (Equation (14)) represents the percentage of all predictions that were correct across both classes.

$$\text{User's Accuracy} = \frac{TP}{TP + FP} \times 100\% \tag{12}$$

$$\text{Producer's Accuracy} = \frac{TP}{TP + FN} \times 100\% \tag{13}$$

$$\text{Overall Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \times 100\% \tag{14}$$

5.2. Generation of Multi-Temporal Validation Data

To rigorously and independently evaluate both the proposed models and the AOPD baseline, we constructed a validation set of 2058 randomly sampled points spanning Indonesia and Malaysia (Figure 4). The validation design enforces three properties: **(i) Independence from training.** The validation points were drawn independently of the patches used for training; their reference labels were never derived from the AOPD product but were assigned by manual visual interpretation, so the assessment does not reuse the noisy supervision. **(ii) Spatial independence.** Points were screened to avoid spatial overlap with training patches and were distributed across the study area to limit spatial autocorrelation between training and evaluation samples. **(iii) Per-year reference labels.** Although the identical set of spatial coordinates was reused across years to ensure temporal comparability, the reference class at each point was re-interpreted *separately* for 2020, 2022, and 2024 using Google Earth historical high-resolution imagery acquired as close as possible to each target year. Each year was thus validated against its own contemporaneous ground truth rather than a single static label. This protocol yields a temporally explicit, spatially independent benchmark, which is especially important given that no 10 m oil palm product exists for this region after 2019. The same 2058 points were used for every target year, so the per-year sample size is held constant across 2020, 2022, and 2024. However, because the reference class at each point is re-interpreted independently for each year, the number of points that are actually oil palm is not identical across years: the reference set contains 728 oil palm and 1330 non-oil-palm points in 2020, 820 oil palm and 1238 non-oil-palm points in 2022, and 841 oil palm and 1217 non-oil-palm points in 2024. Because the assessment is binary (oil palm versus non-oil palm), the class-wise performance of each model in each year is fully summarized by the reported User's Accuracy, Producer's Accuracy, F1 score, and Overall Accuracy (Table 1). To make the accuracy assessment fully transparent, we additionally report the underlying confusion-matrix counts (true positives, false positives,

false negatives, and true negatives) for the U-Net model and the AOPD baseline in each year in Table 2.

We adopted a spatially random sampling design so that the validation set reflects the overall areal proportion of oil palm and non-oil palm across the two countries, which yields a large, well-populated non-oil-palm class (≈1200–1330 points) and a substantial oil palm class (≈728–841 points) in every year (Table 2), sufficient to estimate class-wise accuracies with reasonable precision. We acknowledge, however, two limitations of this design that temper the interpretation of the reported metrics. First, because the sample is allocated proportionally to area rather than stratified, it is dominated by common, spatially extensive landscapes; rare or localized regional patterns (for example, small fragmented smallholder plots, immature replanting blocks, or narrow wetland-edge conversions) are under-represented, so the accuracies reported here are best read as region-wide averages rather than as guarantees of performance in these rarer transitional settings. A stratified or change-focused sampling scheme would be needed to characterize those cases specifically and is left for future work. Second, each validation point was interpreted by a single analyst using Google Earth historical imagery, and we did not carry out an independent multi-interpreter assessment; consequently, inter-interpreter agreement was not quantified, and a component of interpretation uncertainty—particularly for spectrally ambiguous perennial crops such as coconut and rubber, and for immature stands—is not reflected in the reported accuracies. We treat this as an explicit limitation of the accuracy assessment (see also Section 6).

Table 2. Confusion-matrix counts for the U-Net model and the AOPD baseline against the 2058 manually interpreted validation points in each year. TP, FP, FN, and TN denote true positives, false positives, false negatives, and true negatives for the oil palm class. "Oil palm" and "Non-oil palm" give the number of reference (ground-truth) validation points per class, re-interpreted independently for each year. The reported User's, Producer's, and Overall Accuracy in Table 1 are reproduced exactly by these counts.

Year	Method	TP	FP	FN	TN	Oil Palm	Non-Oil Palm
2020	AOPD	633	1223	95	107	728	1330
	U-Net	678	554	50	776	728	1330
2022	U-Net	611	541	209	697	820	1238
2024	U-Net	566	547	275	670	841	1217

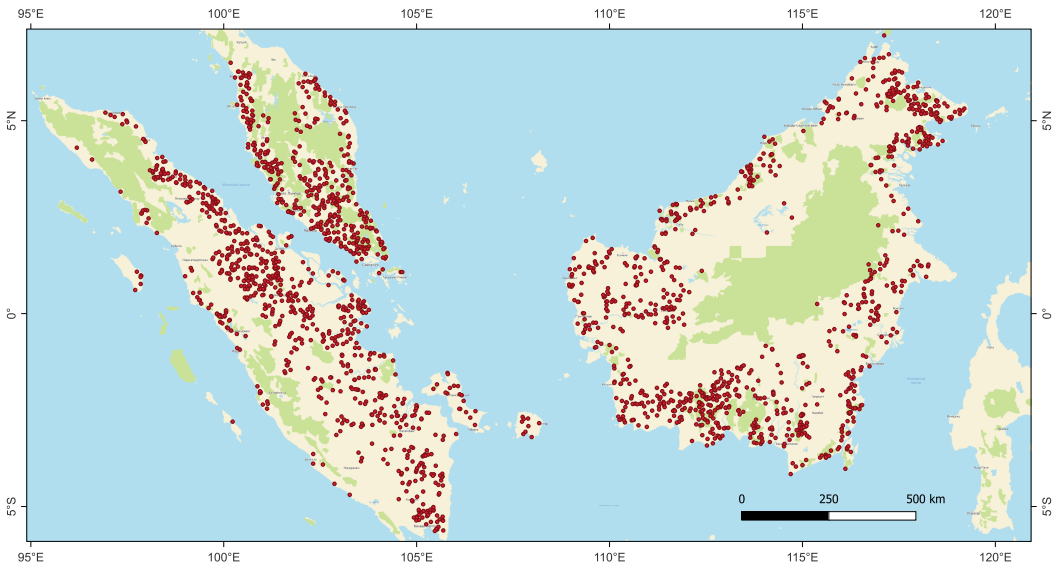


Figure 4. Geographic distribution of the 2058 validation points used for accuracy assessment across the study area.

5.3. Other Comparative Deep Learning-Based Methods

In addition to the U-Net model, this study evaluates several state-of-the-art deep learning methods for semantic segmentation, including MMSFormer [39], DeepLabv3+ [40], UNetFormer [41], HRNet [42], and PyramidMamba [43]. All these methods are recognized for their efficiency in semantic segmentation tasks:

1. MMSFormer [39]: A transformer-based architecture designed for multi-modal fusion, utilizing Mix-Transformer encoders and parallel multi-scale convolutions.
2. DeepLabv3+ [40]: An encoder–decoder network that uses Atrous Spatial Pyramid Pooling (ASPP) to refine segmentation boundaries and capture contextual information.
3. UNetFormer [41]: A hybrid model integrating a ResNet encoder with a transformer decoder to capture both long-range dependencies (via self-attention) and local details.
4. HRNet [42]: A high-resolution network that connects parallel multi-resolution streams to preserve spatial details without recovering resolution from downsampled features.
5. PyramidMamba [43]: A novel architecture incorporating Mamba-based decoders and Dense Spatial Pyramid Pooling to efficiently process spatial contexts with reduced redundancy.

5.4. Classification Accuracy

Visual inspection (Figures 5 and 6) reveals that the U-Net model yields the most precise boundary delineation among all evaluated methods. The model successfully recovers fine-grained morphological details, such as the separation between plantation blocks and adjacent roads or rivers, which are notably absent in the coarse 100 m AOPD training labels. In contrast, while newer architectures like UNetFormer and PyramidMamba produce smoother outputs, they tend to over-smooth boundaries, effectively replicating the blocky artifacts of the noisy labels rather than learning the true underlying features. Furthermore, models such as MMSFormer and DeepLabv3+ exhibited performance inferior to the baseline AOPD data, suggesting a limited capacity to generalize from noisy supervision in this specific domain.

