WorldCereal

SEASONALY UPDATED GLOBAL CROP MAPPING

FINAL REPORT

esa

vito remote sensing · UNIVERSITÉ DE STRASBOURG · CS La force de l'innovation · eleaf · Universitat de València · IIASA · gisat · WAGENINGEN UNIVERSITY & RESEARCH

EXECUTIVE SUMMARY

The WorldCereal Final Report (FR) provides a summarized overview of the results of two-year research on the development of the first global seasonally updated temporary crop land and crop type maps.

This report starts with a high-level overview of the objectives, purpose and scope of the ESA WorldCereal project.

Section two shares more information about the main datasets used by the WorldCereal system (EO, in situ and auxiliary data). WorldCereal uses open and free data sets, one of the main requirements of the user community. Apart from the EU Copernicus Sentinel-1, Sentinel-2 and NASA Landsat 8 EO data, WorldCereal also used in situ data which is the first harmonized global open and free in situ data set for crop mapping.

The used algorithms are described in section three. After a detailed benchmarking activity, considering different input data and the results on accuracy, speed, stability, etc., the algorithms for all products were selected.

More information about the different WorldCereal products, global accuracies and validation methodology can be found in section four and five.

The last sections, section six and seven, show some practical use cases defined by the WorldCereal user community as well as a roadmap of the future steps.

Discover more about WorldCereal at **esa-worldcereal.org** and discover the WordCereal products via the interactive viewer. **https://vdm.esa-worldcereal.org/**

Version: 1.0
Date: May 2023
Author(s): Jeroen Degerickx,
Kristof Van Tricht, Sven Gilliams
ESRIN Contract No. N°4000130569/20/I-NB
Website: https://esa-worldcereal.org/en

TABLE OF CONTENTS

1. INTRODUCTION

1.1 WORLDCEREAL

WorldCereal is a project funded by the European Space Agency (ESA) and aims to develop an efficient, agile and robust EO based system allowing global and detailed agricultural monitoring. More specifically, the open-source WorldCereal system will allow any user to create detailed (10 m resolution), seasonaly updated maps on cropland extent, crop type (specifically for maize and wheat) and irrigation practices, the latter differentiating between actively irrigated and rainfed fields. To this end, the system will build on (1) open-and free satellite data, such as Sentinel-1, -2, and -3 and Landsat 8 and relevant in-situ data, (2) newly developed and tested state-of-the-art classification algorithms, (3) the necessary cloud infrastructure to perform the work and (4) a visualization tool to view and derive useful information from the produced maps.

1.2 PURPOSE AND SCOPE OF THIS DOCUMENT

This Final Report (FR) summarizes at a very high-level the different building blocks of the opensource WorldCereal system in order to produce the first seasonal updated temporary cropland, crop type, active irrigation and active cropland maps.

First, the WorldCereal products are described to ensure that readers fully understand the difference and resemblance between the WorldCereal products and other existing global products.

Secondly, the necessary steps to create these products are described as different building blocks. The main building blocks described within this document are.

- The data set (EO and in-situ)
- The algorithm

More details can be found throughout other approved WorldCereal deliverables. Detailed references to these deliverables are not made to keep the readability of this document.

A couple of examples of the potential use of the products is also given, along with a roadmap for the future and a conclusion.

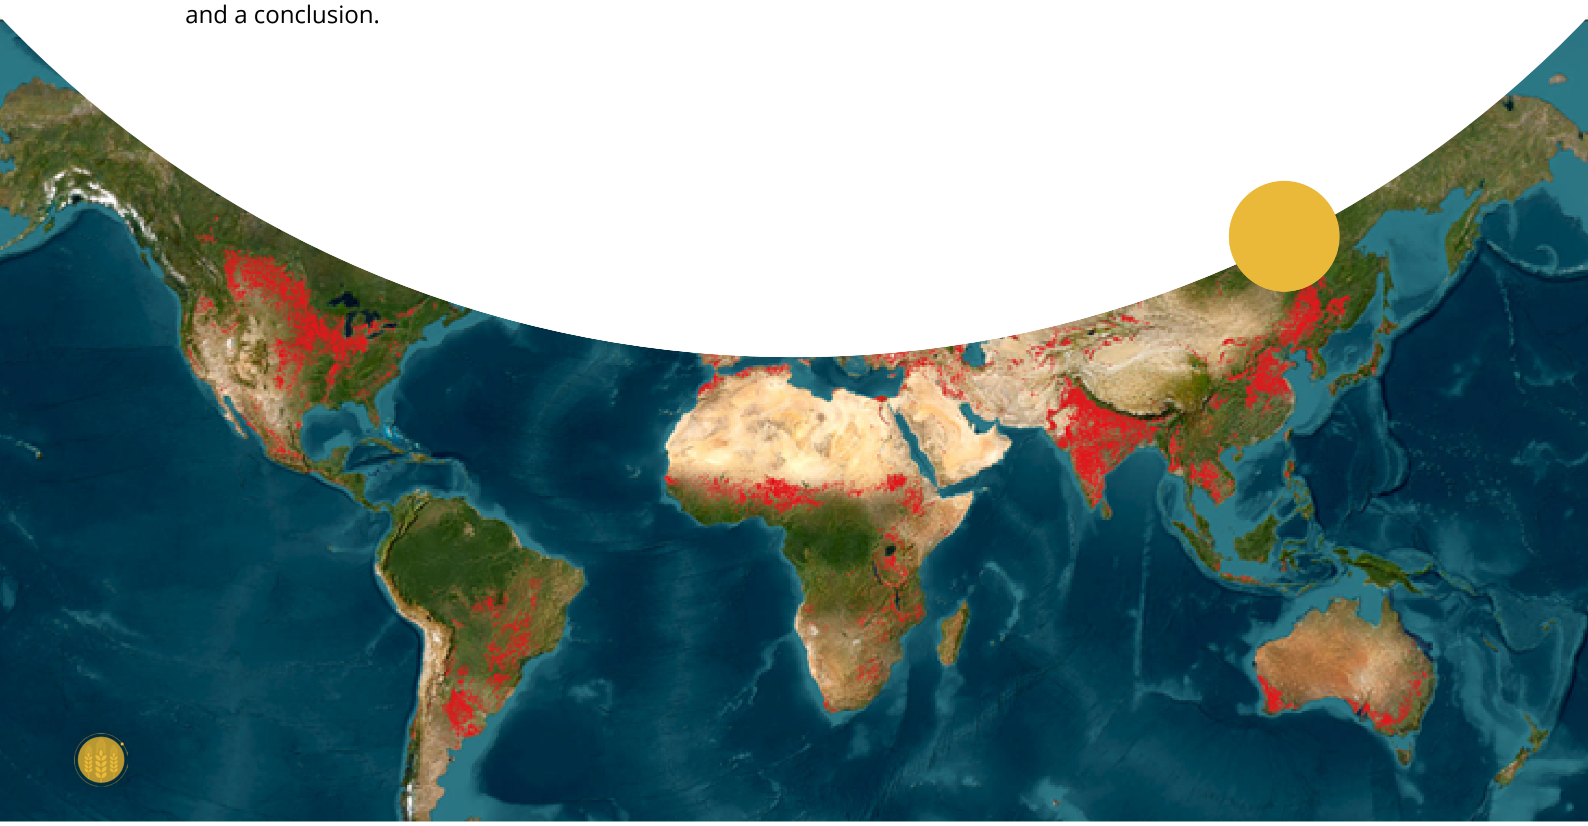

2. WORLDCEREAL DATASETS

A key requirement for WorldCereal was the fact that WorldCereal needed to be an open-source system. Another key requirement was the fact that all of the system code and classification had to be publicly available. This last requirement implies that both the EO data and the in-situ data (used to train and validate algorithms and to create products) had to be available openly and freely.

The section below provides an overview of the main datasets used by the WorldCereal system to create the seasonally updated maps.

2.1 EO-DATA

A key requirement for WorldCereal was the fact that WorldCereal needed to be an open-source system. Another key requirement was the fact that all of the system code and classification had to be publicly available. This last requirement implies that both the EO data and the in-situ data (used to train and validate algorithms and to create products) had to be available openly and freely.

