



Contents lists available at ScienceDirect

# International Journal of Applied Earth Observation and Geoinformation

journal homepage: [www.elsevier.com/locate/jag](http://www.elsevier.com/locate/jag)

## Taking it further: Leveraging pseudo-labels for field delineation across label-scarce smallholder regions

Philippe Rufin<sup>a,b,\*</sup>, Sherrie Wang<sup>c</sup>, Sá Nogueira Lisboa<sup>d,e</sup>, Jan Hemmerling<sup>f</sup>,  
Mirela G. Tulbure<sup>g</sup>, Patrick Meyfroidt<sup>a,h</sup>

<sup>a</sup> Earth and Life Institute, UCLouvain, Place Pasteur 3, 1348 Louvain-la-Neuve, Belgium

<sup>b</sup> Geography Department, Humboldt-Universität zu Berlin, Unter den Linden 6, 10117 Berlin, Germany

<sup>c</sup> Department of Mechanical Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge MA 02139-4307, USA

<sup>d</sup> Nitidae, 500, rue Jean-François Breton, 34 000 Montpellier, France

<sup>e</sup> Department of Forestry Engineering, Universidade Eduardo Mondlane, Av. Julius Nyerere, Maputo, Mozambique

<sup>f</sup> Thünen-Institute of Farm Economics, Bundesallee 63, 38116 Braunschweig, Germany

<sup>g</sup> Center for Geospatial Analytics, North Carolina State University, Jordan Hall, 5112, 2800 Faucette Dr, Raleigh, NC 27695, USA

<sup>h</sup> F.R.S.-FNRS, Rue d'Egmont 5, 1000 Brussels, Belgium

### ARTICLE INFO

#### Keywords:

Mozambique  
Sub-Saharan Africa  
Deep Learning  
Transfer Learning  
Earth Observation  
Cropland

### ABSTRACT

Satellite-based field delineation has entered a quasi-operational stage due to recent advances in machine learning for computer vision. Transfer learning allows for the resource-efficient transfer of pre-trained field delineation models across heterogeneous geographies. However, the scarcity of labeled data for complex and dynamic smallholder landscapes remains a major bottleneck. The key innovation of this study is to overcome this challenge by using pre-trained models to generate sparse (i.e., not fully annotated) field delineation pseudo-labels for fine-tuning models across geographies and sensor characteristics. We build on a FracTAL ResUNet trained for crop field delineation in India (median field size of 0.24 ha) based on multi-spectral imagery at 1.5 m spatial resolution. We use this model to generate pseudo-labels for the use in Northern Mozambique (median field size of 0.06 ha) based on sub-meter resolution true-color satellite imagery. We designed multiple pseudo-label selection strategies based on field-level probability scores and compared the quantities, area properties, seasonal distribution, and spatial agreement of the pseudo-labels against human-annotated training labels ( $n = 1,512$ ). We then used the human-annotated labels and the pseudo-labels for model fine-tuning and compared predictions against human field annotations ( $n = 2,199$ ). We evaluated performance with regards to object-level spatial agreement and site-level field size estimation. Our results indicate i) a good baseline performance of the pre-trained model in both field delineation (mean intersection over union (mIoU) of 0.634) and field size estimation (mean root mean squared error (mRMSE) of 0.071 ha), and ii) the added value of regional fine-tuning with performance improvements in nearly all experiments (mIoU increases of up to 0.060, mRMSE decreases of up to 0.034 ha). Moreover, we found iii) substantial performance increases when using only pseudo-labels (up to 77 % of the mIoU increases and 68 % of the mRMSE decreases obtained by human-annotated labels), and iv) additional performance increases (mIoU+0.008, mRMSE: -0.003 ha) when complementing human annotations with pseudo-labels. Pseudo-labels are architecture-agnostic, can be efficiently generated at scale, and thus facilitate domain adaptation in label-scarce settings. The workflow presented here is a stepping stone for overcoming the persisting challenges in mapping heterogeneous smallholder agriculture.

### 1. Introduction

Smallholder farms operating on less than 2 ha of land produce an estimated 30–34 % of the global food supply on 24 % of the gross

agricultural area (Ricciardi et al., 2018), have higher land productivity (Chiarella et al., 2023), and harbor greater biodiversity as compared to larger farms (Ricciardi et al., 2021). On the other hand, smallholder farming still constitutes a large share of agriculture-driven deforestation

\* Corresponding author.

E-mail address: [philippe.rufin@uclouvain.be](mailto:philippe.rufin@uclouvain.be) (P. Rufin).

<https://doi.org/10.1016/j.jag.2024.104149>

Received 17 November 2023; Received in revised form 25 July 2024; Accepted 4 September 2024

Available online 13 September 2024

1569-8432/© 2024 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

across the tropics (FAO, 2023). The commonly used dichotomy between small and large farms, however, tends to overlook the ongoing structural dynamics related to medium-scale farms (5–100 ha) (Meyfroidt, 2017), which represent an increasingly large portion of agricultural operations in Sub-Saharan Africa (Jayne et al., 2022), causing dynamics of farm consolidation with yet weakly understood impacts and spill-overs affecting productivity, land availability, and labor opportunities at the regional level. As such, a more detailed understanding of land management and change trajectories in smallholder-dominated contexts is required to better balance trade-offs between the social, environmental, and economic impacts of agriculture (Chiarella et al., 2023).

Quantitative and spatially detailed information is limited for tropical smallholder agriculture, and Sub-Saharan Africa in particular (Whitcraft et al., 2019). Numerous empirical blind spots persist, including spatially explicit data on dominant crop types or farming systems (Lambert et al., 2018) or plot-level estimates of productivity (Burke and Lobell, 2017). Moreover, the spatial distribution of agricultural fields and their sizes remain unknown at the sub-national level, hampering a more detailed understanding of the ongoing structural transformations of agriculture in Sub-Saharan Africa (Burke et al., 2020; Jayne et al., 2022; Liverpool-Tasie et al., 2020).

The combination of Earth Observation (EO) data and deep learning offers great potential to close some of the knowledge gaps pertaining to field size distribution by producing accurate, timely, and spatially detailed information in complex smallholder environments (Nakalembe and Kerner, 2023). Challenges related to the spatially fragmented and dynamic nature of smallholder agriculture with region-specific crop portfolios, landscape structures, and sub-plot land management continue to complicate this endeavor. Nevertheless, notable progress has been made in delineating smallholder or generally small fields in EO data in recent years. Amongst the first studies demonstrating the high performance of deep learning approaches for agricultural field delineation, Persello et al. (2019) used a fully convolutional network and WorldView 2/3 data to detect crop field boundaries in Nigeria and Mali, which were processed into field instances using an Oriented Watershed Transform and Single-scale Combinatorial Grouping. Shortly after, Waldner and Diakogiannis (2020) introduced the concept of multi-task semantic segmentation in the domain of field delineation, where models learn to share representation between multiple related tasks. This concept improved generalization and consequently inspired a series of model architectures with high performance when combining pixel-level field boundary predictions with post-processing algorithms such as watershed transforms (Cai et al., 2023; Li et al., 2023; Long et al., 2022; Zhu et al., 2024).

The development of regional to continental scale field delineations in Africa, however, is – amongst limitations in the availability of very-high resolution (VHR) image data – severely challenged by a reference data bottleneck (Nakalembe and Kerner, 2023). Label acquisition must be conducted individually for each target region since cadaster or land parcel information is typically not available in smallholder regions. Further, state-of-the-art multi-tasking architectures require semantically complete field instances as reference data (Cai et al., 2023; Li et al., 2023; Long et al., 2022; Waldner and Diakogiannis, 2020), which can only be collected by well-trained personnel and with repeated expert quality screening, which makes it cost-intensive and often limits the quantities of available labels (Estes et al., 2022).

In recent years, a wealth of training strategies has been proposed to address the paucity of reference data inherent to EO applications (Safonova et al. (2023) provide a comprehensive review of advanced learning strategies in EO applications). Transfer learning, i.e., leveraging pre-trained models, provides a promising avenue for field delineation in data-constrained smallholder settings. Previous studies suggested training models in (source) regions with ubiquitous labels and fine-tuning the models for use in another (target) region using limited reference data (Kerner et al., 2023; Mei et al., 2022; Wang et al., 2022). These studies revealed i) equal or even increased performance when

applying fine-tuned models for use in a new target region compared to training models from scratch, but with ii) significant reductions in training data requirements. For example, Mei et al. (2022) noted a high performance when transferring their model into a region unseen during training. Building on this insight, the authors iteratively corrected and amended predictions in the target region manually and used the corrected predictions for iterative training, which further improved performance. Wang et al. (2022) conducted ablation experiments reducing the number of fields for model fine-tuning and obtained the highest performance with only 500–1,000 labeled fields by combining weak supervision (i.e., partially labeled image chips) with transfer learning. Kerner et al. (2023) used a model fine-tuned for field delineation in South Africa and demonstrated that satisfactory performance (mean intersection over union (mIoU) = 0.64) could be achieved for smallholder field delineation in Kenya without regional fine-tuning.

The potential of transfer learning for task-specific geographic transfer is thus large. The crucial aspect in this regard is domain adaptation between the source and target domain, such as differences in geographic realms (i.e., field sizes, shapes, texture, appearance of field boundaries, or seasonal variations) and image characteristics (i.e., spatial resolution, viewing angle, spectral bands, radiometric properties). A variety of approaches address domain adaptation aiming either at reducing label requirements (e.g. transfer learning, active learning, or semi-supervised learning) or entirely automating domain adaptation without the need for any training labels in the target domain (e.g., self-supervised learning, or self-training; see Toldo et al. (2020) for a review of approaches suitable for dense (i.e., pixel-level) predictions, commonly referred to as semantic segmentation). In the case of semantic segmentation, previous studies approached domain adaptation by optimizing for additional learning objectives during training (“online”), such as image-level test time entropy minimization (Wang et al., 2020), or the use of self-training by complementing source data with pseudo-labels, which represent predictions on unlabeled data in the target domain (Lee, 2013; Prabhu et al., 2022; Zou et al., 2021).

Self-training based on pseudo-labels poses a particularly useful avenue for transfer learning with weak supervision, which may be mobilized without formulating additional learnable objectives in model architectures (i.e., “offline”), and in a source-free setting (i.e., where the training data of the pre-trained model is not readily available). Large quantities of sparse (i.e., partially annotated) labels can be efficiently generated using pre-trained models given unlabeled EO data, which may support domain adaptation. Importantly, however, self-training involves a risk for error propagation with consequent performance loss, e.g., when erroneous pseudo-labels are considered during training. Thus, approaches for selecting accurate and informative labels are essential, and careful performance evaluation and testing for domain-specific biases should be conducted (Mei et al., 2020; Prabhu et al., 2022). The selection of pseudo-labels using constant confidence thresholds typically favors easy cases or classes and generate non-diverse and potentially biased pseudo-labels (He et al., 2021). Consequently, class-balanced approaches were proposed to increase the proportion of pseudo-labels among difficult and rare classes (Zou et al., 2018), or instance-level adaptive thresholds that can be iteratively adjusted to match the prediction confidence present in input data (Mei et al., 2020). In the field of semantic segmentation for EO data, pseudo-labels have been explored in generic land cover mapping tasks (Chen et al., 2022; Liu et al., 2023, 2022; Wu and Prasad, 2018), with most approaches conducting “online” pseudo-label selection reliant on source data. For field delineation in smallholder agriculture, however, conditions differ, as pre-trained models are scarce, and source data is oftentimes not readily available due to the reliance on commercial VHR imagery with prohibitive licensing agreements. Moreover, commonly used pseudo-label selection workflows at the pixel level are unsuitable for field delineation, as many state-of-the-art field delineation architectures require semantically complete field instances (Cai et al., 2023; Li et al., 2023; Long et al., 2022; Waldner and Diakogiannis, 2020).