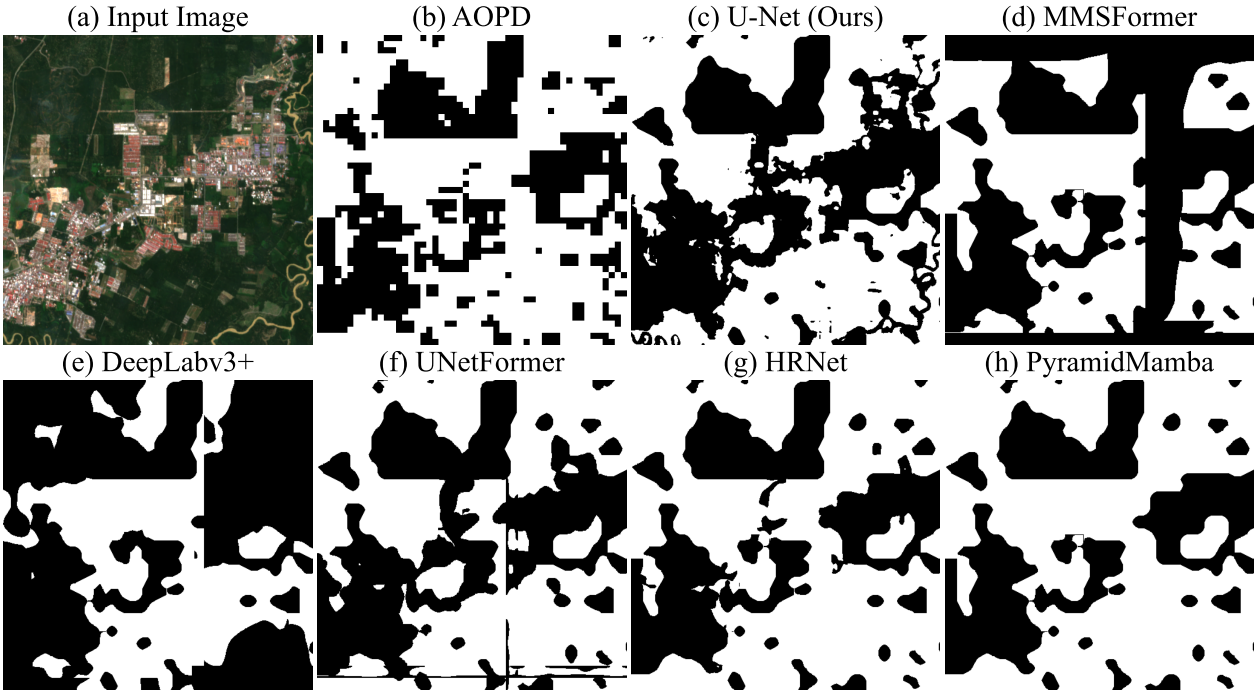


Figure 5. Oil palm detection results in Malaysia across different models: (**a**) Sentinel-2 imagery, (**b**) AOPD noisy label data used for model training, and (**c**–**h**) oil palm detection outputs from respective models.

To ensure a fair and unbiased assessment—given the noisy nature of the training labels—we utilized an independent validation strategy. We conducted a manual visual interpretation of the 2058 sampled pixels (Figure 4) using Google Earth imagery to establish ground truth for 2020, 2022, and 2024. Based on these verified labels, we computed F1 Score, UA, PA, and OA. This method enables a robust comparison of the U-Net and AOPD baselines against the five advanced models (Table 1), effectively bypassing the inaccuracies present in the training ground truth.

The results demonstrate that U-Net consistently achieves the highest performance across most evaluation metrics throughout the study period. Notably, UNetFormer exhibits superior F1 scores in 2022 and 2024, indicating its effectiveness in balancing precision and recall. Furthermore, PyramidMamba consistently maintains the highest Producer's Accuracy values across all three years, demonstrating its exceptional capability in correctly identifying positive instances within the dataset.

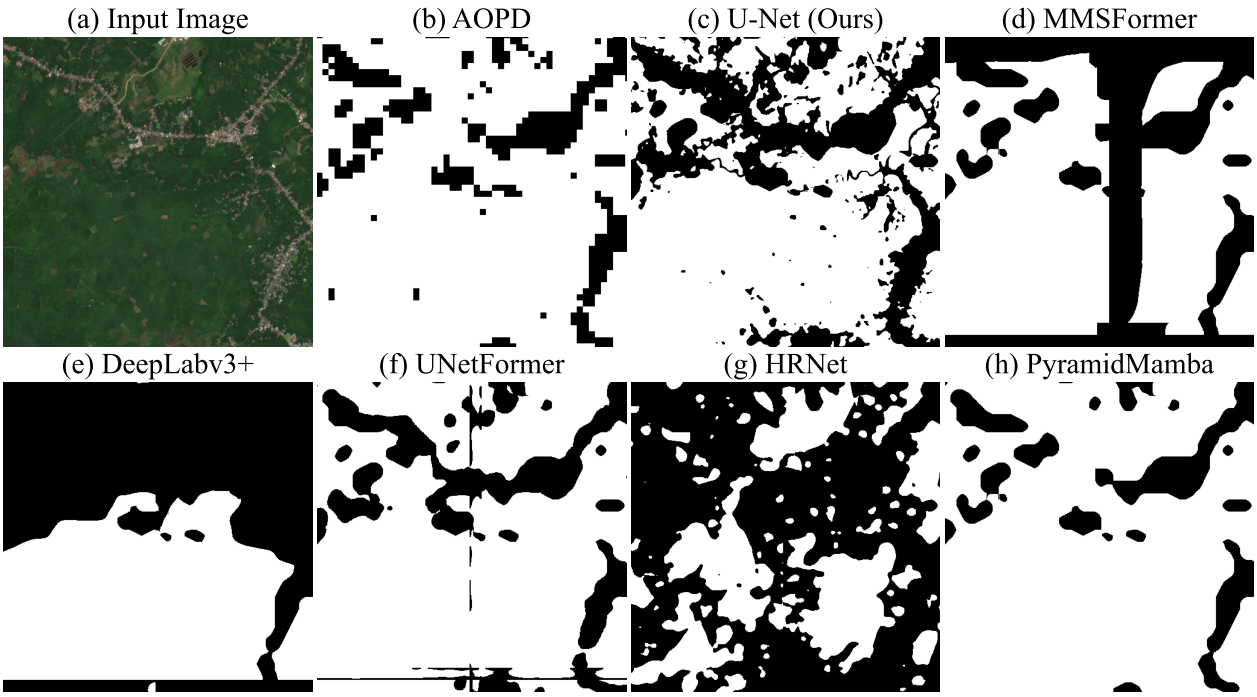


Figure 6. Oil palm detection results in Indonesia across different models: (**a**) Sentinel-2 imagery, (**b**) AOPD noisy label data used for model training, and (**c**–**h**) oil palm detection outputs from respective models.

5.5. Ablation Studies

To illustrate the significant challenge posed by the resolution mismatch between high-resolution imagery (10 m) and spatially coarse labels (100 m), and to demonstrate the efficacy of our learning-with-noisy-labels approach, we conducted an ablation study, i.e., an experiment in which a single component of the method is changed while all others are held fixed, so that the effect of that component can be isolated. Figure 7 presents a qualitative comparison of a U-Net model trained under identical configurations, differing only in the loss function employed (DMI versus BCE).

Panel (c) displays the results using DMI Loss [44], the technique adopted in this study for learning with noisy labels. In contrast, panel (d) shows the results using Binary Cross-Entropy (BCE) loss [51], representing the standard supervised training baseline. The difference in performance is substantial; the U-Net trained with standard BCE loss largely fails to detect oil palm features within the Sentinel-2 imagery. Conversely, the model trained with DMI loss not only detects oil palm with high precision but also recovers

granular spatial details that surpass the coarse resolution of the original ground truth (AOPD [16]).

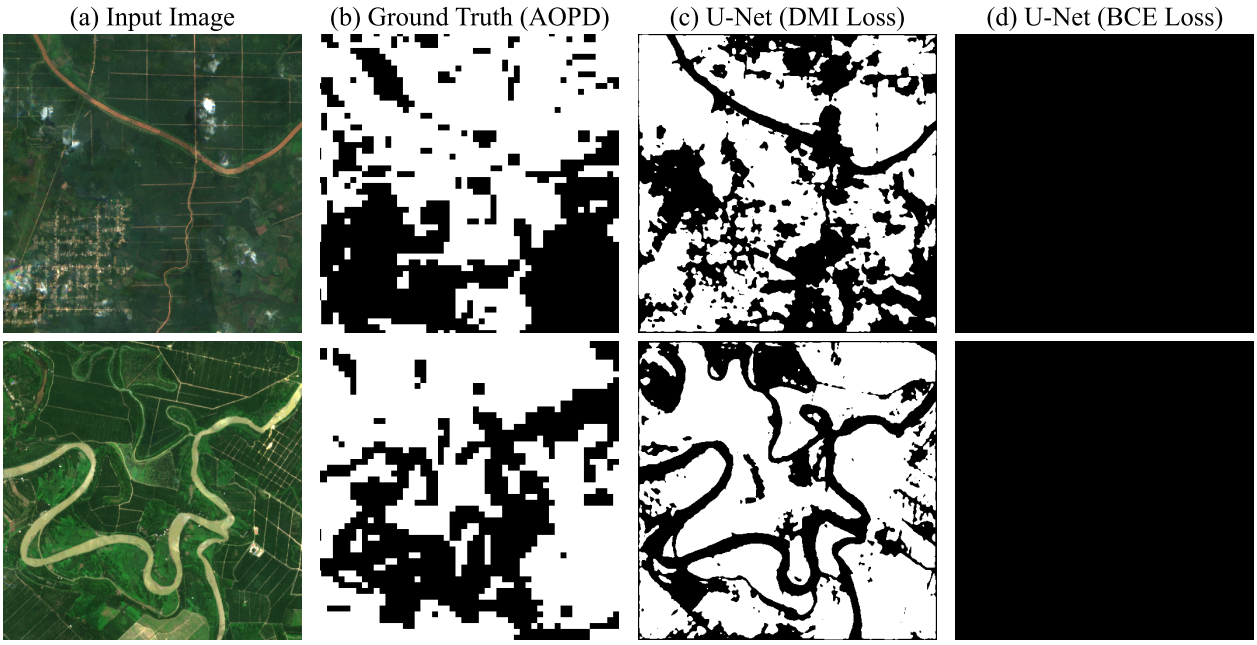


Figure 7. Visual comparison between (**a**) Sentinel-2 satellite imagery, (**b**) Ground Truth (AOPD), (**c**) Results from the U-Net model trained with DMI Loss (learning with noisy labels), and (**d**) Results from the U-Net model trained with Binary Cross-Entropy Loss (standard supervision).

