

# Fields of The World: A Machine Learning Benchmark Dataset for Global Agricultural Field Boundary Segmentation

Hannah Kerner<sup>1\*</sup>, Snehal Chaudhari<sup>1\*</sup>, Aninda Ghosh<sup>1\*</sup>, Caleb Robinson<sup>2\*</sup>,  
Adeel Ahmad<sup>3,4</sup>, Eddie Choi<sup>4</sup>, Nathan Jacobs<sup>4</sup>, Chris Holmes<sup>5</sup>, Matthias Mohr<sup>5</sup>,  
Rahul Dodhia<sup>2</sup>, Juan M. Lavista Ferres<sup>2</sup>, Jennifer Marcus<sup>5</sup>

<sup>1</sup>Arizona State University, Tempe, AZ 85281 USA, hkerner@asu.edu

<sup>2</sup>Microsoft AI for Good Research Lab, Redmond, WA 98052 USA

<sup>3</sup>Taylor Geospatial Institute, St Louis, MO 63108 USA

<sup>4</sup>Washington University in St Louis, St Louis, MO 63130 USA

<sup>5</sup>Taylor Geospatial Engine, St Louis, MO 63130 USA

## Abstract

Crop field boundaries are foundational datasets for agricultural monitoring and assessments but are expensive to collect manually. Machine learning (ML) methods for automatically extracting field boundaries from remotely sensed images could help realize the demand for these datasets at a global scale. However, current ML methods for field instance segmentation lack sufficient geographic coverage, accuracy, and generalization capabilities. Further, research on improving ML methods is restricted by the lack of labeled datasets representing the diversity of global agricultural fields. We present Fields of The World (FTW)—a novel ML benchmark dataset for agricultural field instance segmentation spanning 24 countries on four continents (Europe, Africa, Asia, and South America). FTW is an order of magnitude larger than previous datasets with 70 462 samples, each containing instance and semantic segmentation masks paired with multi-date, multi-spectral Sentinel-2 satellite images. We provide results from baseline models for the new FTW benchmark, show that models trained on FTW have better zero-shot and fine-tuning performance in held-out countries than models that aren't pre-trained with diverse datasets, and show positive qualitative zero-shot results of FTW models in a real-world scenario – running on Sentinel-2 scenes over Ethiopia.

**Code** — <https://github.com/fieldsoftheworld/ftw-baselines>

**Datasets** — <https://beta.source.coop/repositories/kerne-lab/fields-of-the-world/>

## 1 Introduction

Crop field boundary datasets are urgently needed in agricultural monitoring, sustainable agriculture, and development applications (Nakalembe and Kerner 2023). However, these datasets do not exist for most of the world. Automatic field delineation in globally available satellite imagery offers a promising solution, but semantic reasoning about globally diverse agricultural landscapes in satellite imagery remains

\*These authors contributed equally.

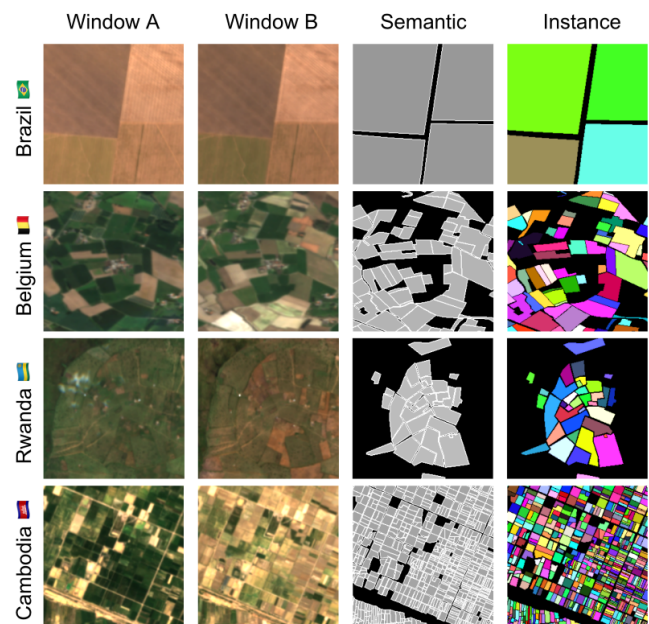


Figure 1: Training samples from four continents, demonstrating the diversity within Fields of The World.

challenging. Field morphologies, agricultural practices, and climate patterns vary greatly across the world. For example, average field sizes range from 1.6 hectares (ha) in Sub-Saharan Africa to 121 ha in North America (Debats et al. 2016). Motivated by this global diversity, we present a novel dataset for agricultural field instance segmentation: Fields of The World (FTW). FTW spans diverse landscapes in 24 countries on 4 continents (Europe, Asia, Africa, and South America). Fields of The World aims to catalyze field instance segmentation research and enable consistent, granular evaluation of different modeling approaches.

Field boundary datasets enable field-scale monitoring of crop conditions, yield, pest/diseases, farming practices, resource utilization, and other agricultural characteris-

tics (Nakalembe and Kerner 2023). They are also in high demand for conservation and climate change policies and programs that require Measurement, Reporting, and Verification (MRV) of greenhouse gas emissions, carbon sequestration, and sustainable land management practices, such as the European Union Deforestation Regulation (European Parliament and Council of the European Union 2023). Field boundaries also simplify challenging tasks like crop-type classification by enabling classification at the object level rather than over individual pixels (Garnot, Landrieu, and Chehata 2022). Statistics agencies use field boundaries for ground-based survey design (Nakalembe and Kerner 2023). Field boundary maps over multiple years enable analyses of environmental and socioeconomic land change dynamics such as aggregation or fragmentation of farm parcels over time (Estes et al. 2022; Sullivan et al. 2023).