Exploring efficient tools for domain adaptation in source-free settings by facilitating the generation of high-quality instance-level labels at scale is a crucial research avenue to leverage the potential of artificial intelligence in smallholder systems, in particular in Sub-Saharan Africa (Nakalembe and Kerner, 2023). This study contributes to the field by exploring the value of pseudo-labels for overcoming domain gaps for field delineation in source-free settings. We used a pre-trained FracTAL ResUNet which has been previously deployed for field delineation in smallholder settings in India (Wang et al., 2022), yet the approach remains independent of the model architecture used for field delineation. We conducted closed-set domain adaptation with geographic, temporal, and sensor-specific domain shifts, where the size, shape, and appearance

of fields in the target region (Mozambique) differed substantially from the source region (India). We used VHR RGB imagery (~0.6 m in target, 1.5 m in source) acquired from various airborne and satellite sensors, with different viewing angles, illumination conditions, and acquired during different times of the year. We tested the generation of semantically complete pseudo-labels using multiple instance-level confidence thresholds and compared them with human-annotated labels to assess their spatial agreement, size, and site-level quantities across seasons. We assessed the value of each respective set of pseudo-labels for overcoming domain gaps in two downstream applications, namely field delineation, and site-level field size estimation. For this, we fine-tuned the pre-trained field delineation model with each set of pseudo-labels and

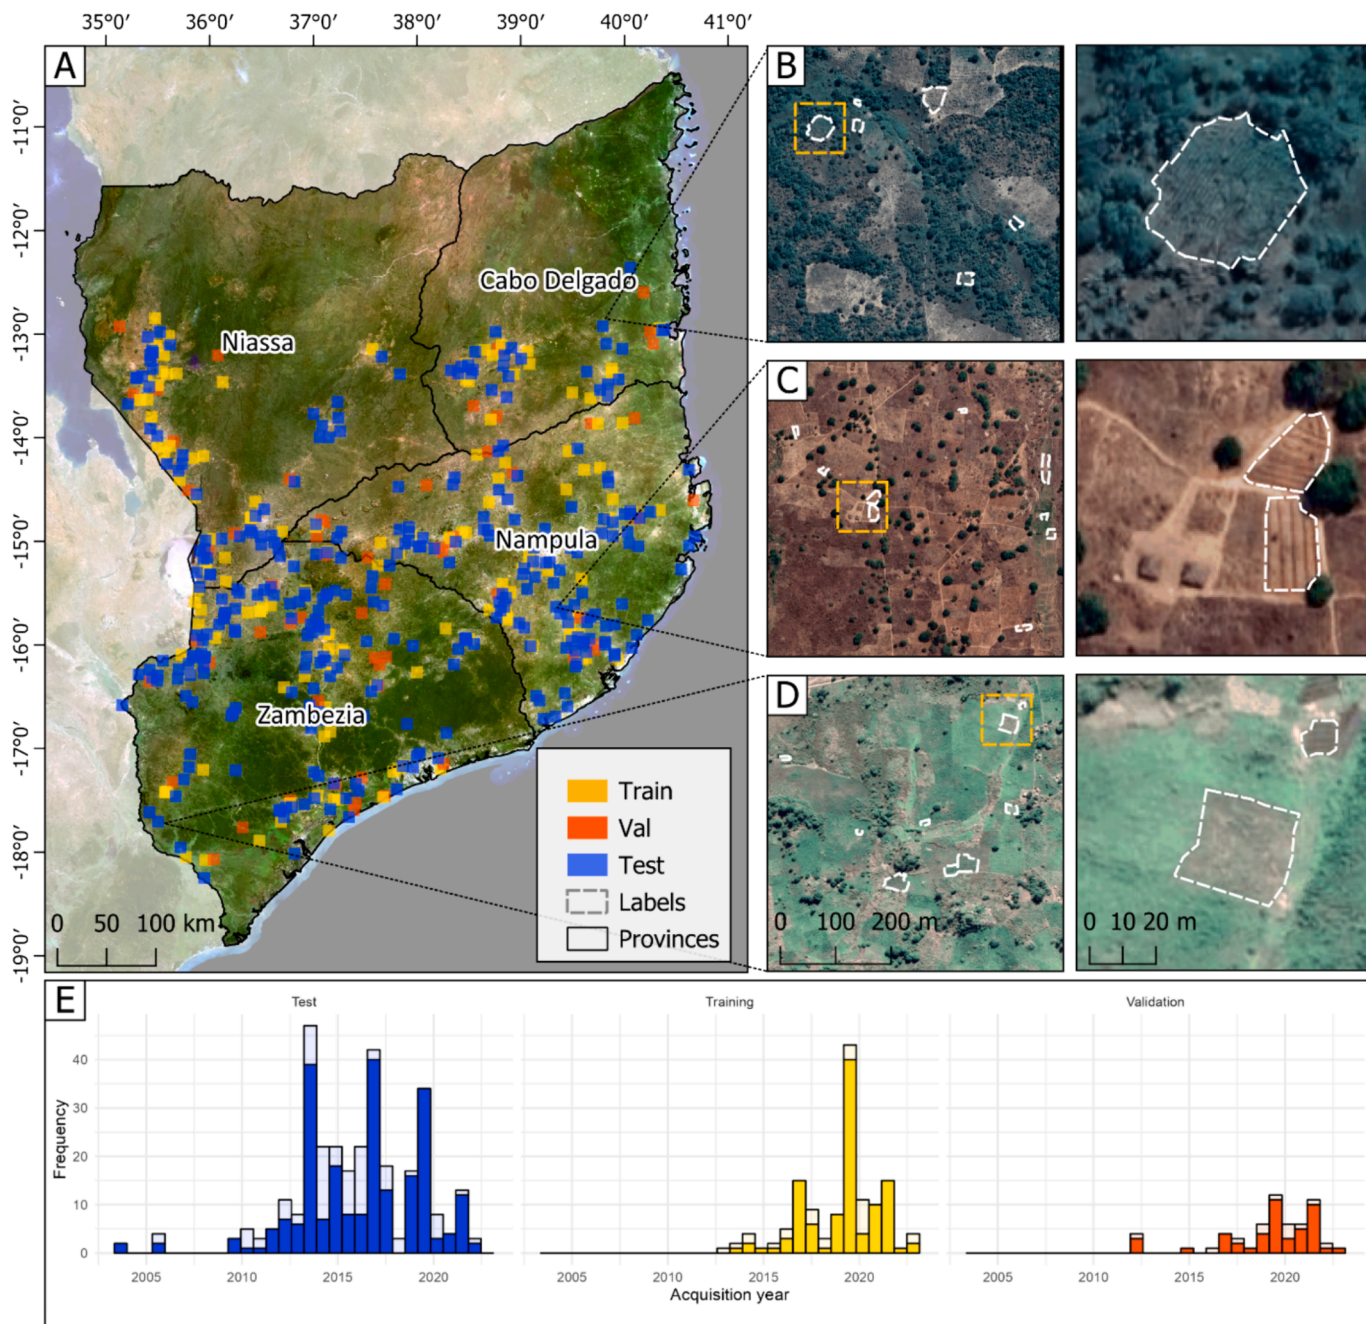


Fig. 1. Study region in Northern Mozambique (A) with site locations for training (yellow), validation (red), and test split (blue). Image data sources: Panel A: PlanetScope mosaic of May 2023, provided through the NICFI data program (Planet Labs Inc., 2023). Panels B-D illustrate example sites with sparse labels and zoom ins. Image data sources: Google Earth Pro VHR imagery (©2023 Maxar Technologies). Panel E shows histogram of image acquisition years in each split, separated by dry season (filled bars) and wet season (hollow bars). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

assessed object-level spatial agreement and site-level field size estimates for our test sites. We compared the performance gains obtained by fine-tuning with human-annotated labels with those obtained by pseudo-labels. While pseudo-labels may present an opportunity for fine-tuning models in the absence of any training labels, we further explored the situation where limited human-annotated labels can be complemented with pseudo-labels. Moreover, we complemented the method proposed by Wang et al. (2022) by using pseudo-labels for learning non-crop land cover types to overcome the requirements for up to date and spatially detailed cropland maps otherwise needed to separate crop fields from other land cover types, which are lacking in many contexts. Our objectives were to: (1) create semantically complete pseudo-labels for field delineation which are similar to human-annotated labels, (2) assess the performance gains achieved by fine-tuning with pseudo-labels relative to the performance gains obtained by using human-annotated labels for fine-tuning in label-scarce settings, and (3) explore the added value of pseudo-labels as a complement to human-annotated labels.

## 2. Methods

### 2.1. Study area

Our study area covers four provinces (Cabo Delgado, Nampula, Niassa, Zambezia) in the North of Mozambique, comprising an area of ~380,000 km<sup>2</sup> (Fig. 1A). The farming system of the region is dominated by semi-subsistence smallholder farming, thus partly cultivating for subsistence while also engaging in commercial activities on local markets, and limited access to capital and a high reliance on manual labor (Fig. 1B-D). Agricultural production is conducted with low inputs, such as fertilizer, pesticides, or herbicides, and access to mechanized equipment is rare. Key crops in the region are maize, beans, sorghum, cassava, soy, or non-food crops such as tobacco and cotton, as well as tree crops such as cashew and macadamia (Ministério da Agricultura e Desenvolvimento Rural, 2017).

Fields in the North of Mozambique are very small overall (Lesiv et al., 2019), but spatially explicit and nuanced data on the distribution of fields is not available. Based on our reference data, we estimated a range of field size between 0.001 ha and 1.4 ha, with a median of 0.06 ha. More than 99.8 % of the fields in our reference data are smaller than one hectare, but substantial local-level variation (standard deviation = 0.16 ha) can be observed. Medium-scale farms, defined as those operating on land holdings of 5 – 100 ha, comprising fields of around 0.5—1 ha, are emerging in many parts of Sub-Saharan Africa (Jayne et al., 2016), including Mozambique (Baumert et al., 2019). Large-scale land acquisitions in the early 2000 s led to the emergence of large fields covering more than one hectare, up to very large fields of more than 50, up to one hundred hectares, although they remain rare in relative terms (Bey et al., 2020).

### 2.2. Data

#### 2.2.1. Very-high resolution imagery

We used VHR image data contained in Google Earth Pro™. The image catalog of Google Earth contains imagery from Maxar Technologies, including WorldView 2 and 3 and GeoEye 1 imagery, as well as imagery acquired during airborne campaigns, all of which have a sub-meter spatial resolution. For site selection, we developed a stratified random sampling scheme to sample from regions with actively used cropland. To identify these regions, we used an existing map of active and fallow cropland for the growing season of September 2020 through August 2021 (Rufin et al., 2022). We aggregated the map to a 1 ha grid and calculated the proportions of active cropland. We sampled 1,000 sites from regions mapped as containing at least 50 % of active cropland within a one-hectare grid cell. We defined a site extent of 600 by 600 m, or 36 ha, in order to assure that a sufficient number of fields can be delineated, even in regions with comparatively large field size. The

selected sites were screened for VHR image quality and acceptable visibility of at least five fields, resulting in 513 sample sites (see Fig. 1A). We requested Google Earth™ imagery through the Google Maps API in R, resampled the RGB imagery to ~ 0.6 m per pixel by extracting 1024 x 1024-pixel images for rectangular sites of 600 x 600 m and stored them in GeoTIFF format. Image acquisition years varied between 2003 and 2022 (Fig. 1E). Approximately 50 % of the images were acquired in 2016 or later, and 2019 was the most frequent year for image acquisition. Over 75 % of the images were acquired during the dry season between June and November.

#### 2.2.2. Labels

Collecting field delineation labels in heterogeneous smallholder landscapes is challenging due to high levels of fragmentation, the presence of very small fields, high intra-field heterogeneity in crop cover, and a high reliance on manual labor, resulting in irregular field boundaries, which partly exert low contrast and are thus not clearly visible. Furthermore, collecting fully annotated labels with high confidence boundaries is in some cases hardly possible for regions where markers of field delineations are not distinct (Estes et al., 2022). In our study region, extensive cultivation in recently cleared forests represents one example of hard-to-interpret field delineations (Fig. 1B). In such contexts, the use of weak supervision adds another benefit as it allows including field delineations from complex regions, where fully annotated labels can either not be created, or would imply sacrificing label quality by including uncertain field delineations.

We followed the routine described in Wang et al. (2022) and tasked human annotators to collect sparse labels (i.e. at least five fields per site). We tasked the interpreters to collect only fields containing non-tree crops by systematically excluding tree crop plantations from our data. While individual trees in the field interior were included in our labels, trees overlapping with the field boundaries were avoided and the tree canopy was considered as the field boundary for completing the labels. All field delineations underwent an iterative quality assessment, where 18 % of the initial field delineations and 7 % of the field delineations in a second iteration were discarded. We compiled a total set of 3,711 field delineations. We split the resulting 513 labeled sites into 200 training and 313 test sites, comprising 1,512, and 2,199 fields, respectively. We determined the number of labels required for training based on the ablation experiments in Wang et al. (2022), which indicated that 500 – 1,000 field labels were sufficient for fine-tuning a model to reach optimal performance. All labels used in this study are made publicly available for download via <https://doi.org/10.5281/zenodo.11488976>.

### 2.3. Model architecture

We based our experiments on the DECODE framework, comprising a Fractal ResUNet architecture for dense predictions of field extent, field boundaries, and within-field distance to the next boundary, which are used to generate semantically complete field instances using hierarchical watershed segmentation. The methods are described in detail in Waldner and Diakogiannis (2020), and Waldner et al. (2021). The Fractal ResUNet resembles the typical UNet encoder-decoder architecture and features residual blocks with a built-in attention mechanism. The multiple model heads are designed for multi-task predictions on crop field extent, field boundaries, and the within-field distance to the nearest field boundary. The model was trained with a fractal Tanimoto loss to evaluate all tasks simultaneously. The Fractal ResUNet outperforms other state-of-the-art model architectures in the task of field delineation (Tetteh et al., 2023), and has been shown to generalize well across regions (Waldner et al., 2021).

The DECODE framework was used by Wang et al. (2022), who reported high performance in smallholder agricultural regions of India using multi-spectral pan-sharpened SPOT6/7 data at 1.5 m spatial resolution. For model training, Wang et al. (2022) processed the four-channel input data (RGB and near-infrared) to a three-channel false-

color representation with channel 1 as the average of the red and green bands, band 2 as the average of the red and NIR bands, and band 3 as the average of the green and blue bands. The authors conducted pre-training using publicly available field delineations from the land parcel information system in France. They fine-tuned the model for field delineation in India using sparse labels created by human annotators across 2,000 sites. Due to the sparse annotation procedure, only five fields per site were included in the training, which allowed for a broader geographic coverage and, thus, a higher diversity of fields in the training data.

#### 2.4. Pseudo-label selection

We developed a workflow for using a pre-trained model for producing semantically complete field instances that can be used as training and validation labels to fine-tune models across geographies (Fig. 2). In this work, we test the workflow using a model pre-trained in France and India for field delineation in Mozambique, which represents a structurally different and more complex target region.