To further substantiate these qualitative observations with quantitative evidence, Table 3 provides a numerical performance comparison. These metrics clearly demonstrate that the model trained with DMI Loss consistently outperforms the BCE Loss baseline across all evaluated years.

Table 3. Accuracy comparison between U-Net trained with DMI loss and BCE loss across different years.

Year	Metric	Method	
		U-Net (DMI Loss)	U-Net (BCE Loss)
2020	F1	**0.69**	0.50
	UA	**55.03**	37.54
	PA	**93.13**	74.05
	OA	**70.64**	47.23
2022	F1	**0.62**	0.54
	UA	**53.04**	42.76
	PA	74.54	**74.70**
	OA	**63.53**	50.03
2024	F1	**0.58**	0.55
	UA	**50.84**	43.41
	PA	67.26	**74.55**
	OA	**60.06**	49.91

5.6. Mapped Oil Palm Area and Regional Patterns

Figure 8 presents the mapped oil palm plantation area across 13 Malaysian regions for 2020, 2022, and 2024. Overall, the mapped areas remained relatively stable during the study period. Sabah and Sarawak consistently ranked as the top-producing states, with mapped areas ranging from 1.1×10^6 ha to 1.5×10^6 ha, reflecting their dominant role in national production [52]. Johor, Pahang, and Perak followed, maintaining substantial and stable extents between 0.5×10^6 ha and 1.0×10^6 ha. While most regions showed

little change, Sabah and Sarawak exhibited a notable decline of approximately 1.0×10^5 ha to 2.5×10^5 ha between 2022 and 2024; possible drivers of this decline are discussed in Section 6.

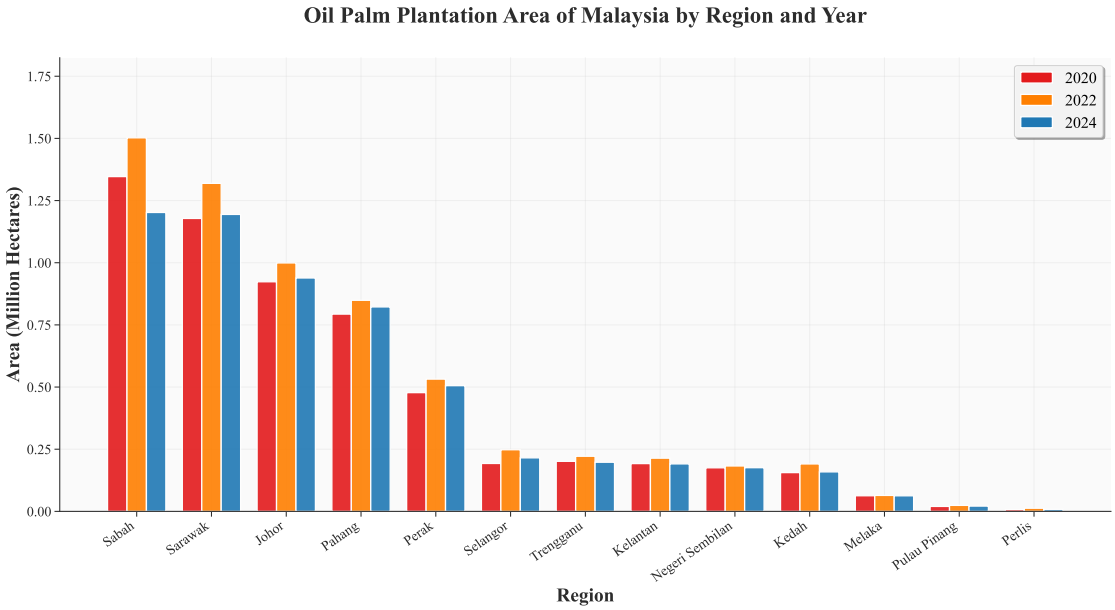


Figure 8. Oil Palm Plantation Area of Malaysia by Region and Year.

Figure 9 further shows that plantations in Malaysia are predominantly concentrated in Malaysian Borneo, with mapped areas ranging from 1.0×10^6 to 2.0×10^6 ha, whereas the Malay Peninsula exhibits a lower mapped extent, averaging approximately 0.65×10^6 ha. The mapped trends also indicate faster land use/cover change rates on the Malay Peninsula (2%–10% annually) compared with Kalimantan (−2.5%–2%).

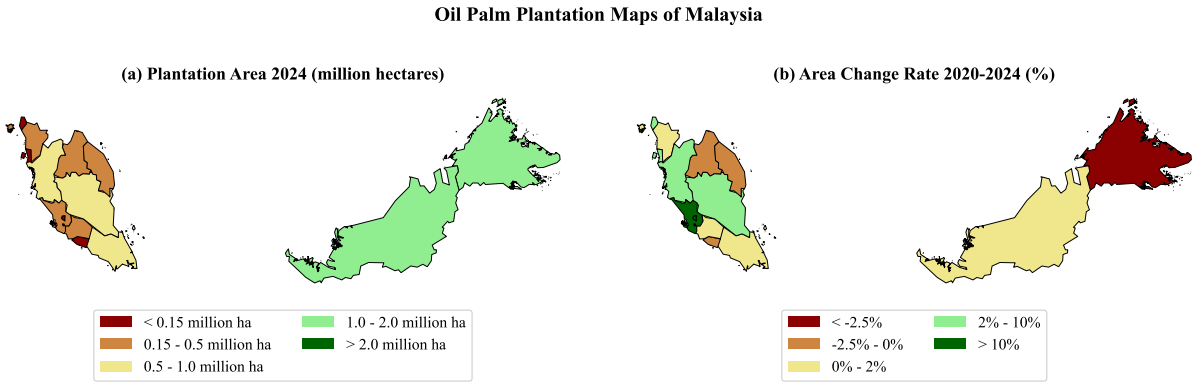


Figure 9. Oil Palm Plantation Maps in Malaysia (**a**) Plantation area by region in 2024 and (**b**) Percentage change in plantation area between 2020 and 2024.

Similarly, Figure 10 shows the mapped oil palm area across 15 Indonesian regions from 2020 to 2024. Riau consistently had the largest mapped area, peaking at 2.54×10^6 ha in 2022 but declining to 1.45×10^6 ha by 2024, a reduction of approximately 43%. Region-wide, the mapped plantation coverage peaked in 2022 and contracted in 2024. We stress that this contraction, and the large Riau reduction in particular, should be read as an indicative regional pattern rather than a robust quantitative estimate of area change: as quantified in Section 6, the mapped absolute areas overestimate official statistics by roughly a factor of two and part of the apparent post-2022 decline may stem from classification uncertainty and the 2016 temporal label bias rather than from genuine land-use change alone.

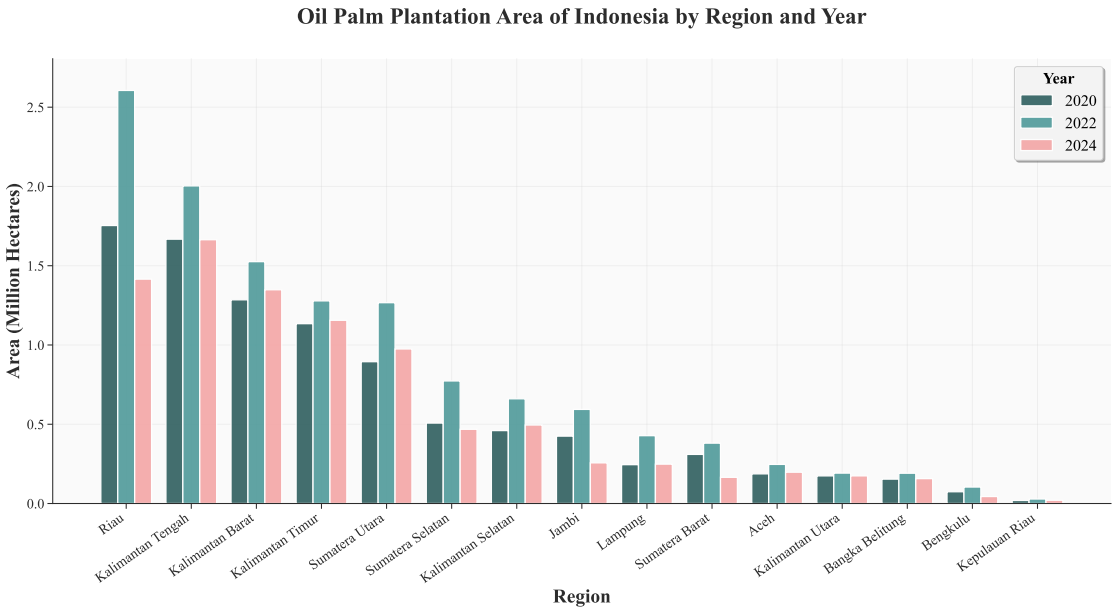


Figure 10. Oil Palm Plantation Area of Indonesia by Region and Year.