The section below provides an overview of the main datasets used by the WorldCereal system to create the seasonally updated maps.

2.1.1 Optical data

The Copernicus Sentinel-2 satellite is the main source of optical data used to generate the WorldCereal products, due to its high spatial and temporal resolution. However, in certain regions of the world, availability of cloud-free Sentinel-2 data remains problematic, due to persistent cloud cover during the majority of the growing season (upper panel of Figure 1).

In these particular regions, the Sentinel-2 timeseries was complemented with additional optical observations derived from Landsat 8 (lower panel of Figure 1). The identification of these regions was done through thresholding the available number of cloud-free Sentinel-2 and Landsat 8 acquisitions on a yearly basis: in the case where an average of less than one cloud-free Sentinel-2 image was available per 10 days and where more than five cloud-free Landsat 8 images were available over a full year, the region was flagged as needing to use Landsat 8 optical imagery.

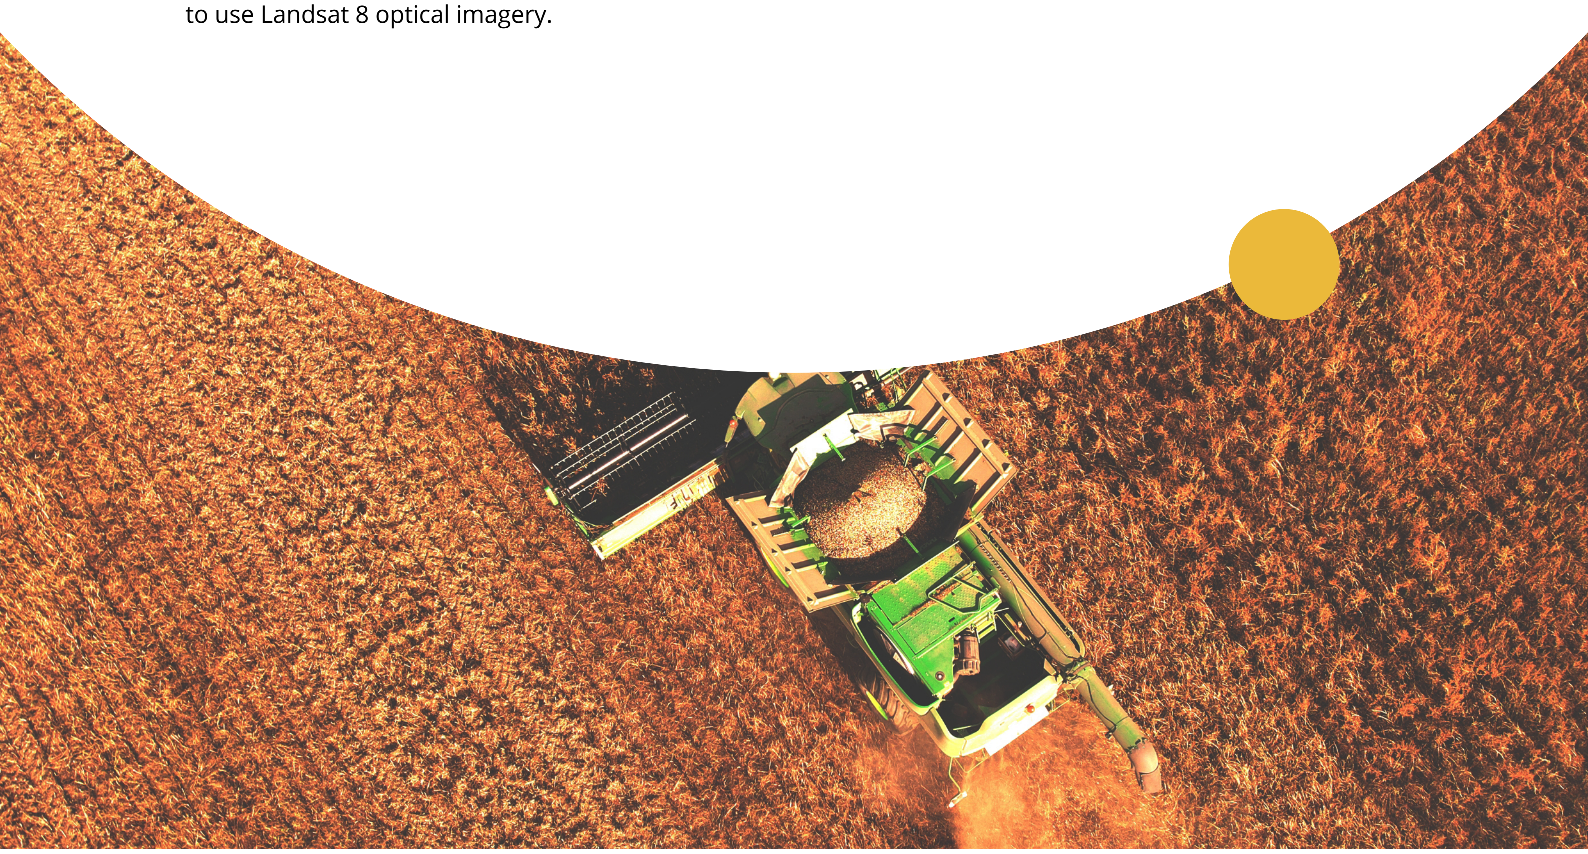

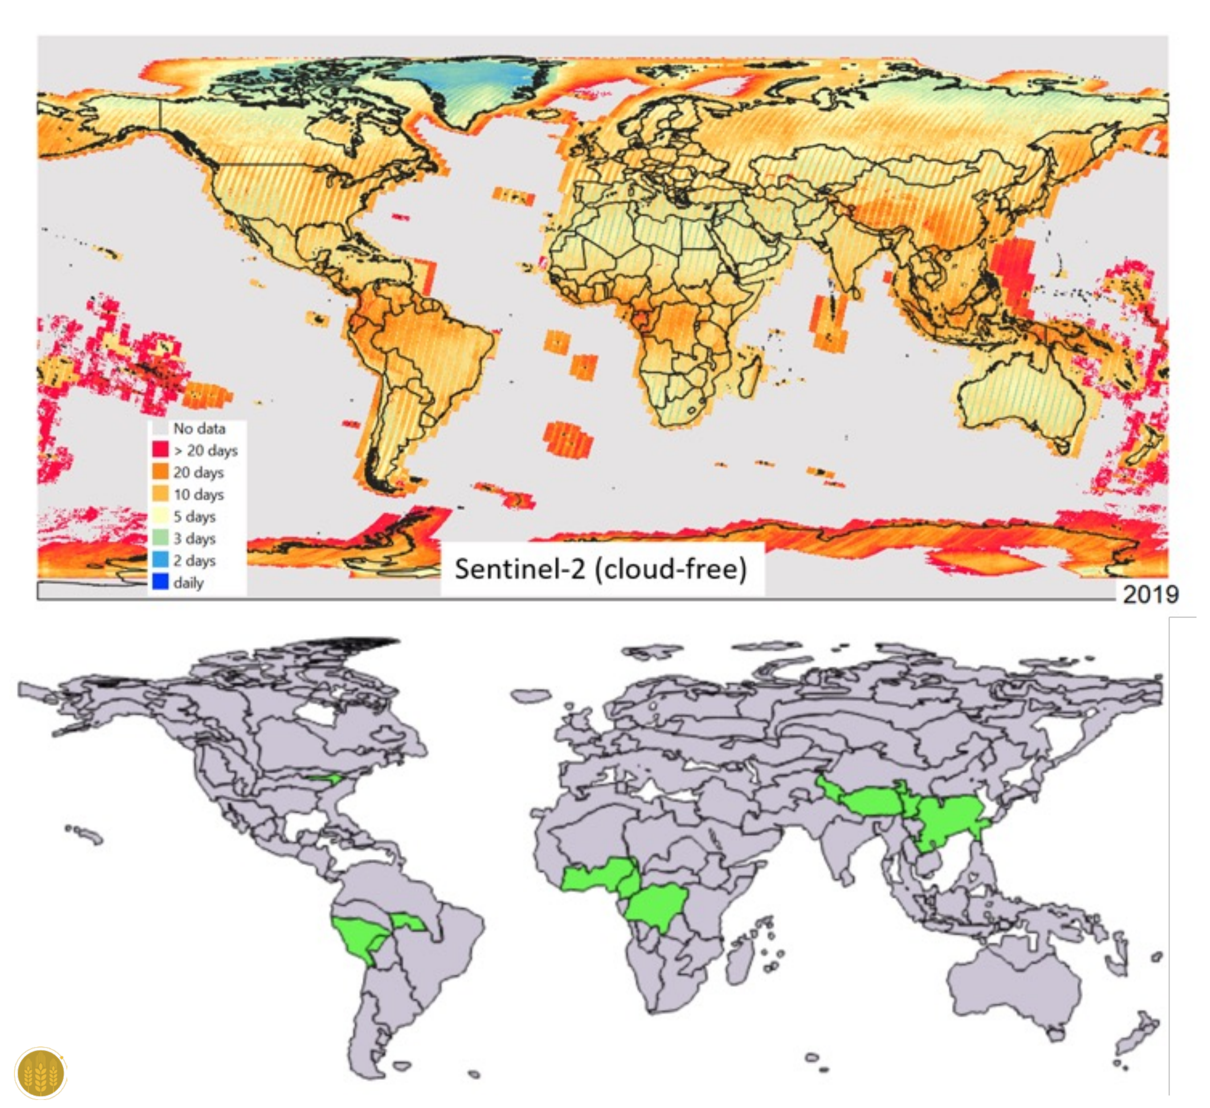

Figure 1: The upper panel shows the availability of cloud-free Sentinel-2 observations expressed as an average revisit frequency over the whole of 2019, whereas the lower panel highlights (in green) the specific Agro-Ecological Zones where Landsat 8 optical observations will be used to complement the Sentinel-2 timeseries.

Spatial resolution

Table 1 shows which optical bands from both Sentinel-2 and Landsat 8 were used within the WorldCereal system and at which spatial resolution. Note that according to Sen2Like convention, B05 of Landsat 8 should be matched with B8A of Sentinel-2. However, we specifically opted to match it with B08, understanding that the calibration would not be as good compared to using B8A. However, due to the fact that (1) all WorldCereal products will be generated at 10m resolution and (2) the NIR spectral band is used as input for the majority of the optical classification features, and (3) Sentinel-2 is the main source of optical information, it was decided that the 10m resolution of Sentinel-2 B08 was more important in this respect.

Wavelength range	Sentinel-21	Landsat 8	WorldCereal	Resolution
Blue	B02	B02	B02	10m
Green	B03	B03	B03	10m
Red	B04	B04	B04	10m
Red edge 1	B05	-	B05	20m
Red edge 2	B06	-	B06	20m
Red edge 3	B07	-	B07	20m
NIR	B08	B05	B08	10m
SWIR 1	B11	B06	B11	20m
SWIR 2	B12	B07	B12	20m

Table 1: Sentinel-2 and Landsat 8 optical bands used in the WorldCereal system along with their spatial resolution.
1 Note that the bands not listed in this table are currently not used by the WorldCereal system.

Cloud & shadow masking