We assessed the performance of the publicly available, pre-trained FracTAL ResUNet by Wang et al. (2022) to predict fields in Mozambique, which can be used as pseudo-labels for fine-tuning the model. Given the assumption that the basic properties of field appearance across geographies are similar, we reason that the pre-trained model can produce accurate field delineations in the known feature space without fine-tuning. We used the model to create predictions of cropland and boundary probabilities for our 200 training sites, and performed instance generation from the predictions based on hierarchical watershed segmentation (Waldner et al., 2021). We kept the hyperparameters of the watershed segmentation constant at boundary threshold = 0.2 and extent threshold = 0.4 following Wang et al. (2022) for better comparison across model runs while acknowledging that optimizing the hyperparameters based on validation labels may further improve performance.

We then calculated a set of instance-level scores using the model output activations (sigmoid-scaled in the case of field boundary predictions) for both cropland field extent and field boundaries. For pseudo-label selection, we derived the semantic confidence  $SemC_i$ , for each instance  $i$  comprising  $n$  pixels as the median of the pixel-level field extent probabilities  $p_{ext}$ :

$$SemC_i = P_{50}(p_{ext})$$

We derived instance confidence  $InsC_i$  relating to the median boundary probabilities  $p_{bnd}$  in all boundary pixels of instance  $i$  as:

$$InsC_i = P_{50}(p_{bnd})$$

Lastly, we calculated the size of each instance as the number of pixels  $n$  comprising the instance:

$$Size_i = n$$

We first tested selection based on thresholding  $SemC_i$  using absolute thresholds  $T_{SemC}$ , so that a pseudo-label was selected if:

$$SemC_i > T_{SemC}$$

We tested four absolute probability thresholds  $T_{SemC} = \{0.925, 0.950, 0.975, 0.990\}$ .

The absolute thresholds constrain pseudo-label selection to the overall most certain predictions, which reduces the likelihood of including erroneous pseudo-labels, but may limit pseudo-label selection to images and settings with comparatively high confidence and may thus reduce the diversity of pseudo-labels. To increase diversity of pseudo-labels, we tested adaptive thresholds based on the site-level distribution of confidence scores. Thresholds were derived using percentiles of all instances-level scores per site  $P_j(SemC_N)$  with  $j$  denoting the respective percentile, and  $N$  being all instances within the site. We considered  $j = 99, j = 98$ , and  $j = 95$ , thus selecting the 1%, and 2%, most confident predictions, respectively. When targeting a pseudo-label selection based on  $P_j(SemC_N)$  and  $P_j(InsC_N)$  simultaneously, we defined  $j = 95$ , i.e., selecting the top 5% from each metric, because instances with both high semantic and high instance confidence were rare:

$$SemC_i > P_{95}(SemC_N) \ \& \ InsC_i > P_{95}(InsC_N)$$

Moreover, we considered a minimum size criterion to avoid the selection of small fragments consisting of a few confident pixels, such as individual trees which were apparent in the predictions. We defined a minimum size of  $T_{Size} = 180 \text{ m}^2$  (i.e., 500 pixels) which corresponds to the extent of a large tree canopy, as often present on or nearby fields in the region. We tested the size criteria in combination with  $P_{99}(SemC_{image})$  only. Following this logic, an instance was selected from the pool of pseudo-label candidates if

$$SemC_i > P_{99}(SemC_N) \ \& \ Size_i > T_{Size}$$

Following this procedure, we compiled a total of eight sets of pseudo-labels, with four selected based on absolute and four based on adaptive thresholds. We then compared human-annotated and pseudo-labels in terms of label quantity, area properties, and spatial agreement. Besides testing the use of pseudo-labels for fully replacing human-annotated data, we also tested the value of pseudo-labels as a complement to human-annotated labels. For this, we generated pseudo-labels for an additional set of images ( $n = 400$ ) following the percentile-based routine described above and complemented the pseudo-labels with human-annotated data for the 200 training sites, thereby providing 600 training sites to the model. The code used to produce pseudo-labels are publicly accessible via GitHub (<https://github.com/philipperufin/pseudo-fields/>).

#### 2.5. Non-cropland labels

One caveat of the weak supervision strategy employed by Wang et al. (2022) is that the model did not learn to discriminate cropland from non-cropland, as reference data for non-cropland classes were not included during training. Reliable attribution of instances to cropland or other land cover or land use types thus requires a detailed, timely, and accurate cropland mask or additional updating of the non-cropland classes. We leveraged the uncertain predictions of the pre-trained model to add patches of non-cropland samples to training. We complemented non-cropland pseudo-labels for the absolute thresholds by selecting instances if:

$$SemC_i < T_{SemC} = 0.75,$$

For the adaptive strategy, we defined  $j = 10$ , selecting non-cropland labels if:

$$SemC_i < P_{10}(SemC_N),$$

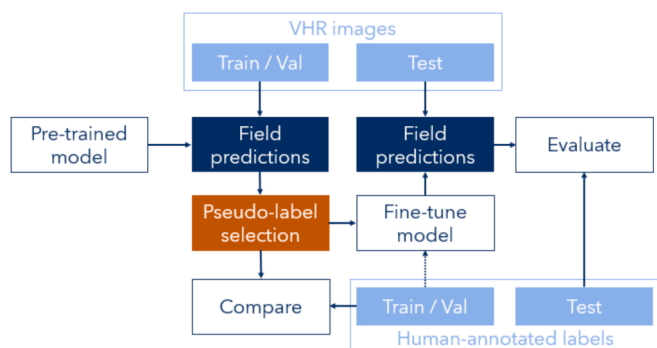


Fig. 2. Workflow of this study.

For the combined selection based on semantic and instance confidence, we selected an instance of non-cropland if:

$$SemC_i < P_{25}(SemC_N) \& InsC_i < P_{25}(InsC_N).$$

We did not account for a minimum instance size during negative pseudo-label selection due to the complexity of the landscapes with highly variable instance sizes depending on the underlying non-cropland land cover type. An example for field and non-cropland pseudo-label selection is illustrated in Fig. 3.

### 2.6. Model fine-tuning

We used the training data obtained from human annotators and the five sets of machine-generated training data to update the pre-trained FracTAL ResUNet in ten experiments, namely 1) fine-tuning using the sparse labels generated from human annotators, 2) – 5) fine-tuning using pseudo-labels with  $SemC_i > T_{SemC}$  with  $T_{SemC} = 0.925, 0.950, 0.975, 0.990$ , 6)  $SemC_i > P_{99}(SemC_N)$ , vii)  $SemC_i > P_{99}(SemC_N) \&$

$Size_i > 180m^2$ , viii)  $SemC_i > P_{98}(SemC_N)$ , ix)  $SemC_i > P_{95}(SemC_N) \& InsC_i > P_{95}(InsC_N)$ , and x) human-annotated labels combined with pseudo-labels with  $SemC_i > P_{99}(SemC_N)$  for 400 additional sites.

From the pseudo-labels, we generated multi-task labels representing field boundaries, field extent, and within-field distance to the nearest boundary using Python’s OpenCV module (Bradski, 2000). For model updating, we cropped the 1024 x 1024-pixel images into sixteen 256 x 256-pixel image chips, resulting in 2,240 image chips. We removed chips without any field labels. This resulted in a varying number of image chips for the different fine-tuning experiments. The 200 training sites were split into 70 % (n = 140) for training and 30 % (n = 60) for validation. We ran the fine-tuning for 100 epochs with a batch size of 4 and a start learning rate of 0.0001 using Adam optimizer in all experiments to obtain comparable insights on training time and performance. We selected the best model based on the maximum Matthews correlation coefficient in the validation split and registered the number of epochs used for training. In the absence of the commercial source imagery used to pre-train the model in India, we are not able to investigate domain shifts in the feature space but are constrained to exploring the

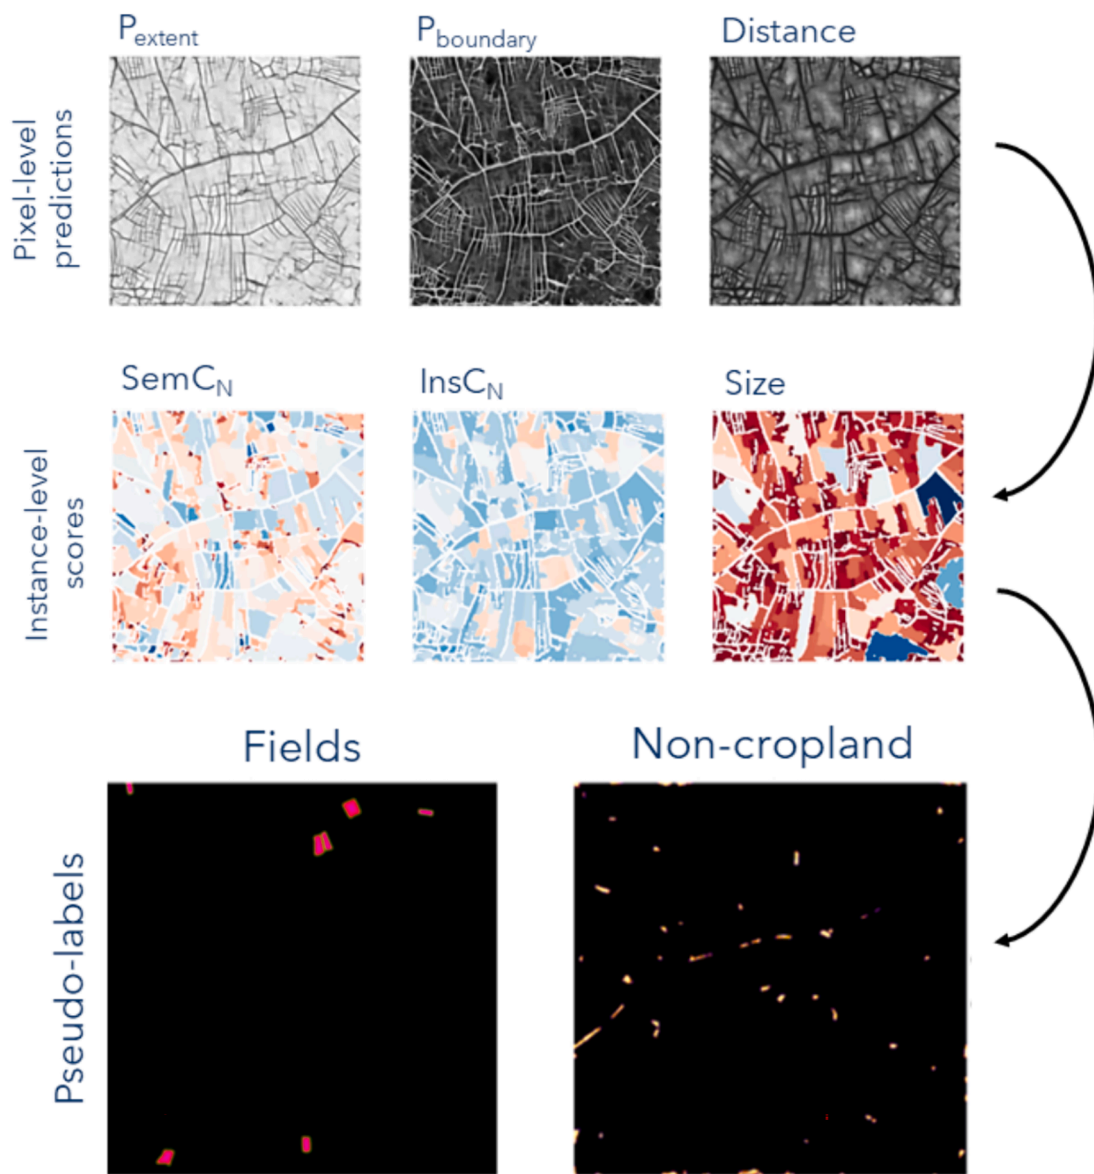


Fig. 3. Example for pseudo-label generation using  $SemC_i > P_{99}(SemC_N)$  for field labels and  $SemC_i < P_{10}(SemC_N)$  for non-cropland labels. Predictions of the input image (top row) are used to generate instances, for which scores are computed (middle row). Pseudo-labels for fields and non-cropland are selected from all instances based on the respective selection strategy, here  $SemC_i < P_{10}(SemC_N)$ (bottom row).

effect of fine-tuning with pseudo-labels through exhaustive evaluation of the prediction performance.