Figure 11 shows that mapped plantations in Indonesia are heavily concentrated in Kalimantan, averaging 1.8×10^6 ha and forming the core of national production. Kalimantan is also the main mapped driver of expansion, with an average growth rate of 5%, while other areas declined by approximately 1%. In contrast, Sumatra has extensive plantation coverage—exceeding 2.0×10^6 ha in some provinces—but shows stagnant or declining mapped trends (−2.5% to 0%). Java and the eastern islands have minimal coverage (typically $< 0.15 \times 10^6$ ha) and mixed growth.

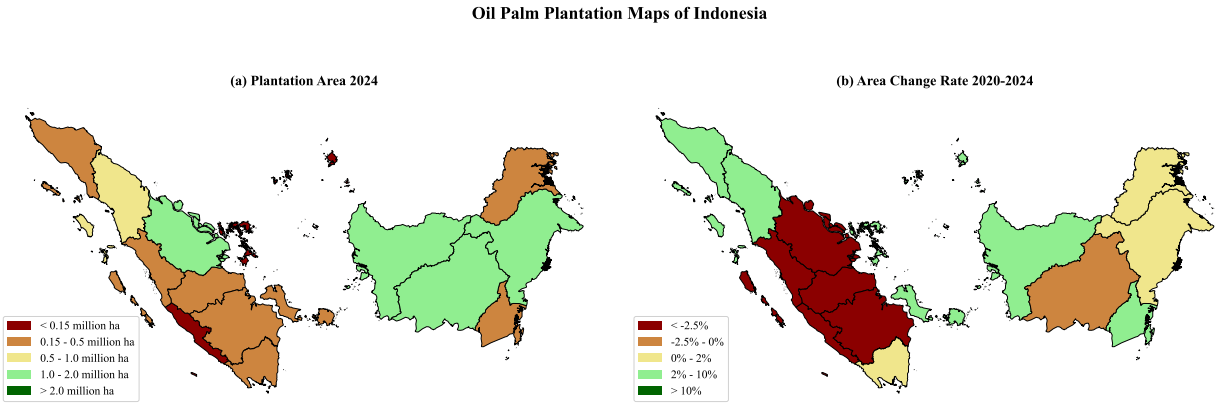


Figure 11. Oil Palm Plantation Maps in Indonesia (**a**) Plantation area by region in 2024 and (**b**) Percentage change in plantation area between 2020 and 2024.

5.7. Land-Cover Transitions Between 2020 and 2024

Based on the generated 10 m resolution oil palm plantation maps, we analyzed the spatiotemporal dynamics of oil palm cultivation from 2020 to 2024. To characterize the land cover types interacting with oil palm, we utilized the Esri Sentinel-2 10 m Land Use/Land Cover time series [20]. We use this product only for contextual analysis of the surrounding land-cover classes (built area, crops, flooded vegetation, rangeland, and water) that exchange with oil palm; because its broad class taxonomy does not explicitly identify oil palm plantations, the oil palm class itself is taken from our own 10 m maps rather than from the Esri product. In the land cover flow analysis presented in Figure 12, all transition magnitudes are reported in km^2 (one 10 m pixel corresponds to 1.0×10^{-4} km^2).

Because the persistent (unchanged) oil palm area is roughly two orders of magnitude larger than the individual transition flows and would otherwise visually dominate a single Sankey diagram, we complement the Sankey view (Figure 12, top) with per-class trend panels (Figure 12, bottom) that plot each directional flow on its own scale, so that the smaller conversions can be compared directly. The total mapped oil palm area initially expanded from approximately 1.13×10^5 km^2 in 2020 to 1.17×10^5 km^2 in 2022, and subsequently contracted to about 1.10×10^5 km^2 in 2024.

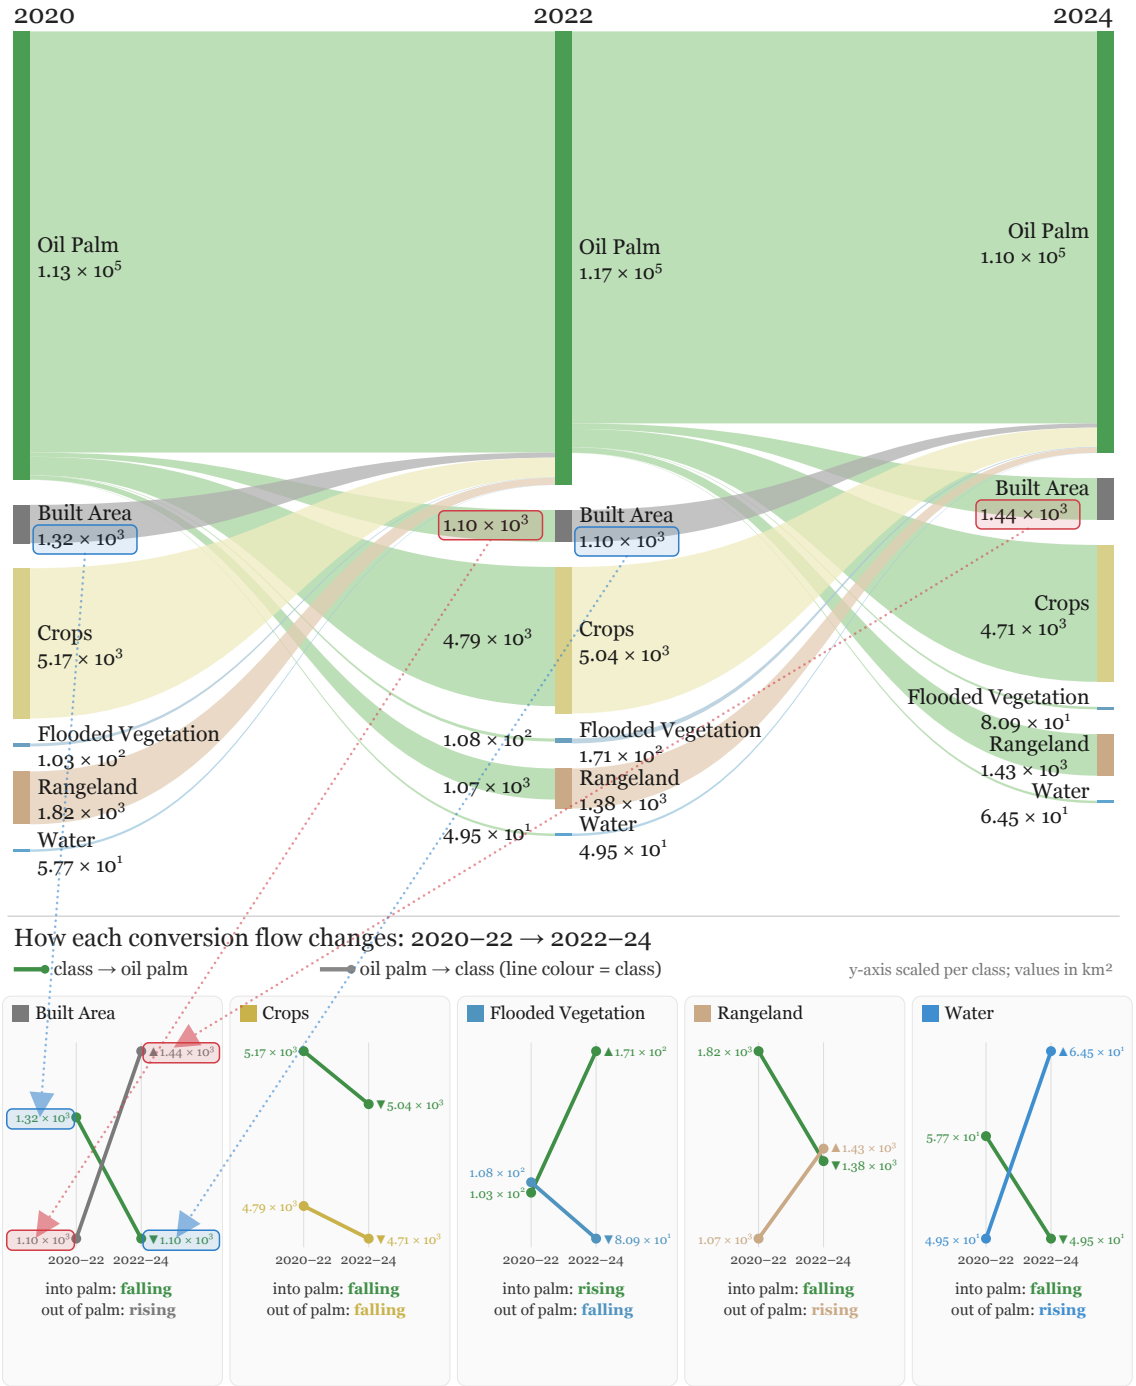


Figure 12. Land-cover exchange between oil palm and other classes (built area, crops, flooded vegetation, rangeland, and water), with all magnitudes in km^2. (**Top**) Sankey diagram of the oil palm state and its gross exchanges with each class at 2020, 2022, and 2024. (**Bottom**) Per-class trend panels that separate the two transition directions so that the smaller but ecologically important flows remain

legible despite the dominance of the persistent (unchanged) oil palm area: green lines denote conversion *into* oil palm (class → oil palm) and grey lines denote loss *out of* oil palm (oil palm → class) over the 2020–2022 and 2022–2024 intervals. Each panel is scaled independently to its own class magnitude, and the up/down markers summarize whether each directional flow is rising or falling.