Previous work has demonstrated good performance for field instance segmentation in European countries, enabled by a combination of novel algorithms and labeled datasets (e.g., Wang, Waldner, and Lobell (2022); Sainte Fare Garnot and Landrieu (2021)). Research progress has been driven by benchmark datasets such as PASTIS (Sainte Fare Garnot and Landrieu 2021; Garnot, Landrieu, and Chehata 2022), AI4Boundaries (d’Andrimont et al. 2023), and AI4FoodSecurity (Planet, TUM, DLR and Radiant Earth 2021). While these datasets have been critical research catalysts, they do not fully capture the diversity and complexity of global agricultural landscapes. Existing datasets have limited geographic diversity, with labels concentrated in a handful of (usually European) countries (Table 2).

Fields of The World captures greater geographic diversity, morphological diversity, and agro-climatic diversity than any previous dataset. It includes fields of various sizes (Figure 3), shape, and orientation (Figure 2). FTW is also an order of magnitude larger than previous datasets, with 70 462 samples covering a total geographic area of 166 293 km<sup>2</sup>.

Fields of The World provides harmonized ML-ready inputs from the optical Sentinel-2 satellite. Each example in Fields of The World includes four spectral channels (red, green, blue, and near-infrared) from two contrasting dates. Labels include instance and semantic segmentation masks. We provide the polygon label annotations in a standardized format using the *fiboa* (field boundaries for agriculture) specification (fiboa contributors 2024). This makes previously siloed datasets interoperable and enables users to obtain custom satellite data or other inputs corresponding to geo-referenced labels. The provided metadata also allows users to easily subset the dataset depending on their needs (e.g., commercial license or specific location).

We include training, validation, and test sets for each country, using existing splits when possible to maximize compatibility with existing work. We propose benchmark tasks that mimic real-world scenarios relevant to downstream users of field boundary datasets (e.g., region-specific evaluation, transfer learning, and zero-shot generalization). Finally, we perform experiments to demonstrate the value of the FTW dataset and provide baseline results for benchmark tasks. We release code via Github, data via Source Cooperative, and data loaders and pre-trained models via TorchGeo.

## 2 Dataset Description

### Annotations

**Field Boundary Representations** Field boundary annotations are typically in the form of geo-referenced polygons. Since field boundaries at the same location may change across growing seasons, these polygons should also be temporally referenced to specify when the boundary is valid. Field polygons may be farmer-reported, manually drawn on high-resolution satellite images with GIS software, or recorded by walking the field perimeter with a handheld location-recording device. Polygons can then be paired with satellite data from the same location and time.

We conducted a comprehensive search for field polygons from government databases, published literature, and other websites. We looked for datasets with diverse geographic coverage, high-quality and trustworthy polygon annotations, and licenses that permit reuse. We included all datasets meeting these criteria in FTW. We considered author-reported quality assessment in each dataset’s documentation, previous use of the dataset in ML analyses, and visual inspection (e.g., closed polygons, polygons consistent with satellite images from reported dates, etc). Table 1 lists the 24 source datasets selected for Fields of The World.

**Presence/absence Labels** *Presence/absence labels* are binary labels providing information about both the occurrence and non-occurrence of a phenomenon across sampled locations or time periods—for example, the presence of a field boundary or its absence. Most of the datasets in Fields of The World have presence/absence labels. However, some have *presence-only* labels, i.e., they are partially labeled. These indicate the presence of some, but not necessarily all, field boundaries in the sampled locations. Some pixels in presence-only label masks have unknown labels that might be labeled as background. Partial labels are a common challenge in field boundary segmentation (Wang, Waldner, and Lobell 2022). The Rwanda example in Figure 1 (row 3) illustrates presence-only labels while the Cambodia example (row 4) illustrates presence/absence labels.

**Semantic Filtering** We focused Fields of The World on field boundaries for annual crops. Annual crops, also called temporary crops, are planted, grown, and harvested within a single growing season or year. Common examples include wheat, rice, maize, soybeans, and barley. This does not include permanent or perennial crops, which are cultivated for longer than one year and are not replanted annually, such as fruit trees, nut trees, and some grasses. We also excluded parcels used for pasture, fallow land, or other non-crop agricultural activities, such as grazing, orchards, vineyards, and forestry. If a dataset included parcels that were not active annual crops, we filtered them out (details in supplement).

**Sample Grids** Many datasets in Table 1, particularly those sourced from the European Union government websites and EuroCrops (Schneider and Körner 2022), have millions of dense annotations spanning the entire country. Including all of these annotations in Fields of The World would bias the dataset toward large European countries. We sub-sampled these datasets by: 1) creating a bounding box enclosing the

Country													Presence/absence labels		
	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Subsampled	# train, val, test	Source
Austria				█	█	█	█						✓	5303, 637, 745	Schneider and Körner (2022)
Belgium		█	█	█	█	█	█	█	█	█			✓	1554, 189, 198	Schneider and Körner (2022)
Cambodia			█	█	█	█	█	█	█	█	█	█	✗	245, 27, 25	Persello et al. (2023)
Croatia				█	█	█	█	█	█	█			✓	2778, 351, 353	Agency for Payments in Agriculture, Fisheries & Rural Development (2024)
Denmark	█	█	█	█	█	█	█	█	█	█			✓	2868, 360, 332	Ministry of Food, Agriculture and Fisheries of Denmark (2021)
Estonia				█	█	█	█	█	█	█			✓	5348, 681, 684	Schneider and Körner (2022)
Finland					█	█	█	█	█	█			✓	4527, 550, 588	Finnish Food Authority (2021)
Corsica				█	█	█	█	█	█	█			✓	1974, 240, 258	The Service and Payment Agency (ASP) (2024)
France				█	█	█	█	█	█	█			✓	2773, 339, 396	The Service and Payment Agency (ASP) (2024)
Germany				█	█	█	█	█	█	█			✗	306, 30, 350	Kondmann et al. (2021)
Latvia					█	█	█	█	█	█			✓	5529, 668, 741	Schneider and Körner (2022)
Lithuania				█	█	█	█	█	█	█			✓	4208, 522, 528	Schneider and Körner (2022)
Luxembourg						█	█	█	█	█	█		✓	643, 81, 84	Administration of Technical Agricultural Services (2024)
Netherlands			█	█	█	█	█	█	█	█	█	█	✓	3110, 381, 388	Netherlands National GeoRegistry (2021)
Portugal			█	█	█	█	█	█	█	█	█	█	✓	47, 9, 10	Instituto de Financiamento da Agricultura e Pescas (2024)
Slovakia				█	█	█	█	█	█	█	█	█	✓	3275, 390, 408	Slovak Republic Government (2024)
Slovenia					█	█	█	█	█	█	█	█	✓	1733, 216, 228	Schneider and Körner (2022)
South Africa													✗	590, 72, 85	Planet et al. (2021)
Spain				█	█	█	█	█	█	█	█	█	✓	2015, 201, 216	Schneider and Körner (2022)
Sweden				█	█	█	█	█	█	█	█	█	✓	3802, 442, 516	Swedish Board of Agriculture (2024)
Vietnam													✗	228, 36, 23	Persello et al. (2023)