## 2.7. Performance evaluation

### 2.7.1. Field-level spatial agreement

For each resulting model, we predicted field boundaries for the 313 test sites and created field instances based on hierarchical watershed segmentation documented in Waldner et al. (2021) using cropland extent and boundary thresholds of 0.4 and 0.2, as described in Wang et al. (2022). We vectorized the predicted instances using GDAL/OGR contributors, (2023), matched predicted fields with reference fields using the centroid of the reference fields, and calculated the Intersection over Union (IoU), Precision, and Recall of each field instance using the segmetric package in R (Simoes et al., 2022). IoU is defined as the area of overlap between reference and predicted field divided by the area of their union. Precision refers to the summed area of the true positives over the area of true and false positives combined. Recall refers to the summed area of the true positives over the area of true positive and false negatives combined. For details on the implementation of the performance metrics in the segmetric package, we advise readers to consult (Rezatofighi et al., 2019) for IoU, and (Zhang et al., 2015) for Precision and Recall. We aggregated all field-level IoU scores using mean (mIoU), median (medIoU), the share of IoU scores above 0.5 (IoU<sub>50</sub>) and 0.8 (IoU<sub>80</sub>) conducted in-depth error diagnoses by assessing instance-level IoU scores by field size, the season of image acquisition, and province. The resulting insights on the error distribution were used to test for performance declines in particular subsets of the target domain when considering the different sets of pseudo-labels. We quantified pseudo-label performance gains as percent relative to the performance obtained with human-annotated labels to assess to what degree domain adaptation can be achieved without the need for any human-annotated labels:

$$\text{relativegain} = \left( \frac{\text{SCORE}_{\text{pseudo}} - \text{SCORE}_{\text{baseline}}}{\text{SCORE}_{\text{human}} - \text{SCORE}_{\text{baseline}}} \right) * 100$$

### 2.7.2. Site-level field size estimation

We further evaluated the suitability of the respective models for estimating field size at the site level. The assumption here was that while object-level performance scores may correlate with field size estimation performance, the latter relaxes the constraints on the spatial alignment of the individual field delineations while maintaining focus on the overall distribution of field size present in the predictions. We calculated field size for the reference and predicted fields, computed the residuals between reference and prediction, and derived root mean squared error (RMSE), mean absolute error (MAE), and mean error (ME) at the site level (for formulas please see e.g. Okujeni et al. (2018)). We then averaged the site-level scores across the test sites ( $n = 313$ ) to obtain mean RMSE (mRMSE), mean MAE (mMAE), and mean bias (mME). We further investigated the spatial distribution of the RMSE scores across provinces and calculated global Moran's I to test for spatial autocorrelation in the model errors. Additionally, we fitted linear regression models between observed and predicted field sizes and recorded the goodness-of-fit as  $R^2$ , slope, and intercept of the regression. Following the object-level assessment, we quantified performance gains from pseudo-labels relative to performance gains obtained from human-annotated labels as percent of the performance increase (or error reductions, respectively) of using human-annotated labels against the pre-trained model.

### 2.7.3. Non-cropland identification

We evaluated the discrimination of cropland against other types of land cover and land use that was only trained during fine-tuning based on non-cropland pseudo-labels. While we acknowledge that a robust evaluation of the separation between cropland and non-cropland is

ideally conducted on fully annotated labels, the landscape complexity did not allow for compiling fully segmented labels on non-cropland. We therefore compiled sparsely annotated non-cropland labels for the test sites, which we could derive from an existing map product indicating the presence of cropland at approximately 5 m spatial resolution for the growing season 2020/2021 (Rufin et al., 2022). We considered all regions mapped as open woodlands, closed woodlands, non-vegetated surfaces, and water with an area comprising at least four pixels (or  $\sim 91 \text{ m}^2$ ) as candidates for evaluating the ability of the model to separate non-cropland. Due to the mismatches between spatial resolutions and acquisition dates for most images in the test dataset, we manually screened all candidate regions to remove all non-cropland patches that were not representing the non-cropland domain accurately. We thereby compiled a final set of 2,143 reference data patches representing the non-cropland domain in our study region. We evaluated the non-cropland identification in the region using pixel-level overall accuracy, Precision, Recall, and F1-score per image and as average across all images in the test set. Moreover, we conducted an in-depth qualitative evaluation of the non-cropland identification to support the empirical data with qualitative insights.

## 3. Results

### 3.1. Pseudo-labels for field delineation

We compared the quantities and area properties of the pseudo-labels with human-annotated labels for the 200 training sites (Table 1, Fig. 4). When using absolute thresholds, the number of selected instances in the pseudo-labels as well as the average number of labels per site decreased with increasing confidence thresholds, with values ranging from 3,472 field instances for the  $T_{\text{SemC}} = 0.925$ , to 267 field instances for  $T_{\text{SemC}} = 0.990$ , against 1,517 field instances in our human-annotated data. A minimum of 5 fields per site was obtained at 163 sites for the human-annotated data, and between 167 ( $T_{\text{SemC}} = 0.925$ ) down to only 12 ( $T_{\text{SemC}} = 0.990$ ) sites for the pseudo-labels. Generally, the spread in the number of fields per site was higher for the pseudo-labels than the human-annotated data, where in some cases, up to 70 fields ( $T_{\text{SemC}} = 0.925$ ) or 38 fields ( $T_{\text{SemC}} = 0.990$ ) were generated. For several sites, no pseudo-labels could be generated, with the number of sites without pseudo-labels ranging from 5 ( $T_{\text{SemC}} = 0.925$ ) up to 133 when using the most conservative threshold ( $T_{\text{SemC}} = 0.990$ ). For the adaptive thresholds, we observed a lower variation in the number of selected fields compared to the absolute thresholds. Selecting the top 1 % predictions ( $P_{99}(\text{SemC}_N)$ ) yielded 766 fields, or 695 when including with the minimum size criterion. Relaxing the threshold to  $P_{98}(\text{SemC}_N)$  yielded 1,334 field instances, while the selection based on  $P_{95}(\text{SemC}_N)$  and  $P_{95}(\text{InsC}_N)$  combined yielded only 342 field instances (Table 1, Fig. 4).

The area properties of both the human-annotated and pseudo-labels selected with absolute thresholds were more similar when applying conservative thresholds. Discrepancies increased when relaxing the threshold, and the number of large fields increased substantially when using lower confidence thresholds. Mean and median field sizes in the human-annotated labels (0.120 ha, and 0.060 ha) and  $T_{\text{SemC}} = 0.990$  pseudo-labels (0.120 ha, and 0.080 ha) agreed well (Table 1). The standard deviation for the human-annotated labels was slightly higher (0.164 ha) as compared to the pseudo-labels (0.114 ha). Using the adaptive thresholds, the mean area was closest to the human-annotated labels with  $P_{95}(\text{SemC}_N)$  and  $P_{95}(\text{InsC}_N)$  combined, followed by  $P_{99}(\text{SemC}_N)$ . We found that using  $P_{95}(\text{SemC}_N)$  and  $P_{95}(\text{InsC}_N)$  yielded a range of field sizes comparable to our human-annotated labels but with an overrepresentation of small fields with a size of approximately 0.02 ha and less (Fig. 4).

We found high spatial agreement between the human-annotated labels and the pseudo-labels with increasing agreement at higher confidence thresholds (Table 2). The pseudo-labels selected at  $T_{\text{SemC}} = 0.990$ , reached a mIoU of 0.802, medIoU of 0.868, IoU<sub>50</sub> of 0.914, IoU<sub>80</sub> of

**Table 1**  
Descriptive statistics of human-annotated and pseudo-labels.

	Human-annotated labels	Absolute thresholds				Adaptive thresholds			
		$T_{SemC} = 0.925$	$T_{SemC} = 0.950$	$T_{SemC} = 0.975$	$T_{SemC} = 0.990$	$P_{99}(SemC_N)$	$P_{99}(SemC_N)$ & Size	$P_{98}(SemC_N)$	$P_{95}(SemC_N)$ & $P_{95}(InsC_N)$
Total N fields	1,517	3,472	2,406	935	267	766	695	1,334	342
Mean N fields	7.74	17.36	12.03	4.68	1.34	3.83	3.475	6.67	1.71
Max N fields	24	70	61	64	38	13	10	19	12
N sites $\geq 5$ fields	163	167	137	55	12	70	54	152	15
N sites 0 fields	4	5	18	71	133	3	3	0	68
Mean (ha)	0.1187	0.6856	0.4052	0.1727	0.1172	0.1285	0.1326	0.2114	0.1161
Median (ha)	0.0600	0.1162	0.1113	0.0950	0.0806	0.0741	0.0771	0.0797	0.0594
SD (ha)	0.1642	2.7494	1.5031	0.3466	0.1144	0.2149	0.2176	1.2423	0.2305

0.743, and a high Precision of 0.871 and Recall of 0.904. Importantly, Precision and Recall were also balanced in this case.

We observed slight differences in the number of pseudo-labels per site between the dry season, here defined as June through November, and the wet season, lasting from December through May (Fig. 5). Differences between seasons were statistically insignificant (two-tailed  $t$ -test) for the human-annotated labels and absolute confidence thresholds ( $p > 0.1$ ), but significant differences ( $p < 0.05$ ) were observed for three adaptive selection strategies, with a higher average number of labels for wet season observations as compared to dry season observations for  $P_{99}(SemC_N)$  (dry season: 3.4, wet season: 4.4,  $p = 0.015$ ),  $P_{98}(SemC_N)$  (dry season: 6.4, wet season: 8.0,  $p = 0.022$ ), and  $P_{95}(InsC_N)$  and  $P_{95}(InsC_N)$  (dry season: 1.6, wet season: 2.5,  $p = 0.020$ ).

### 3.1.1. Fine-tuning & learning curves

The learning curves for all fine-tuning runs indicate stable learning behavior (Fig. 6). Maximum MCC scores on the validation split were reached in 20 ( $T_{SemC} = 0.990$ ) to 66 (Human-annotated labels) epochs. While training loss (blue) further decreases after the selected number of epochs, increasing validation loss suggests that the models start to overfit on the training data and fine-tuning performance saturates. Variations in absolute loss and metric values between the different training runs arise from differing number of image chips containing labels and the characteristics of the fields therein.

### 3.1.2. Object-level performance

We assessed object-level agreement between predictions of the fine-tuned models and our independent reference data using mIoU, IoU<sub>50</sub>, Precision (or User's accuracy), and Recall (or Producer's accuracy) (Fig. 7). We used the pre-trained model as a baseline to assess relative performance improvements, while the model trained based on human-annotated labels represented the target (Table 3).

When using absolute thresholds, performance generally declined with lower threshold values. A notable exception was  $T_{SemC} = 0.950$  with high mIoU and IoU<sub>50</sub> scores, yet these came at the expense of below-baseline precision. We found the lowest performance of all pseudo-labels at  $T_{SemC} = 0.925$ , except for a substantial increase in Recall, which was, however, outweighed by a decline in Precision, indicating substantial undersegmentation. The high absolute thresholds  $T_{SemC} = 0.975$  and  $T_{SemC} = 0.990$  yielded moderate performance gains across all metrics.

Adaptive thresholds in pseudo-label selection led to overall higher performance gains. The combination of  $P_{95}(SemC_N)$  and  $P_{95}(InsC_N)$  performed the worst, followed by the selection using  $P_{98}(SemC_N)$ .  $P_{99}(SemC_N)$  outperformed all other approaches at the object level and consistently led to substantial performance increases. Further constraining pseudo-label selection using size thresholds degraded the performance gains. Fine-tuning with pseudo-labels based on  $P_{99}(SemC_N)$  allowed for achieving 77.4 %, and 85.6 % of the human-annotated label performance gains in mIoU, and IoU<sub>50</sub>, while consistently reducing errors across field sizes and sub-regions.

Combining human-annotated data with pseudo-labels for additional sites further improved the performance of the model across all scores. At the object level, mIoU was 0.686 using human-annotated labels and 0.694 when complemented with pseudo-labels, and IoU<sub>50</sub> increased from 0.801 to 0.816, exceeding the performance gains obtained from human-annotated labels.

In-depth error diagnoses reveal the distribution of performance gains across the observed range of field size and across different parts of the study region in more detail (Fig. 8). The overall distribution of IoU scores improved for all fine-tuning experiments relative to the baseline model (Fig. 8, left panel). Correctly identifying very fields smaller than 0.01 ha was challenging in all models (Fig. 8, center panel). For fields above 0.01 ha, all fine-tuned models had substantially improved performance over the pre-trained baseline model. For fields above 0.05 ha, the models involving human-annotated labels outperformed the models trained exclusively with pseudo-labels. Comparing human-annotated labels only against combined human-annotated and pseudo-labels reveals similar performance across the range of field size, with combined labels slightly outperforming human-annotated labels across the entire range. Assessing median model performance across provinces further attested that combining pseudo-labels with human-annotated labels performs at least similarly well, or better than human-annotated labels only. Using only pseudo-labels was on par with using only human-annotated labels in Nampula province, and pseudo-labels even outperformed human-annotated labels in Niassa province. Differences in IoU between predictions in dry season versus wet season imagery were statistically insignificant ( $p > 0.1$ ) in all models.

A qualitative assessment of the field predictions (Fig. 9) attested to the accurate field delineation results and how regional fine-tuning helps to overcome context-specific issues of under- (Fig. 9A, C, E), or over-segmentation (Fig. 9D). The pre-trained model performed well in some cases (Fig. 9B, D), attesting to the hypothesis that accurate delineations can be created for individual (but not all) fields when transferring models across geographies. However, there were cases where neither the pre-trained nor the fine-tuned models produced optimal field delineations (Fig. 9F).

### 3.1.3. Field size estimation performance

We assessed site-level performance based on mean RMSE (mRMSE), median RMSE (P50RMSE) and mean MAE (mMAE) scores, as well as adjusted  $R^2$  of observed against predicted field size (Fig. 10). Again, we assessed performance gains of pseudo-labels relative to using only pseudo-labels (Table 4).

Similar to the object-level analyses, using liberal absolute thresholds ( $T_{SemC} = 0.925$  and  $T_{SemC} = 0.950$ ) led to a performance below baseline for mRMSE, mMAE, and  $R^2$ . More conservative thresholds ( $T_{SemC} = 0.975$  and  $T_{SemC} = 0.990$ ) yielded performance increases in mRMSE and mMAE, but only  $T_{SemC} = 0.990$  led to increases in  $R^2$ . Again,  $P_{99}(SemC_N)$  consistently outperformed all other approaches.