Between the two intervals (2020–2022 vs. 2022–2024), the rate of land conversion into oil palm from built area, rangeland, and water decreased; for example, built-area conversion dropped from 1.32×10^3 km^2 to 1.10×10^3 km^2. Conversely, the loss of oil palm to these classes accelerated, with conversion to built area rising from 1.10×10^3 km^2 to 1.44×10^3 km^2.

Interactions with flooded vegetation and crops followed distinct patterns. Conversion of flooded vegetation into oil palm increased from 1.03×10^2 km^2 to 1.71×10^2 km^2, while the reverse flow declined. Meanwhile, the exchange between oil palm and crops showed a reduction in bidirectional flow; despite this reduction, the crop–palm exchange remained the most substantial, accounting for the highest proportion of land-cover change among all analyzed classes. Because these transition magnitudes inherit the per-year classification uncertainty of both our oil palm maps and the Esri land-cover product, and are far smaller than the persistent oil palm area, the absolute flow values should be regarded as indicative regional patterns rather than robust quantitative transition estimates; we therefore emphasize the *direction* and *relative* ordering of these flows rather than their exact magnitudes. Interpretations of these transition patterns are provided in Section 6.

6. Discussion

6.1. Comparison of Our Results with Statistics and Other Products

To assess whether our predicted plantation areas reproduce the ordering of official statistics, we used Spearman's rank correlation coefficient ρ, a non-parametric measure that quantifies the strength of a monotonic relationship between two variables by comparing the ranks of their values rather than the raw magnitudes. Here, the variables being ranked are the per-region oil palm areas: for each region we compare its rank in our predicted areas against its rank in the reference statistics, so that ρ measures how well our maps preserve the relative regional ordering of plantation extent. The general formula, which also handles tied ranks, is

$$\rho = \frac{\frac{1}{n}\sum_{i=1}^{n}\left(R(x_i) - \overline{R(x)}\right) \cdot \left(R(y_i) - \overline{R(y)}\right)}{\sqrt{\left(\frac{1}{n}\sum_{i=1}^{n}\left(R(x_i) - \overline{R(x)}\right)^2\right) \cdot \left(\frac{1}{n}\sum_{i=1}^{n}\left(R(y_i) - \overline{R(y)}\right)^2\right)}} \tag{15}$$

where x_i and y_i are the i-th paired observations, n denotes the total number of observations, and $\overline{R(x)}$ and $\overline{R(y)}$ are the average ranks of variables x and y, respectively. For data without tied ranks, both mean ranks equal $\frac{n+1}{2}$.

We assessed the regional rank-order consistency of our mapping results by comparing the predicted plantation areas against the CEIC Global Database [53] using Spearman's rank correlation coefficient (Equation (15)). Because the CEIC database provides state/province-level oil palm planted-area statistics for Malaysia but not for Indonesia, this comparison is restricted to the 13 Malaysian states within the study area. As illustrated in Figure 13, our predictions demonstrate a strong positive rank correlation with the Malaysian reference statistics across all three years, yielding coefficients of 0.95, 0.92, and 0.94 for 2020, 2022, and 2024, respectively. Furthermore, the calculated q-values (6.4×10^{-7}, 9.9×10^{-6}, and 1.9×10^{-6}) fall well below the standard significance threshold ($\alpha = 0.05$). These results support the regional ordering of mapped plantation extent in Malaysia, but they should not be interpreted as a direct validation of absolute area estimates or temporal change

magnitudes, and an equivalent statistical cross-check for Indonesia is not currently possible owing to the lack of a comparable region-level statistical source.

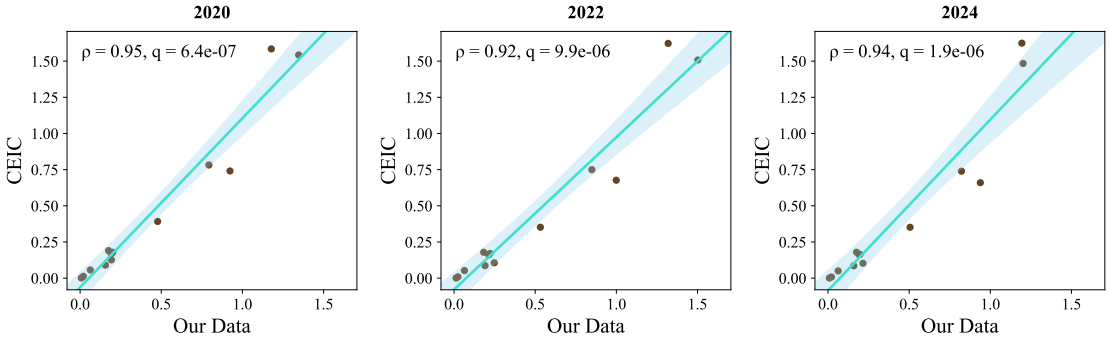


Figure 13. Comparison of mapped oil palm area with CEIC regional statistics for Malaysia (2020–2024). The CEIC database provides region-level oil palm statistics for Malaysia only; no equivalent region-level source is available for Indonesia.

To make the limitation on absolute magnitude explicit, we additionally compared our mapped areas against the CEIC statistics for Malaysia, where region-level statistics are available for the same administrative units used in our maps (Table 4). At the national level, our maps overestimate the reported planted area by roughly a factor of two (relative error of +109%, +150%, and +127% for 2020, 2022, and 2024), and the mean absolute per-region relative error is correspondingly large (MAPE of 100%, 153%, and 110%). This systematic overestimation is consistent with the moderate User's Accuracy (approximately 50%–55%) reported in Table 1: the high commission error of the binary classifier inflates absolute area, particularly in states such as Perlis, Selangor, and Pulau Pinang where the officially reported area is small and spectral confusion with other perennial vegetation is proportionally more damaging. Conversely, in the two largest producing states, Sabah and Sarawak, the maps slightly underestimate the CEIC area (by 12%–27%). Importantly, despite these large absolute discrepancies, the relative ranking of regions is preserved almost exactly (Spearman's $\rho \geq 0.92$), which is why we restrict our quantitative claims to regional ordering and relative trajectories rather than absolute planted area. We therefore recommend that the absolute areas reported here be treated as upper-bound indicators and be rescaled against official statistics before use in area-accounting applications.

Table 4. Comparison of mapped oil palm area (this study) with CEIC regional statistics for Malaysia (in 10^3 ha), with relative error RE = (mapped − CEIC)/CEIC. Only states inside the study area with reported CEIC statistics are shown.

State	2020			2022			2024		
	Mapped	CEIC	RE%	Mapped	CEIC	RE%	Mapped	CEIC	RE%
Johor	923	741	+25	999	677	+48	939	660	+42
Kedah	156	90	+74	191	86	+121	159	86	+86
Kelantan	192	168	+15	214	162	+32	191	164	+17
Melaka	63	56	+11	64	52	+23	63	50	+26
Negeri Sembilan	175	190	−8	183	179	+2	176	179	−2
Pahang	793	782	+1	849	750	+13	822	739	+11
Perak	478	392	+22	532	352	+51	506	351	+44
Perlis	8	0.7	+988	13	0.9	+1334	8	0.9	+857
Pulau Pinang	20	13	+56	24	9	+182	21	8	+175
Sabah	1346	1543	−13	1502	1508	−0.4	1202	1484	−19
Sarawak	1178	1585	−26	1319	1622	−19	1194	1624	−27
Selangor	193	127	+52	248	106	+134	215	103	+109
Trengganu	201	179	+13	222	171	+30	198	164	+20
National	**5727**	**2738**	**+109**	**6360**	**2544**	**+150**	**5693**	**2505**	**+127**

6.2. Interpretation of Regional Trends and Potential Drivers

The regional patterns reported in Section 5.6 admit several plausible, non-exclusive explanations. In Malaysia, the overall stability of the mapped areas is consistent with the maturity of the national oil palm sector, and the concentration of plantations in Malaysian Borneo reflects favorable agro-ecological conditions and extensive land availability [54,55], whereas the lower extent on the Malay Peninsula is likely related to more limited expansion space and less suitable soil conditions. This spatial disparity highlights the combined influence of climate, soil characteristics, and historical land-use patterns. The faster land use/cover change rates observed on the peninsula are potentially driven by the expansion of oil palm processing capacity and encroachment into previously undeveloped areas following land saturation [56], while the decline mapped in Sabah and Sarawak between 2022 and 2024 is potentially due to policy adjustments, land-use reallocation, or market-driven replanting cycles [57,58].

In Indonesia, the contraction of the mapped area after the 2022 peak—most pronounced in Riau—may reflect environmental degradation, policy restrictions on expansion, shifts in market demand, or economic factors such as commodity price fluctuations and trade regulations [57,58]. The divergence between continued expansion in Kalimantan and stagnation or decline in Sumatra suggests land saturation or stricter regulations in the older production zones, and highlights the influence of environmental, economic, and policy factors shaping Indonesia's oil palm landscape. Part of the apparent contraction may also reflect classification uncertainty and the temporal label bias discussed below rather than genuine land-use change alone, so further investigation would be necessary to determine the primary drivers behind this reduction in mapped oil palm cultivation.