Country													Presence-only labels		
	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Subsampled	# train, val, test	Source
Brazil	█	█											✗	1289, 130, 188	Oldoni et al. (2020)
India				█	█	█	█	█	█	█	█	█	✗	1261, 300, 399	Wang, Waldner, and Lobell (2023)
Kenya					█	█	█	█	█	█	█	█	✗	316, 20, 55	Pula Advisors (2022)
Rwanda													✗	57, 6, 7	NASA Harvest and Radiant Earth Foundation (2024)

Table 1: Key dataset details for each country in FTW, with green and brown cells indicating Windows A and B respectively.test

entire dataset, 2) splitting the bounding box into a grid where each grid cell covered an area between 3300 to 5000 km<sup>2</sup>, and 3) selecting 2-4 grid cells per country that captured a mixture of high-density and low-density agricultural areas.

For the eight datasets in Table 1 that were not sub-sampled, we used all field boundaries provided by the source dataset. Four of these datasets were published with predefined grids, which we used without modification: Germany (Kondmann et al. 2021), South Africa (Planet et al. 2021), Cambodia (Persello et al. 2023), and Vietnam (Persello et al. 2023). The Kenya (Pula Advisors 2022) and Brazil (Oldoni et al. 2020) label polygons were highly clustered but did not have predefined grids or clusters. We used *k*-means clustering to cluster the label polygons (using the center latitude/longitude of each polygon as the features). We defined a rectangular grid spanning the bounds of each cluster. We chose *k* by visually inspecting each dataset and balancing between over-clustering (resulting in high overlap between cluster grids) or under-clustering (resulting in sparse grids with large unlabeled areas). In the India (Wang, Waldner, and Lobell 2023) and Rwanda (NASA

Harvest and Radiant Earth Foundation 2024) datasets, polygon labels were in small clusters (e.g., 5 fields per cluster for India). We did not define grids for these datasets because we created sample chips directly from each small cluster.

**Sample Chip ROIs** We tiled each sample grid into 1536m × 1536m sample patch ROIs (regions of interest), which we call ‘chips’. For India and Rwanda, we created a 1536m × 1536m chip around the center of each label cluster.

**Metadata Standardization** We converted the label polygon datasets to the fiboa specification (fiboa contributors 2024). If a dataset was sub-sampled from a larger dataset, we only converted the sub-sampled version. Per the fiboa core specification, each dataset has per-polygon attributes *id* (unique field identifier), *determination\_datetime* (last timestamp at which the field boundary existed/was observed), *area* (field area in hectares), and *geometry* (field polygon geometry; we use the WGS84/EPSG:4326 coordinate reference system). We included crop type or other attributes when available. Converted GeoParquet files are available on Source Cooperative at <https://beta.source.coop/>




Figure 2: Field orientation histograms for selected countries.

repositories/kerner-lab/fields-of-the-world/. The README for each dataset provides details including the source dataset license (extends to the FTW subset using it) and link. We provide GeoParquet files for the sample grids and chip ROIs.

**Label Masks** Previous work explored several approaches to convert (“rasterize”) field label polygons to label masks, which are the ML prediction targets. Approaches include binary field extent masks (field interiors vs background), binary field boundary masks (field boundaries vs background), 3-class masks (field interiors, field boundaries, and background) (Taravat et al. 2021), and distance masks (to field centroids) (d’Andrimont et al. 2023). We provide binary field extent and 3-class semantic masks and instance masks.

### Satellite Data

We obtained multispectral Sentinel-2 satellite images using Microsoft Planetary Computer (Microsoft Open Source et al. 2022). Images in this catalog are processed to Level 2A (bottom-of-atmosphere) and stored in cloud-optimized GeoTIFF (COG) format. We used the red (B04), green (B03), blue (B02), and near-infrared (B08) spectral bands, all of which have spatial resolution of 10 m per pixel. We used Sentinel-2 because it is the highest-resolution optical satellite dataset that is freely accessible. In the FTW Github repository, we provide a CSV file containing the Sentinel-2 scene ID, cloud percentage, and date ranges for each sample.

**Dates** Previous work showed that contrasting images from different times during the same year improved crop field segmentation by highlighting the intra-annual variation characteristic of active crop fields (Estes et al. 2022; Debats et al. 2016). This contrast can help models rule out potential false positives such as fallow fields or forest stands.

For each country, we collected images from two date ranges (Table 1). Growing seasons vary greatly globally (and even within countries). To choose the date ranges, we looked at each country’s crop calendar which specifies the planting, mid-season, harvesting, and off-season months for the main crops (USDA Foreign Agricultural Service 2024). We inspected the available satellite images in the sample area(s) in two date ranges spanning the planting/mid-season and harvesting/off-season months. If a country had multiple growing seasons (e.g., winter and summer crops), we visualized both seasons and chose the one appearing most active. We then iteratively adjusted the date ranges to account for good contrast between images and cloud cover. After adjustment, the date ranges do not necessarily match the growing season stages, so we call them Window A and B.




Figure 3: Field area distribution across four countries.