Combining human-annotated data with pseudo-labels for additional sites improved the performance across all metrics. We observed an

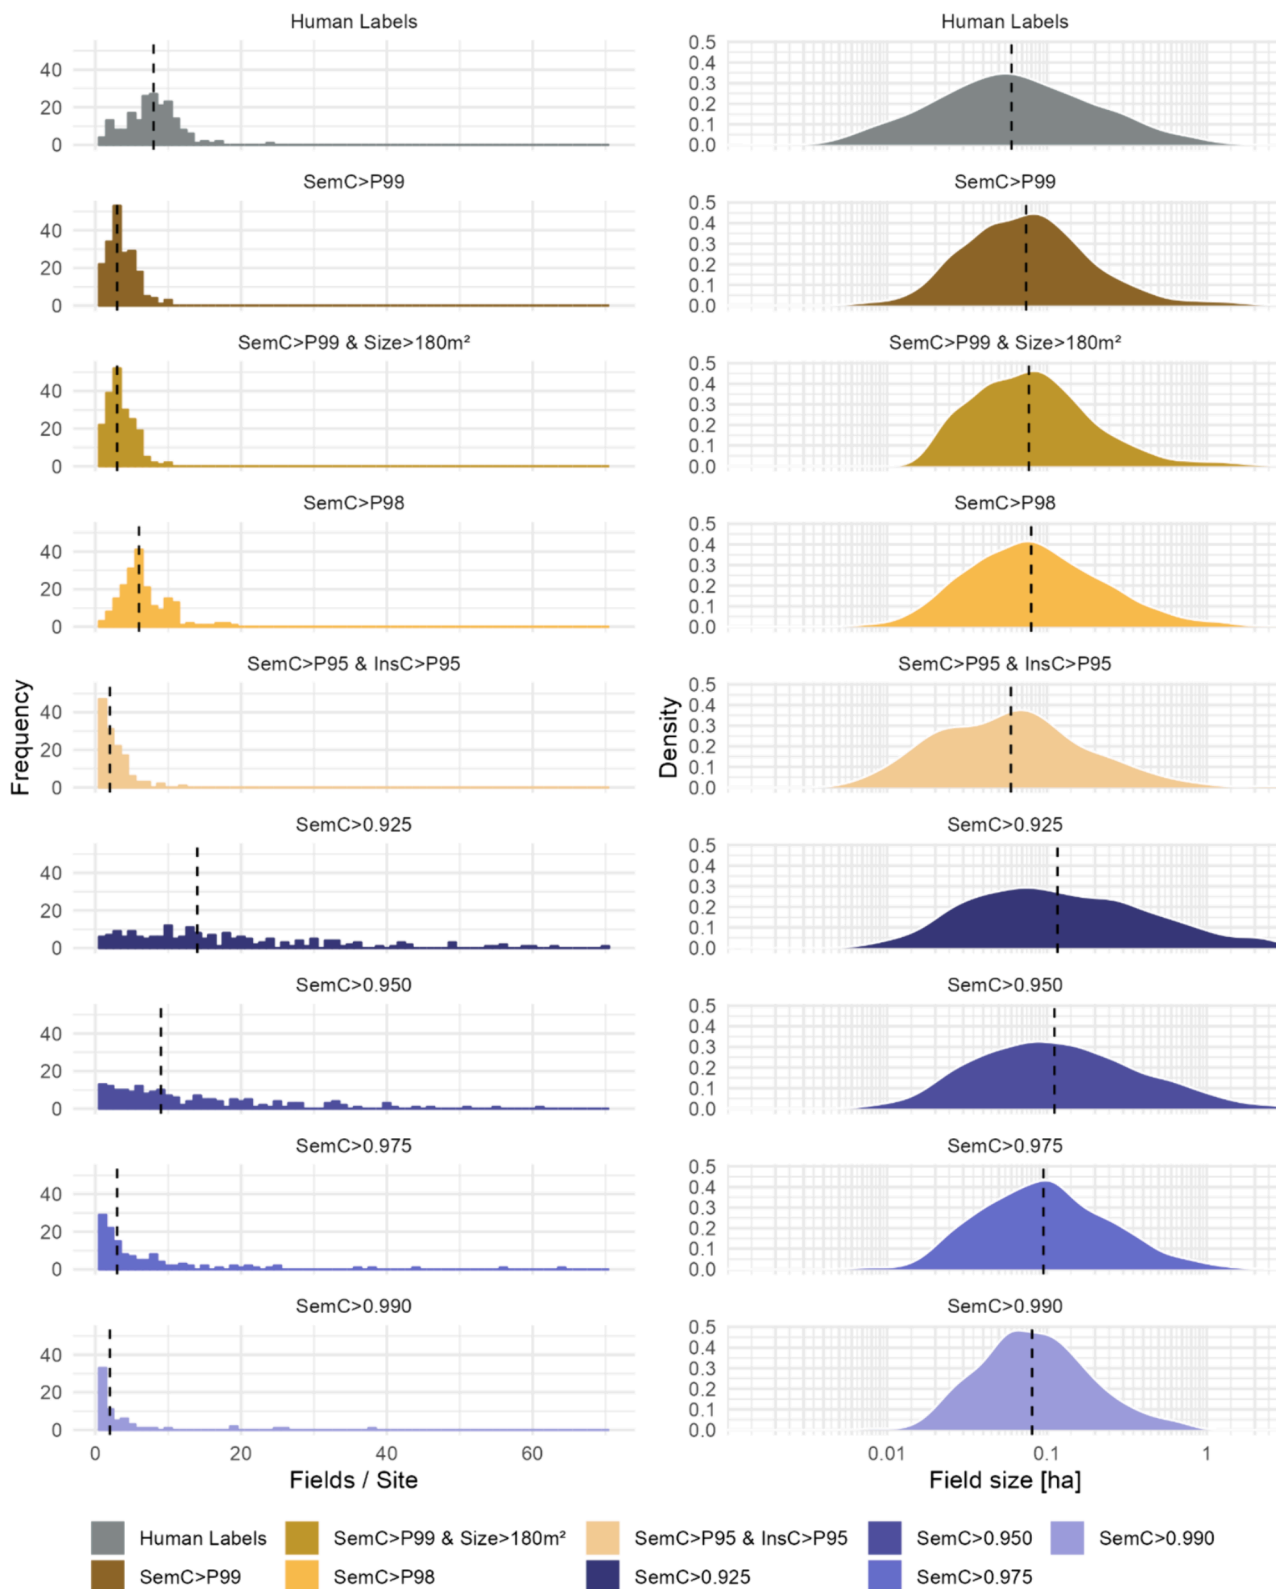


Fig. 4. Number of fields per site (left column) and density curves of field size in hectares (right column) of human-annotated labels (top row) and pseudo-labels with varying thresholds (rows 2–9). Note that x-axis of field size is log-scaled.

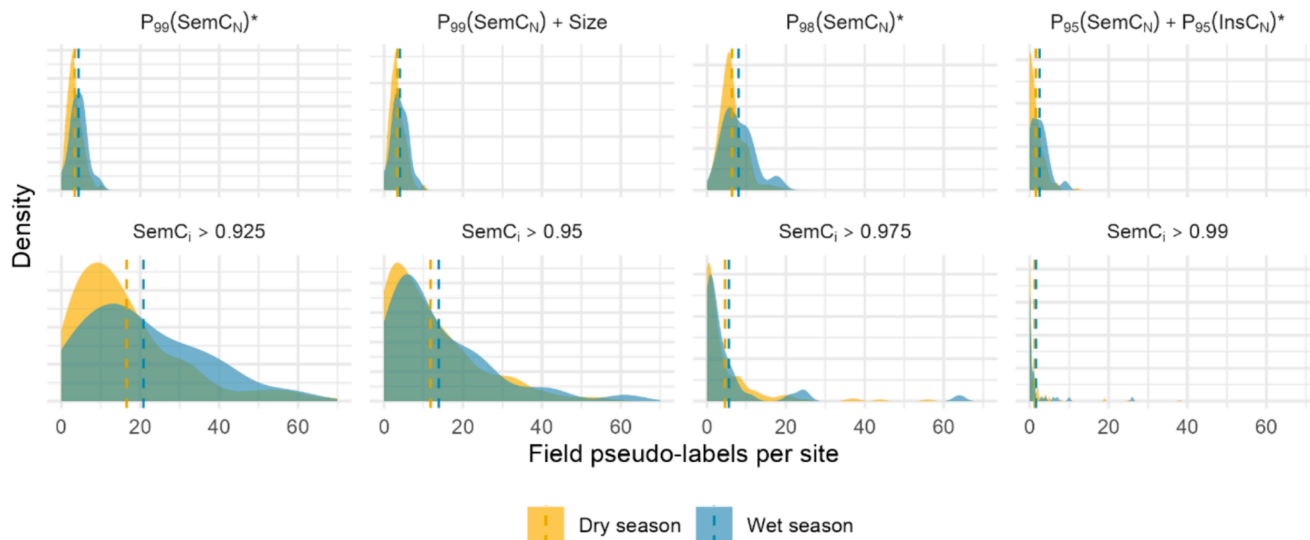
mRMSE decrease from 0.040 ha using only human-annotated labels to 0.037 ha when complementing with pseudo-labels, and  $R^2$  increased from 0.799 to 0.831, again exceeding the performance gains obtained from human labels only.

Error diagnoses revealed the high agreement between observed and

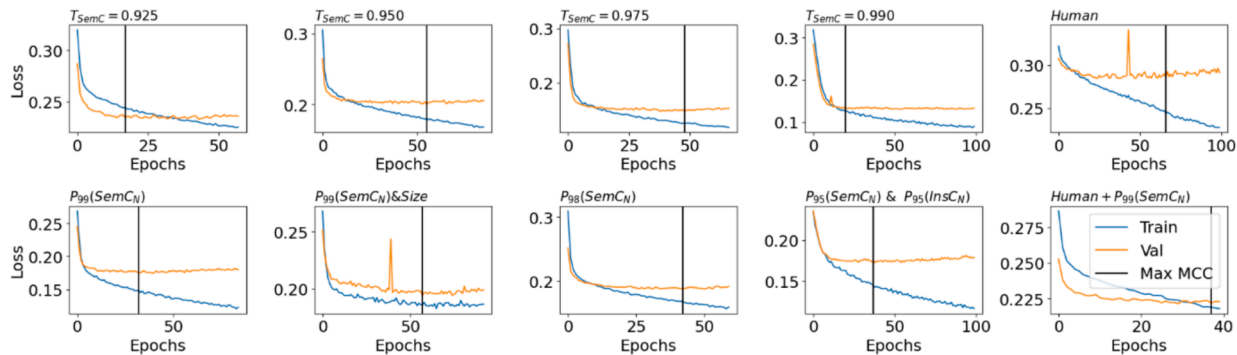
predicted field size when using human-annotated labels or combined labels and how distributions of RMSE shift towards lower RMSE (Fig. 11). Using human-annotated labels while consistently reducing median errors across the four sub-regions. Combining pseudo-labels with human-annotated labels outperformed human-annotated labels

**Table 2**  
Performance metrics comparing pseudo-labels with human-annotated labels for training sites.

	$T_{SemC} = 0.925$	$T_{SemC} = 0.950$	$T_{SemC} = 0.975$	$T_{SemC} = 0.990$	$P_{99}(SemC_N)$	$P_{99}(SemC_N) \& Size$	$P_{98}(SemC_N)$	$P_{95}(SemC_N) \& P_{95}(InsC_N)$
mIoU	0.629	0.663	0.743	0.802	0.741	0.741	0.735	0.732
medIoU	0.741	0.777	0.821	0.868	0.804	0.807	0.804	0.811
IoU <sub>50</sub>	0.729	0.773	0.857	0.914	0.855	0.853	0.847	0.845
IoU <sub>80</sub>	0.396	0.439	0.571	0.743	0.503	0.509	0.510	0.563
Dice	0.561	0.554	0.830	0.887	0.828	0.827	0.816	0.797
Precision	0.422	0.406	0.769	0.871	0.779	0.779	0.763	0.709
Recall	0.838	0.868	0.902	0.904	0.882	0.882	0.876	0.911



**Fig. 5.** Number of machine-generated reference fields per site by season of image acquisition. Pseudo-labels with significant ( $p < 0.05$ ) differences between dry season (yellow) and wet season (blue) are marked with asterisk. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)



**Fig. 6.** Learning curves for all experiments. The number of epochs for training is determined using the maximum MCC score of the validation loss.

only. To account for spatial autocorrelation in model errors, we assessed global Moran's I statistics of the site-level RMSE scores. We did not find a significant ( $p < 0.1$ ) spatial autocorrelation in the field size residuals.

### 3.1.4. Non-cropland identification performance

The quantitative assessment of the non-cropland identification revealed a general tendency to underestimate non-cropland. While errors of commission were rare in the non-cropland domain, errors of omission were common. This resulted in high Precision (mean Precision = 0.822), and low Recall (mean Recall = 0.511) scores for the non-cropland identification (Fig. 12). The mean accuracy (i.e. combining the cropland and non-cropland classes) across all test sites was 0.652, which is not satisfactory but a large step forward given that the initial

model was not able to identify non-cropland at all, and given that the improvement has been achieved only by including non-cropland pseudo-labels in fine-tuning.

