Integrating radar and multi-spectral data to detect cocoa crops
A deep learning approach

Therias, Adele; Rafiee, Azarakhsh; Lhermitte, Stef; van der Lugt, Philip; Lindenbergh, Roderik

DOI
10.1016/j.rsase.2025.101652

Publication date
2025

Document Version
Final published version

Published in
Remote Sensing Applications: Society and Environment

Remote Sensing Applications: Society and Environment 39 (2025) 101652

Contents lists available at ScienceDirect

Remote Sensing Applications: Society and Environment

journal homepage: www.elsevier.com/locate/rsase

Integrating radar and multi-spectral data to detect cocoa crops: a deep learning approach

Adele Therias [a], Azarakhsh Rafiee [a,*], Stef Lhermitte [b,c], Philip van der Lugt [d], Roderik Lindenbergh [c]

[a] Architectural Engineering + Technology Department, Delft University of Technology, Delft, the Netherlands
[b] Department Earth & Environmental Sciences, KU Leuven, Leuven, Belgium
[c] Department Geoscience & Remote Sensing, Delft University of Technology, Delft, the Netherlands
[d] LugtAarde Consult, the Netherlands

ARTICLE INFO

Keywords:
cocoa farms
SAR
MSI
U-NET
Semantic segmentation
Temporality
Polarization

ABSTRACT

The production of cocoa beans contributes to 7.5 % of European Union (EU) driven deforestation. As a result, the recent European Union Deforestation-free Regulation (EUDR) mandates producers to track cocoa farm extents comprehensively. While Remote Sensing has enormous capacity in dynamic crop monitoring, cocoa crop detection shows challenges due to cocoa complex canopy structure, spectral similarity to forest, variable farming methods, and location in frequently cloudy regions. Previous research on cocoa crop detection has mainly focused on pixel-based classification, disregarding spatial context. In this research we have performed a semantic segmentation approach to incorporate spatial configuration and enhance cocoa crop detection. We have applied Convolutional Neural Network (CNN) for the to semantic segmentation of cocoa parcels, considering both spectral and spatial characteristics. Additionally, we have evaluated the impact of combining Synthetic Aperture RADAR (SAR) and MSI (Multi-Spectral Imagery) data in the training of a CNN to demonstrate the importance of texture, moisture, and canopy characteristics in identifying cocoa canopies. The impact of MSI dataset stack with different SAR polarizations, seasons and temporality has been evaluated. The methodology is tested on Sentinel 1 and 2 data over an area of 100 × 100 km in Ghana for which an extensive ground truth data set of almost 90,000 polygons was available for training and validation. The results show that the addition of single-day and temporal SAR to a single-day MSI image can improve the predictions, reaching an F1 score of 86.62 %. This research demonstrates the influence of SAR measurements, seasons, polarization, and ground truth classes on the semantic segmentation of cocoa.

1. Introduction

Forests play a key role in the functioning of ecosystems at local and global scales. They provide habitat for over 80 % of terrestrial biodiversity and the sequestration of around 289 Gt of carbon (Kumar et al., 2022). However, forests are under significant threat from human activities such as agricultural expansion, infrastructure development, and illegal logging (Bastin et al., 2019). The loss of forests also has severe and long-term impacts on ecological and human well-being by contributing to degradation of ecosystems, rising

* Corresponding author.
E-mail address: a.rafiee@tudelft.nl (A. Rafiee).

https://doi.org/10.1016/j.rsase.2025.101652
Received 30 September 2024; Received in revised form 15 April 2025; Accepted 30 June 2025
Available online 1 July 2025

temperatures, habitat destruction, pollution, and soil degradation (Kumar et al., 2022; Houghton, 2012). Deforestation remains a significant global concern, with tropical forests being particularly vulnerable. According to the Food and Agriculture Organization (FAO), the world loses approximately 10 million hectares of forest annually, contributing to biodiversity loss, carbon emissions, and disruptions in local climate patterns (FAO, 2020). Between 2015 and 2020, Africa had the highest rate of deforestation globally. A total of 4.41 million hectares of land were cleared, including 1.90 million hectares in Central and Western Africa (Global forest resources assessment, 2020). West Africa is one of the main producers of cocoa beans (Abu et al., 2021), which are estimated to cause 7.54 % of EU-driven deforestation(European Commission, 2021). In December 2022 the European Union (EU) approved the European Union Deforestation-free Regulation (EUDR). The Regulation aims to reduce the impact of EU consumption on global deforestation by banning the import of products, including cattle, wood, palm oil, soy, coffee and cocoa, issued from deforested areas (European Commission, 2021).

To preserve forests and promote responsible farming, satellite data is used by policymakers to aid in decision-making and enforce laws, wherein crop detection is an essential task in proper forest monitoring. Using optical spaceborne Multispectral Imagery (MSI) various crops can be detected via automated classification (Abu et al., 2021). However, the detection cocoa presents unique challenges. First, West Africa has frequent cloud cover due to its semi-equatorial climate. This limits the availability of cloud-free MSI and thereby reduces its temporal resolution (Ashiagbor et al., 2020). Second, agroforest cocoa crops, a frequent practice which integrates shade trees and other crops to improve growing conditions, has a spectral signature and canopy structure like nearby forest (Ashiagbor et al., 2020). Third, the canopy structure of cocoa can vary widely, e.g. due to differences in farming practices (Kaba et al., 2020; Asitoakor et al., 2022; Sonwa et al., 2019). Whereas the application of MSI for cocoa detection is limited by atmospheric and spectral characteristics, Synthetic Aperture Radar (SAR) backscatter can provide insight into the texture, geometry, canopy density, and water content of vegetation (Jansing, 2021) without interruption due to cloud cover (Evans et al., 1993). The polarization of a radar pulse influences the types of surface characteristics that are detected (Jansing, 2021). For instance, co-polarized Vertical-Vertical (VV) has the strongest backscatter when interacting with a rough and dry soil surface. Cross-polarized Vertical-Horizontal (VH) has the greatest intensity because of the volume scattering created by a larger amount of biomass (Numbisi and Van Coillie, 2020; Ruetschi et al., 2018; Jansing, 2021). Furthermore, when considered as a time-series, SAR can capture temporal variation in vegetated land cover (Vreugdenhil et al., 2018). Previous work has shown that a land classification of >85 % accuracy can be achieved when using imagery in two polarizations from four days over the course of one year (Thiel et al., 2009).

While there has been several studies on cocoa crop detection, the majority of existing research involves pixel-based classification using conventional Machine Learning (ML) algorithms. These studies consider only MSI, such as Landsat or Sentinel-2 (S2) data, to which they apply classification algorithms such as Maximum Likelihood Algorithm (Overall Accuracy (OA) = 82.6 %) (Benefoh et al., 2018), Random Forest (OA = 89.8 %) (Ashiagbor et al., 2022), and XGBoost (OA = 95.17 %) (Batista et al., 2022). Other studies use SAR-only datasets for classification via Supervised Maximum-likelihood Classifier (OA = 89 %) (Saatchi et al., 2001), Random Forest combined with Grey-Level Co-occurrence Matrix (GLCM) (OA = 88.1 %) (Numbisi et al., 2019), and Multi-Layer Perceptron Neural Networks Regression (Root Mean Squared Error (RMSE) = 7.18 %) (Numbisi and Van Coillie, 2020). There are two studies which combine SAR and MSI data to perform pixel-based classification using multi-feature Random Forest classifiers (Producer Accuracy (PA) = 82.9 % and User Accuracy (UA) = 62.2 % (Abu et al., 2021), and PA = 88 % and UA = 91 % (Tamga et al., 2022). The effectiveness of pixel-based classification is limited by its lacking consideration of spatial context, and that it does not reflect the patch-based composition of landscapes (Ashiagbor et al., 2020). Another two studies apply object-detection algorithms prior to cocoa detection using an Object-based Nearest Neighbour classifier (OA = 89 %) (Erasmi and Twele, 2009) and Random Forest classification (OA = 89.76 %) (Ashiagbor et al., 2020), both of which combine MSI and SAR. While this approach improves the spatial consideration in cocoa detection, object-based approaches are highly dependent on user-defined parameters and may lead to a higher number of misclassified pixels if entire patches are poorly detected (Ashiagbor et al., 2020).

Convolutional Neural Network (CNN)s enable the detection and classification of cocoa parcels by considering both spectral and spatial characteristics in a deep learning architecture (Kalischek et al., 2022). Two studies that demonstrate the potential for CNNs to detect the extent of cocoa farms in West Africa based on temporal MSI data are (Filella, 2018), using a U-NET architecture (Intersection over Union (IoU) = 58.2 %), and (Kalischek et al., 2022), using a custom-trained CNN with additional height prediction input (F1 score = 87.3 %). Despite the progress made with these methods, neither of these studies incorporate both MSI and SAR datasets as input for their deep learning network to explore their synergistic potential for improved accuracy in cocoa detection. In line with this, the objective of our research is to build upon existing work on cocoa detection via ML, by advancing Deep Learning (DL) approaches that integrate SAR data. Specifically, the study explores to what extent a CNN trained with multispectral *and* SAR datasets can enable the automated detection of cocoa crops in Ghana, while considering the impact of temporality, seasonality, and polarization on the segmentation results. Our proposed method involves the use of a U-NET, as an established state-of-the-art deep learning architecture, that combines MSI and SAR data from different seasons and polarizations to carry out binary semantic segmentation. Our study yields 10m resolution maps indicating the probability of cocoa crop presence. Finally, we analyze the influence of SAR addition on segmentation with different dataset stacks (i.e. single-day, temporal, dry and wet seasons) and evaluate the influence of various parameters on the results, including background class characteristics and the removal of clouds.

1.1. Cocoa characteristics

Cocoa is a small evergreen tree that produces pods containing cocoa beans. In Ghana, cocoa is cultivated under various systems that differ in canopy structure and associated vegetation. These systems are broadly classified into non-shaded (monoculture) cocoa systems, shaded cocoa systems (agroforestry cocoa), and intercropped cocoa systems, particularly during the early establishment phase.