The pre-processing chains of both Sentinel-2 and Landsat 8 come with a cloud mask which is based on either the Sen2Cor scene classification or on Landsat's FMask algorithm. Usually, these cloud masks tend to be conservative, with many cloud and shadow-related impurities remaining unflagged, especially near cloud edges. As pointed out by Baetens et al. (2019), both Sen2Cor and FMask cloud/shadow masks should be dilated to reduce these impurities and drastically boost the quality of the mask. At the same time, pixel-based cloud masks tend to be sensitive to bright surfaces such as certain greenhouse roofs, which are marked as cloudy in the cloud mask.

When applying a dilation to these masks, wrongly labelled bright surfaces are smeared out and will mask otherwise clear areas. Therefore, the improved cloud and shadow masking deployed in the WorldCereal system consists of two consecutive steps:

1. Erosion of the original mask by 30m to reduce bright surfaces wrongly labelled as clouds
2. Dilation of the original mask by 200m to reduce impurities near cloud edges wrongly labelled as clear

The project's benchmarking experiments have shown that this approach is an excellent trade-off between computational expensiveness (as opposed to for example multi-temporal approaches such as the MAJA cloud mask) and mask quality.

2.1.2 SAR data

Only one source of SAR data is used in the project, and that is Copernicus Sentinel-1. At the start of the project S-1A and S-1B were available. So, all training is done based on the availability of these two sensors.

The following common processing routines are performed on the Sentinel-1 data. Note that the first step (orbit direction selection) happens before the actual backscatter data is produced during pre-processing, while the next steps take place on the pre-processed inputs.

Spatial resolution
Sentinel-1 data in WorldCereal is by default processed at a spatial resolution of 20m and is handled by the features processor as such to minimize resource requirements. Only when the SAR features have been computed, an upsampling to 10m resolution is performed to match the 10m product requirement.

Orbit direction selection
Figure 2 shows the Sentinel-1 observation scenario that was in place at the start of the project. This demonstrates that in many countries of the world, either the ascending or descending orbit is available. At the same time, ascending and descending observations of the same area cannot easily be combined, due their different viewing angles having an impact on the backscatter time series. The best way to cope with this is to keep these time series apart for analysis.

Due to the globally robust nature of the WorldCereal system, for any location on the globe only one orbit direction is considered, even when two orbit passes are available. This ensures consistency of the resulting backscatter time series as well as the robustness and general applicability of the WorldCereal algorithm workflows throughout the world.

Sentinel-1 Constellation Observation Scenario: Revisit & Coverage Frequency

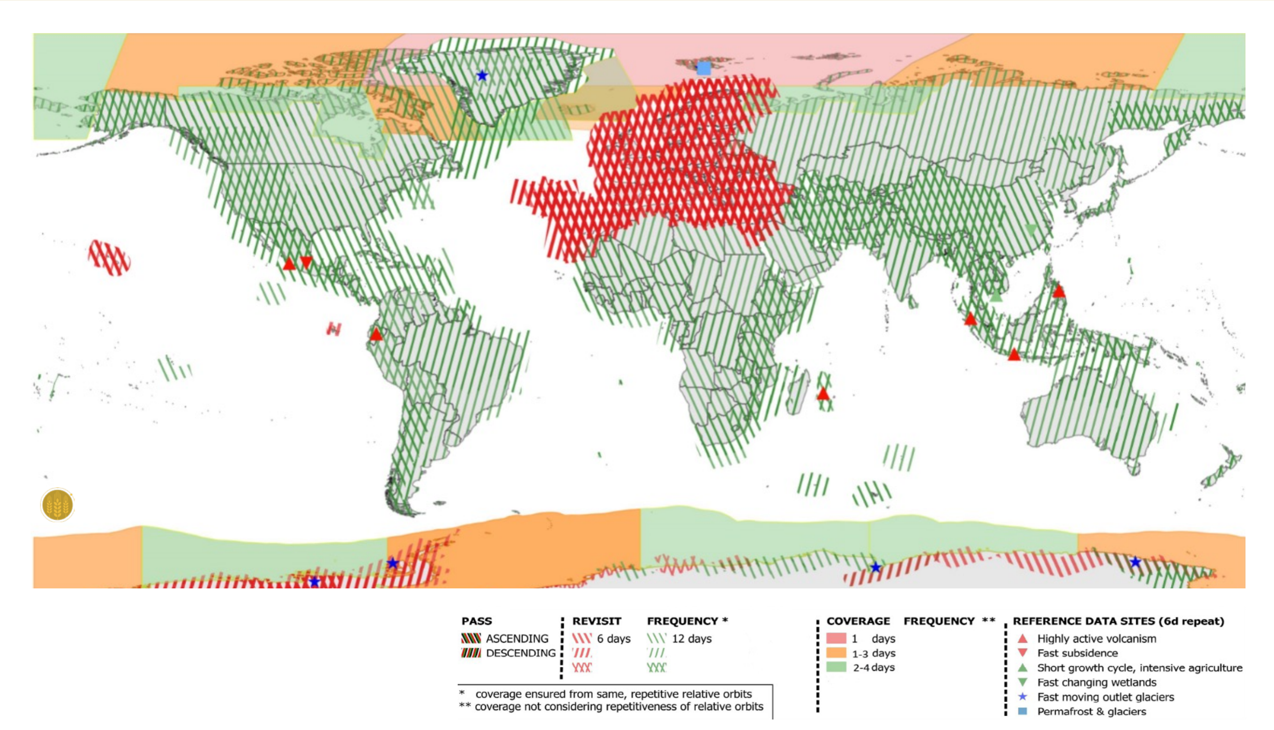

Figure 2: Sentinel-1 constellation observation scenario currently in place as of 05/2019.