A qualitative evaluation of the field delineation and the non-cropland discrimination when using non-cropland pseudo-labels confirmed the qualitative findings of lacking discriminative capacities regarding the non-cropland identification. Larger continuous patches of open or closed tree cover were correctly mapped as non-cropland in most cases, although smaller patches therein were partly wrongly mapped as fields (Fig. 13A). In some cases, the improved recognition of non-cropland, relative to the pre-trained model, helped in discriminating isolated fields surrounded by non-cropland (Fig. 13B). Individual trees were in many cases correctly identified as non-cropland (Fig. 13C),

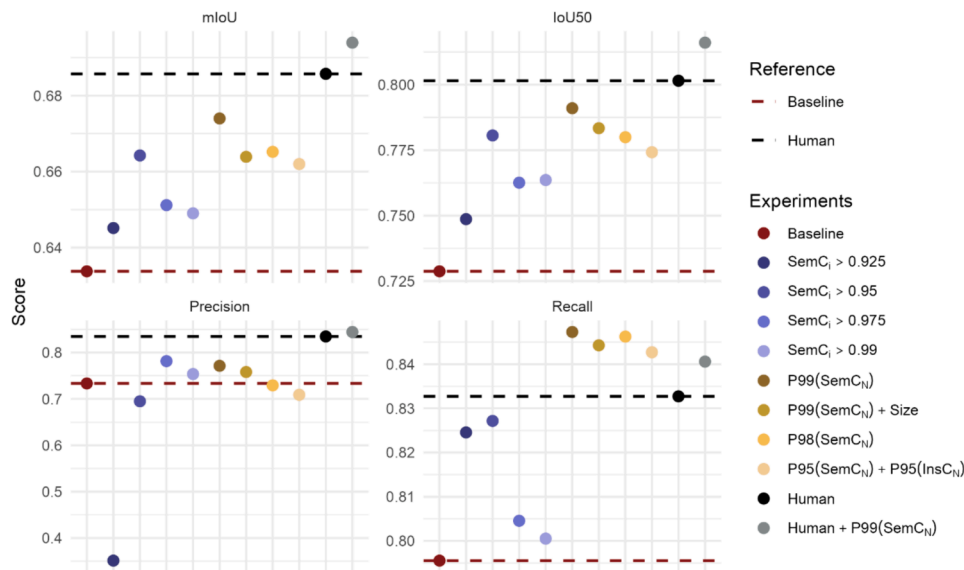


Fig. 7. Comparison of mIoU, IoU50, Precision, and Recall across all experiments. Dashed lines represent baseline model (red) and model trained with human-annotated labels (black). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3

Absolute performance scores for Baseline & Human-annotated labels, performance gain from Human-annotated labels, and relative gain for object-level. Best results using only pseudo-labels and combined human and pseudo-labels highlighted in bold.

	Experiment	mIoU	IoU <sub>50</sub>	Precision	Recall
Performance	Baseline	0.634	0.729	0.733	0.796
	Human-annotated	0.686	0.801	0.835	0.833
Gain from Human-annotated labels	Human-annotated – Baseline	0.052	0.073	0.101	0.037
	T <sub>SemC=0.925</sub>	21.9	27.4	–377.1	78.0
Gain relative to gain from Human labels [%]	T <sub>SemC=0.950</sub>	58.7	71.3	–38.0	85.0
	T <sub>SemC=0.975</sub>	33.5	46.5	<b>47.5</b>	24.2
Human-annotated + P <sub>99</sub> (SemC <sub>N</sub> )	T <sub>SemC=0.990</sub>	29.4	47.9	20.0	13.4
	P <sub>99</sub> (SemC <sub>N</sub> )	<b>77.4</b>	<b>85.6</b>	37.5	<b>139.2</b>
	P <sub>99</sub> (SemC <sub>N</sub> ) & Size	58.0	75.1	24.3	131.1
	P <sub>98</sub> (SemC <sub>N</sub> )	60.5	70.4	–4.1	136.4
	P <sub>95</sub> (SemC <sub>N</sub> ) & P <sub>95</sub> (InsC <sub>N</sub> )	54.3	62.5	–24.2	126.9
	Human-annotated + P <sub>99</sub> (SemC <sub>N</sub> )	<b>115.8</b>	<b>120.0</b>	<b>109.5</b>	<b>121.2</b>

but were falsely identified as fields in other cases, including for clusters of trees (Fig. 13D) or large trees with unevenly shaped canopies (Fig. 13E). Fallow lands were often correctly excluded from the cropland domain, although the discrimination was sensitive to the structure and texture and not consistent in all cases (Fig. 13F).

#### 4. Discussion

We demonstrated how pre-trained field delineation models can be leveraged to generate informative and diverse pseudo-labels which facilitate domain adaptation and help fine-tuning models for use in a new target region without the need to train models from scratch. We discuss our findings concerning 1) the high baseline performance of the pre-trained model and the added value of fine-tuning with human-annotated labels, 2) opportunities for mobilizing the value of pseudo-labels for geographic domain adaptation, 3) and challenges and

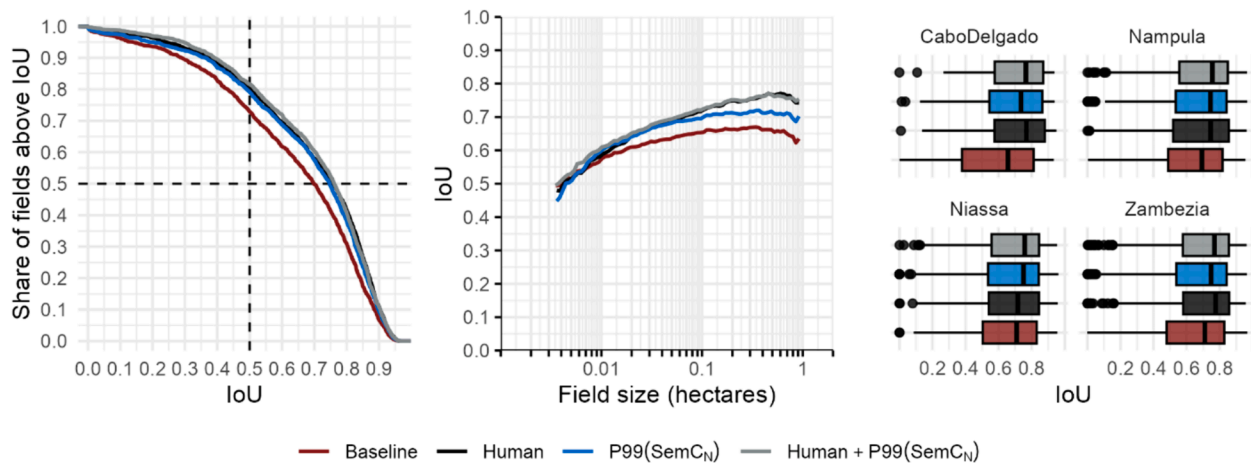
opportunities for large-area field delineation in smallholder agriculture.

##### 4.1. Baseline performance & added value of model fine-tuning

Our experiments are based on a pre-trained FracTAL ResUNet which has been shown to generalize well across regions and outperforms other state-of-the-art deep learning architectures for field delineation (Tetteh et al., 2023). In the heterogeneous setting of the North of Mozambique, the FracTAL ResUNet achieved a good baseline performance in both field delineation and field size estimation (mIoU=0.634, mRMSE=0.071). The performance of the pre-trained model was in line, or even exceeding the performance of previous studies in Sub-Saharan Africa, including field delineation in Kenya using Sentinel-2 and PlanetScope data with mIoU scores of up to 0.64 (Kerner et al., 2023), field delineation in Nigeria and Mali using 0.5 m World-View 3 imagery with F1 scores ranging between 0.6 and 0.7 (Persello et al., 2019), or field size estimation in Ghana, where predicted field size was on average 2.91 ha (141 %) above reference field size (2.06 ha) (Estes et al., 2022). It is important to note that the absence of a “good practices” protocol outlining key metrics for field delineation hampers direct comparisons of field delineation performance across studies.

While the good baseline performance of the pre-trained model may suggest that satisfying results in increasingly complex landscapes can be achieved without fine-tuning, as also reported by (Wang et al., 2022), this finding should be interpreted as context-specific to our study rather than generalized across all landscapes. The performance of field delineation is strongly dependent on the input imagery and the regional characteristics of the agricultural system, most notably the size of the fields present (Kerner et al., 2023; Nakalembe and Kerner, 2023). Here, we performed inference on images with a spatial resolution of approximately one-third of the training data (0.6 m used here as compared to 1.5 m in the training data), while median field size was around one-quarter of that present in the India data (0.06 ha as compared to 0.24 ha). Wang et al. (2022) highlighted the relevance of spatial resolution for smallholder field delineation by systematically degrading the spatial resolution of Airbus data, finding that performance decreased sharply, particularly affecting the lower end of the field size distribution. Consequently, inference on 1.5 m SPOT data in Mozambique may yield reduced baseline performance for smaller fields as compared to the higher resolution images used here.

Similar to Wang et al. (2022), fine-tuning the India model using data



**Fig. 8.** Comparison of for pre-trained baseline model (red), pseudo-labels selected with  $P99(SemC_N)$  (blue), human-annotated labels (black), and combined human-annotated labels and pseudo-labels (azure). IoU curves (left), mIoU across field size (middle), and province-level IoU distribution with median in black (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

from Mozambique consistently increased model performance across metrics in a comparable range of IoU (our mIoU: +0.05, median IoU in Wang et al.: +0.11) and IoU50 (ours: +0.07, Wang et al. 2022: +0.20), noting that the pre-trained model already yielded satisfactory performance.

#### 4.2. Mobilizing the potential of pseudo-labels

We found that when using pseudo-labels, IoU50 increased by up to 86 % and the decrease in mRMSE by up to 65 % compared to human-annotated labels. We explored a variety of methods for filtering pseudo-labels using instance-level scores involving semantic confidence, instance confidence, and instance size and tested absolute as well as image-adaptive thresholds for pseudo-label selection. The best selection strategy was based on site-level adaptive thresholds for semantic confidence, by which the 1 % most confident predictions were selected for each site. Compared to the selection based on absolute thresholds, adaptive selection likely increases the diversity of pseudo-labels and thus leads to better generalization (Prabhu et al., 2022). Simultaneously, the adaptive thresholds generated a higher number of pseudo-labels as compared to the most conservative absolute threshold. Due to the observed properties in terms of quantities, high spatial agreement, and similar area properties as compared to human-annotated labels, a selection based on adaptive thresholds is a suggested option when employing the confidence metrics used here.

The conservative pseudo-label selection used here represents a trade-off in terms of quantity, which may be mitigated by running pseudo-label selection across larger quantities of unlabeled image chips when these are available. Diversity may further be enhanced by adding sampling strata in site selection to account for agroecological zones (FAO, 2021). In applications where multi-temporal observations or image acquisition dates are accessible and can be included during sampling, the timing of image acquisition may be further accounted for to enhance the temporal generalization of the model. As such, we advise users to use conservative (e.g. 99th percentile) adaptive thresholds for pseudo-label selection and control the number and characteristics of the sites included appropriately to obtain optimal model performance. In cases when the pre-trained model has a poor baseline performance in the target region, or in cases where fields are numerous, even higher percentiles (e.g., 99.5 %) may be considered, or the adaptive thresholds may be combined with absolute confidence threshold to avoid the inclusion of highly uncertain predictions.

Self-training using pseudo-labels may lead to error propagation in model performance due to the self-referential nature of the approach


(Toldo et al., 2020). The introduction of biases in the model is a potential risk in self-learning approaches, which is hard to monitor. Therefore, we recommend a comparison with human-annotated data or a quality screening of pseudo-labels for downstream applications. We here observed a tendency of the models to undersegment fields in the region when including pseudo-labels of lower confidence, as expressed in substantial declines in Recall. We, therefore, advise against the use of liberal selection criteria relying on absolute thresholds below 0.975, and adaptive thresholds selecting beyond the 1 % most confident instances. Moreover, manual screening (i.e., a simple site-level quality assessment of pseudo-labels) can be conducted to effectively minimize the risk for erroneous pseudo-labels in the training dataset and thus represents a cost-efficient middle ground to produce labels at scale without the need for costly annotation.

Our experiments on complementing human-annotated data with pseudo-labels for additional sites confirmed that performance increases can be easily achieved by providing diverse, but highly confident pseudo-labels for complementary sites. We thus advise the combined use of human-annotated labels and quality-screened pseudo-labels wherever possible to obtain optimal performance and reduce the risk of error propagation. Moreover, the use of iterative training may help to overcome the relatively poor discrimination of cropland and non-cropland in the region, which was dominated by overpredictions of cropland in our results.

#### 4.3. Challenges and opportunities for wall-to-wall field delineation

Two key barriers to wall-to-wall field delineation persist in Sub-Saharan Africa and beyond, namely a lack of labeled data to train deep learning models, and a lack of accessible VHR imagery (Nakalembe and Kerner, 2023). Regarding labels, a compilation of existing labeled datasets in smallholder agriculture from in Rwanda (NASA Harvest, 2022), Kenya (Kerner et al., 2023), and Ghana (Estes et al., 2022), may allow for generating a standardized benchmark for testing field delineation approaches across heterogeneous geographies and image datasets. Publicly sharing existing labeled data and model weights, as done by Wang et al. (2022) and in this study, can help to alleviate the existing label bottleneck. Additionally, we propose to consider pseudo-labels for producing or complementing training data in label-free or label-scarce settings.