Non-shaded cocoa systems—also referred to as monoculture or full-sun cocoa—are characterized by the absence or minimal presence of shade trees, typically with canopy cover below 15 %, and a relatively uniform cocoa canopy height of approximately 5.5 m (Numbisi and Van Coillie, 2020). In contrast, shaded cocoa systems incorporate shade trees that form a stratified canopy, where cocoa trees (~5 m) comprise the lower stratum and taller shade trees (>5 m) form an upper layer, often resulting in total canopy cover above 15 % (Blaser-Hart et al., 2021; Kaba et al., 2020).Cocoa can also be intercropped with food crops such as yam or plantain, a practice commonly observed during the early stages of farm development. However, even in mature shaded or non-shaded cocoa farms, food crops may occasionally be present beneath the cocoa canopy (Kaba et al., 2020). Although shaded cocoa systems may exhibit some structural similarities with open-canopy forests, several characteristics can help differentiate them. Natural forests generally have a denser and more vertically complex canopy—often exceeding 20 m in height—and exhibit greater structural stability over time due to the persistence of mature evergreen species (Numbisi and Van Coillie, 2020). In comparison, cocoa agroforests tend to have more open canopies with gaps, especially in non-shaded systems, and may show seasonal fluctuations such as the loss of shade tree foliage during the dry season. Vegetation height also provides insight, as shade trees within cocoa farms vary in crown shape and stature, forming distinct upper and lower canopy layers (Blaser-Hart et al., 2021; Asitoakor et al., 2022). The combination of tree height, canopy density, and stratification can aid in distinguishing cocoa systems—especially shaded ones—from natural forest cover (Asitoakor et al., 2022; Kaba et al., 2020). In the case of agroforestry, the height difference between vegetation types can be significant: cocoa grows to a maximum of 8 m tall (Kalischek et al., 2022), while tropical forest trees in Ghana have a mean height of around 20 m (Laurin et al., 2019).

2. Study context

2.1. Study area

The study area is located in Ghana, West Africa. Ghana's climate is tropical, with a distinct wet and dry season. Rainfall varies across regions, with the southern parts receiving up to 2000 mm annually, while the northern areas experience drier conditions. The country's topography is diverse, ranging from coastal plains along the Gulf of Guinea in the south to the mountainous areas in the Ashanti Region and the Northern savannah plains. Ghana's geology is dominated by Precambrian rock formations, with fertile soils in the forested regions ideal for agriculture, particularly cocoa cultivation. The population is over 30 million, with a significant proportion living in rural areas where agriculture, including cocoa farming, is the primary livelihood.

The study area is delineated by a 110 × 110km S2 tile located in central south Ghana shown in Fig. 1. The region has a semi-equatorial climate consisting of dry and rainy seasons, including amajor rainy season (from April–July), and a minor rainy season (from September–October) (Asitoakor et al., 2022). Considering that cocoa requires a high amount of precipitation in order to grow, these rainy seasons are respectively referred to "main crop" and "light crop" cocoa seasons (Asitoakor et al., 2022). During the main crop season, the high levels of rainfall, high relative humidity and low temperatures lead to a higher production of cocoa. During the light crop season, the semi-drought conditions lead to a lower productivity and a change in canopy structures as deciduous trees lose their foliage (Asitoakor et al., 2022).

2.2. Data sources

2.2.1. Ground truth data

Meridia B.V. (Meridia) provided access to cocoa and non-cocoa crop polygons to be used as labels. The data was collected by field agents in Ghana, who mapped the farms between 2017 and 2022 alongside farmers, farm owners, or their representatives with an average accuracy of 2m. The 88,975 cocoa polygons range in size from <1 ha to >10 ha, although most polygons have an area under 3 ha. Each polygon dataset contains attributes collected during field visits, such as cultivation type. Only cocoa polygons labelled as monocrop or intercrop, located within the study area, were used in the research. The non-cocoa crop polygons in the study included tree crops (e.g. cashew, oil palm, mango) and smaller crops (e.g. maize, cassava, beans). Meridia B.V. also provided access to a 2019 forest reserve dataset issued by the Ghana Forestry Commission. The layers were prepared by removing any ambiguous or overlapping polygons and applying a 500m internal buffer to forest reserve polygons to remove any land cover ambiguity alongside the edges. The vector datasets were converted to raster layers indicating cocoa, non-cocoa and unknown pixels.

2.2.2. Sentinel-1/-2 data

Sentinel-1 and -2 datasets were retrieved from the WEkEO JupyterHub (Earth Observation Server) using the WEkEO Harmonized Data Access API (Wekeo harmonized data access). Four datasets of each were downloaded, each within the time period of January 2020–January 2021 to overlap with the ground truth dataset collection dates. The Sentinel-2 MSI datasets (Level 2A) were filtered to contain less than 15 % cloud cover. In line with (Kalischek et al., 2022), the S2 bands of 10m and 20m resolution MSI were used for this analysis, which include bands 2–8, band 8a, 11, and 12; covering parts of the visible, near infrared, and shortwave infrared wavelengths.

The Sentinel-1 SAR datasets (Level 1C) were filtered to Interferometric Wide (IW) mode and Ground Range Detected (GRD) products in ascending orbit. These datasets were pre-processed, multi-looked, and converted to WGS84 by WekEO with a pixel resolution of 10m and a spatial resolution of 20m × 22m (Reiche et al., 2021). In line with (Reiche et al., 2018), the SAR values used in this research are gamma0 backscatter in decibels, to calibrate the values to the logarithmic scale which normalizes the values and increases contrast between features on the ground (Numbisi and Van Coillie, 2020).

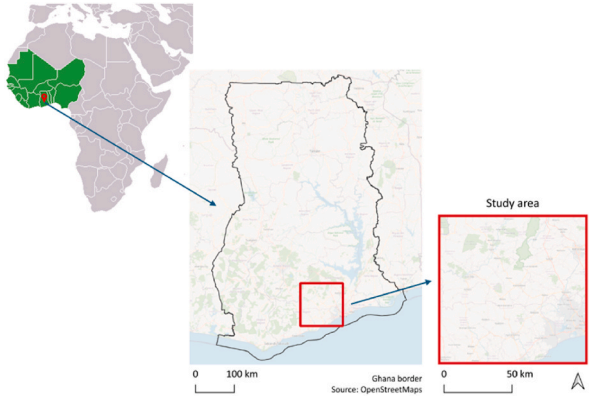

Fig. 1. Study area in Ghana, West Africa.[11].

All downloaded images were resampled to 10m resolution using bilinear interpolation, projected to UTM zone 30N, clipped to study area, and stacked as virtual rasters. Each virtual raster was a 3-dimensional matrix where each pixel had a vector containing reflectance and backscatter values from each of the dataset's bands, in line with (Kalischek et al., 2022) and (Ashiagbor et al., 2020). The virtual rasters and corresponding class labels were subdivided into patches of 128x128 pixels for input to the network. Data was

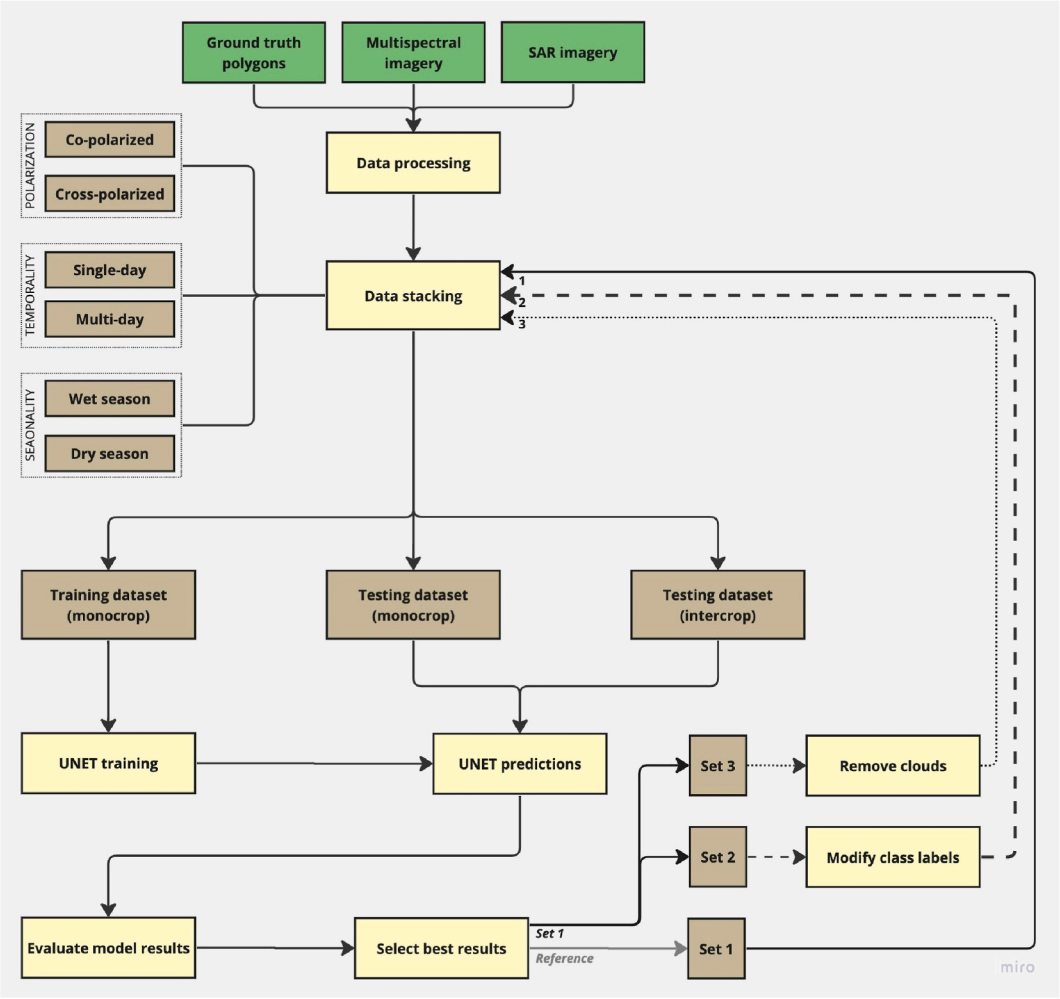

Fig. 2. Methodological framework

linearly scaled using the following equation (KaduncNika Oman, 2022):

$$val_{out} = (val_{in} - c)\left(\frac{b-a}{d-c}\right) + a \tag{2.1}$$

where val_{out} is the normalized pixel value, val_{in} the original pixel index, a and b are the lower and upper values of the desired range (0 and 1), and c and d are the minimum and maximum values of the original range.

3. Method

3.1. Methodological framework

For this study a methodological framework was created, as depicted in Fig. 2. First, ground truth polygons, MSI data, and SAR data are retrieved and processed. Second, satellite imagery is stacked in various combinations of polarization, temporality, and seasonality. Training and testing datasets are formed to observe the impact of different stacking scenarios on the training of a U-NET. The testing data is composed of a set of images containing monocrop cocoa and another containing intercrop cocoa, and each is used separately to evaluate the training results.

3.1.1. Temporal data stacks

Considering the variability in canopy cover within cocoa fields across the wet and dry seasons, efforts were made to select MSI datasets that are evenly distributed over the course of one year and most likely to capture seasonal changes. Further effort was taken to select imagery with minimal cloud cover. The MSI dates selected for the temporal stack were: January 2020 (middle of dry season), March 2020 (end of dry season), December 2020 (beginning of dry season), and January 2021 (middle of dry season). A SAR time-series was created from images selected as close as possible to the dates of MSI data time-series. An additional SAR time-series was formed from four images distributed evenly over the course of the wet and dry seasons (May 2020, August 2020, November 2020, and January 2021). The distribution of both time-series stacks is shown in Fig. 3.