The transition flows reported in Section 5.7 can be interpreted in the same cautious spirit. The decelerating conversion into oil palm from built area, rangeland, and water, combined with the accelerating losses to these classes, suggests that oil palm expansion is slowing in these sectors in favor of urbanization and reversion to rangeland or water bodies. The increasing conversion of flooded vegetation into oil palm, together with the declining reverse flow, is indicative of encroachment into wetland areas and may signal ecological risks that merit closer monitoring, particularly in Kalimantan. The reduction in bidirectional exchange between oil palm and crops suggests a stabilization of agricultural rotation, although this exchange remains the largest single component of the observed land-cover change.

6.3. Reliability of the Multi-Temporal Change Analysis

A potential concern regarding our change analysis is that, given per-year overall accuracies in the 60%–70% range, the naive product of two single-year accuracies (e.g., $0.63 \times 0.63 \approx 0.40$) would suggest a per-pixel change-detection reliability below 40%. We argue that this lower bound substantially understates the reliability of our *aggregate* analysis, for three reasons. First, the multiplicative bound assumes that errors in different years are statistically independent. In practice, classification errors in remote sensing are strongly spatially and temporally correlated: a pixel misclassified in 2020 because of terrain, mixed land cover, or persistent spectral confusion (e.g., with coconut or rubber) is likely to be misclassified in 2022 and 2024 as well, because the underlying cause is stationary. Such correlated errors largely cancel when computing *net change*, so the realized joint accuracy of a change estimate is considerably higher than the independence-based product implies. Second, our analysis is conducted at the level of regional and national aggregates rather than individual pixels. Under spatial aggregation, zero-mean random classification errors partially offset one another and the variance of the area estimate decreases as the number of aggregated pixels grows; consequently, regional area trends

are far more robust to per-pixel error than pixel-level change maps. Third, the temporal consistency of our aggregate estimates is corroborated by an independent reference: for Malaysia, where region-level statistics are available, the predicted plantation areas correlate strongly with official CEIC economic statistics (Spearman's $\rho = 0.95$, 0.92, and 0.94 for 2020, 2022, and 2024; Figure 13), which would not be expected if the year-to-year trends were dominated by random error. We therefore interpret our results as reliable indicators of *regional-scale trajectories* rather than pixel-accurate change maps, and we frame all change statements accordingly. We emphasize, however, that the Spearman comparison assesses only the rank-order agreement between mapped and reported areas; it does not validate the absolute magnitude of the mapped areas or of the year-to-year changes. As quantified in Table 4, the mapped areas in fact overestimate the official CEIC statistics by roughly a factor of two at the national level, consistent with the moderate User's Accuracy of the classifier; the strong rank correlation is therefore a statement about relative regional ordering rather than absolute area, and absolute planted-area values should be rescaled against official statistics before being used quantitatively.

6.4. Uncertainty of Our Results

Although our proposed method significantly improves spatial resolution from 100 m to 10 m, several sources of uncertainty remain in the final oil palm plantation maps. The primary source of uncertainty stems from the propagation of errors inherent in the training labels. As noted in the AOPD documentation [16], the original 100 m datasets contain confusion between oil palm and other vegetation types, particularly in impervious areas and mixed-crop plantations. While the DMI loss function [44] effectively mitigates the impact of noisy labels during training, it cannot fully eliminate the bias introduced when the ground truth itself contains systematic classification errors.

Secondly, the exclusive reliance on optical Sentinel-2 imagery introduces susceptibility to atmospheric interference, which is a pervasive challenge in the tropical climates of Indonesia and Malaysia. Although we utilized surface reflectance products, persistent cloud cover, haze, and cloud shadows can distort the spectral reflectance values of the bands. In the training phase, this spectral noise may prevent the model from learning a consistent feature representation for oil palm. In the prediction phase, residual thin clouds or unmasked shadows can alter the radiometric integrity of pixels, potentially leading to errors in persistently cloudy regions.

Thirdly, spectral confusion remains a challenge in optical remote sensing [7]. Mature oil palm trees share similar spectral signatures with other tropical perennial crops, most notably coconut palms and rubber plantations. In regions like coastal Indonesia, where coconut and oil palm plantations frequently coexist, the model may generate false positives. This uncertainty is exacerbated in smallholder plantations where crops are often intercropped rather than planted in distinct, large-scale monocultures.

Finally, uncertainty is introduced through the resolution mismatch during the validation process. Our model predicts boundaries at a 10 m granularity, capturing small roads and drainage canals within plantations. However, comparing these fine-grained predictions against 100 m resolution historical data or statistical aggregates often results in a "mixed pixel" problem, where the coarse resolution of the reference data fails to capture the fragmented nature of plantation edges, potentially skewing the accuracy metrics in transition zones.

A further and systematic source of uncertainty is temporal bias in the training data. Our model is supervised exclusively with the 2016 AOPD product, the most recent year available, yet it is applied to imagery from 2020, 2022, and 2024. Because plantation extent and structure evolve through replanting, expansion, and abandonment, the mismatch

between the 2016 supervision and the prediction year grows with temporal distance. This mechanism offers a parsimonious explanation for the monotonic decline in overall accuracy that we observe over time (70.64% in 2020, 63.53% in 2022, and 60.06% in 2024): the further a target year is from 2016, the less the learned 2016 spatial distribution reflects the true contemporary landscape. This temporal bias should be considered when interpreting absolute accuracies for the later years, and it motivates future updates of the reference product or self-training schemes that propagate labels forward in time.

We also acknowledge a limitation in our characterization of label noise. Both the Sentinel-2 imagery and the AOPD reference product contain inherent errors, and the true noise level in our training data is therefore neither known nor directly controllable. We treat the coarse 100 m AOPD product as a form of noisy supervision, but we did not conduct controlled noise-injection experiments (e.g., systematically flipping or spatially shifting a known fraction of training samples) that would directly quantify the robustness of the DMI loss to increasing noise relative to standard cross-entropy. While our ablation against the BCE baseline (Table 3) provides indirect evidence of this robustness, a controlled noise-sensitivity study remains an important direction for future work and would more rigorously establish the advantage of determinant-based mutual information under quantified label corruption. We note, however, that a synthetic noise-injection experiment would necessarily add controlled noise *on top of* the already-unknown noise inherent in the 2016 AOPD supervision, so its absolute noise levels could not be calibrated to the real operating point of this study; we therefore rely on the direct DMI-versus-BCE comparison under the identical, real coarse-label supervision as the most representative evidence available for our resolution-mismatch setting.

We also acknowledge that our ablation analysis is deliberately focused and does not exhaustively isolate every component of the pipeline. In particular, we report the DMI-versus-BCE loss comparison (Table 3) as the central ablation, but we did not run additional controlled ablations such as removing the Bayesian AOPD fusion, disabling the class-balancing (random undersampling) step, or substituting alternative noise-robust losses (e.g., GCE or forward/backward loss correction). As a result, the individual contribution of the Bayesian fusion and the class-balancing step to the final accuracy is not separately quantified here, and our claim of superiority and robustness should be understood as established primarily against the BCE baseline under identical coarse-label supervision rather than against the full space of noise-robust alternatives. Regarding the five-epoch training budget, we did not include formal learning curves; instead, as described in Section 5, we observed that validation accuracy saturated early under the coarse 100 m supervision and that longer training degraded 10 m boundary quality by memorizing blocky label artifacts, and we verified that the heavier transformer- and Mamba-based baselines had likewise reached their validation plateau within this budget. A more comprehensive ablation suite, together with explicit per-architecture learning curves, would further strengthen these conclusions and is left for future work.

Finally, the magnitude of area change observed over a five-year window (2020–2024) may appear large relative to the roughly 25-year productive lifespan of an oil palm stand. Several factors reconcile this. First, the observed dynamics largely reflect the replanting cycle rather than the establishment of entirely new stands: aging or low-yielding blocks are cleared and replanted, and during the immature phase their canopy signature differs enough from mature palm that they may not be detected as oil palm, producing apparent short-term losses that are in fact part of normal estate management. Second, policy interventions (e.g., moratoria on new concessions and sustainability-certification pressure) and market forces (commodity price volatility and shifting trade regulations) can drive measurable reallocation of land over short intervals. Third, part of the observed variability

is attributable to the temporal label bias and per-year classification uncertainty discussed above. Disentangling genuine land-use change from replanting-phase detection gaps is an important avenue for future work, ideally supported by stand-age information or radar time series that are less sensitive to canopy maturity.

7. Conclusions

This study addressed an important gap in high-resolution, up-to-date monitoring of oil palm plantations across Malaysia and Indonesia. By leveraging a noise-robust deep learning framework, we successfully generated 10 m resolution maps for the years 2020, 2022, and 2024, overcoming the significant challenge of supervising high-resolution Sentinel-2 imagery with coarse 100 m historical labels. Our experimental results demonstrated that the U-Net architecture, trained with DMI loss, exhibited greater robustness to label noise than more complex state-of-the-art models such as MMSFormer and UNetFormer, achieving a maximum overall accuracy of 70.64% in 2020.

The spatiotemporal analysis derived from these maps provided useful insights into regional plantation dynamics. At the regional scale, the mapped area suggests a peak in 2022 followed by a lower mapped extent in 2024. Regionally, the mapped extent remains concentrated in Kalimantan, while Sumatra exhibits signs of saturation or decline. Our land cover transition analysis indicates that, while the exchange between oil palm and general crops constitutes the largest volume of change, the conversion of flooded vegetation (wetlands) into oil palm plantations appears to have increased between 2022 and 2024, which may signal ecological risks and merits closer monitoring. Given the moderate per-year accuracies and the temporal mismatch between the 2016 supervision and the later prediction years, these results should be interpreted as regional-scale indicators rather than pixel-level change maps. The generated maps can support regional monitoring, sustainability assessment, and the prioritization of areas for further validation in support of zero-deforestation commitments and sustainable agricultural development.