Critically, wall-to-wall image data at appropriate spatial resolution is needed to conduct mapping across large areas. In the case of Northern Mozambique, half of the fields in our reference data were smaller than 0.06 ha, which translates into the need for VHR data at appropriate



**Fig. 9.** Selected test sites with Google Earth RGB images (©Maxar Technologies 2019) (left column, yellow dashed line is the area zoomed in to in columns 2–6), only labels (second column), predictions (black outlines, white transparent fill) and labels (red) (columns 3–6) using the pre-trained model, pseudo-labels, human-annotated labels, and human-annotated labels combined with pseudo-labels. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

spatial resolution. Wang et al. (2022) revealed that IoU scores for fields below 0.06 ha size (affecting 11 % of the fields in their sample) often had IoU scores below 0.8 when using SPOT data, indicating that ideally, sub-meter spatial resolution data would be needed to capture these small fields. Accessing commercial VHR imagery remains a costly exercise, and local-level or sample-based studies (Lesiv et al., 2019) will remain common trade-offs as long as these cost barriers persist. Alternatively, strengthening public-private initiatives similar to the NICFI data

program (Planet Labs Inc., 2023) enabling public access to monthly PlanetScope mosaics at ~ 5 m spatial resolution across the global tropics – can aid in overcoming this issue (Estes et al., 2022; Wang et al., 2022). Alternatively, airborne imaging campaigns are available for selected countries and can be leveraged for mapping efforts (e.g. Yin et al. (2023) for Benin), but these remain rare, particularly in Sub-Saharan Africa.

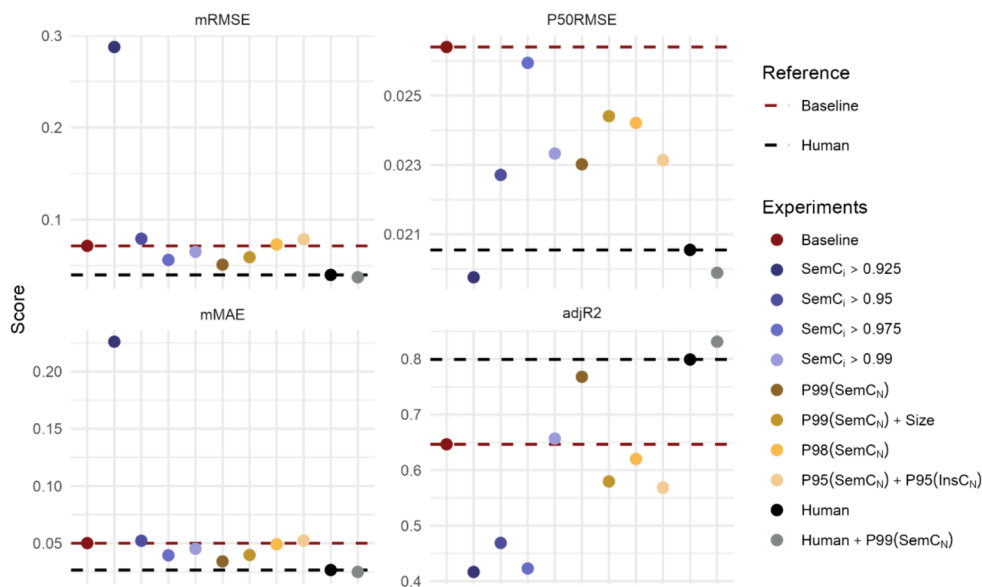


Fig. 10. Comparison of mRMSE, P50RMSE, mMAE, and  $R^2$  across experiments. Dashed lines represent baseline model (red) and model trained with human-annotated labels (black). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 4

Absolute performance scores for Baseline & Human-annotated labels, performance gain from Human-annotated labels, and relative gain across field size estimation experiments. Best results using only pseudo-labels and combined human-annotated with pseudo-labels highlighted in bold.

	Experiment	mRMSE	P <sub>50</sub> RMSE	MAE	adj. R <sup>2</sup>
Performance	Baseline	0.071	0.026	0.050	0.647
	Human-annotated	0.040	0.021	0.027	0.799
Gain from Human-annotated labels	Human-annotated – Baseline	-0.032	-0.006	-0.023	0.153
Gain relative to gain from Human-annotated labels [%]	T <sub>SemC</sub> =0.925	-685.0	113.5	-749.8	-150.5
	T <sub>SemC</sub> =0.950	-24.7	<b>63.1</b>	-8.8	-116.3
	T <sub>SemC</sub> =0.975	48.3	7.8	45.3	-146.4
	T <sub>SemC</sub> =0.990	20.3	52.5	20.8	6.6
	<b>P<sub>99</sub>(SemC<sub>N</sub>)</b>	<b>64.7</b>	57.8	<b>67.9</b>	<b>79.6</b>
	P <sub>99</sub> (SemC <sub>N</sub> ) & Size	39.0	34.1	44.1	-43.8
	P <sub>98</sub> (SemC <sub>N</sub> )	-4.9	37.4	4.2	-17.3
	P <sub>95</sub> (SemC <sub>N</sub> ) & P <sub>95</sub> (InsC <sub>N</sub> )	-22.6	55.7	-9.7	-51.1
	<b>Human-annotated + P<sub>99</sub>(SemC<sub>N</sub>)</b>	<b>108.1</b>	<b>111.3</b>	<b>107.3</b>	<b>121.0</b>

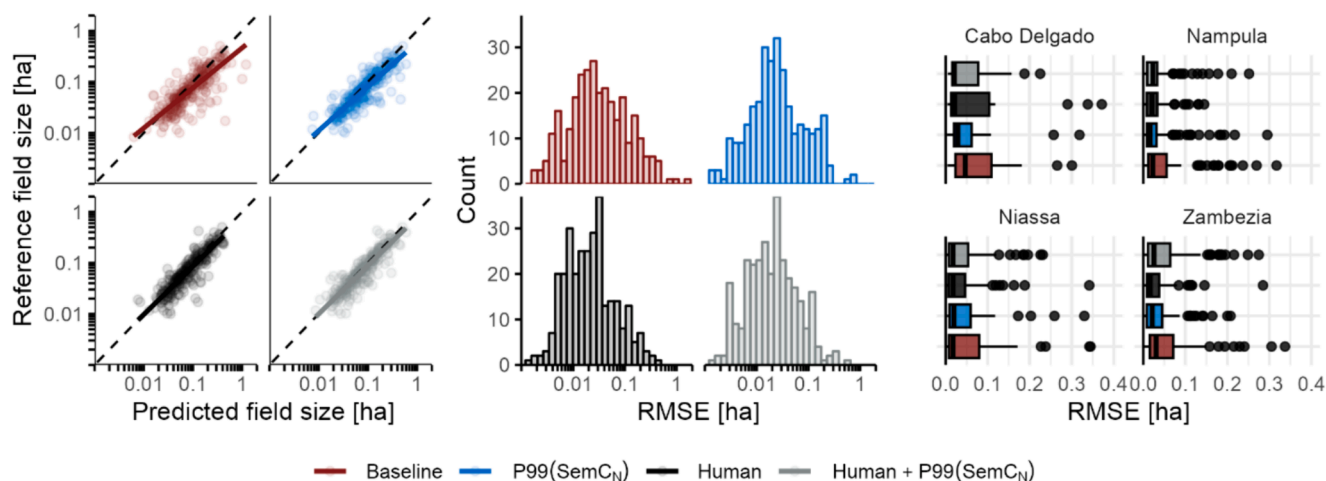


Fig. 11. Comparison of for pre-trained baseline model (red), pseudo-labels selected with  $P99(SemC_N)$  (blue), human-annotated labels (black), and combined human-annotated labels and pseudo-labels (azure). Plots show linear regression lines between reference field size and predicted field size (left), histogram of site-level RMSE scores (middle), and site-level RMSE across the four different provinces with median in black (right). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

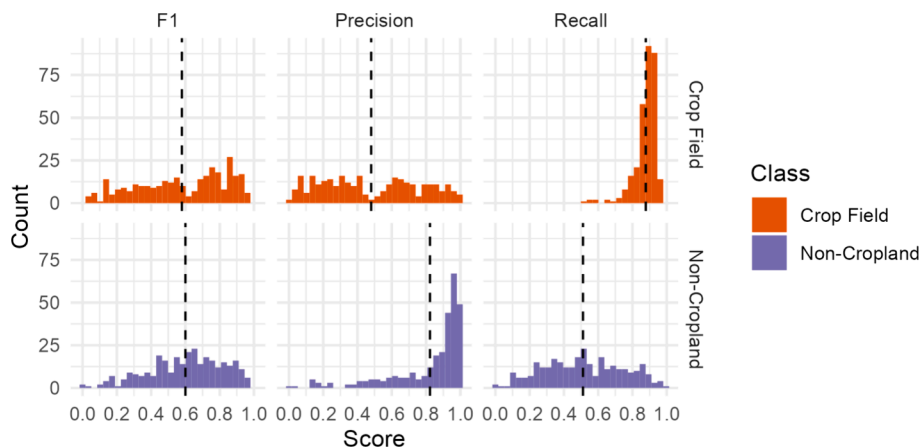


Fig. 12. Quantitative evaluation of non-cropland discrimination.



Fig. 13. Example input imagery with human-annotated labels (first column), predictions of the baseline model (second column), and pseudo-labels (third column). Boxes indicate examples of high (green), medium (orange), and low (red) performance of non-cropland identification, relating to the detection of (A) large patches of woody vegetation, (B) isolated fields in non-cropland surroundings, (C) isolated trees, (D) clusters of trees, (E) large trees with irregular canopies, and (F) short-term fallow cropland. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

### 5. Conclusion

This study explores the fully automated generation of pseudo-labels for smallholder field delineation, which can be used for fine-tuning deep convolutional neural networks across geographies. The workflow is based on publicly available model weights with good capabilities for generating predictions of field boundaries but is architecture-agnostic and independent of source data and can thus be implemented in a variety of settings. We demonstrated how the “known knows” can be leveraged to inform on the “known unknowns” and thereby assist in geographic domain adaptation for field delineation. Our findings highlight the potential of pseudo-labels for overcoming domain gaps free of reference data, or their use as a complement to limited human-

annotated data, which can further increase model performance.

While our study opens a pathway for achieving accurate field delineation in smallholder contexts with reduced label requirements, we identified several open questions and limitations that warrant further research. These include performance assessments across a larger and more diverse set of geographies, testing the production of pseudo-labels with coarser spatial resolution imagery, exploring alternative instance-level metrics, optimizing selection thresholds, and investigating the optimal quantitative balance between human-annotated labels and pseudo-labels. Further mobilizing the potential of pseudo-labels for closing geographic domain gaps will be a stepping stone for large-area field delineation in smallholder contexts without the need for costly label collection.

## CRedit authorship contribution statement

**Philippe Rufin:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Sherrie Wang:** Writing – review & editing, Software, Methodology. **Sá Nogueira Lisboa:** Writing – review & editing. **Jan Hemmerling:** Writing – review & editing. **Mirela G. Tulbure:** Writing – review & editing. **Patrick Meyfroidt:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Data availability

Links for data and code will be made available after acceptance of the manuscript.

## Acknowledgements

This work was supported by the F.R.S.-FNRS, grant no. T.0154.21 and grant no. 1.B422.24 and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant agreement No 677140 MIDLAND). This research contributes to the Global Land Programme (<http://glp.earth>). The authors would like to thank the labeling team at Humans in the Loop for organizing and performing label collection.

## References

- Baumert, S., Fisher, J., Ryan, C., Woollen, E., Vollmer, F., Artur, L., Zorrilla-Miras, P., Mahamane, M., 2019. Forgone opportunities of large-scale agricultural investment: A comparison of three models of soya production in Central Mozambique. *World Development Perspectives* 16. <https://doi.org/10.1016/j.wdp.2019.100145>.
- Bey, A., Jetimane, J., Lisboa, S.N., Ribeiro, N., Sitoe, A., Meyfroidt, P., 2020. Mapping smallholder and large-scale cropland dynamics with a flexible classification system and pixel-based composites in an emerging frontier of Mozambique. *Remote Sens. Environ.* 239, 111611 <https://doi.org/10.1016/j.rse.2019.111611>.
- Bradski, G., 2000. The openCV library. *Dr. Dobb's Journal: Software Tools for the Professional Programmer* 25, 120–123.
- Burke, W.J., Jayne, T.S., Sitko, N.J., 2020. Do Medium-scale Farms Improve Market Access Conditions for Zambian Smallholders? *J. Agric. Econ.* 71, 517–533. <https://doi.org/10.1111/1477-9552.12360>.
- Burke, M., Lobell, D.B., 2017. Satellite-based assessment of yield variation and its determinants in smallholder African systems. *Proc. Natl. Acad. Sci. U.S.A.* 114, 2189–2194. <https://doi.org/10.1073/pnas.1616919114>.
- Cai, Z., Hu, Q., Zhang, X., Yang, J., Wei, H., Wang, J., Zeng, Y., Yin, G., Li, W., You, L., Xu, B., Shi, Z., 2023. Improving agricultural field parcel delineation with a dual branch spatiotemporal fusion network by integrating multimodal satellite data. *ISPRS J. Photogramm. Remote Sens.* 205, 34–49. <https://doi.org/10.1016/j.isprsjprs.2023.09.021>.
- Chen, J., Sun, B., Wang, L., Fang, B., Chang, Y., Li, Y., Zhang, J., Lyu, X., Chen, G., 2022. Semi-supervised semantic segmentation framework with pseudo supervisions for land-use/land-cover mapping in coastal areas. *Int. J. Appl. Earth Obs. Geoinf.* 112, 102881 <https://doi.org/10.1016/j.jag.2022.102881>.
- Chiarella, C., Meyfroidt, P., Abeygunawardane, D., Conforti, P., 2023. Balancing the trade-offs between land productivity, labor productivity and labor intensity. *Ambio.* <https://doi.org/10.1007/s13280-023-01887-4>.
- Estes, L.D., Ye, S., Song, L., Luo, B., Eastman, J.R., Meng, Z., Zhang, Q., McRitchie, D., Debats, S.R., Muhando, J., Amukoa, A.H., Kaloo, B.W., Makuru, J., Mbatia, B.K., Muasa, I.M., Mucha, J., Mugami, A.M., Mugami, J.M., Muinde, F.W., Mwawaza, F.M., Ochieng, J., Oduol, C.J., Oduor, P., Wanjiku, T., Wanyoike, J.G., Avery, R.B., Caylor, K.K., 2022. High Resolution, Annual Maps of Field Boundaries for Smallholder-Dominated Croplands at National Scales. *Front. Artif. Intell.* 4, 744863 <https://doi.org/10.3389/frai.2021.744863>.
- Fao, 2021. Global Agro-Ecological Zones (GAEZ v4) - Data Portal user's guide. FAO and IIASA. <https://doi.org/10.4060/cb5167en>.
- Fao, 2023. How much do large-scale and small-scale farming contribute to global deforestation? FAO. <https://doi.org/10.4060/cc5723en>.
- GDAL/OGR contributors, 2023. GDAL/OGR Geospatial Data Abstraction Software Library. Open Source Geospatial Foundation.