3.1.2. Polarization data stack

Combining both polarizations of SAR is expected to yield better results than a single polarization as it will capture the effects of volume scattering, surface scattering and double bounce. Fig. 4 illustrates the expected SAR backscatter variation on monocrop, intercrop, and forest area.

Fig. 5 presents an example of SAR VV and VH polarization.

The combined impact of SAR polarization and temporality was measured by stacking (1) single-day MSI with single-day SAR, (2) multi-day MSI with single-day SAR, (3) single-day MSI with multi-day SAR, and (4) multi-day MSI with multi-day SAR (Fig. 6).

3.2. Deep learning

3.2.1. Training and testing datasets

Three training datasets were created to carry out each of the three experimental sets (see Table 1). Each dataset had a unique definition of the "non-cocoa" label and was composed of 128x128 pixel patches containing a minimum of 10 % monocrop cocoa and/or non-cocoa labels. Any other pixels were masked by assigning the value "unknown." 10 % of each training dataset was set aside as the test dataset. To obtain a more robust estimate of each model's performance and determine whether it was consistent across different splits of the data, a 10-fold cross-validation approach was used during training. To evaluate experimental results for intercrop cocoa detection, an additional test dataset was created containing 56 images with intercrop-only labels and was used for all experiment sets. The procedure of training/testing data setup is illustrated in Fig. 7.

3.2.2. U-NET implementation and experimental sets

U-NET is a convolutional neural network well suited for land cover classification (Clark et al., 2023) that has previously been used for cocoa detection (Filella, 2018). The U-NET architecture is composed of five encoder blocks and five decoder blocks using Rectified Linear Unit (ReLU) as its activation function, with the "same" padding and max pooling layers between each block.

Our U-NET architecture was adapted from previous cocoa segmentation work (Filella, 2018) and an open source Jupyter notebook U-NET implementation (Bhatia, 2021). It was trained on ground truth data via AWS Studio Lab to identify cocoa crop areas at the field level. The random weights were initialized using the He normal distribution and Adam optimizer, alongside a 0.3 dropout probability to prevent overfitting. In line with (Kalischek et al., 2022), the mini-batch size was set to 32, the learning rate to 10^{-5}, and an L2 regularization rate of 10^{-2} was used. For each experiment, the model with the lowest validation loss value was saved using checkpoints and consecutively used for testing. To prevent overfitting, early stopping was used to automatically stop training when there had not been validation loss improvement for 10 epochs.

During the "Data impact" experimental set, the U-NET was trained with single-and multi-day stacks of MSI data. It was also trained

[1] Image sources: Open Street Maps & https://en.wikipedia.org/wiki/West_Africa#/media/File:Africa-countries-WAFU-UFOA.png

Fig. 3. Seasons of southern Ghana and distribution of datasets over one year.

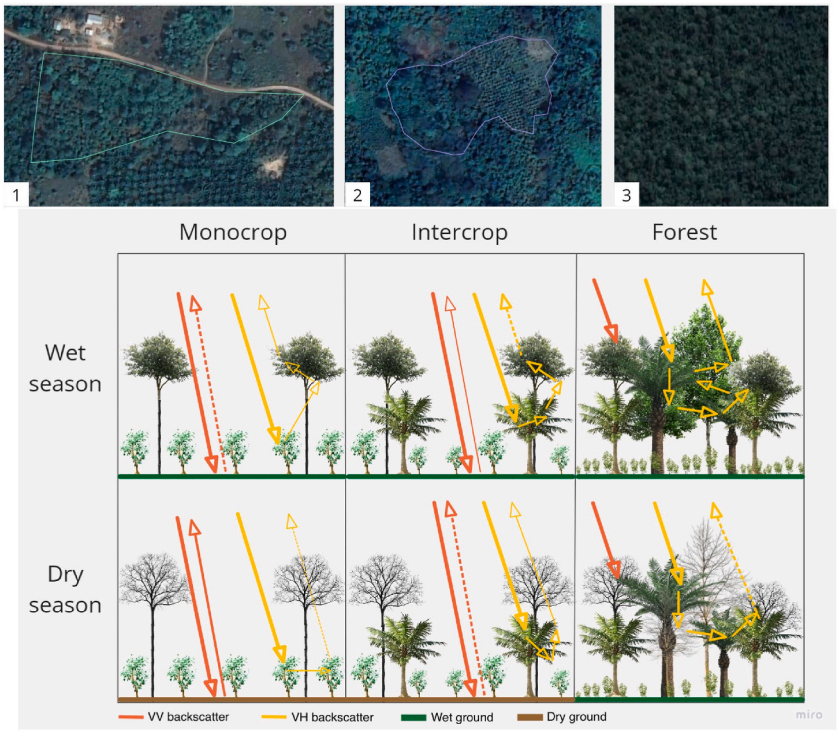

Fig. 4. Satellite imagery and expected SAR backscatter variability in cocoa agriculture showing (1) Monocrop cocoa, (2) Intercrop cocoa and (3) Forest reserve Satellite data: Google Earth accessed via QGIS XYZ tile layer connection. Polygon data: Meridia B.V.

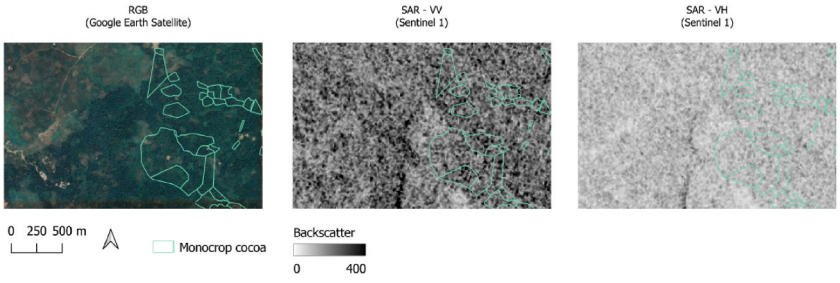

Fig. 5. SAR imagery in different polarizations over study area

separately with single- and multi-day SAR datasets. The "Fusion" experimental set involved stacking SAR data (single- and multi-day rasters) onto MSI datasets to observe the impacts on evaluation metrics and visual output. The experimental set "Label sensitivity" involved repeating some experiments from "Fusion" with a modified labelling approach in which the non-cocoa class includes both forest and non-cocoa crops. The experimental set "Cloud cover" involved repeating some experiments from "Fusion" with the removal of clouded regions. A summary of the experimental sets is shown in Fig. 8 and a list of all experiments is depicted in the supplementary material.

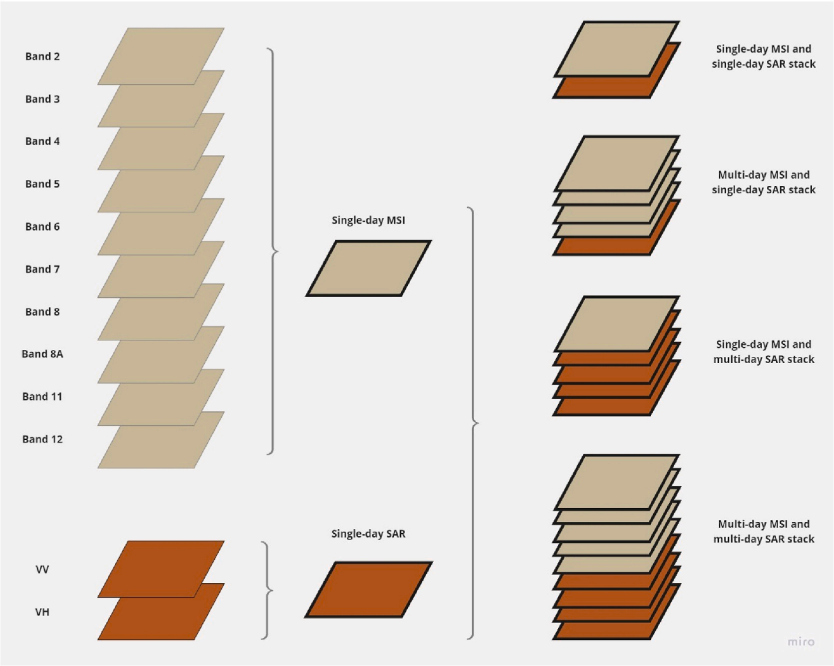

Fig. 6. Possible stacking combinations of MSI and SAR datasets.

Table 1
Training datasets per set. *The "Cloud cover" set dataset involves cloud removal.

Set	Experiment name	Non-cocoa label	Total number of patches	% cocoa	% non-cocoa
Reference	Data Impact	Forest	496	13.38	86.62
1	Fusion	Forest	496	13.38	86.62
2	Label Sensitivity	Non-cocoa crop + forest	556	18.44	82.46
3*	Cloud Cover	Non-cocoa crop + forest	537	20.30	79.70

3.2.3. Weighted loss function for class imbalance

In each of the labelling scenarios, shown in Table 2, the cocoa class is underrepresented compared to the non-cocoa class. To avoid relying on undersampling, the class imbalance within the dataset was rectified via a weighted loss function. The weights of the cocoa and non-cocoa classes were determined based on the "inverse proportion of class frequencies" (Balanced weights for imbalanced classification, 2022): the proportion of labels from each class was calculated to increase the influence of the minority class (i.e. cocoa) on the loss computation, as depicted in Equation (3.1). Furthermore, there was uncertainty in the reference polygon dataset, as the dataset contained sparse ground truth labels (Schutera et al., 2022) and was known to be missing labels for the cocoa and non-cocoa classes because it was not exhaustive for the study area. In this research, any pixel without a ground truth label was assigned a weight of zero as to prevent it from influencing the training. The weights for each class were integrated into the categorical cross entropy function, as detailed in Equation (3.2).

$$w_c = \frac{1}{\frac{n_c}{N}} \tag{3.1}$$

where c is a given class, nc is the number of labels in this class and N is the total number of labels in the dataset.

$$L_{cross-entropy}(\hat{y}, y) = -\frac{1}{N}\sum_{j}^{N}\sum_{c}^{M} w_c y_{c,j} \ln\left(\hat{y}_{c,j}\right) \tag{3.2}$$

where N is the number of labelled samples, and M is the number of classes, $y_{c,j}$ is a binary indicator (0 or 1) of a pixel belonging to a certain class c and $\hat{y}_{c,j}$ is the predicted probability of pixel belonging to a certain class c.