**Cloud Filtering** For each chip ROI, we searched for Sentinel-2 scenes with  $< 90\%$  scene level cloud cover in the two date ranges. For the resulting scenes, we cropped each scene to the sample chip ROI and computed the cloud percentage in the patch using the Sentinel-2 scene classification layer (SCL) “Cloud medium probability” and “Cloud high probability” classes. We selected the chip with the lowest cloud percentage. If there were no chips with cloud percentage  $< 10\%$ , we discarded the chip. Table 1 gives the resulting number of Sentinel-2 chips created for each country.

**Resizing and Normalization** We resized each chip to  $256 \times 256$  pixels. Each chip is stored as a GeoTIFF with EPSG:4326. We normalize images during training by dividing each channel by 3000 (an approximate mean value).

### Dataset Splits

FTW defines training, validation, and test sets for each country to facilitate evaluation of test metrics at the country scale. Many sample chips are spatially adjacent since they were tiled from large grids. Spatial autocorrelation between adjacent chips may cause leakage between data subsets if chips are randomly split into subsets (Rolf 2023). To reduce the impact of spatial autocorrelation, we implemented a blocked random splitting strategy. We grouped chips into  $3 \times 3$  blocks and randomly assigned 80% to training, 10% validation, and 10% test. To ensure  $3 \times 3$  blocks were large enough, we performed an experiment to quantify the sensitivity of test performance to the block size and distance from each test patch to the nearest training patch (see supplement).

## 3 Dataset Analysis and Related Work

The dramatic differences in field morphology across the globe motivated the construction of FTW (Figure 1). FTW has significant advantages over previous field instance segmentation datasets in terms of (i) geographic representation and extent, (ii) annotation volume, and (iii) annotation/scene complexity. The most relevant datasets for comparison are AI4Boundaries (d’Andrimont et al. 2023), AI4SmallFarms (Persello et al. 2023), PASTIS (Sainte Fare Garnot and Landrieu 2021), and PASTIS-R (Garnot, Landrieu, and Chehata 2022). We only include datasets that explicitly label individual field instances, excluding semantic segmentation labels since field instances enable a broader range of applications. We also exclude datasets providing field polygons but no imagery, such as the field boundary dataset created by the French Land Parcel Identification

Dataset	# countries	Sensor(s)	Input dim ( $H, W, C, T$ )	# field polygons (million)	# total samples	Sample dim. ( $m \times m$ )	Total area ( $km^2$ )
Fields of The World (FTW)	24	Sentinel-2	[256, 256, 4, 2]	1.63	70,484	1,536	166,293
AI4Boundaries (d’Andrimont et al. 2023)	7	Sentinel-2, Aerial	[256, 256, 4, 12] [512, 512, 3, 1]	1.07 <sup>3</sup>	7,831 7,598	2,560 512	51,321 1,992
AI4SmallFarms (Persello et al. 2023)	2	Sentinel-2	[256, 256, 4, 1]	0.44	62	5,000	1,550
PASTIS (Sainte Fare Garnot and Landrieu 2021)	1	Sentinel-2	[128, 128, 10, $N$ ] <sup>1</sup>	0.12	2,433	1,280	3,986
PASTIS-R (Garnot, Landrieu, and Chehata 2022)	1	Sentinel-2, Sentinel-1	[128, 128, 10, $N$ ] <sup>2</sup> [128, 128, 3, 70]	0.12	2,433	1,280	3,986

<sup>1</sup> Varying observations (38-61) taken between September 2018 and November 2019.

<sup>2</sup> All available observations for the 2019 season.

<sup>3</sup> d’Andrimont et al. (2023) reports 2.5M parcels contained in 7,831 4-km samples, however the number of polygons included in dataset sample masks is smaller.

Table 2: Key attributes of Fields of The World and previous field boundary segmentation datasets.

System (The Service and Payment Agency (ASP) 2024). The lack of standardized satellite imagery for polygon-only datasets hampers the creation and comparison of automated algorithms for field boundary delineation and related tasks.

Table 2 compares the key attributes for all datasets (see geographic distributions in the supplement). FTW has a significantly broader geographic distribution than all previous datasets, including fields from 24 countries spanning four continents (Europe, Asia, Africa, and South America). Previous datasets include at most 7 countries and are mostly concentrated in Europe, except AI4SmallFarms which only has images from Cambodia and Vietnam. FTW is the largest dataset in terms of number of samples, total area, and total annotations. FTW has an order of magnitude more samples and area covered than the next-largest dataset (AI4Boundaries (d’Andrimont et al. 2023)), with 70 462 sample chips spanning 166,293  $km^2$ . It also has the highest annotation volume with 1.63M field polygons, compared to the next-largest volume of 1.07M in AI4Boundaries.

The FTW dataset captures greater morphological diversity of agricultural field instances than any other dataset. Figure 1 shows example field boundaries from four countries. The diverse shapes and sizes reflect the unique topographical, environmental, and historical factors that influenced the development of field boundaries in each region. Figure 3 shows the dramatic difference between the field areas in different countries. There is little overlap between the distributions of Vietnam (small fields) and Brazil (large fields). Figure 9 shows the distribution field areas for AI4SmallFarms, AI4Boundaries, and FTW (note that we did not include PASTIS/PASTIS-R since both provide field boundaries in raster, not vector, format so we could not compute morphological statistics). AI4SmallFarms consists mostly of small-area fields, AI4Boundaries consists mostly of medium-area fields, and FTW has a broader distribution of field areas. There is also significant diversity in field shape complexity (see visualizations in supplement). For example, Estonia and Slovakia have complex field shapes, with 111.9 and 95.4 polygon vertices on average, respectively. In contrast, Kenya and Rwanda have simpler field shapes, with 4.5 and 5.6 vertices on average, respectively.

FTW provides a more complete representation of the diversity and complexity of agricultural landscapes across the globe than previous datasets. We hope that FTW will lead to

more broadly applicable models for field boundary segmentation by providing a large and diverse training dataset and enabling region-specific evaluation and error analysis.