Optional precipitation masking

Despite the fact that SAR can make surface observations through clouds, the signal might still be affected by meteorological events. Especially strong precipitation events may alter the radar signal directly, while also indirectly impacting the backscattered energy by altered dielectric properties of soil and vegetation due to a strong rainfall event.

The three most common practices to minimize this impact are:

- **Filtering outliers** in the backscatter timeseries: the WorldCereal system takes care of this approach during its compositing step (see below).
- Regularization of the impact by means of **band combinations**: the WorldCereal system deploys this step during feature computation, where the VH/VV ratio is computed minimizing the humidity impact on the time series.
- **Filtering out strong precipitation events**

Also, this last option is deployed optionally in the WorldCereal system by connecting the Sentinel-1 processing component to the AgERA5 global reanalysis data source which includes daily precipitation rates. The goal is to identify Sentinel-1 acquisitions that happened on rainy days, so that these observations can be eliminated from the time series. According to Tamm et al. (2016), a precipitation of 0.25 mm over 3 hours was a threshold above which Sentinel-1 acquisitions started to show significant impact. As the **AgERA5** reanalysis data only holds daily precipitation rates, we extrapolated this threshold to 10 mm per day as the threshold to decide whether or not to keep a Sentinel-1 acquisition. When this optional precipitation masking step is activated, all observations occurring on a day with more than 10 mm of precipitation are removed from the timeseries.

Multi-temporal speckle filtering

SAR backscatter acquisitions are prone to speckle noise due to the random nature of the scatterers in a radar resolution cell leading to either addition or cancellation of the individual contribution of these scatterers to the return signal. The WorldCereal system deploys an efficient multitemporal speckle filtering approach on both VV and VH backscatter time series which is implemented at the pixel level, preserving the original spatial resolution while significantly reducing speckle noise in the signal. The parameters of the filter are provided in Table 2.

Parameter	Value
Kernel type	Gamma-MAP
Filter window size	7
ENL	3

Table 2: Parameter settings of the multi-temporal speckle filtering routine

2.1.3 Thermal data

Within the WorldCereal project, the Landsat 8 Level 2 surface temperature product (ST_B10) is used as the source for high-resolution thermal data.

Spatial resolution

The original 30 m resolution ST_B10 product is resampled to 10 m using nearest neighbor prior to any further pre-processing and feature computation.

Cloud & shadow masking

Cloud and shadow masking for the thermal data is based on the cloud mask originally delivered together with the data itself, in turn based on the CFMask algorithm. Similarly to the optical data (section 2.1.1.2), this mask is additionally eroded by 30 m and dilated by 200 m to correct for obvious errors caused by bright surfaces or near cloud edges respectively.

2.2 IN-SITU DATA

A huge task throughout the project was the collection and harmonization of the crop land, crop type and irrigation in-situ data for validation and calibration.

This effort resulted in the first open, free and harmonized in-situ data set for crop land and crop type classification. The harmonized reference data used in the WorldCereal system can be accessed in two ways. See section 6.1.2 for more information on how to access the WorldCereal in-situ data.

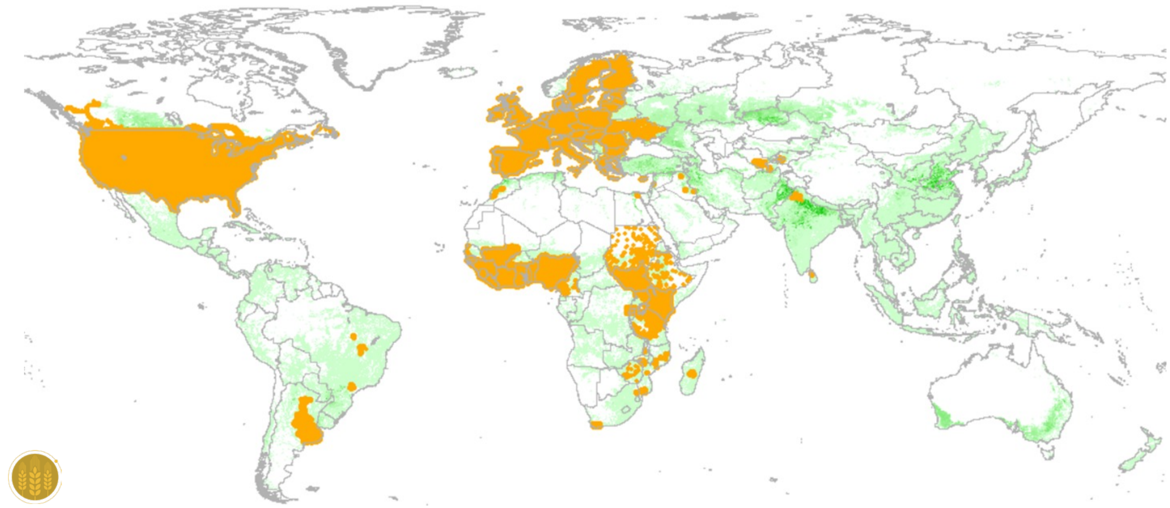

Figure 11: Overview of the locations of the 75 million harmonized observations with standardized metadata available in the WorldCereal in-situ database. (Orange areas indicate were in-situ data is available, and green areas indicate maize and cereals areas according to SPAM 2010)

These data sets are also a huge contribution to the newly started GEOGLAM in-situ data working group. The idea is that the data set through GEOGLAM will be kept available and will grow even after the end of WorldCereal.

2.3 AUXILIARY DATA

Next to the main satellite inputs and in-situ data, the workflow also makes use of ancillary data sources.

2.3.1 DEM data

The DEM that is integrated in the WorldCereal classification workflows to provide auxiliary features to the algorithms is the Copernicus DEM - Global and European Digital Elevation Model (COP-DEM) at approximately 30m spatial resolution ("GLO-30"). The original 30m data is resampled to 20m spatial resolution to align with the Sentinel-2 tile grid and to be compatible with the WorldCereal classification workflow.

2.3.2 Biome data

In addition to the DEM data, one additional layer of auxiliary information used by some of the WorldCereal classification algorithms is based on biome membership. These biomes are based on the 846 ecoregions of the Ecoregions 2017 map (Dinerstein et al., 2017). Biome membership allows classification algorithms to become "spatially aware" and hence gain knowledge on where training or inference data is originating from. Biome membership was used for implicit grouping of training and inference data based on shared characteristics as described by their biome.

2.3.3 Meteorological data

Within WorldCereal, the AgERA5 dataset provided by ECMWF is used as the source for meteo information. The AgERA5 dataset provides daily surface meteorological data for the period from 1979 to present, and is based on the hourly ECMWF ERA5 data at surface level. This data was used for crop type mapping, a crop-specific Growing Degree Days (GDD) normalization step was performed on the original time series using mean daily temperature data from the global AgERA5 reanalysis dataset. It was also used to extract features for the active irrigation mapping and clean SAR signals.

Spatial resolution

The ECWMF data are aggregated to daily time steps at the local time zone and corrected towards a finer topography at a 0.1° spatial resolution. For use in WorldCereal, the data is projected on the fly from lat/lon to the local UTM projection system, and further resampled to a 1km spatial resolution using nearest neighbour.