3.2.4. Network output and evaluation metrics

The output of the U-NET is a n-dimensional matrix in which n is the number of classes, containing the probability that a pixel belongs to the given class. The Softmax activation function was applied to normalize values to be between 0 and 1 (Filella, 2018;

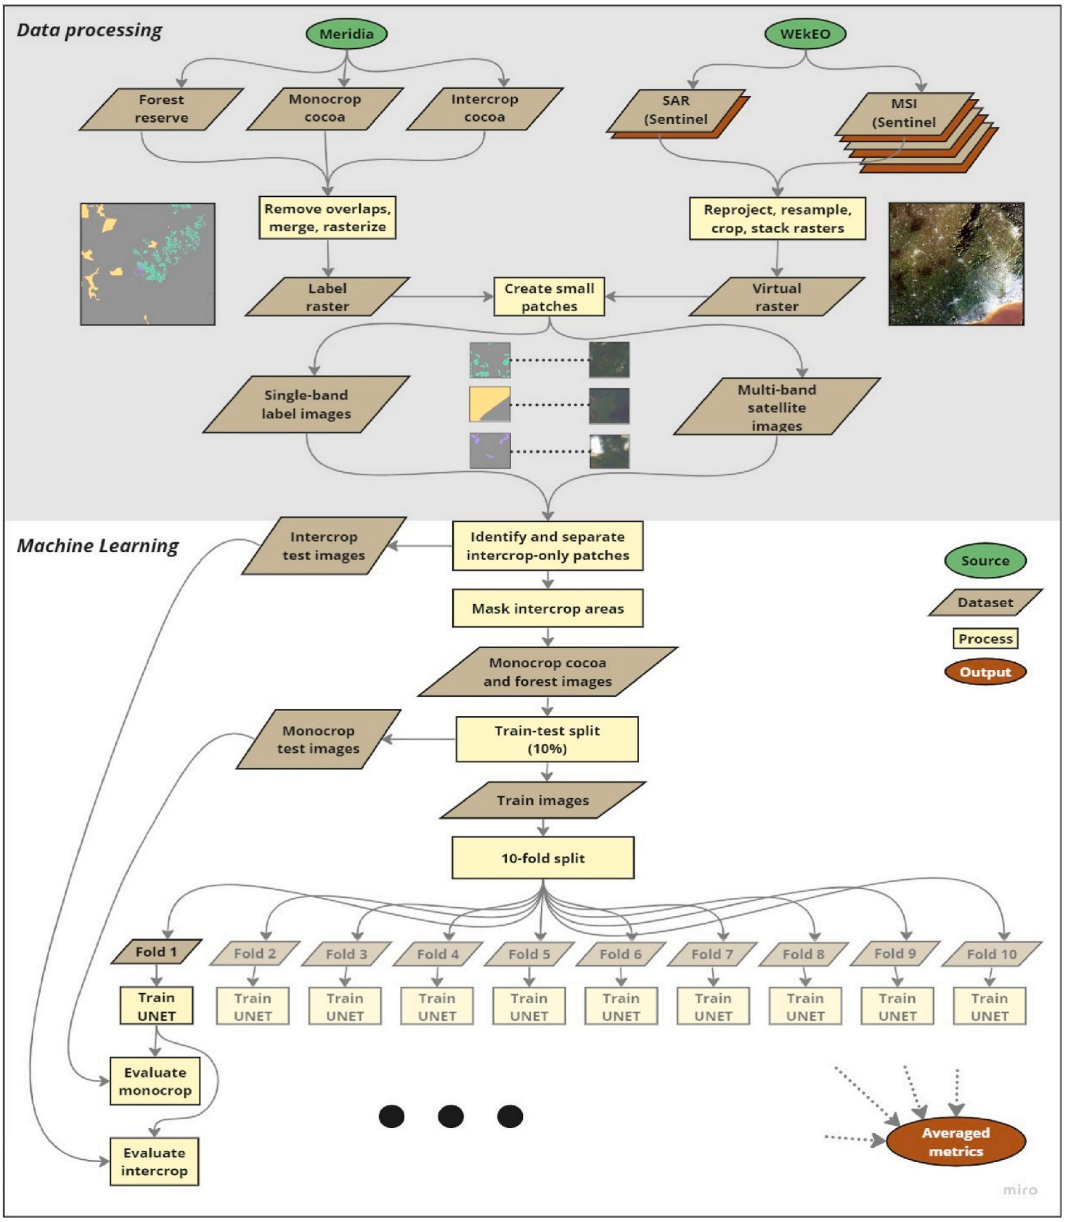

Fig. 7. Workflow for setting up training and testing data

Alonso et al., 2019). This output enabled the visualization of per-pixel cocoa probability maps. The U-NET's Argmax function also generated a prediction mask which for each pixel identifies the class with the highest probability. In order to evaluate the segmentation outputs for each of the experiments, the prediction masks were compared to the ground truth labels excluding the "unknown" regions. Finally, the following metrics were computed and averaged over all experiment folds.

1. Loss: measures the difference between the ground truth of each pixel and the probabilities predicted for each class. The loss values are averaged for all labelled pixels in all test images.
2. IoU: measures the degree of similarity between the predicted mask and ground truth and is computed only for the Cocoa class

$$IOU = \frac{AreaOfOverlap}{AreaOfUnion} = \frac{TP}{TP + FP + FN} \tag{3.3}$$

3. Recall: calculates the proportion of correctly predicted ground truth cocoa, also referred to as Producer's Accuracy

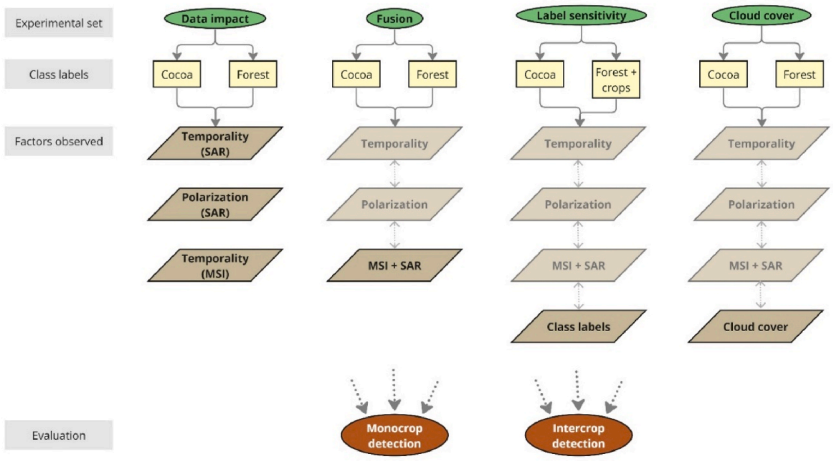

Fig. 8. Overview of experiments

$$PA = \frac{AreaOfOverlap}{AreaOfUnion} = \frac{TP}{TP + FN} \tag{3.4}$$

4. Precision: calculates the proportion of true cocoa predictions, also referred to as User's accuracy

$$UA = \frac{TP}{TP + FP} \tag{3.5}$$

5. F1: the harmonic mean of Precision and Recall

$$F1 = \frac{2 \times UA \times PA}{UA + PA} \tag{3.6}$$

These metrics are commonly used to evaluate the results of semantic segmentation as they provide a comprehensive understanding of the model's performance in terms of both accuracy and reliability. These metrics assess how well the model identifies and classifies the image in relation to the ground truth, highlighting the trade-offs between false positives and false negatives. By using a combination of these metrics, researchers can gain insights into the model's ability to generalize, detect objects accurately, and minimize errors, which is crucial for applications such as cocoa crop detection.

Table 2
"Data impact" training results: average metrics across all folds. DS = Dry Season, MS = Multi-season.

#	label	time (s)	loss	accuracy	IoU*	precision*	recall*	F1*
1	Dec MSI	624	1.03	0.92	0.71	0.89	0.80	0.84
2	March MSI	797	0.76	0.94	0.77	0.95	0.81	0.88
3	Temp MSI	735	0.52	0.95	**0.78**	0.98	0.83	0.90
4	May VV	1216	1.66	0.51	0.11	0.49	0.21	0.29
5	May VH	1148	1.59	0.60	0.22	0.59	0.31	0.41
6	May VV VH	957	1.61	0.56	0.15	0.53	0.24	0.33
7	Jan VV	1038	1.53	0.67	0.31	0.66	0.38	0.48
8	Jan VH	938	1.61	0.72	**0.68**	0.68	0.44	0.53
9	Jan VV VH	893	1.40	0.72	0.36	0.71	0.44	0.54
10	Temp VV (DS)	689	1.55	0.53	0.26	0.74	0.32	0.45
11	Temp VH (DS)	774	1.34	0.71	0.39	0.76	0.46	0.57
12	Temp VV VH (DS)	757	1.30	0.76	0.41	0.73	0.49	0.59
13	Temp VV (MS)	779	1.46	0.69	0.33	0.67	0.41	0.51
14	Temp VH (MS)	802	1.25	0.78	**0.44**	0.77	0.52	0.62
15	Temp VV VH (MS)	666	1.22	0.78	0.44	0.78	0.51	0.61

4. Results and discussion

4.1. Experimental results: data impact

The resulting metrics from the Data impact experimental set are shown in Table 2. When comparing the training results of the single-day MSI from March (end of the Dry season) and December (beginning of the dry season), the March imagery leads to improved results. Results are further improved when training the U-NET with a temporal stack of MSI selected from four days across the Dry season (IoU = 0.78) compared to both single-day datasets (IoU = 0.71 and 0.77).

Overall, the results from the January dataset in the Dry season are improved (IoUs = 0.31, 0.68 and 0.36) compared to those trained with May data in the Wet season (IoUs = 0.11, 0.22, and 0.15). When comparing the results from different polarizations of data captured on the same day, VH offers the best results (IoU May = 0.22, IoU Jan = 0.68), followed by combined VV and VH (IoU May = 0.15, IoU Jan = 0.36), and with VV leading to the poorest metrics (IoU May = 0.11, IoU Jan = 0.31). Overall, the training results for the multi-season stacks (e.g. VH IoU = 0.44) are better than Dry season only (VH IoU = 0.39). When comparing polarization metrics, the use of VV leads to considerably lower performance in both temporal stacks compared to VH or VV + VH, which have similar results.

4.2. Experimental results: fusion

The metrics of the "Fusion" experimental set are summarized in Table 3. The box plot (Fig. 9) shows a slightly improved loss values and reduced spread when single-day December MSI results (loss = 1.03) is combined with the addition of single day SAR (loss = 0.78). This suggests that the model is less sensitive to the split of the training and validation data and is therefore more robust. The spread, as shown in the box plot, is further reduced with the use of multi-season SAR (loss = 0.69), and the median loss value is improved in a statistically significant way. The addition of *multi-season* VV data to *single-day* MSI was found to cause the most significant improvement, decreasing the average loss value from 1.03 to 0.69, as well as reducing the spread of the 10-fold loss values and the vulnerability of the model to imbalanced data splits. This improvement is in line with the findings presented in (Numbisi and Van Coillie, 2020) which identify multi-temporal VV SAR as having the greatest importance in cocoa detection via Random Forest regression. In the case of multi-day MSI, the results are less conclusive. While they are indeed better than any single-day experiment, the addition of SAR from a single day or from across the seasons does not significantly improve the loss results.