## 4 Baseline Experiments

**Setup and Metrics** We follow the common approach to field instance segmentation of segmentation then polygonization of predicted raster masks (Persello et al. 2023). Unless specified otherwise, we used a U-net with EfficientNet-b3 backbone with inputs consisting of concatenated 4-channel RGB-NIR images from Window A and B. We found that this architecture performed well compared to other architectures and backbones (see **Architectures** paragraph in this section and Table 7 in supplement). We also found that concatenating both temporal windows and using four spectral bands performed best compared to other configurations (see **Multi-temporal and Multispectral Channels** paragraph and Table 4). We initialized the RGB channels using ImageNet and NIR channels using random weights. We optimized cross-entropy loss with class weights inversely proportional to each class’s frequency in the training set. We trained all models for 100 epochs. We used a fixed random seed for all experiments (randomly chosen). We did not perform hyperparameter tuning. Experiments required 4 A100 and 8 V100 GPUs for approximately one week.

We used semantic (pixel-level) and instance (object-level) segmentation metrics: pixel-level intersection over union (IoU), precision, and recall, and object-level precision and recall (functions in the FTW code repository). We converted segmentation masks to polygons using `rasterio` then computed object-level metrics with IoU threshold of 0.5.

**Modeling Configuration** Fields of The World includes multiple target mask formats and satellite images from two dates and four spectral channels, giving users many modeling choices. In semantic segmentation followed by polygonization, there are other choices such as which architecture to choose. We performed several experiments to assess the impact of these choices on model performance. In each experiment, we evaluate performance using the test set for each country. We report results for countries with presence/absence labels chosen to span a range of field sizes: Slovenia - SVN (average field size 0.64 ha), France - FRA (average 5.7 ha), and South Africa - ZAF (average 13.8 ha). We provide results for all countries in the supplementary material. An

Mask type	Pixel IoU			Pixel precision			Pixel recall			Object precision			Object recall		
	SVN	FRA	ZAF	SVN	FRA	ZAF	SVN	FRA	ZAF	SVN	FRA	ZAF	SVN	FRA	ZAF
2-class	0.66	<b>0.83</b>	<b>0.83</b>	0.84	0.87	0.87	0.76	0.95	0.94	0.52	0.64	0.54	0.06	0.29	0.24
2-class (ignore presence-only bg)	<b>0.69</b>	0.82	0.81	0.78	0.86	0.84	<b>0.86</b>	<b>0.95</b>	<b>0.96</b>	0.30	0.38	0.44	0.08	0.15	0.19
3-class	0.67	<b>0.83</b>	<b>0.83</b>	0.87	0.88	0.88	0.75	0.94	0.94	<b>0.60</b>	<b>0.71</b>	<b>0.63</b>	0.07	0.45	0.35
3-class (ignore presence-only bg)	0.59	0.79	0.79	<b>0.90</b>	<b>0.89</b>	<b>0.89</b>	0.63	0.88	0.88	0.33	0.55	0.51	<b>0.20</b>	<b>0.58</b>	<b>0.55</b>

Table 3: Performance metrics for different target mask formats in Slovenia (SVN), France (FRA), and South Africa (ZAF). We compared 2-class field extent and 3-class masks with or without ignoring background (bg) pixels for presence-only samples.

Channels	Pixel IoU			Pixel precision			Pixel recall			Object precision			Object recall		
	SVN	FRA	ZAF	SVN	FRA	ZAF	SVN	FRA	ZAF	SVN	FRA	ZAF	SVN	FRA	ZAF
Stacked Windows A and B	<b>0.58</b>	<b>0.79</b>	<b>0.80</b>	<b>0.91</b>	<b>0.89</b>	<b>0.90</b>	0.61	<b>0.87</b>	0.87	<b>0.30</b>	<b>0.54</b>	<b>0.55</b>	<b>0.18</b>	<b>0.58</b>	<b>0.54</b>
Stacked Windows A and B (RGB only)	<b>0.58</b>	0.78	0.79	0.90	<b>0.89</b>	0.89	<b>0.62</b>	0.86	<b>0.88</b>	0.27	0.51	0.53	<b>0.18</b>	0.56	<b>0.54</b>
Mean of Windows A and B	0.54	0.77	0.78	0.88	<b>0.89</b>	0.88	0.59	0.86	<b>0.88</b>	0.27	0.50	0.49	0.16	0.55	0.53
Window A only	0.55	0.77	0.78	0.88	0.88	0.88	0.59	0.86	0.87	0.27	0.47	0.47	0.17	0.54	0.52
Window B only	0.52	0.78	0.79	0.87	<b>0.89</b>	0.89	0.57	0.86	0.87	0.24	0.49	0.52	0.15	0.54	0.53

Table 4: Ablation results for multispectral (RGB-NIR vs. RGB only) and multi-temporal (Window A and Window B) input channels in Slovenia (SVN), France (FRA), and South Africa (ZAF).

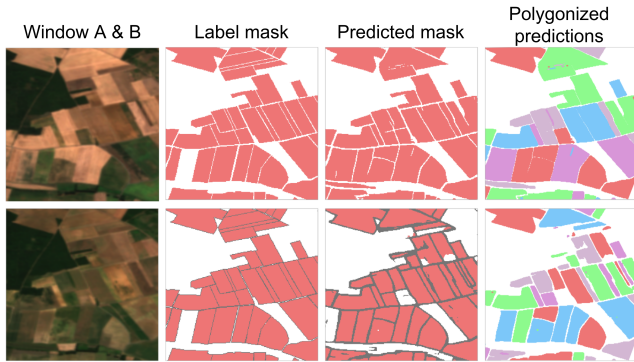


Figure 4: Example France predictions for 2-class (top row) and 3-class (bottom row) models in rows 2 and 4 of Table 3.

extensive search of modeling configurations is beyond the scope of this work, but we hope future studies will explore a greater range of options using the FTW dataset.