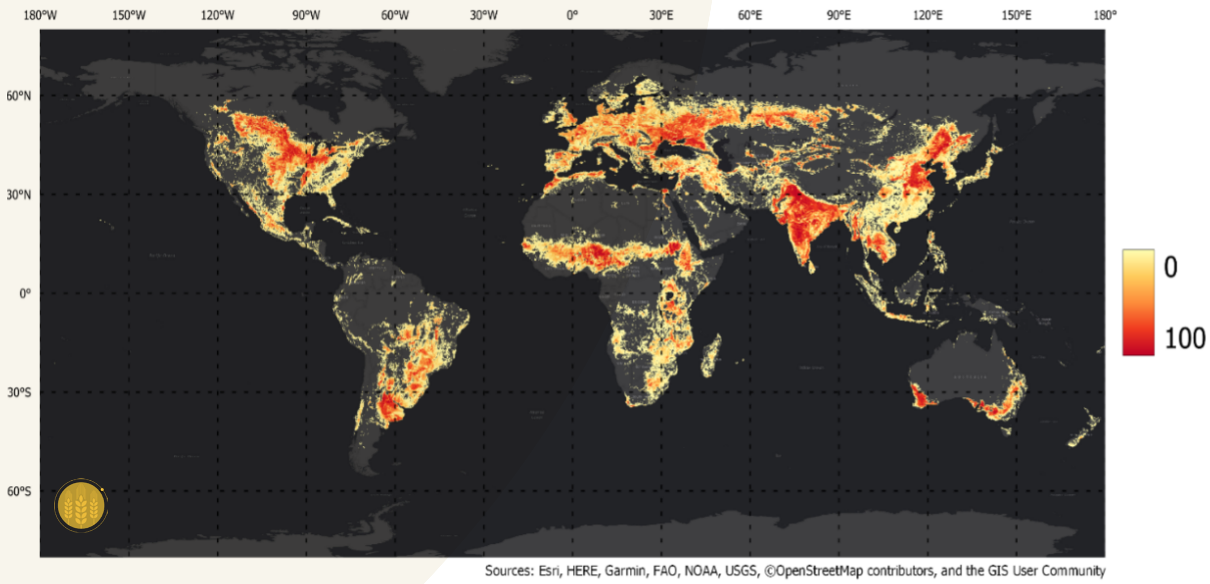

Figure 4: Biogeographical realms to localize models as used within the WorldCereal system.

3. WORLDCEREAL ALGORITHMS

Based on the above-mentioned data, the consortium developed the necessary open-source algorithms to produce the seasonally updated WorldCereal products.
The first phase of the project consisted of a benchmarking of different methodologies, in-situ data combination etc. Tests were done on scalability, transferability, production speed etc. The algorithms that were selected after benchmarking are described below.
The classification algorithms are based on a CatBoost model, which is a high-performance open-source library for gradient boosting on decision trees. Separate temporary crop mapping models were trained for each realm (see figure 4). Crop type and irrigation models were trained at the global level because of a lack of sufficient training samples at all individual realm levels. The models were trained on their respective seasonal training features.

3.1 FEATURE EXTRACTION

Classification features were derived from five data sources i.e. optical, radar, thermal infrared, DEM and biomes. Aside from the DEM and biomes, feature extraction always starts from pre-processed time series, either directly derived from the data source (e.g. Sentinel-2 reflectance bands) or a derived time series using a combination of multiple input variables (e.g. a spectral index). The exact timing and length of the time series was determined by the pixel-based crop calendars for the respective products. In the remainder of this section, we describe in more detail the specific features that were computed for generating the different WorldCereal products.

3.1.1 Temporary crop mapping features

When mapping temporary crops at global scale, using satellite data remains challenging in many regions, due to the variability in agricultural landscapes, to spectral similarity with other land cover classes, to fallow practices and to cloud obstruction during the growing season. Defining a robust and characteristic set of features that separates temporary crops from all other land cover classes is therefore key. From the Sentinel-2 optical pre-processed inputs, the following vegetation indices were computed, which have a proven record for mapping cropland: normalized difference vegetation index (NDVI), normalized difference water index (NDWI), normalized difference greenness index (NDGI), angle on near infrared (ANIR), normalized difference moisture index (NDMI), and two normalized difference red edge indices (NDRE85 and NDRE75).

Together with the shortwave-infrared bands B11 and B12, these time series were summarized using the 10th, 50th and 90th percentiles as well as the interquartile range (IQR). For NDVI in specific, also six equidistant time series descriptors (ts0-ts5) were added to explicitly capture the temporal profile, as well as 12 of the temporal features describing the evolution of a crop. As for Sentinel-1 SAR features, three time series were used as the basis for feature computation, i.e. VV, VH backscatter and the radar vegetation index (RVI), all of which have proven their use in crop mapping studies. These time series were summarized using the p10, p50 and p90 percentiles and the IQR. DEM altitude and slope and biome membership features were included as well. Finally, positional features latitude/longitude were also added and are called localization features. Localization features allow classification algorithms to become "spatially aware" and hence gain knowledge on where training or inference data is originating from. To avoid overfitting on exact combinations of latitude and longitude and at the same time reduce the risk of inferior product quality in data sparse regions, a random perturbation of up to 2.5° and 10° was added during training to latitude and longitude, respectively.
The complete list of features used for temporary crop mapping is provided in Table 3.

Data source	Time series	Features
OPTICAL	NDVI	p10, p50, p90, IQR, ts0, ts1, ts2, ts3, ts4, ts5, maxdif, mindif, difminmax, peak, lengthpeak, areapeak, ascarea, asclength, ascratio, descarea, desclength, descratio,
	NDWI, NDGI, ANIR, NDMI, NDRE85, NDRE75, B11, B12	p10, p50, p90, IQR
SAR	VV, VH, RVI	p10, p50, p90, IQR
DEM		Altitude, slope
BIOMES		13 biome fuzzy membership features
LOCALIZATION	latitude/longitude	Center latitude/longitude for a pixel

Table 3: Selected features for temporary crop mapping

3.1.2 Crop type mapping features

Specific crop type identification within the temporary crop mask started from a similar collection of features as for temporary crop mapping. To further enrich the feature set for distinguishing between different crop types, the standard deviation (STD) temporal statistic was added, in addition to the Sentinel-2 RGB bands (B02, B03, B04). The enhanced vegetation index (EVI) was also computed, and it was used to automatically detect the growing seasons in a time series. Outputs of the season detection algorithm include the number of detected growing seasons and for each season the date of its start, peak and end. Based on these outputs, minimum, median and maximum of both the length and EVI amplitude for all detected seasons were derived and added as classification features. Biome and localization features, in turn, were not included because insufficient global coverage of training data was available to cover all possible biome and localization combinations. The full feature set is described in Table 4. The seasonality detection is also used in the WorldCereal system to generate a seasonal active cropland layer (see section 4.4).

Data source	Time series	Features
OPTICAL	NDVI	p10, p50, p90, STD, ts0, ts1, ts2, ts3, ts4, ts5, maxdif, mindif, difminmax, peak, lengthpeak, areapeak, ascarea, asclength, ascratio, descarea, desclength, descratio,
		p10, p50, p90, STD
	NDWI, NDGI, ANIR, NDRE1, NDRE5,	p10, p50, p90, STD
	B2, B3, B4, B12	p10, p50, p90, STD
SAR	VV, VH, RVI	p10, p50, p90, IQR
DEM		Altitude, slope

Table 4: Selected features for crop type mapping

3.1.3 Irrigation mapping features

The feature collection of the WorldCereal irrigation model focuses on optical and thermal satellite observations from Sentinel-2 and Landsat 8, respectively, in combination with meteorological data from AgERA5. The basic features of the algorithm are pure Sentinel-2 based indices, such as NDVI, NDWI, modified normalized difference water index (MNDWI), EVI, and global vegetation moisture index (GVMI). These features can explain the health of a crop or if a crop is experiencing any form of stress. To prevent overfitting of the model, only the p90 and STD were calculated for these indices and added as features to the model. The p90 explains if a crop was able to flourish, potentially because of irrigation, or if a crop showed clear signs of stress. The STD is used to understand how dynamic the growing season of a crop was. A more advanced feature based on multiple Sentinel-2 bands is the spectral (cosine) median absolute deviation (SMAD). This feature highlights the temporal variation of multiple optical bands and has a positive impact on the detection of irrigation. Finally, also the geomedian (GM) calculated for the near-infrared and shortwave-infrared bands of Sentinel-2 were added to the model to emphasize the absorption patterns of chlorophyll and water.