4.3. Experimental results: label sensitivity

In this experiment the "non-cocoa" class that previously included only forest reserves, is supplemented with non-cocoa crops, bringing more variation in this class. The "Label sensitivity" set metrics are summarized in Table 4.

A comparison of the Fusion and Label sensitivity results is depicted in Fig. 10. In the case of single-day MSI data (original loss = 1.03), the loss value is improved with the modified labels (loss = 0.73), and further improved with the addition of single-day VH SAR (0.69). As shown in Fig. 10, unlike the Fusion experiments, the Label sensitivity experiments results indicate an outlier in the single day MSI + single day SAR plot. This can be explained by the fact that the non-cocoa crops made up a relatively small portion of the total ground truth data. As a consequence, the predictions may be negatively affected if the random split excludes any of these areas from training.

4.4. Experiment results: the "cloud cover" set

The "Cloud cover" set metrics are summarized in Table 5. The removal of cloudy training data leads to overall lower loss values and higher IoUs. These results suggests that the relevance of SAR in improving segmentation results is more important in the case of

Table 3
The "Fusion" set experiment results: average metrics across all folds.

#	label	time (s)	loss	accuracy	IoU*	precision*	recall*	F1*
16	Dec MSI + Jan VV	841	**0.79**	0.93	0.73	0.92	0.80	0.86
17	Dec MSI + Jan VH	820	0.97	0.79	0.51	0.93	0.53	0.67
18	Dec MSI + Jan VV VH	749	0.88	0.93	0.73	0.93	0.79	0.85
19	Dec MSI + May VV	738	0.86	0.92	0.71	0.93	0.77	0.84
20	Dec MSI + May VH	793	**0.78**	0.92	0.72	0.93	0.79	0.85
21	Dec MSI + May VV VH	761	0.83	0.92	0.72	0.94	0.78	0.85
22	Dec MSI + Temp VV	800	**0.69**	0.93	0.74	0.93	0.80	0.86
23	Dec MSI + Temp VH	724	0.80	0.94	0.75	0.92	0.82	0.87
24	Dec MSI + Temp VV VH	756	0.77	0.93	0.73	0.92	0.80	0.85
25	Temp MSI + May VV	723	0.54	0.95	0.77	0.98	0.82	0.89
26	Temp MSI + May VH	771	**0.50**	0.95	0.79	0.98	0.83	0.90
27	Temp MSI + May VV VH	743	0.51	0.95	0.78	0.98	0.83	0.90
28	Temp MSI + Temp VV	731	0.51	0.95	0.78	0.98	0.83	0.90

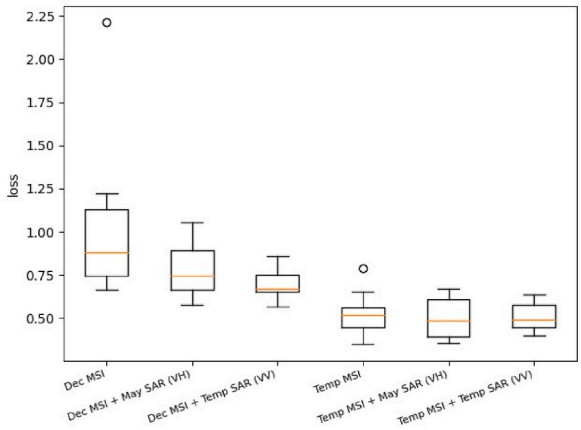

Fig. 9. Box plot comparison of the best loss metric values from the Data impact and Fusion experiments.

Table 4
"Label sensitivity" experiment results.

#	label	time (s)	loss	accuracy	IoU*	precision*	recall*	F1*
29	Dec MSI	671	0.73	0.93	0.65	0.93	0.68	0.79
30	Dec MSI + May VH	712	0.69	0.94	0.68	0.94	0.70	0.79
31	Dec MSI + Temp VV	502	0.78	0.93	0.64	0.88	0.72	0.74

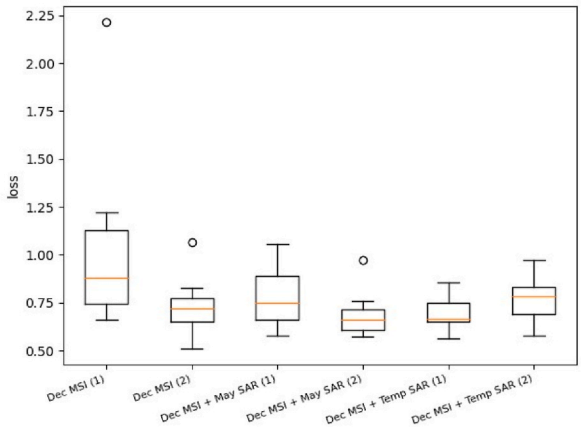

Fig. 10. Comparison of loss box plot for "Fusion" (1) and "Label sensitivity" (2) experiments

Table 5
Set 3 experiment results: original labels with clouds removed.

#	label	time (s)	loss	accuracy	IoU*	precision*	recall*	F1*
32	Dec MSI	875	0.52	0.95	0.78	0.97	0.79	0.86
33	Dec MSI + May VH	705	0.57	0.94	0.74	0.96	0.76	0.84
34	Dec MSI + Temp VV	688	0.56	0.95	0.77	0.95	0.80	0.85

partially cloudy MSI.

In the Cloud cover experiments, the removal of cloudy training data leads to overall lower loss values and higher IoUs. However, the metrics of the single-day MSI are not improved with the addition of any SAR (single-day or temporal) data. These results suggests that the relevance of SAR in improving segmentation results is more important in the case of partially cloudy MSI.

4.5. Overview of all results

Fig. 11 depicts the metrics from each of the experimental sets. It shows that the best overall results were generated when combining temporal MSI with SAR in experimental set "Fusion". While the worst results involved the SAR-only "Data impact" experiments.

4.6. Additional observation

Following are some observations regarding the key factors tested in the experiments, including temporality, polarization, class labels, and the inter-relations between them. We also discuss and evaluate intercrop cocoa detection. Lastly, an overview of the research limitations is provided.

4.6.1. Temporality

The training results for the multi-season SAR-only stacks are better than Dry season only stacks (Table 2). This is likely due to the additional insight Wet season imagery provides into seasonal farm changes, such as the loss of foliage of cocoa shade trees compared to the more seasonally stable forest canopies. Furthermore, the seasonality of datasets affects the effectiveness of different SAR polarizations in cocoa crop detection. Stacking VV data with MSI, imagery from the dry season (the same season as the MSI imagery) leads to better metrics. This may be because during the dry season, cocoa is likely to have more canopy gaps and therefore has a higher contrast in ground backscatter compared to forest backscatter. In contrast, stacking VH data leads to better results when working with Wet season SAR data. This may be because combining a Dry season MSI image with a Wet season SAR image provides temporal information about the canopy characteristics that is not available when both datasets originate from the same season.

Overall, depending on the seasonality and polarization of the dataset, the results show that the addition of SAR to *single-day* MSI can improve the accuracy of the model. In the "Fusion" experiments (Table 3), the addition of *single-day* VH Wet season imagery to a *single-day* image from December (Dry season) slightly improved the loss metrics. It was found that less cocoa was predicted in forested areas and in visibly identifiable oil palm fields, suggesting an improved distinction between vegetation canopies. This insight suggests that more detailed ground truth data and fine-tuning of the model could potentially lay the foundation for cocoa predictions at a higher temporal frequency than is possible with MSI-only dataset analyses.

4.6.2. Polarization

In the "Data impact" experimental set, when comparing the results from different polarizations of single-day data, VH offers the best results, followed by combined VV and VH, and with VV leading to the poorest metrics (Table 2). These results may be due to high vegetation density in some areas of cocoa fields, which prevents co-polarized ground backscatter. When comparing the results from different polarizations of multi-day data, the use of VV leads to significantly worse results in both temporal stacks compared to VH or VV + VH, which have similar results. It may be that seasonal changes in ground backscatter alone are not sufficient to distinguish between different vegetation types. In The "Fusion" set experiments, the polarization with the greatest impact depends on the temporality and seasonality. When stacking single-day MSI with single-day SAR, VV is found to be most effective in the dry season, and VH in the wet season. When stacking single-day MSI with multi-day SAR, VV is found to be most effective.

Combining both co- and cross-polarized SAR datasets was expected to yield better results than a single polarization by capturing the different surface characteristics. In the case of the "Data impact" training set (both single-day and multi-day), the combined polarizations did indeed generate the highest cocoa F1 score and the lowest loss value, indicating the best results. This is in line with the results presented in (Numbisi et al., 2018, 2019), which indicates that the use of combined co- and cross-polarized SAR data led to the best classification of land cover types. However, in the "Fusion" set, there is no clear tendency for the combined polarizations to improve predictions when combined with MSI. In the case of single-day MSI + single-day SAR, single-day MSI + temporal SAR, and temporal MSI + temporal SAR; the loss values of combined polarizations were consistently between that of the best result and the worst result. The reason for this difference in outcome may be attributable to the higher volume of data, which may require adjustments to the deep learning structure.

4.6.3. Class labels

The "Label sensitivity" set approach (Table 4), in which the non-cocoa class contained forest and non-cocoa crop land cover, led to improved results compared to the "Fusion" set labels. This is to be expected, as this additional training data should enable the network to more effectively differentiate between cocoa and other crops. In addition, the results indicate that the labelling approach influences the effectiveness of different SAR polarizations and the influence of dataset temporality in cocoa crop detection. One key example is the differentiation between oil palm and cocoa which have many similarities. First, both plant species are evergreen. However, unlike palm, cocoa has deciduous shade trees above the canopy. Second, both are likely to have some ground scattering. Yet, during the dry season, cocoa will experience a decrease in shade foliage therefore an increased intensity of VV. Whereas palm will exhibit a small increase in VV due to a dryer soil, in contrast to forest which will have a consistently low intensity of VV backscatter. Finally, both cocoa and palm have a rather high amount of volume scattering. Moreover palm appears to have a lusher canopy with fewer gaps, therefore the VH intensity of palm is closer to that of forest and higher than that of cocoa.