**Target Mask** We evaluated the two types of target masks provided in FTW: 2-class (field interior vs. background) and 3-class (field interior, boundary, and background). In presence-only countries, pixels with an unknown class are labeled as background. We evaluated two scenarios: 1) when computing the loss, ignore all pixels labeled background for presence-only countries, and 2) compute the loss for all pixels treating unknown labels as background. Before computing test metrics, we converted outputs from all models to binary field extent masks to ensure a common evaluation basis.

Rows 1 and 3 of Table 3 show that almost all metrics are higher with 3-class masks than binary masks. Rows 2 and 4 compare when unknown-label pixels are counted as background or ignored for presence-only samples when training with 3-class masks. Although object precision is lower when ignoring unknown pixels, object recall is significantly

higher, especially for Slovenia. Figure 4 shows that the segmentation masks predicted by both models are good, but 3-class masks improve the delineation of contiguous fields and thus object recall. From these results, we concluded that training with 3-class masks and ignoring presence-only background is most likely to give good performance across all regions and used this setup for subsequent experiments.

**Multi-temporal and Multispectral Channels** We did an ablation experiment to evaluate the benefit of the two contrasting image dates (Window A and B) in FTW. We also evaluated a mean of both windows. Finally, we evaluated with/without the NIR channel. Table 4 shows the best performance comes from both time windows and all spectral channels. Overall, removing one of the time windows causes a greater drop in metrics than removing the NIR channel. This is consistent with previous work that showed performance improvements were greater when adding more timesteps compared to more spectral channels (Debats et al. 2016).

**Architecture** We evaluated U-net (Ronneberger, Fischer, and Brox 2015) and DeepLabv3+ (Chen et al. 2018) models with 5 different backbones: ResNet-18, ResNet-50, ResNeXt-50, EfficientNet-b3, and EfficientNet-b4. Overall performance is similar across different architectures and backbones, though U-net models tend to outperform DeepLabv3+ models (results in supplement).

**Transfer Learning** Some countries have many labels while others have few (Table 1). Prior work showed that performance on a data-scarce region could be improved by pre-training models on a country with a large labeled dataset and then fine-tuning on the target region. Wang, Waldner, and Lobell (2022) fine-tuned a model for India after pre-training on France. Persello et al. (2023) fine-tuned for Cambodia and Vietnam after pre-training on the Netherlands.

Direct comparisons to these works are not possible due to differences in data formats. Instead, we performed three

Analogous work	Fine-tune/test	Pre-train	Pixel IoU	Pixel precision	Pixel recall	Object precision	Object recall
Wang, Waldner, and Lobell (2022)	India	FRA	0.51 / 0.50	-	0.60 / 0.55	-	0.03 / 0.13
		AUT, ESP, FRA, LUX, NLD, SVN, SWE	0.16 / 0.54	-	0.16 / 0.59	-	0.05 / 0.16
		FTW - {India}	<b>0.57 / 0.55</b>	-	<b>0.63 / 0.60</b>	-	<b>0.14 / 0.19</b>
Persello et al. (2023)	Cambodia Vietnam	NLD	0.04 / 0.48	0.77 / 0.90	0.04 / 0.51	0.02 / 0.22	0.00 / 0.18
		AUT, ESP, FRA, LUX, NLD, SVN, SWE	0.16 / 0.51	0.90 / 0.90	0.16 / 0.54	0.12 / 0.27	0.04 / 0.22
		FTW - {Cambodia, Vietnam}	<b>0.43 / 0.55</b>	<b>0.92 / 0.91</b>	<b>0.44 / 0.58</b>	<b>0.21 / 0.29</b>	<b>0.16 / 0.24</b>

Table 5: Transfer learning results for models pre-trained on France (FRA) or Netherlands (NLD), AI4Boundaries countries (Austria/AUT, Spain/ESP, FRA, Luxembourg/LUX, NLD, Slovenia/SVN, and Sweden/SWE), or FTW minus the target region. Models are fine-tuned and tested on the target region. We report recall metrics only for India since it has presence-only labels. Each cell gives two results: no fine-tuning / after fine-tuning using the target training set.

analogous experiments to evaluate the improvement of pre-training on FTW compared to smaller, more geographically limited datasets as in prior work: 1) pre-training on one data-rich country (France or Netherlands), 2) pre-training on the countries included in AI4Boundaries (to emulate pre-training on AI4Boundaries), and 3) pre-training on FTW with the target country held-out. We then fine-tuned each model for 200 epochs using the target country (India or Cambodia+Vietnam) training set and evaluated on its test set.

Table 5 shows that models pre-trained with FTW outperform models trained on more geographically limited subsets in both target regions. The performance of FTW models without any fine-tuning is especially impressive. FTW models with no fine-tuning perform similarly or better than fully fine-tuned versions of compared models.

**Deployment Readiness** Motivated by the impressive performance of FTW pre-trained models without fine-tuning in Table 5, we used a FTW pre-trained model to predict field boundaries in Ethiopia, a challenging region not in FTW (Figure 5). The results show good qualitative performance that could be improved with local fine-tuning and post-processing. This shows the high potential of FTW to be immediately used in practice with little adaptation effort.

## 5 Discussion and Conclusion

ML research on automatic extraction of agricultural field boundaries from remotely sensed imagery is limited by a lack of ML-ready datasets to train and evaluate models on the global diversity of crop fields. These datasets are urgently needed in many applications for agriculture, climate change, and development. We designed Fields of The World to improve ML model performance for field boundary segmentation in diverse global agricultural landscapes and enable granular country-scale evaluation for more countries than any prior dataset. Our experiments established a performance baseline for the new FTW benchmark and showed that FTW-trained models perform better than more geographically limited datasets analogous to existing benchmarks.