Second, the relation between the air temperature (Tair) and land surface temperature (LST) is used to further understand the stress conditions of a crop. Under well-watered conditions, the difference between Tair¬ and LST is minimal, because the crop is cooling itself through transpiration processes. An increasing difference between Tair¬ and LST indicates that the crop is unable to transpire to its maximum potential and that stomata are being closed. Additional water is necessary for the crop to continue growing.

The third feature set focuses on the impact of irrigation on evapotranspiration (ET). Similarly to the proposed Sentinel-2-based indices, ET indicates if a crop can thrive or not. Modeling the actual ET using remote sensing data is complex and requires many inputs. To increase the computational efficiency of the model, a simple relation between the reference evapotranspiration (ET0) and the NDVI is used to calculate the actual ET. The

ET0 is calculated using AgERA5 data and relies on the penman-monteith equations. Since ET only explains if a crop is thriving and cannot help making a distinction between a rainfed crop in a humid climate or an irrigated crop in a more arid climate, precipitation data were added. The resulting precipitation deficit (Pdef) explains the difference between evapotranspiration and precipitation, where a large Pdef can be the result of extensive irrigation. From the Pdef time series, multiple features were calculated. The basic features are p10, p50, and p90, followed by the STD of the ET data. Additionally, the cumulative Pdef was calculated to understand the trend, longevity, and severity of the precipitation deficit. From this cumulative Pdef, the maximum and minimum were calculated, together with the maximum duration of a positive Pdef and the maximum slope of the cumulative Pdef curve. To conclude, also the sum divided by the length of the growing season of the ET0, ETact, and P were added as features to the algorithm. The duration of a positive cumulative Pdef, the maximum cumulative Pdef, and the minimum cumulative Pdef were also divided by the length of the season. These divisions were made to ensure that there is no bias toward regions with longer growing seasons.

Finally, to also include the relation between soil moisture and irrigation, the optical trapezoid model (OPTRAM) was used. This model focuses on the relationship between shortwave infrared reflectance and the NDVI. In this model, the shortwave infrared reflectance is converted into surface-transformed reflectance (STR). A trapezoidal model relies on a predefined wet and dry edge. These edges explain at which NDVI and STR value the soil is saturated or at its wilting point. In contrast to the original OPTRAM model, the edges are defined by grouping the STR data of one growing season by discrete NDVI steps. The dry edge of each step is represented by the minimum STR value within that specific step. The wet edge is calculated by adding the median STR with the standard deviation of the STR to prevent the model to become too sensitive to oversaturated conditions, which is a known issue. The final wet and dry edges are calculated by applying a linear regression through all the individual wet and dry edge values. For this soil moisture data, p50, p90 and STD were calculated and used as features. The final feature set was based on the correlation between the precipitation and the OPTRAM-based soil moisture content. A high correlation indicates that an increase in soil moisture is primarily driven by precipitation. A low correlation, on the other hand, might indicate that other factors, like irrigation, could have caused the increase in soil moisture content. For this dataset also p50, p90 and STD were calculated. Table 5 shows an overview of all the features used in the irrigation classification algorithm.

Data source	Time series	Features
OPTICAL	NDVI, EVI, MNDWI, NDWIveg, GVMI B08, B11 B02 + B03 + B04 + B08 + B11 + B12	p90, std p90, std GM SMAD
TIR	LST - Ta	p50, p90, STD
METEO	Precipitation, ETa, ET0 Precipitation deficit SSM, SSM_adj	Sum p10, p50, p90, std, cum_max, cum_min, cum_maxdur, cum_maxslope p50, p90, STD

Table 5: Selected features for active irrigation mapping

3.1.4 Classification

The output of each model is a binary classification of the inputs into the class of interest vs. all other classes. For the temporary crop map, this means a binary classification of temporary crops vs. all other land cover types. For seasonal crop type maps, this means maize or cereals vs. all other crops. For irrigation maps, actively irrigated crops are mapped against rainfed crops. Separate temporary crop mapping models were trained for each realm. Crop type and irrigation models were trained at the global level because of a lack of sufficient training samples in each individual realm. The models were trained on their respective seasonal training features: we trained temporary crop models based on annual features; a winter cereals model was trained based on the main wheat season features; a spring cereals model was trained using the features from the maize season in selected northern zones that are known to grow spring cereals; and a maize model was trained on the combined features of up to two maize seasons. For each model, the training data was randomly divided into 70% calibration, 20% validation and 10% test samples. During training, only calibration and validation samples were used, while test samples were retained for performance assessment.

Each model was trained with a maximum of 4000 iterations, a depth of 8, a learning rate of 0.05 and early stopping activated after 40 rounds without improvement. The distribution of the binarized training samples is imbalanced, the degree by which depends on the availability of different sources of reference datasets. To cope with this imbalance, we computed the class weights that balance the distribution, which we then used for loss weighting to eliminate the imbalance problem. In addition to these class weights, sample specific weights were also adjusted based on the confidence score of the respective reference dataset they were originating from. It is important to note that the models were trained on the combined training data from the available years (2017-2021), without providing year-specific information to the model. The aim was to train generalized models across multiple years that do not specifically require new training data in unseen years.

As a complementary product of the binary prediction, the models also provide binary class probabilities which we used to assess the pixel-based model's confidence in its prediction. Unconfident model predictions are characterized by binary probabilities close to 0.5, while confident model predictions are close to 0 or 1. Therefore, we defined classification confidence as a value between 0 and 100.

4. WORLDCEREAL PRODUCTS

The WorldCereal products are hierarchical and are generated based on global crop calendars. A requirement for WorldCereal was the fact that products needed to be ready within one month after the season. This implied that the trigger for the WorldCereal system had to be the growing season. The first step was to define these growing seasons, which was done based on information coming from FAO, JRC-ASAP, Crop Monitor and NASA-Harvest. Where no information was available, crop calendars were simulated (see Franch et al., 2022). These crop calendars were then stratified, based on start and end of season, in 203 WorldCereal production zones (AEZ) (see Figure 5).

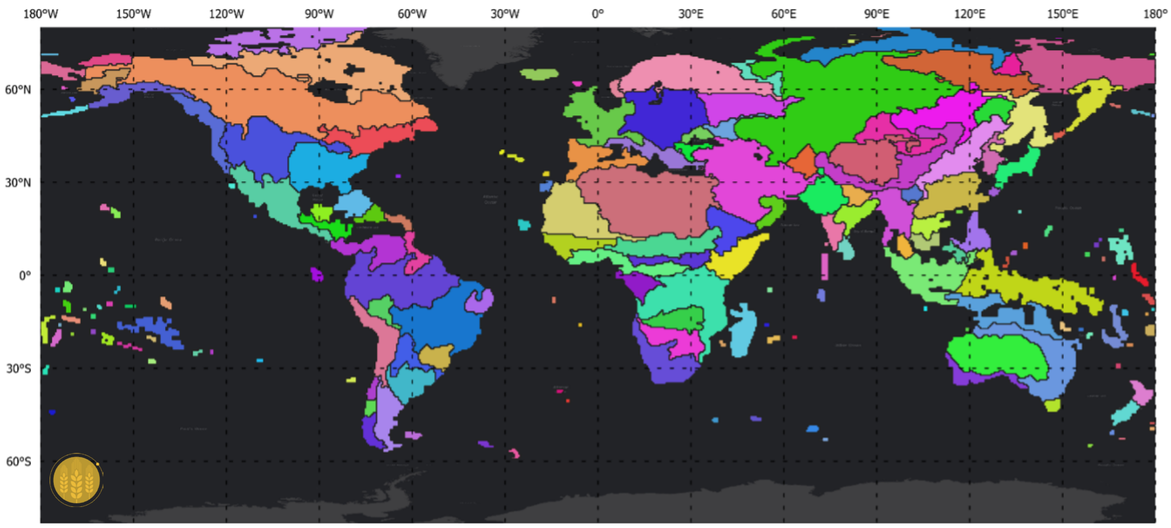

Figure 5: Global stratification based on crop calendar similarity. Each resulting zone is represented by a color in the map, and serves as a WorldCereal map trigger to generate products based on local seasonality.