In the "Fusion" set experiments, the seasonal variation in palm ground backscatter may appear more like that of cocoa than that of forest, causing better results with the temporal VV stack. In contrast, the single day VH groups together forest and palm, possibly because they both have higher amounts of volume backscatter due to canopy structure. When training the model for the "Label sensitivity" set experiments, the seasonal variation in palm appears less likely to be confused for cocoa, and instead the model appears

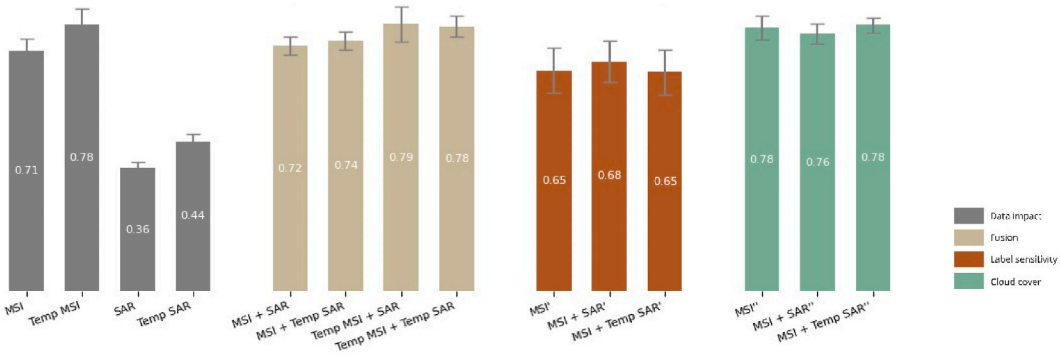

(a) IoU

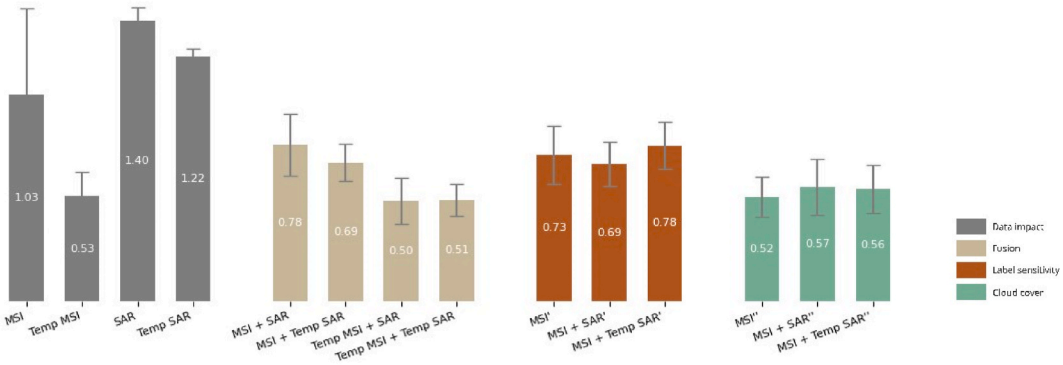

(b) Loss

Fig. 11. Comparison of metrics for all experiments

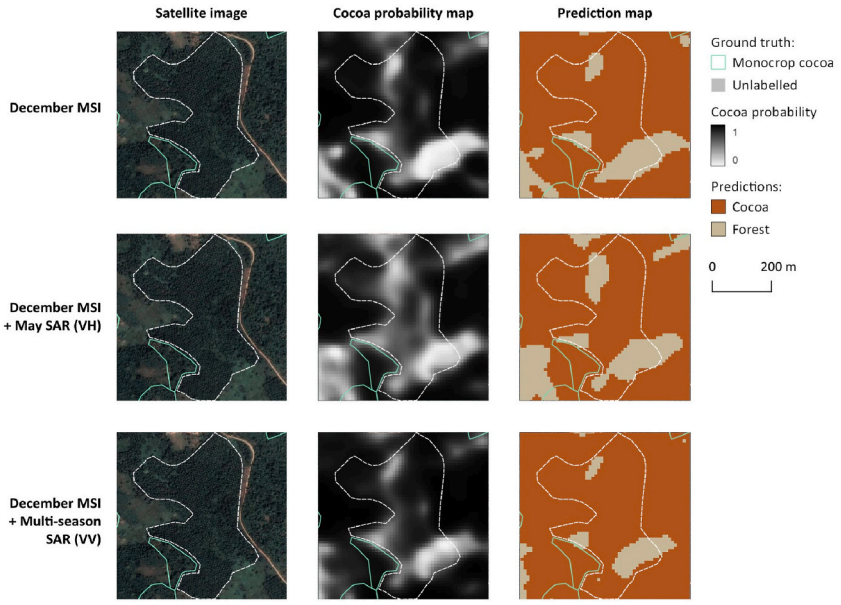

Fig. 12. Single-day MSI (December, dry season) stacked with single and multi-season SAR datasets showing impact on palm crop differentiation (white dotted polygon) when training with non-cocoa forest labels.

to better differentiate between their seasonal patterns.

Fig. 12 compares the "Fusion" set predictions, and indicates that the single-day MSI + single-day VH SAR experiment has better results compared to single-day MSI + multi-day VV SAR as cocoa has the lowest probability in the palm field. One possible reason is that palm and cocoa may behave more similarly across the seasons in contrast with forest, which makes the co-polarized backscatter less effective at differentiating them, but palm and forest have similarly dense canopies in the wet season, and therefore similar levels of cross-polarized backscatter.

In contrast, the results from the "Label sensitivity" set experiments show that the single-day MSI + single-day VH SAR is labelling oil palm crops as cocoa (see Fig. 13), whereas the single-day MSI + multi-day SAR identifies these areas as "not cocoa." In this case, the seasonal variation in palm is less likely to be confused for cocoa. However, the wide range of crop types and canopies in the non-cocoa class may create a less clear distinction between cocoa canopy and other canopies, therefore explaining why the single day VH predicts most vegetated areas as cocoa. Overall, these observations bring attention to the importance of considering land cover classes when selecting SAR polarization and temporality.

4.6.4. *Intercrop detection*

For each of the experimental sets, the intercrop-only test dataset was used to observe the influence of dataset selection on cocoa detection in more complex and mixed canopies. Considering the uncertainty of intercrop labels, the use of quantitative metrics does not provide meaningful insight into the effectiveness of the models. Efforts were made to observe visually whether certain experiments led to more accurate intercrop detection. In the case of the "Fusion" set experiments, a comparison of predictions overlaid on an intercrop polygon does offer some insight, see Fig. 14. The stacked multi-day MSI and single-day SAR, which has the highest loss metric (least effective monocrop cocoa prediction) predicts cocoa in the entire polygon with a high certainty and uniformity. This prediction is unlikely considering that the defining characteristic of "intercrop" is the combination of different crop types. In contrast, the multi-day MSI stack predicts a lower cocoa probability in the top half of the polygon where a denser tree canopy is visible on the satellite image. Finally, the multi-day MSI and multi-day SAR stack has a lower cocoa probability in the top half and lower right corner of the polygon, where a line of large trees is visible. While it is not possible to determine the exact location of cocoa in the imagery at this resolution, it can be observed that the addition of multi-seasonal SAR enables the model to be sensitive to some vegetation differences that are not detected by MSI only.

The findings of this study aligns with the previous research in several aspects. On polarization effects, the findings of our study on the enhancement of canopy structure detection through VH polarization combined with dry-season MSI imagery, while VV polarization combined with temporal SAR data improves ground texture differentiation aligns with the findings by Numbisi et al. (2019), who also noted the importance of polarization in distinguishing vegetation types. On seasonality and temporality effects, the advantage of multi-season SAR combinations over dry-season-only stacks is consistent with the observations of Numbisi et al. (2019), who found that incorporating seasonal variations in SAR data enhances classification accuracy. While this study aligns with previous research in many ways, it also introduces several notable distinctions that advance the current understanding of SAR/MSI integration for cocoa detection. Most previous studies generally acknowledged that SAR polarization plays a role in vegetation classification, but

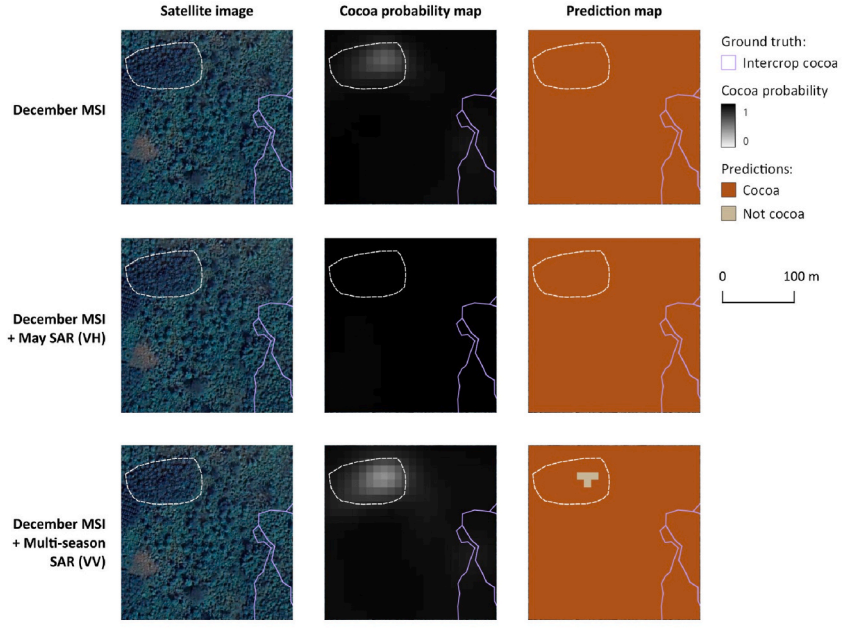

Fig. 13. Single-day MSI (December, dry season) stacked with single and multi-season SAR datasets showing impact on palm crop differentiation (white dotted polygon) when training with non-cocoa crop labels.