Future work could build on these baselines by testing more model architectures, including instance segmentation architectures (e.g., Mask-RCNN (He et al. 2017) or SAM (Kirillov et al. 2023)) and geospatial foundation models (e.g., SatMAE (Cong et al. 2022) or Presto (Tseng et al. 2023)). Future work could also explore other methods of constructing target masks, motivated by our result that train-

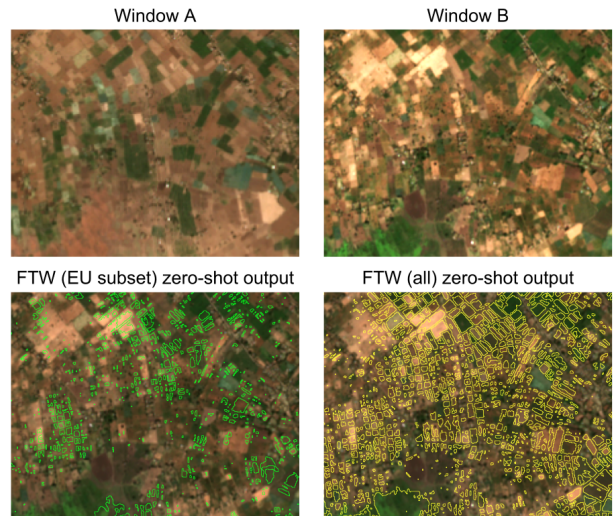


Figure 5: Zero-shot predictions with no post-processing for a 20 sq km region in Ethiopia ( $8^{\circ}05'N$ ,  $38^{\circ}51'E$ ). A FTW pre-trained model achieves (qualitatively) good performance, even though Ethiopia is not in the FTW training set and is a challenging region to delineate fields.

ing with 3-class masks performed better than 2-class masks.

We provide complete metadata for sample grids, sample chips, and field boundary polygons to enable future extensions of Fields of The World. For example, future work could add more spectral channels or sensors, timesteps, or sample locations. FTW can be extended as field polygons become available for more countries. We hope the community will build on FTW as research on this important task grows.

**Benchmarking on FTW** We hope this study will inspire researchers to develop new methods for field boundary segmentation and measure improvement using the FTW benchmark. We suggest benchmarking performance on the per-country test sets and reporting individual country results, the mean across all countries, or the minimum across countries (worst-case performance). Supplement Table 13 reports these metrics for the best model evaluated in this paper.

## Ethics statement

Researchers, practitioners, and other users of Fields of The World must be aware of important ethical considerations

raised by the digitization of field boundaries and other information from publicly accessible satellite data. Digitized field boundary data could inadvertently expose the practices and characteristics of individual land parcels, which could infringe on the privacy of local landowners who may be unaware of this digitization or its implications. There are also risks that private or public entities may use digitized field boundary data in a way that marginalizes vulnerable individuals such as smallholder farmers. Rolf et al. (2024) summarized distinct ethical concerns of machine learning applied to satellite data. In line with the recommendations of Rolf et al. (2024), we suggest that users of FTW work with local organizations and communities to build and release responsible field boundary datasets and ensure their project goals and practices align with local needs and regulations.

### Acknowledgments

This project was supported by funding from the Taylor Geospatial Engine and a NASA Supplemental Open Source Software Award.

### References

- Administration of Technical Agricultural Services. 2024. Flik Plot Repository - Open Data. <https://data.public.lu/en/datasets/referentiel-des-parcelles-flik/#resources>. Accessed: 2024-08-15.
- Agency for Payments in Agriculture, Fisheries & Rural Development. 2024. ARKOD. <https://www.aprrr.hr/prostorni-podaci-servisi/>. Accessed: 2024-08-15.
- Beck, H. E.; Zimmermann, N. E.; McVicar, T. R.; Vergopolan, N.; Berg, A.; and Wood, E. F. 2018. Present and future Köppen-Geiger climate classification maps at 1-km resolution. *Scientific data*, 5(1): 1–12.
- Chen, L.-C.; Zhu, Y.; Papandreou, G.; Schroff, F.; and Adam, H. 2018. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 801–818.
- Cong, Y.; Khanna, S.; Meng, C.; Liu, P.; Rozi, E.; He, Y.; Burke, M.; Lobell, D.; and Ermon, S. 2022. SatMAE: Pre-training transformers for temporal and multi-spectral satellite imagery. *Advances in Neural Information Processing Systems (NeurIPS)*, 35: 197–211.
- d'Andrimont, R.; Claverie, M.; Kempeneers, P.; Muraro, D.; Yordanov, M.; Peressutti, D.; Batič, M.; and Waldner, F. 2023. AI4Boundaries: an open AI-ready dataset to map field boundaries with Sentinel-2 and aerial photography. *Earth System Science Data*, 15(1): 317–329.
- Debats, S. R.; Luo, D.; Estes, L. D.; Fuchs, T. J.; and Caylor, K. K. 2016. A generalized computer vision approach to mapping crop fields in heterogeneous agricultural landscapes. *Remote Sensing of Environment*, 179: 210–221.
- Estes, L. D.; Ye, S.; Song, L.; Luo, B.; Eastman, J. R.; Meng, Z.; Zhang, Q.; McRitchie, D.; Debats, S. R.; Muhando, J.; et al. 2022. High resolution, annual maps of field boundaries for smallholder-dominated croplands at national scales. *Frontiers in Artificial Intelligence*, 4: 744863.
- European Parliament; and Council of the European Union. 2023. Regulation (EU) 2023/1115 on Deforestation-free Products. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32023R1115>. Accessed: 2024-08-08.
- fiboa contributors. 2024. Field Boundaries for Agriculture (fiboa) specification. <https://github.com/fiboa>. Accessed: 2024-08-15.
- Finnish Food Authority. 2021. Agricultural parcel containing spatial data. <https://www.ruokavirasto.fi/en/about-us/open-information/inspire/>. Accessed: 2024-08-15.
- Garnot, V. S. F.; Landrieu, L.; and Chehata, N. 2022. Multi-modal temporal attention models for crop mapping from satellite time series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 187: 294–305.
- Hall, J. V.; Argueta, F.; and Giglio, L. 2024. GloCAB cropland field boundary dataset. *Data in Brief*, 55: 110739.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask R-CNN. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2961–2969.
- Instituto de Financiamento da Agricultura e Pescas. 2024. Sistema de Informação de Parcelas.
- Jung, S.; Rasmussen, L. V.; Watkins, C.; Newton, P.; and Agrawal, A. 2017. Brazil's national environmental registry of rural properties: implications for livelihoods. *Ecological Economics*, 136: 53–61.
- Kehs, A.; McCloskey, P.; Chelal, J.; Morr, D.; Amakove, S.; Plimo, B.; Mayieka, J.; Ntango, G.; Nyongesa, K.; Pamba, L.; et al. 2021. From village to globe: A dynamic real-time map of African fields through PlantVillage. *Frontiers in Sustainable Food Systems*, 5: 514785.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; Dollár, P.; and Girshick, R. 2023. Segment Anything. *arXiv preprint arXiv:2304.02643*.
- Kondmann, L.; Toker, A.; Rußwurm, M.; Camero, A.; Peressutti, D.; Milcinski, G.; Mathieu, P.-P.; Longépé, N.; Davis, T.; Marchisio, G.; et al. 2021. DENETHOR: The DynamicEarthNET dataset for Harmonized, interoperable, analysis-Ready, daily crop monitoring from space. *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*.
- Microsoft Open Source; McFarland, M.; Emanuele, R.; Morris, D.; and Augspurger, T. 2022. microsoft/PlanetaryComputer: October 2022. <https://doi.org/10.5281/zenodo.7261897>.
- Ministry of Food, Agriculture and Fisheries of Denmark. 2021. LandbrugsGIS. <https://landbrugsgeodata.fvm.dk/>. Accessed: 2024-08-15.
- Nakalembe, C.; and Kerner, H. 2023. Considerations for AI-EO for agriculture in Sub-Saharan Africa. *Environmental Research Letters*, 18(4): 041002.
- NASA Harvest and Radiant Earth Foundation. 2024. Rwanda Field Boundary Competition. Accessed: 2024-08-07.
- Netherlands National GeoRegistry. 2021. Dataset: Basic registration Crop Parcels (BRP). <https://www.pdok.nl/atom->