- NASA Harvest, Radiant Earth Foundation, TaQadam, 2022. Rwanda Field Boundary Competition Dataset (Version 1.0). Radiant MLHub.
- He, R., Yang, J., Qi, X., 2021. Re-distributing Biased Pseudo Labels for Semi-supervised Semantic Segmentation: A Baseline Investigation. in: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). Presented at the 2021 IEEE/CVF International Conference on Computer Vision (ICCV), IEEE, Montreal, QC, Canada, pp. 6910–6920. <https://doi.org/10.1109/ICCV48922.2021.00685>.
- Planet Labs Inc., 2023. NICFI Data Program [WWW Document]. URL <https://www.planet.com/nicfi/>.
- Jayne, T.S., Chamberlin, J., Traub, L., Sitko, N., Muyanga, M., Yeboah, F.K., Anseuw, W., Chapoto, A., Wineman, A., Nkonde, C., Kachule, R., 2016. Africa's changing farm size distribution patterns: the rise of medium-scale farms. *Agricultural Economics (United Kingdom)* 47, 197–214. <https://doi.org/10.1111/agec.12308>.
- Jayne, T.S., Wineman, A., Chamberlin, J., Muyanga, M., Yeboah, F.K., 2022. Changing Farm Size Distributions and Agricultural Transformation in Sub-Saharan Africa. *Ann. Rev. Resour. Econ.* 14, 109–130. <https://doi.org/10.1146/annurev-resource-111220-025657>.
- Kerner, H., Sundar, S., Satish, M., 2023. Multi-Region Transfer Learning for Segmentation of Crop Field Boundaries in Satellite Images with Limited Labels.
- Lambert, M.-J., Traoré, P.C.S., Blaes, X., Baret, P., Defourny, P., 2018. Estimating smallholder crops production at village level from Sentinel-2 time series in Mali's cotton belt. *Remote Sens. Environ.* 216, 647–657. <https://doi.org/10.1016/j.rse.2018.06.036>.
- Lee, D.-H., 2013. Pseudo-Label: The Simple and Efficient Semi-Supervised Learning Method for Deep Neural Networks.
- Lesiv, M., Laso Bayas, J.C., See, L., Duerauer, M., Dahlia, D., Durando, N., Hazarika, R., Kumar Sahariah, P., Vakolyuk, M., Blyshchik, V., Bilous, A., Perez-Hoyos, A., Gengler, S., Prestele, R., Bilous, S., Akhtar, I. ul H., Singha, K., Choudhury, S.B., Chetri, T., Malek, Z., Bungnamei, K., Saikia, A., Sahariah, D., Narzary, W., Danylo, O., Sturn, T., Karner, M., McCallum, I., Schepaschenko, D., Moltchanova, E., Fraisl, D., Moorthy, I., Fritz, S., 2019. Estimating the global distribution of field size using crowdsourcing. *Glob. Change Biol.* 25, 174–186. <https://doi.org/10.1111/gcb.14492>.
- Li, M., Long, J., Stein, A., Wang, X., 2023. Using a semantic edge-aware multi-task neural network to delineate agricultural parcels from remote sensing images. *ISPRS J. Photogramm. Remote Sens.* 200, 24–40. <https://doi.org/10.1016/j.isprsjprs.2023.04.019>.
- Liu, W., Liu, J., Luo, Z., Zhang, H., Gao, K., Li, J., 2022. Weakly supervised high spatial resolution land cover mapping based on self-training with weighted pseudo-labels. *Int. J. Appl. Earth Obs. Geoinf.* 112, 102931 <https://doi.org/10.1016/j.jag.2022.102931>.
- Liu, Q., Peng, J., Zhang, G., Sun, W., Du, Q., 2023. Deep Contrastive Learning Network for Small-Sample Hyperspectral Image Classification. *J. Remote Sens.* 3, 0025. <https://doi.org/10.34133/remotesensing.0025>.
- Liverpool-Tasie, L.S.O., Nuhu, A.S., Awokuse, T., Jayne, T., Muyanga, M., Aromolaran, A., Adelaja, A., 2020. Spillover effects of medium-scale farms on smallholder behaviour and welfare: evidence from Nigeria.
- Long, J., Li, M., Wang, X., Stein, A., 2022. Delineation of agricultural fields using multi-task BsiNet from high-resolution satellite images. *Int. J. Appl. Earth Obs. Geoinf.* 112, 102871 <https://doi.org/10.1016/j.jag.2022.102871>.
- Mei, K., Zhu, C., Zou, J., Zhang, S., 2020. Instance Adaptive Self-training for Unsupervised Domain Adaptation, in: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (Eds.), *Computer Vision – ECCV 2020, Lecture Notes in Computer Science*. Springer International Publishing, Cham, pp. 415–430. [https://doi.org/10.1007/978-3-030-58574-7\\_25](https://doi.org/10.1007/978-3-030-58574-7_25).
- Mei, W., Wang, H., Fouhey, D., Zhou, W., Hinks, I., Gray, J.M., Van Berkel, D., Jain, M., 2022. Using Deep Learning and Very-High-Resolution Imagery to Map Smallholder Field Boundaries. *Remote Sens. (Basel)* 14, 3046. <https://doi.org/10.3390/rs14133046>.
- Meyfroidt, P., 2017. Mapping farm size globally: benchmarking the smallholders debate. *Environ. Res. Lett.* 12, 031002 <https://doi.org/10.1088/1748-9326/aa5ef6>.
- Ministério da Agricultura e Desenvolvimento Rural, 2017. Inquérito Agrícola Integrado. Instituto Nacional de Estatística.
- Nakalembe, C., Kerner, H., 2023. Considerations for AI-EO for agriculture in Sub-Saharan Africa. *Environ. Res. Lett.* 18, 041002 <https://doi.org/10.1088/1748-9326/acc476>.
- Okujeni, A., Canters, F., Cooper, S.D., Degerickx, J., Heiden, U., Hostert, P., Priem, F., Roberts, D.A., Somers, B., Van Der Linden, S., 2018. Generalizing machine learning regression models using multi-site spectral libraries for mapping vegetation-impervious-soil fractions across multiple cities. *Remote Sens. Environ.* 216, 482–496. <https://doi.org/10.1016/j.rse.2018.07.011>.
- Persello, C., Tolpekin, V.A., Bergado, J.R., de By, R.A., 2019. Delineation of agricultural fields in smallholder farms from satellite images using fully convolutional networks and combinatorial grouping. *Remote Sens. Environ.* 231, 111253 <https://doi.org/10.1016/j.rse.2019.111253>.
- Prabhu, V., Khare, S., Kartik, D., Hoffman, J., 2022. AUGCO: Augmentation Consistency-guided Self-training for Source-free Domain Adaptive Semantic Segmentation.
- Rezatofighi, H., Tsai, N., Gwak, J., Sadeghian, A., Reid, I., Savares, S., 2019. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). Presented at the 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE, Long Beach, CA, USA, pp. 658–666. <https://doi.org/10.1109/CVPR.2019.00075>.
- Ricciardi, V., Ramankutty, N., Mehrabi, Z., Jarvis, L., Chookolingo, B., 2018. How much of the world's food do smallholders produce? *Glob. Food Sec.* 17, 64–72. <https://doi.org/10.1016/j.gfs.2018.05.002>.

- Ricciardi, V., Mehrabi, Z., Wittman, H., James, D., Ramankutty, N., 2021. Higher yields and more biodiversity on smaller farms. *Nat Sustain* 4, 651–657. <https://doi.org/10.1038/s41893-021-00699-2>.
- Rufin, P., Bey, A., Picoli, M., Meyfroidt, P., 2022. Large-area mapping of active cropland and short-term fallows in smallholder landscapes using PlanetScope data. *Int. J. Appl. Earth Obs. Geoinf.* 112, 102937 <https://doi.org/10.1016/j.jag.2022.102937>.
- Safonova, A., Ghazaryan, G., Stiller, S., Main-Knorn, M., Nendel, C., Ryo, M., 2023. Ten deep learning techniques to address small data problems with remote sensing. *Int. J. Appl. Earth Obs. Geoinf.* 125, 103569 <https://doi.org/10.1016/j.jag.2023.103569>.
- Simoes, R., Sanchez, A., Picoli, M.C.A., 2022. segmetric: Metrics for Assessing Segmentation Accuracy for Geospatial Data.
- Tetteh, G.O., Schwieder, M., Erasmi, S., Conrad, C., Gocht, A., 2023. Comparison of an Optimised Multiresolution Segmentation Approach with Deep Neural Networks for Delineating Agricultural Fields from Sentinel-2 Images. *PFG*. <https://doi.org/10.1007/s41064-023-00247-x>.
- Toldo, M., Maracani, A., Michieli, U., Zanuttigh, P., 2020. Unsupervised Domain Adaptation in Semantic Segmentation: A Review. *Technologies* 8, 35. <https://doi.org/10.3390/technologies8020035>.
- Waldner, F., Diakogiannis, F.I., 2020. Deep learning on edge: Extracting field boundaries from satellite images with a convolutional neural network. *Remote Sens. Environ.* 245, 111741 <https://doi.org/10.1016/j.rse.2020.111741>.
- Waldner, F., Diakogiannis, F.I., Batchelor, K., Ciccotosto-Camp, M., Cooper-Williams, E., Herrmann, C., Mata, G., Toovey, A., 2021. Detect, Consolidate, Delineate: Scalable Mapping of Field Boundaries Using Satellite Images. *Remote Sens. (basel)* 13, 2197. <https://doi.org/10.3390/rs13112197>.
- Wang, D., Shelhamer, E., Liu, S., Olshausen, B., Darrell, T., 2020. Tent: Fully Test-time Adaptation by Entropy Minimization. <https://doi.org/10.48550/ARXIV.2006.10726>.
- Wang, S., Waldner, F., Lobell, D.B., 2022. Unlocking large-scale crop field delineation in smallholder farming systems with transfer learning and weak supervision. *Remote Sens. (basel)* 14, 5738. <https://doi.org/10.3390/rs14225738>.
- Whitcraft, A.K., Becker-Reshef, I., Justice, C.O., Gifford, L., Kavvada, A., Jarvis, I., 2019. No pixel left behind: Toward integrating Earth Observations for agriculture into the United Nations Sustainable Development Goals framework. *Remote Sens. Environ.* 235, 111470 <https://doi.org/10.1016/j.rse.2019.111470>.
- Wu, H., Prasad, S., 2018. Semi-Supervised Deep Learning Using Pseudo Labels for Hyperspectral Image Classification. *IEEE Trans. on Image Process.* 27, 1259–1270. <https://doi.org/10.1109/TIP.2017.2772836>.
- Yin, L., Ghosh, R., Lin, C., Hale, D., Weigl, C., Obarowski, J., Zhou, J., Till, J., Jia, X., You, N., Mao, T., Kumar, V., Jin, Z., 2023. Mapping smallholder cashew plantations to inform sustainable tree crop expansion in Benin. *Remote Sens. Environ.* 295, 113695 <https://doi.org/10.1016/j.rse.2023.113695>.
- Zhang, X., Feng, X., Xiao, P., He, G., Zhu, L., 2015. Segmentation quality evaluation using region-based precision and recall measures for remote sensing images. *ISPRS J. Photogramm. Remote Sens.* 102, 73–84. <https://doi.org/10.1016/j.isprsjprs.2015.01.009>.
- Zhu, Y., Pan, Y., Hu, T., Zhang, D., Zhao, C., Gao, Y., 2024. A generalized framework for agricultural field delineation from high-resolution satellite imageries. *Int. J. Digital Earth* 17, 2297947. <https://doi.org/10.1080/17538947.2023.2297947>.
- Zou, Y., Yu, Z., Kumar, B.V.K.V., Wang, J., 2018. Domain Adaptation for Semantic Segmentation via Class-Balanced Self-Training.
- Zou, Y., Zhang, Z., Zhang, H., Li, C.-L., Bian, X., Huang, J.-B., Pfister, T., 2021. PseudoSeg: Designing Pseudo Labels for Semantic Segmentation.