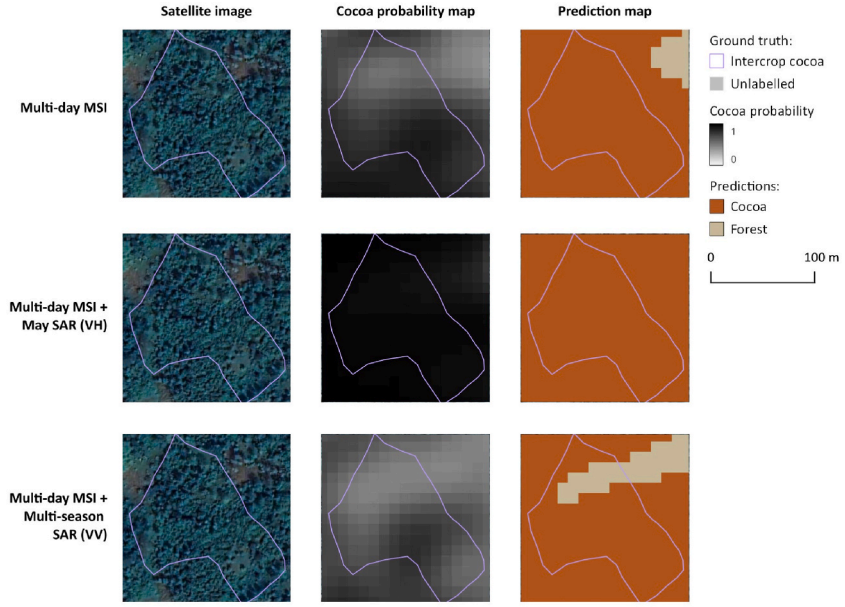

Fig. 14. Multi-day MSI stacked with single and multi-season SAR datasets showing impact on cocoa detection in intercrop labels.

our study explicitly compares VH and VV polarizations in relation to MSI and SAR seasonality. Furthermore, previous studies mostly emphasized the utility of multi-season SAR or longer time-series analysis, but our research uniquely focuses on the strategic pairing of SAR and MSI from different seasons (e.g., dry MSI with wet SAR), rather than just stacking same-season or long-term imagery. This approach showed that even single-day MSI, when paired optimally with seasonal SAR, can reach high classification performance, which is particularly relevant for operational scalability in resource-limited contexts.

The findings of this research have significant socio-economic implications for cocoa-growing regions in Ghana, particularly for smallholder farmers who form the backbone of the industry. By improving the accuracy of cocoa farm detection through the integration of SAR and MSI data, the study enables more efficient and cost-effective mapping, which can support land tenure documentation, improve access to agricultural financing, and enhance delivery of extension services. Additionally, the ability to distinguish cocoa from forest and other crops with seasonal sensitivity aids in monitoring deforestation and enforcing sustainable land-use policies, which is crucial in balancing cocoa-driven economic development with environmental conservation. This technology can also support certification efforts for sustainable cocoa production by providing verifiable spatial data, thereby opening up premium market opportunities for farmers. Furthermore, it empowers policymakers and development agencies to allocate resources more effectively, respond to seasonal changes in cocoa productivity, and monitor the impacts of climate change or disease outbreaks. Overall, the research contributes to a more equitable, sustainable, and data-driven cocoa sector in Ghana, with potential ripple effects across livelihoods, environmental health, and market access.

4.6.5. Limitations

Label classes. For clarity, two classes are used, cocoa and non-cocoa. The non-cocoa class is composed of forest in the "Fusion" set, or forest and non-cocoa crops in the "Label sensitivity" set. However, this approach excludes many other land cover types that are present in the study area, such as water bodies and urban areas. As a result, the network is limited in its ability to differentiate between other classes.

Label uncertainty. In the cocoa polygon data collection, farmers could indicate whether their farms were "Mostly mono-cropped cocoa" or "Mostly intercropped cocoa." This wording suggests that the categories are fuzzy, e.g. that a mostly cocoa-producing farm may contain a small section of other crop(s), therefore it is important to note that there is a degree of labelling uncertainty in the dataset. Furthermore, the forest labels are reserves, which does not guarantee that such areas are undisturbed forest. Cocoa producers are increasingly encroaching on protected areas due to environmental and social pressures (Abu et al., 2021; Kalischek et al., 2022). As can be seen in Fig. 15, numerous cocoa polygons in the Meridia datasets overlap with protected area polygons. This is an important indication of the limitation of using forest reserves as forest ground truth labels. Therefore, the network may be mistakenly trained with illegal cocoa or other crops.

Metrics. The labelling challenges described above limit the relevance of the metrics used to evaluate cocoa crop predictions. For one, the metrics are computed after the application of the Argmax function, meaning that predictions with cocoa probabilities above 50 % are considered "cocoa", which does not reflect the difference in levels of certainty between different experiments. Computing metrics from predictions above a certain threshold (e.g. 80 %) would offer more meaningful comparisons between experiments. Moreover, two challenges arose from the exclusion of unknown areas from the metrics. One challenge was that a model which overpredicted cocoa (e.g. high probability cocoa predicted across large regions) could still lead to good metrics. Another challenge was

Fig. 15. Examples of overlap between cocoa and forest ground truth.

that there was missing information from many image regions where the prediction accuracy is not quantified or estimated. These two challenges were addressed via visual observation of the predicted masks and probabilities. By viewing the predicted maps, it could be identified whether cocoa is predicted across most of the image. Furthermore, by comparing the predictions to satellite imagery, it was possible to qualitatively estimate the likelihood that certain predictions were correct (e.g. by visually identifying crops, urban areas, water, and other features visible to the human eye), and therefore contribute to the validation of the results.

Satellite data. The 10m resolution of Sentinel 1 (S1) and S2 datasets is likely not sufficient to detect many important details that differentiate between intercrop cocoa and forest. For instance, low spatial resolutions can increase the uncertainty along boundaries between forested and non-forested areas (Hansen et al., 2016).

Temporal resolution. The temporal stacks were limited to four images per year; a higher temporal resolution would increase the amount of information available for training. Furthermore, the temporal resolution of the MSI is limited to the dry season in order to avoid the use of imagery with >15 % cloud cover. This limits the effectiveness of training with the temporal MSI stack. Potentially this limitation could be mitigated by spreading out the data, evenly across the year, while masking cloudy areas.

Variety of data. Additional steps in data preparation could maximize the use of datasets for training. Data augmentation, such as the rotating and translating of imagery, can increase the ability of the U-NET to identify land cover classes from different angles. Likewise, expanding the study area would also allow for a more diverse representation of cocoa and forest from different regions.

5. Conclusions

This research aimed to evaluate the influence of SAR and MSI datasets integration on the performance of U-NET to detect cocoa in Ghana. A major focus of the study was on exploring the impact of SAR/MSI temporality, seasonality, and polarization on the semantic segmentation performance. Four experimental sets were carried out to observe the influence of data, polarization and seasonality/temporality. Regarding polarization impact, the result of research demonstrated that the combination of single-day dry season MSI imagery with SAR data from the wet season has the most significant positive impact when considering VH polarization. This combination provides the best insight into canopy structure and volume backscatter during the productive season. In contrast, adding temporal SAR data to MSI yields the best results when using VV polarization, likely because it offers a clearer view of canopy gaps, as well as ground texture and moisture, which show more distinct differences between forest and cocoa areas than biomass does. Regarding seasonality/temporality impact, the results demonstrated that multi-season SAR-only combinations outperform those for dry season-only stacks. This improvement is probably attributed to the extra information provided by wet season imagery, which captures seasonal changes on farms, such as the shedding of cocoa shade tree foliage, in contrast to the more stable forest canopies that experience less seasonal variation. More specifically, the results showed that the addition of single-day and temporal SAR to a single-day MSI image can improve predictions when the relevant seasons and polarizations are considered, thereby reaching a highest F1 score of 86.62 %.

Overall, the findings from this research contribute to a more in-depth understanding of the potential, limitations, and complexity of combining Sentinel 1 and Sentinel 2 datasets for cocoa crop detection. First, the paper provides an exploration into the importance of temporal datasets that capture changes in reflectance, moisture, and biomass of different land cover types over the seasons. For instance, it suggests ways to combine temporal SAR data with single day MSI (and vice-versa) to improve predictions and potentially reduce the amount of data needed to identify cocoa. Second, the research demonstrates the importance of carefully choosing SAR polarizations, when stacking with MSI data, by showing the outcome of different dataset stacks depending on the scenario (e.g. SAR alone, temporal SAR and MSI, single-day SAR and MSI). Third, the research has explored the impact of temporality and polarization with respect to the seasonal changes in cocoa crops. Finally, the research has demonstrated the need for high quality labels and a careful choice of classes to generate results that are relevant for the purpose of the cocoa map. Whereas this research offers insight into the effect of SAR in differentiating between the cocoa, forest, and “other crops” classes; it is expected that the conclusions drawn in this research are relevant to cocoa classification with a greater number of “non-cocoa” ground truth classes.

A key limitations of this study results from the nature of the training data used. The polygon dataset merged shaded and unshaded cocoa farms into a single class, resulting in high intra-class variability. This introduced considerable spectral heterogeneity in the

Sentinel-2 MSI training pixels, as well as high variability in backscatter values in the Sentinel-1 SAR training pixels. These inconsistencies are particularly problematic when using satellite datasets with 10-m spatial resolution, as the spectral and structural diversity within the cocoa class complicates the training and performance of machine learning classifiers. This variability also increases the risk of misclassification, particularly between shaded cocoa systems and open-canopy forests, which often exhibit similar spectral and backscatter characteristics. The lack of dedicated training data for open-canopy forest areas further exacerbates this issue, limiting the classifier's ability to reliably distinguish cocoa from certain forest types. Additionally, the use of relatively large training polygons—such as 1-ha samples (approximately 100 pixels at 10-m resolution)—limits the purity of class signatures. Future efforts should aim to refine the training dataset by separating cocoa classes based on shading characteristics, incorporating more representative forest samples, and using smaller, more homogeneous training polygons better suited to the spatial resolution of the satellite imagery.

As the reporting requirements for cocoa producers continue to grow in response to the threat of deforestation, further work on the use of ML and remotely sensed data will play an important role in producing more reliable cocoa maps. Future research could explore the potential of semantic segmentation of cocoa using more precise remotely sensed data. Longer wavelengths for SAR (i.e. L- and P-bands), or higher resolution MSI and/or hyperspectral imagery offer the potential to detect cocoa in complex environments, such as that of agroforestry.

CRediT authorship contribution statement

Adele Therias: Visualization, Investigation, Conceptualization, Methodology, Data curation, Writing – original draft, Validation, Formal analysis. **Azarakhsh Rafiee:** Writing – review & editing, Conceptualization, Supervision, Methodology. **Stef Lhermitte:** Supervision, Methodology, Writing – review & editing, Conceptualization. **Philip van der Lugt:** Data curation, Writing – review & editing, Conceptualization, Supervision. **Roderik Lindenbergh:** Writing – review & editing.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:Philip van der Lugt was a Meridia representative and supervisor for the duration of the thesis on a freelance basis. Furthermore, Meridia was involved in initiating and facilitating the thesis, but it did not fund the research.

Data availability

The authors do not have permission to share data.

References

Abu, Itohan-Osa, Szantoi, Zoltan, Brink, Andreas, Marine, Robuchon, Thiel, Michael, 2021. Detecting cocoa plantations in Cote d'ivoire and Ghana and their implicationsˆ on protected areas. Ecol. Indic. 129, 107863. ISSN 1470-160X.

Alonso, Inigo, Yuval, Matan, Eyal, Gal, Treibitz, Tali, Murillo, Ana C., 2019. ~ Coralseg: learning coral segmentation from sparse annotations. J. Field Robot. 36 (8), 1456–1477.