- downloads/services/-/article/basisregistratie-gewaspercelen-brp-. Accessed: 2024-08-15.
- Oldoni, L. V.; Sanches, I. D.; Picoli, M. C. A.; Covre, R. M.; and Fronza, J. G. 2020. LEM+ dataset: for agricultural remote sensing applications. Mendeley Data, V1.
- Persello, C.; Grift, J.; Fan, X.; Paris, C.; Hänsch, R.; Koeva, M.; and Nelson, A. 2023. AI4SmallFarms: A Data Set for Crop Field Delineation in Southeast Asian Smallholder Farms. *IEEE Geoscience and Remote Sensing Letters*.
- Planet; Foundation, R. E.; of Agriculture, W. C. D.; and (DLR), G. A. C. 2021. A Fusion Dataset for Crop Type Classification in Western Cape, South Africa (Version 1.0). <https://doi.org/10.34911/rdnt.gqy868>. Radiant MLHub. [Date Accessed].
- Planet, TUM, DLR and Radiant Earth. 2021. AI4FoodSecurity Challenge. <https://platform.ai4eo.eu/ai4food-security-south-africa/data>. Accessed: 2024-08-15.
- Pula Advisors. 2022. Bird's Eye. [https://ecass-project-documentation.readthedocs.io/en/latest/modules/about\\_Ecass.html](https://ecass-project-documentation.readthedocs.io/en/latest/modules/about_Ecass.html). Accessed: 2024-08-15.
- Rolf, E. 2023. Evaluation challenges for geospatial ML. *arXiv preprint arXiv:2303.18087*.
- Rolf, E.; Klemmer, K.; Robinson, C.; and Kerner, H. 2024. Position: Mission Critical—Satellite Data is a Distinct Modality in Machine Learning. In *Forty-first International Conference on Machine Learning*.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 234–241. Springer.
- Sainte Fare Garnot, V.; and Landrieu, L. 2021. Panoptic Segmentation of Satellite Image Time Series with Convolutional Temporal Attention Networks. *ICCV*.
- Schneider, M.; and Körner, M. 2022. EuroCrops.
- Slovak Republic Government. 2024. Spatial Data and Services. Accessed: 2024-08-07.
- Sullivan, J. A.; Samii, C.; Brown, D. G.; Moyo, F.; and Agrawal, A. 2023. Large-scale land acquisitions exacerbate local farmland inequalities in Tanzania. *Proceedings of the National Academy of Sciences*, 120(32): e2207398120.
- Swedish Board of Agriculture. 2024. Agricultural block database. <https://www.geodata.se/geodataportalen/srv/swe/catalog.search;jsessionid=6C2D281619D69AC2356E1BD4C1923A3A#/metadata/df439ba5-014e-44ec-86cb-ddb9e5ba306c>. Accessed: 2024-08-15.
- Taravat, A.; Wagner, M. P.; Bonifacio, R.; and Petit, D. 2021. Advanced fully convolutional networks for agricultural field boundary detection. *Remote Sensing*, 13(4): 722.
- The Service and Payment Agency (ASP). 2024. The Graphical Parcel Register (Registre parcellaire graphique (RPG)). <https://geoservices.ign.fr/rpg#telechargementrpg2021>.
- Tseng, G.; Cartuyvels, R.; Zvonkov, I.; Purohit, M.; Rolnick, D.; and Kerner, H. 2023. Lightweight, pre-trained transformers for remote sensing timeseries. *arXiv preprint arXiv:2304.14065*.
- USDA Foreign Agricultural Service. 2024. Crop Calendar Charts. <https://ipad.fas.usda.gov/ogamaps/cropcalendar.aspx>. Accessed: 2024-08-15.
- Wang, S.; Waldner, F.; and Lobell, D. B. 2022. Unlocking large-scale crop field delineation in smallholder farming systems with transfer learning and weak supervision. *Remote Sensing*, 14(22): 5738.
- Wang, S.; Waldner, F.; and Lobell, D. B. 2023. 10,000 crop field boundaries across India.