Author Contributions: Conceptualization, J.Z. and Z.Z.; methodology, N.K. and Z.Z.; software, N.K., B.S. and J.L.; validation, N.K., J.L., J.J., J.H. and S.T.; formal analysis, N.K. and Z.Z.; investigation, N.K. and Z.Z.; data curation, N.K., J.L. and J.J.; writing—original draft preparation, N.K. and Z.Z.; writing—review and editing, B.S., J.L., J.J., J.H., S.T. and J.Z.; visualization, N.K. and Z.Z.; supervision, J.Z.; project administration, J.Z.; funding acquisition, J.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Natural Science Foundation of China (Grant No. 42401415), and in part by the Shenzhen Science and Technology Program under Grant KCXFZ20240903093759004 and Grant KJZD20230923115106012.

Data Availability Statement: Sentinel-2 multispectral imagery used in this study is publicly available from the Copernicus Data Space Ecosystem: https://dataspace.copernicus.eu/data-collections/copernicus-sentinel-data/sentinel-2 (accessed on 8 July 2026). The nationwide 10 m oil palm mapping dataset for Malaysia and Indonesia (2020–2024) is publicly available on Zenodo under a CC-BY 4.0 license at https://doi.org/10.5281/zenodo.17768444. Additional data and materials are available from the corresponding author upon reasonable request.

Acknowledgments: The authors would like to thank the Copernicus Data Space Ecosystem and Google Earth Engine for providing access to Sentinel-2 data. The nationwide 10 m oil palm maps (2020–2024) are publicly available on Zenodo under a CC-BY 4.0 license.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; the collection, analysis, or interpretation of data; the writing of the manuscript; or the decision to publish the results.

References

1. Ritchie, H. Palm Oil. *Our World in Data*. 2021. Available online: https://ourworldindata.org/palm-oil (accessed on 8 July 2026).
2. Gatti, R.C.; Velichevskaya, A. Certified "sustainable" palm oil took the place of endangered Bornean and Sumatran large mammals habitat and tropical forests in the last 30 years. *Sci. Total Environ.* **2020**, *742*, 140712. [CrossRef]
3. Hanafiah, K.M.; Mutalib, A.H.A.; Miard, P.; Goh, C.S.; Sah, S.A.M.; Ruppert, N. Impact of Malaysian palm oil on sustainable development goals: Co-benefits and trade-offs across mitigation strategies. *Sustain. Sci.* **2022**, *17*, 1639–1661.
4. Dhandapani, S.; Ritz, K.; Evers, S.; Sjögersten, S. GHG emission under different cropping systems in some Histosols of Malaysia. *Geoderma Reg.* **2019**, *18*, e00229. [CrossRef]
5. Dohong, A.; Aziz, A.A.; Dargusch, P. Carbon emissions from oil palm development on deep peat soil in Central Kalimantan Indonesia. *Anthropocene* **2018**, *22*, 31–39. [CrossRef]
6. Danylo, O.; Pirker, J.; Lemoine, G.; Ceccherini, G.; See, L.; McCallum, I.; Hadi; Kraxner, F.; Achard, F.; Fritz, S. A map of the extent and year of detection of oil palm plantations in Indonesia, Malaysia and Thailand. *Sci. Data* **2021**, *8*, 96. [CrossRef] [PubMed]
7. Nurmasari, Y.; Wijayanto, A.W. Oil palm plantation detection in indonesia using sentinel-2 and landsat-8 optical satellite imagery (case study: Rokan Hulu Regency, Riau Province). *Int. J. Remote Sens. Earth Sci.* **2021**, *18*, 1–18. [CrossRef]
8. Descals, A.; Wich, S.; Meijaard, E.; Gaveau, D.L.A.; Peedell, S.; Szantoi, Z. High-resolution global map of smallholder and industrial closed-canopy oil palm plantations. *Earth Syst. Sci. Data* **2021**, 13, 1211–1231. [CrossRef]
9. Rodríguez, A.C.; D'Aronco, S.; Schindler, K.; Wegner, J.D. Mapping oil palm density at country scale: An active learning approach. *Remote Sens. Environ.* **2021**, *261*, 112479. [CrossRef]
10. Putra, Y.C.; Wijayanto, A.W. Automatic detection and counting of oil palm trees using remote sensing and object-based deep learning. *Remote Sens. Appl. Soc. Environ.* **2023**, *29*, 100914. [CrossRef]
11. Mirpulatov, I.; Illarionova, S.; Shadrin, D.; Burnaev, E. Pseudo-Labeling Approach for Land Cover Classification Through Remote Sensing Observations With Noisy Labels. *IEEE Access* **2023**, *11*, 82570–82583. [CrossRef]
12. Hell, M.; Brandmeier, M. Identifying Plausible Labels from Noisy Training Data for a Land Use and Land Cover Classification Application in Amazônia Legal. *Remote Sens.* **2024**, *16*, 2080. [CrossRef]
13. Xu, G.; Fang, Y.; Deng, M.; Sun, G.; Chen, J. Remote Sensing Mapping of Build-Up Land with Noisy Label via Fault-Tolerant Learning. *Remote Sens.* **2022**, *14*, 2263. [CrossRef]
14. Zhang, Z.; Ye, Z.; Wen, Y.; Yuan, S.; Fu, H.; Jianxi, H.; Zheng, J. GTPBD: A Fine-Grained Global Terraced Parcel and Boundary Dataset. In Proceedings of the Advances in Neural Information Processing Systems 38 (NeurIPS 2025), San Diego, CA, USA, 2–7 December 2025.
15. Zheng, J.; Yuan, S.; Li, W.; Fu, H.; Yu, L.; Huang, J. A Review of Individual Tree Crown Detection and Delineation From Optical Remote Sensing Images: Current progress and future. *IEEE Geosci. Remote Sens. Mag.* **2024**, *13*, 209–236. [CrossRef]
16. Xu, Y.; Yu, L.; Li, W.; Ciais, P.; Cheng, Y.; Gong, P. Annual oil palm plantation maps in Malaysia and Indonesia from 2001 to 2016. *Earth Syst. Sci. Data* **2020**, *12*, 847–867. [CrossRef]
17. Descals, A.; Gaveau, D.L.A.; Wich, S.; Szantoi, Z.; Meijaard, E. Global mapping of oil palm planting year from 1990 to 2021. *Earth Syst. Sci. Data* **2024**, *16*, 5111–5129. [CrossRef]
18. Zheng, J.; Fu, H.; Li, W.; Wu, W.; Zhao, Y.; Dong, R.; Yu, L. Cross-regional oil palm tree counting and detection via a multi-level attention domain adaptation network. *ISPRS J. Photogramm. Remote Sens.* **2020**, *167*, 154–177. [CrossRef]
19. Zhang, Z.; Zeng, X.; Kong, X.; Zhang, K.; Liang, H.; Shi, B.; Zheng, J.; Huang, J.; Lu, Y.; Fu, H. GTPBD-MM: A Global Terraced Parcel and Boundary Dataset with Multi-Modality. *arXiv* **2026**, arXiv:2604.12315.
20. Esri; ImpactObservatory; Microsoft. Sentinel-2 10m Land Use/Land Cover Time Series, 2024. Data Derived from European Space Agency (ESA) Sentinel-2 Imagery. Accessed via the Sentinel-2 Land Cover Explorer Application. Available online: https://www.pacificgeoportal.com/datasets/cfcb7609de5f478eb7666240902d4d3d/about (accessed on 8 July 2026).
21. Zheng, J.; Wu, W.; Yu, L.; Fu, H. Coconut trees detection on the tenarunga using high-resolution satellite images and deep learning. In *Proceedings of the 2021 IEEE International Geoscience and Remote Sensing Symposium IGARSS*; IEEE: Piscataway, NJ, USA, 2021; pp. 6512–6515.
22. Daliman, S.; Abu-Bakar, S.A.R.; Busu, I. Oil Palm Tree Enumeration Based on Template Matching. In Proceedings of the World Virtual Conference on Applied Sciences and Engineering Applications (WVCASEA2015), Johor Bahru, Malaysia, 27–29 March 2015.
23. Shafri, H.Z.M.; Humdan, N.; Saripan, M.I. Semi-automatic detection and counting of oil palm trees from high spatial resolution airborne imagery. *Int. J. Remote Sens.* **2011**, *32*, 2095–2115. [CrossRef]
24. Wang, M.; Mao, D.; Wang, Y.; Xiao, X.; Xiang, H.; Feng, K.; Luo, L.; Jia, M.; Song, K.; Wang, Z. Wetland mapping in East Asia by two-stage object-based Random Forest and hierarchical decision tree algorithms on Sentinel-1/2 images. *Remote Sens. Environ.* **2023**, *297*, 113793. [CrossRef]
25. Wijayanto, A.W.; Afira, N.; Nurkarim, W. Machine Learning Approaches using Satellite Data for Oil Palm Area Detection in Pekanbaru City, Riau. In Proceedings of the 2022 IEEE International Conference on Cybernetics and Computational Intelligence (CyberneticsCom), Malang, Indonesia, 16–18 June 2022; pp. 84–89.