For each of these zones, based on the crop calendars, different products are created, as shown in Table 6 below.

Season	Product	Remarks
TC-ANNUAL	Temporary crops map	
TC-WINTERCEREALS	Winter cereals map Active cropland map Active irrigation map	Considerd as the main cereals season
TC-MAIZE-MAIN TC-SPRINGCEREALS	Maize map Spring cereals map Active cropland map Active irrigation map	 Only in parts of Northern hemisphere
TC-MAIZE-SECOND	Maize map Active cropland map Active irrigation map	

Table 6: Overview of products delivered per season for an AEZ.

Each AEZ can potentially have 4 seasons for which products are calculated. The first is always the tc-annual season, where the temporary crop map is made. This map is the base layer for all other seasons. Other products are calculated only for those areas where temporary crops are detected.
Next, depending on the AEZ, the products for the tc-wintercereals season are calculated, ie. winter cereals, active cropland and active irrigation. This season is actually considered as the main cereal season for an AEZ. In the northern hemisphere, some AEZs have an additional cereals season defined, i.e. tc-springcereals.

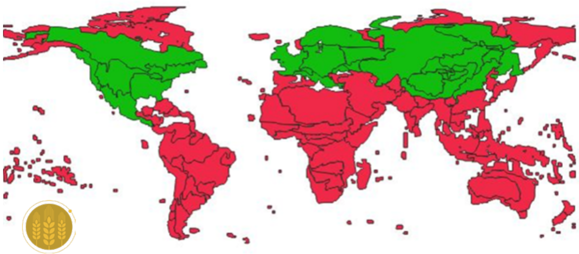

Figure 6: Overview of AEZs where tc-springcereals are produced (green areas)

Figure 7: Overview of AEZs where second maize season maps are produced (green areas)

The tc-springcereals season is in most cases combined with the main maize season (tc-maize-main). For the tc-maize-main growing period, an active cropland and active irrigation maps are also calculated. Finally, there are some AEZs with a second maize season, and for them, a second maize map is created, as well as an extra active cropland map and an active irrigation map.

The definition of these products is detailed in the following sections.

4.1 ANNUAL TEMPORARY CROP MAP

This is the base product generated by the WorldCereal system. The annual temporary crop map consists on a binary map identifying land used for crops with a growing cycle of less than one year, which must be newly sown or planted for further production after the harvest (FAO, 2023).

Sugar cane, asparagus and cassava are considered as temporary crops, despite the fact that they remain in the field for more than one year. The WorldCereal temporary crop maps excludes perennial crops as well as (temporary) pastures.

These annual temporary crop maps are generated once a year, and, for a certain region, the period is defined using the end of the last growing season considered by the system.

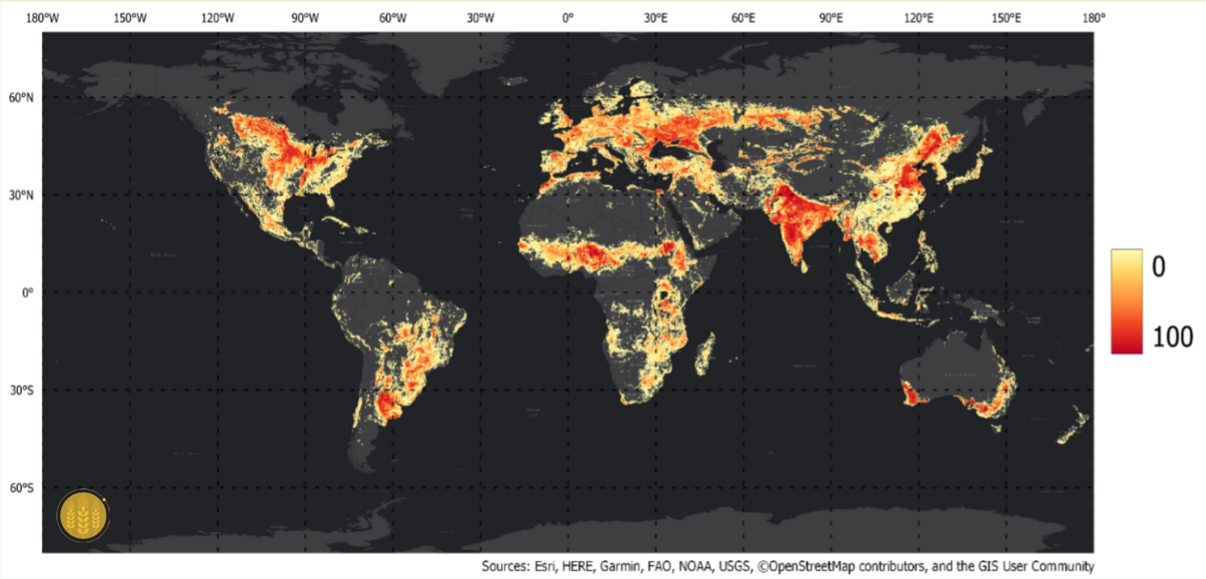

Figure 8: WorldCereal 2021 temporary crop extent map. The original 10m product was resampled to ~0.004° resolution, showing the fraction of the original 10m pixels that were labelled as temporary crops (0 being no crops to 100 all crops). This global overview consists of a mosaic of the individual zones for which the product was generated.

4.2 SEASONAL CROP TYPE MAPS (MAIZE AND CEREALS)

The WorldCereal crop type products provide binary maps for the maize and wheat growing seasons, as defined by the global crop calendars, showing where maize and cereals are grown. Cereals include wheat, barley and rye, all of which belong to the Triticeae tribe. These crops were grouped together because their spectral signatures and growing seasons were too similar to reliably distinguish them at a global scale. The WorldCereal crop type maps are generated within the respective annual temporary crop mask.