Ashiagbor, George, Forkuo, Eric K., Asante, Winston A., Acheampong, Emmanuel, Quaye-Ballard, Jonathan A., Boamah, Prince, Mohammed, Yakubu, Foli, Ernest, 2020. Pixel based and object-oriented approaches in segregating cocoa from forest in the Juabesobia landscape of Ghana. Remote Sens. Appl.: Society and Environment 19, 100349. ISSN 2352-9385.

Ashiagbor, George, Adams Asante, Winston, Forkuo, Eric Kwabena, Acheampong, Emmanuel, Foli, Ernest, 2022. Monitoring cocoa-driven deforestation: the contexts of encroachment and land use policy implications for deforestation free cocoa supply chains in Ghana. Appl. Geogr. 147, 102788. ISSN 0143-6228.

Asitoakor, Bismark Kwesi, Vaast, Philippe, Ræbild, Anders, Ravn, Hans Peter, Eziah, Vincent Yao, Owusu, Kwadwo, Mensah, Eric Opoku, Asare, Richard, 2022. Selected shade tree species improved cocoa yields in low-input agroforestry systems in Ghana. Agric. Syst. 202, 103476. ISSN 0308-521X.

Balanced weights for imbalanced classification, 2022. Medium. URL. https://medium.com/grabngoinfo/balanced-weights-for-imbalanced-classification-465f0e13c5ad#:~:text=Theweightsarecalculatedusing,theminorityclassis97.919.

Batista, Joao E., Rodrigues, Nuno M., Cabral, Ana I.R., Vasconcelos, Maria J.P., Venturieri, Adrianõ, Silva, Luiz G.T., Silva, Sara, 2022. Optical time series for the separation of land cover types with similar spectral signatures: cocoa agroforest and forest. Int. J. Rem. Sens. 43 (9), 3298–3319.

Bastin, J.F., Finegold, Y., Garcia, C., Mollicone, D., Rezende, M., Routh, D., et al., 2019. The global tree restoration potential. Science 365 (6448), 76–79.

Benefoh, Daniel Tutu, Villamor, Grace B., van Noordwijk, Meine, Borgemeister, Christian, Asante, Winston A., Asubonteng, Kwabena O., 2018. Assessing land-use typologies and change intensities in a structurally complex Ghanaian cocoa landscape. Appl. Geogr. 99, 109–119. ISSN 0143-6228.

Bhatia, Vidushi, 2021. U-net implementation from scratch using tensorflow. Medium. URL. https://medium.com/geekculture/u-net-implementation-from-scratch-using-tensorflow-b4342266e406.

Blaser-Hart, W.J., Hart, S.P., Oppong, J., Kyereh, D., Yeboah, E., Six, J., 2021. The effectiveness of cocoa agroforests depends on shade-tree canopy height. Agric. Ecosyst. Environ. 322, 107676. ISSN 0167-8809.

Clark, Andrew, Phinn, Stuart, Scarth, Peter, 2023. Optimised u-net for land use–land cover classification using aerial photography. PFG–Journal of Photogrammetry, Remote Sensing and Geoinformation Science 1–23.

Erasmi, Stefan, Twele, Andre, 2009. Regional land cover mapping in the humid tropics using combined optical and SARsatellite data—a case study from central sulawesi, Indonesia. Int. J. Rem. Sens. 30 (10), 2465–2478.

European Commission, 2021. Proposal for a Regulation of the European Parliament and of the Council on the Making Available on the Union Market as Well as Export from the Union of Certain Commodities and Products Associated with Deforestation and Forest Degradation and Repealing Regulation (Eu) No 995/2010. URL. https://environment.ec.europa.eu/system/files/2021-11/COM_2021_706_1_EN_ACT_part1_v6.pdf.

Evans, D.L., Elachi, C., Stofan, E.R., Holt, B., Way, J.B., Kobrick, M., Vogt, M., Wall, S., van Zyl, J., Schier, M., Ottl, H., Pampaloni, P., 1993. The shuttle imaging radar-c and x- SARmission. EOS, Transactions American Geophysical Union 74 (13), 145–158. https://doi.org/10.1029/93EO00097.

FAO, 2020. Global Forest Resources Assessment 2020. Food and Agriculture Organization of the United Nations.

Filella, Guillem Bonet, 2018. Cocoa Segmentation in Satellite Images with Deep Learning. Master's thesis, ETH Zurich. URL. https://ethz.ch/content/dam/ethz/special-interest/baug/igp/photogrammetry-remote-sensing-dam/documents/pdf/Student_Theses/BA_BonetFilella.pdf.

Global forest resources assessment, 2020. Food and Agriculture Organization of the United Nations.
Hansen, Matthew C., Krylov, Alexander, Tyukavina, Alexandra, Potapov, Peter V., Turubanova, Svetlana, Zutta, Bryan, Ifo, Suspense, Margono, Belinda, Stolle, Fred, Moore, Rebecca, 2016. Humid tropical forest disturbance alerts using landsat data. Environ. Res. Lett. 11 (3), 034008.
Houghton, R.A., 2012. Carbon emissions and deforestation. Annu. Rev. Environ. Resour. 37, 79–106.
Jansing, E. David, 2021. Introduction to Synthetic Aperture Radar: Concepts and Practice. McGraw-Hill Education.
Kaba, James S., Otu-Nyanteh, Alexander, Abunyewa, Akwasi A., 2020. The role of shade trees in influencing farmers' adoption of cocoa agroforestry systems: insight from semideciduous rain forest agroecological zone of Ghana. NJAS - Wageningen J. Life Sci. 92, 100332. ISSN 1573-5214.
Kadunc, Nika Oman, 2022. How to normalize satellite images for deep learning. Medium. URL. https://medium.com/sentinel-hub/how-to-normalize-satellite-images-for-deep-learning-d5b668c885af.
Kalischek, Nikolai, Lang, Nico, Renier, Cecile, Daudt, Rodrigo, Addoah, Thomas, Thompson, William, Blaser-Hart, Wilma, Garrett, Rachael, Wegner, Jan, 2022. Satellite-based highresolution maps of cocoa planted area for Cote d'ivoire and Ghana. arXiv preprint arXiv:2206.06119.
Kumar, Rahul, Kumar, Amit, Saikia, Purabi, 2022. Deforestation and forests degradation impacts on the environment. Chapter Deforestation and Forests Degradation Impacts on the Environment. Springer International Publishing, Cham, ISBN 978-3-030-95542-7, pp. 19–39.
Laurin, Gaia Vaglio, Ding, Jianqi, Disney, Mathias, Bartholomeus, Harm, Herold, Martin, Papale, Dario, Valentini, Riccardo, 2019. Tree height in tropical forest as measured by different ground, proximal, and remote sensing instruments, and impacts on above ground biomass estimates. Int. J. Appl. Earth Obs. Geoinf. 82, 101899. ISSN 1569-8432.899.
Meridia. Field data solutions for smallholder supply chains, n.d. URL https://www.meridia.land/.
Numbisi, Frederick N., Vancoillie, Frieke, De Wulf, Robert, 2018. Multi-date sentinel1 sar image textures discriminate perennial agroforests in a tropical forest-savannah transition landscape. Int. Arch. Photogram. Rem. Sens. Spatial Inf. Sci. XLII-1, 339–346.
Numbisi, Frederick N., Van Coillie, Frieke M.B., De Wulf, Robert, 2019. Delineation of cocoa agroforests using multiseason Sentinel-1 SARimages: a low grey level range reduces uncertainties in GLCM texture-based mapping. ISPRS Int. J. GeoInf. 8 (4). ISSN 2220-9964.
Numbisi, Frederick N., Van Coillie, Frieke, 2020. Does Sentinel-1a backscatter capture the spatial variability in canopy gaps of tropical agroforests? A proof-of-concept in cocoa landscapes in Cameroon. Remote Sens. 12 (24). ISSN 2072-4292.
Reiche, Johannes, Hamunyela, Eliakim, Verbesselt, Jan, Hoekman, Dirk, Herold, Martin, 2018. Improving near-real time deforestation monitoring in tropical dry forests by combining dense sentinel-1 time series with landsat and alos-2 palsar-2. Rem. Sens. Environ. 204, 147–161. ISSN 0034-4257.
Reiche, Johannes, Mullissa, Adugna, Slagter, Bart, Gou, Yaqing, Tsendbazar, Nandin-Erdene, Odongo-Braun, Christelle, Vollrath, Andreas, Weisse, Mikaela J., Stolle, Fred, Pickens, Amy, et al., 2021. Forest disturbance alerts for the congo basin using sentinel-1. Environ. Res. Lett. 16 (2), 024005.
Ruetschi, Marius, Schaepman, Michael E., Small, David, 2018. Using multitemporal sentinel-1 c-band backscatter to monitor phenology and classify deciduous and coniferous forests in northern Switzerland. Remote Sens. 10 (1). ISSN 2072-4292.
Saatchi, S., Agosti, D., Alger, K., Delabie, J., Musinsky, J., 2001. Examining fragmentation and loss of primary forest in the southern Bahian atlantic forest of Brazil with radar imagery. Conserv. Biol. 15 (4), 867–875.
Schutera, Mark, Rettenberger, Luca, Pylatiuk, Christian, Reischl, Markus, 2022. Methods for the frugal labeler: Multi-class semantic segmentation on heterogeneous labels. PLoS One 17 (2), 1–14. https://doi.org/10.1371/journal.pone.0263656.
Sonwa, Denis J., Weise, Stephan F., Schroth, Goetz, Janssens, Marc J.J., Shapiro, Howard-Yana, 2019. Structure of cocoa farming systems in west and central Africa: a review. Agrofor. Syst. 93, 2009–2025.
Tamga, Dan Kanmegne, Latifi, Hooman, Ullmann, Tobias, Baumhauer, Roland, Thiel, Michael, Bayala, Jules, 2022. Modelling the spatial distribution of the classification error of remote sensing data in cocoa agroforestry systems. Agrofor. Syst. 97 (1), 109–119. https://doi.org/10.1007/s10457-022-00791-2.
Thiel, Carolin, Cartus, Oliver, Eckardt, Robert, Richter, Nicole, Thiel, Christian, Schmullius, Christiane, 2009. Analysis of multi-temporal land observation at c-band. 2009 IEEE International Geoscience and Remote Sensing Symposium 3, 318–321.
Vreugdenhil, Mariette, Wagner, Wolfgang, Bauer-Marschallinger, Bernhard, Pfeil, Isabella, Teubner, Irene, Rudiger, Christoph, Strauss, Peter, 2018. Sensitivity of sentinel-1 backscat- ter to vegetation dynamics: an austrian case study. Remote Sens. 10 (9). ISSN 2072-4292.
Wekeo harmonized data access. URL https://help.wekeo.eu/en/collections/353.0725-wekeo-harmonized-data-access.