26. Crasto, N.; Hopkinson, C.; Forbes, D.; Lesack, L.; Marsh, P.; Spooner, I.; van der Sanden, J. A LiDAR-based decision-tree classification of open water surfaces in an Arctic delta. *Remote Sens. Environ.* **2015**, *164*, 90–102. [CrossRef]
27. Frenay, B.; Verleysen, M. Classification in the Presence of Label Noise: A Survey. *IEEE Trans. Neural Netw. Learn. Syst.* **2014**, *25*, 845–869. [CrossRef] [PubMed]
28. Wright, N.; Duncan, J.M.; Callow, J.N.; Thompson, S.E.; George, R.J. Training sensor-agnostic deep learning models for remote sensing: Achieving state-of-the-art cloud and cloud shadow identification with OmniCloudMask. *Remote Sens. Environ.* **2025**, *322*, 114694. [CrossRef]
29. Radman, A.; Mahdianpari, M.; Varon, D.J.; Mohammadimanesh, F. S2MetNet: A novel dataset and deep learning benchmark for methane point source quantification using Sentinel-2 satellite imagery. *Remote Sens. Environ.* **2023**, *295*, 113708. [CrossRef]
30. Dixon, D.J.; Zhu, Y.; Brown, C.F.; Jin, Y. Satellite detection of canopy-scale tree mortality and survival from California wildfires with spatio-temporal deep learning. *Remote Sens. Environ.* **2023**, *298*, 113842. [CrossRef]
31. Robinson, C.; Hou, L.; Malkin, K.; Soobitsky, R.; Czawlytko, J.; Dilkina, B.; Jojic, N. Large Scale High-Resolution Land Cover Mapping With Multi-Resolution Data. In Proceedings of the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, 15–20 June 2019; pp. 12718–12727.
32. Malkin, K.; Robinson, C.; Hou, L.; Soobitsky, R.; Czawlytko, J.; Samaras, D.; Saltz, J.; Joppa, L.; Jojic, N. Label Super-Resolution Networks. In Proceedings of the International Conference on Learning Representations, New Orleans, LA, USA, 6–9 May 2019.
33. Zheng, J.; Fu, H.; Li, W.; Wu, W.; Yu, L.; Yuan, S.; Tao, W.Y.W.; Pang, T.K.; Kanniah, K.D. Growing status observation for oil palm trees using Unmanned Aerial Vehicle (UAV) images. *ISPRS J. Photogramm. Remote Sens.* **2021**, *173*, 95–121. [CrossRef]
34. Li, Z.; Zhang, H.; Lu, F.; Xue, R.; Yang, G.; Zhang, L. Breaking the resolution barrier: A low-to-high network for large-scale high-resolution land-cover mapping using low-resolution labels. *ISPRS J. Photogramm. Remote Sens.* **2022**, *192*, 244–267. [CrossRef]
35. Chen, Y.; Zhang, G.; Cui, H.; Li, X.; Hou, S.; Ma, J.; Li, Z.; Li, H.; Wang, H. A novel weakly supervised semantic segmentation framework to improve the resolution of land cover product. *ISPRS J. Photogramm. Remote Sens.* **2023**, *196*, 73–92. [CrossRef]
36. Li, Z.; He, W.; Cheng, M.; Hu, J.; Yang, G.; Zhang, H. SinoLC-1: The first 1 m resolution national-scale land-cover map of China created with a deep learning framework and open-access data. *Earth Syst. Sci. Data* **2023**, *15*, 4749–4780. [CrossRef]
37. Li, Z.; He, W.; Li, J.; Lu, F.; Zhang, H. Learning without Exact Guidance: Updating Large-Scale High-Resolution Land Cover Maps from Low-Resolution Historical Labels. In Proceedings of the 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Seattle, WA, USA, 17–21 June 2024; pp. 27717–27727.
38. Ronneberger, O.; Fischer, P.; Brox, T. U-Net: Convolutional Networks for Biomedical Image Segmentation. *arXiv* **2015**, arXiv:1505.04597.
39. Reza, M.K.; Prater-Bennette, A.; Asif, M.S. MMSFormer: Multimodal Transformer for Material and Semantic Segmentation. *arXiv* **2023**, arXiv:2309.04001.
40. Chen, L.C.; Zhu, Y.; Papandreou, G.; Schroff, F.; Adam, H. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. *arXiv* **2018**, arXiv:1802.02611.
41. Wang, L.; Li, R.; Zhang, C.; Fang, S.; Duan, C.; Meng, X.; Atkinson, P.M. UNetFormer: A UNet-like transformer for efficient semantic segmentation of remote sensing urban scene imagery. *ISPRS J. Photogramm. Remote Sens.* **2022**, *190*, 196–214. [CrossRef]
42. Wang, J.; Sun, K.; Cheng, T.; Jiang, B.; Deng, C.; Zhao, Y.; Liu, D.; Mu, Y.; Tan, M.; Wang, X.; et al. Deep High-Resolution Representation Learning for Visual Recognition. *arXiv* **2019**, arXiv:1908.07919.
43. Wang, L.; Li, D.; Dong, S.; Meng, X.; Zhang, X.; Hong, D. PyramidMamba: Rethinking Pyramid Feature Fusion with Selective Space State Model for Semantic Segmentation of Remote Sensing Imagery. *arXiv* **2024**, arXiv:2406.10828.
44. Xu, Y.; Cao, P.; Kong, Y.; Wang, Y. L_dmi: A novel information-theoretic loss function for training deep nets robust to label noise. In Proceedings of the Advances in Neural Information Processing Systems 32 (NeurIPS 2019), Vancouver, BC, Canada, 8–14 December 2019.
45. Ioffe, S.; Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Proceedings of the 32nd International Conference on International Conference on Machine Learning-JMLR.org , Lille, France, 7–9 July 2015; ICML'15, Volume 37, pp. 448–456.
46. Kong, Y. Dominantly truthful multi-task peer prediction with a constant number of tasks. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms*; Society for Industrial and Applied Mathematics: Philadelphia, PA, USA, 2020; SODA'20, pp. 2398–2411.
47. Zhang, Z.; li, Y.; Li, Y.; Qin, Y. Making Deep Neural Networks Robust to Label Noise: A Reweighting Loss and Data Filtration. In Proceedings of the 2019 IEEE 4th International Conference on Signal and Image Processing (ICSIP), Wuxi, China, 19–21 July 2019; pp. 289–293. [CrossRef]
48. Zhang, Z.; Sabuncu, M.R. Generalized Cross Entropy Loss for Training Deep Neural Networks with Noisy Labels. *arXiv* **2018**, arXiv:1805.07836.
49. Bottou, L. Large-Scale Machine Learning with Stochastic Gradient Descent. In *Proceedings of COMPSTAT'2010*; Lechevallier, Y., Saporta, G., Eds.; Physica-Verlag HD : Heidelberg, Germany, 2010; pp. 177–186.

50. Zhao, W.; Chellappa, R. Chapter 21–Multimodal Biometrics: Augmenting Face with Other Cues. In *Face Processing*; Academic Press: Burlington, NJ, USA, 2006; pp. 679–705.
51. Baum, E.B.; Wilczek, F. Supervised Learning of Probability Distributions by Neural Networks. In Proceedings of the Neural Information Processing Systems, Denver, CO, USA, 28 November–1 December 1988; pp. 52–61.
52. Zhao, Q.; Yu, L.; Li, X.; Xu, Y.; Du, Z.; Kanniah, K.; Li, C.; Cai, W.; Lin, H.; Peng, D.; et al. The expansion and remaining suitable areas of global oil palm plantations. *Glob. Sustain.* **2024**, *7*, e9. [CrossRef]
53. CEIC Data. CEIC Global Database: Malaysia and Indonesia Oil Palm Planted Area [Data Set]. 2023. Accessed via the CEIC Flex Student Plan. Available online: https://www.ceicdata.com/en/malaysia/planted-area-oil-palm/oil-palm-planted-area (accessed on 2 September 2025).
54. Kuok Ho, D.T.; Qahtani, H. Sustainability of oil palm plantations in Malaysia. *Environ. Dev. Sustain.* **2020**, *22*, 4999–5023.
55. Shevade, V.S.; Loboda, T.V. Oil palm plantations in Peninsular Malaysia: Determinants and constraints on expansion. *PLoS ONE* **2019**, *14*, e0210628. [CrossRef] [PubMed]
56. Alam, A.S.A.F.; Er, A.C.; Begum, H. Malaysian oil palm industry: Prospect and problem. *J. Food Agric. Environ.* **2015**, *1313*, 143–148.
57. Gaveau, D.L.A.; Locatelli, B.; Salim, M.A.; Husnayaen.; Manurung, T.; Descals, A.; Angelsen, A.; Meijaard, E.; Sheil, D. Slowing deforestation in Indonesia follows declining oil palm expansion and lower oil prices. *PLoS ONE* **2022**, *17*, e0266178. [CrossRef] [PubMed]
58. Astari, A.J.; Lovett, J.C.; Wasesa, M. Sustainable pathways in Indonesia's palm oil industry through historical institutionalism. *World Dev. Sustain.* **2025**, *6*, 100200. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.