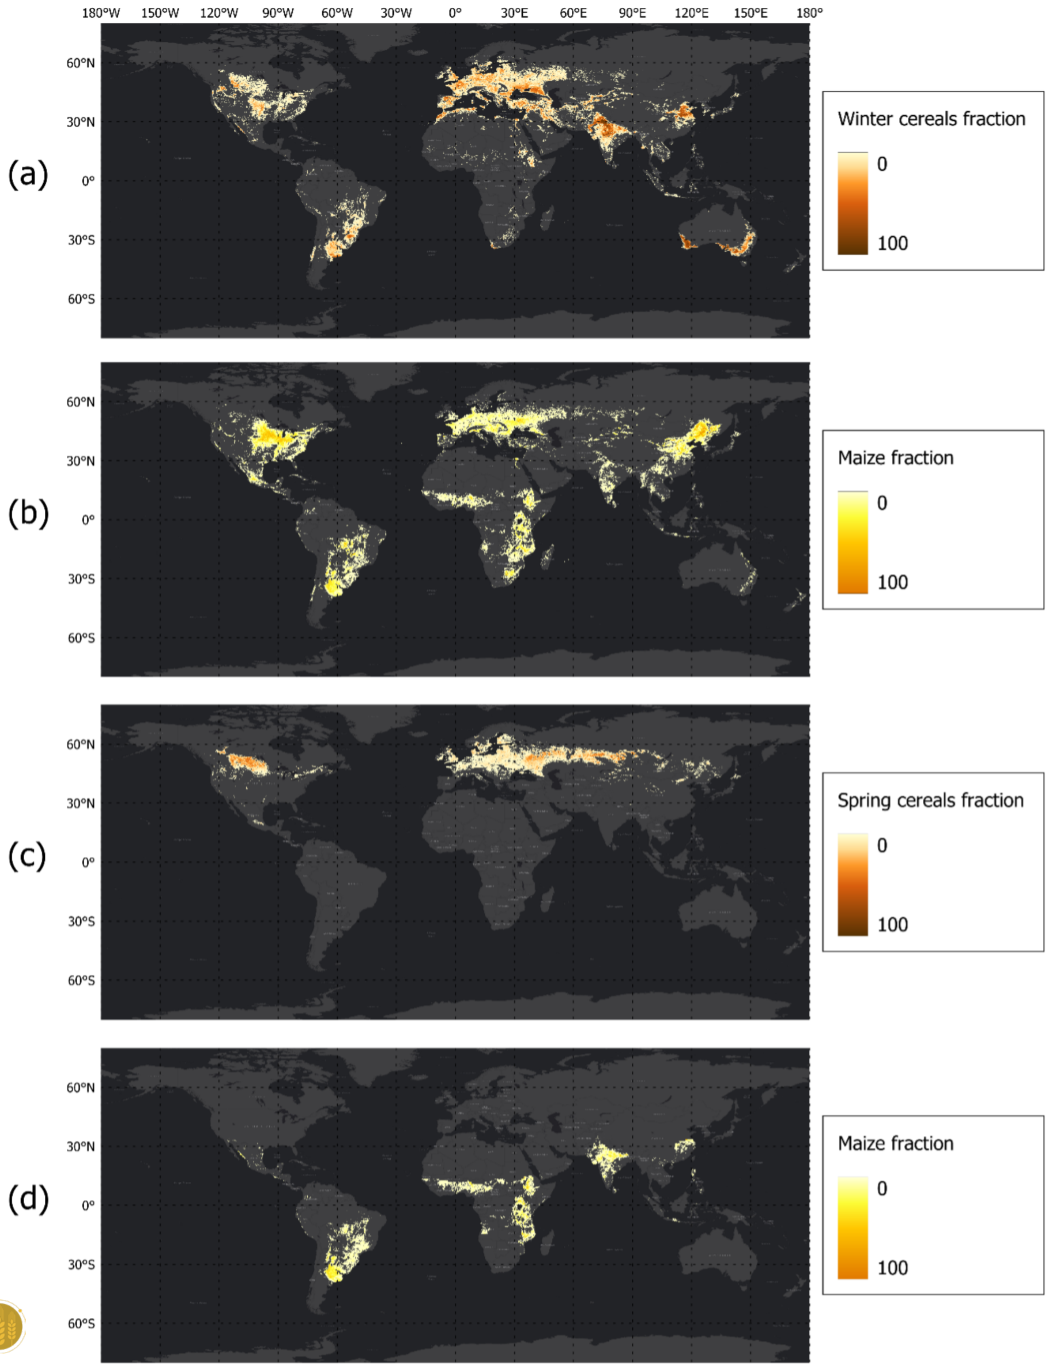

Figure 9: WorldCereal 2021 seasonal crop type products. The original 10m products were resampled to ~0.004° resolution, showing the fraction of land covered by each crop type. These global overviews consist of a mosaic of the individual AEZs for which the product was generated. (a) Winter cereals fraction in the tc-wintercereals season. (b) Maize fraction in the tc-maize-main season. (c) Spring cereals fraction in the tc-springcereals season. (d) Maize fraction in the tc-maize-second season.

4.3 SEASONAL CROP TYPE MAPS (MAIZE AND CEREALS)

The seasonal actively irrigated cropland map is defined by the WorldCereal system as a piece of land that is extensively irrigated during a specific growing season, and where, without irrigation applied at regular intervals, crop growth would be significantly reduced or impossible. Incidental irrigation, such as irrigation that has been applied only during the sowing period of a crop, is not considered in the determination of actively irrigated cropland.

A pixel can only be classified as being irrigated when it is included inside the annual temporary crop mask.

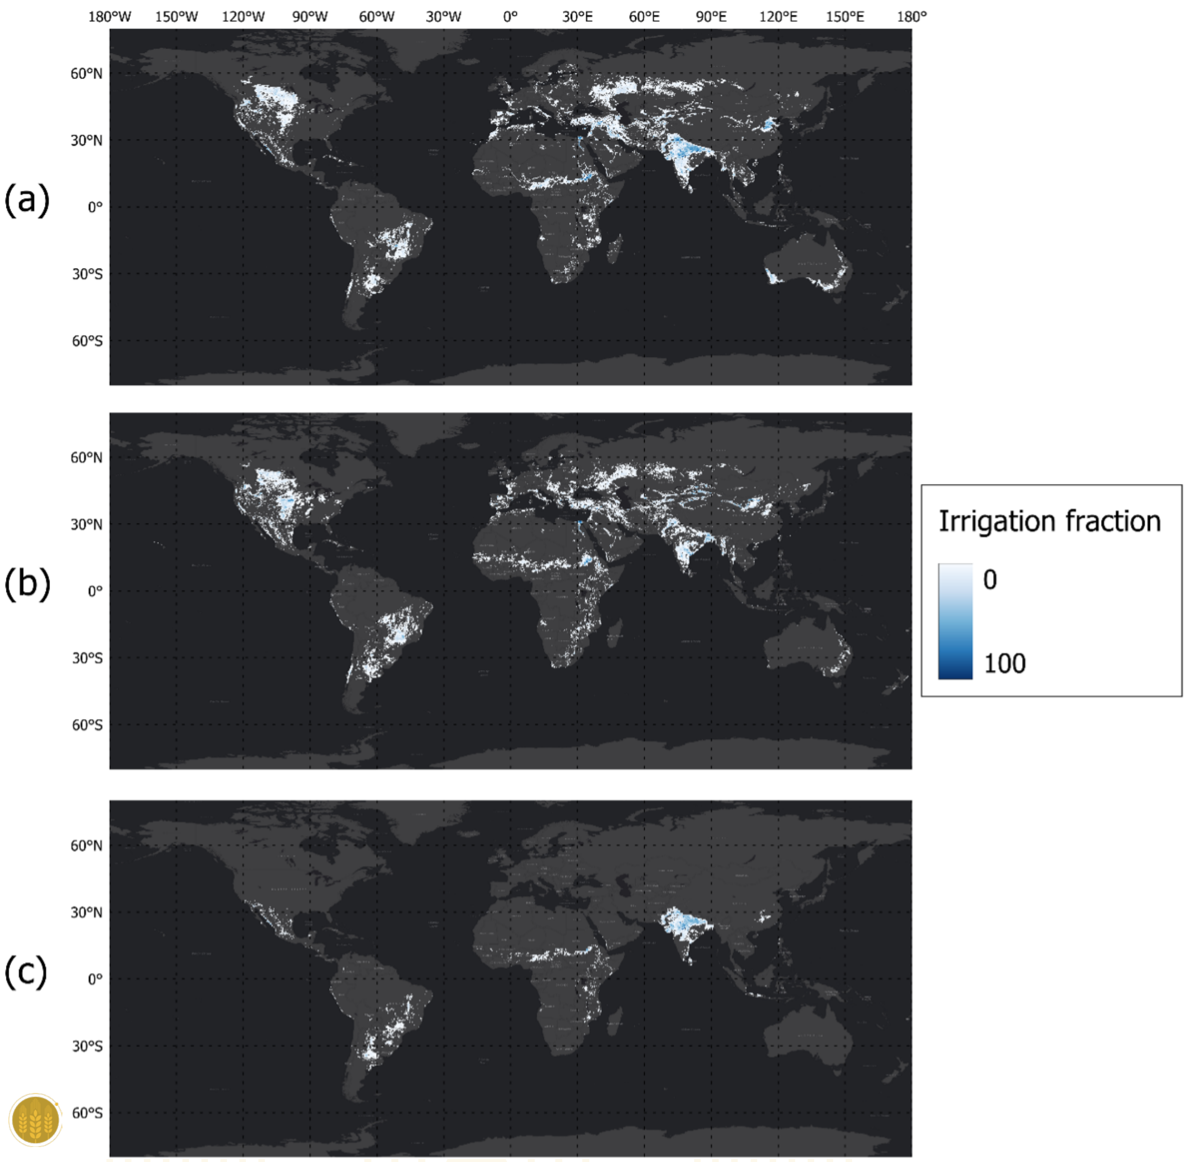

Figure 10: WorldCereal 2021 seasonal active irrigation products. The original 10m products were resampled to ~0.004° resolution, showing the fraction of irrigated land. These global overviews consist of a mosaic of the individual AEZs for which the product was generated. (a) Fraction of irrigated land in the tc-wintercereals season. (b) Fraction of irrigated land in the tc-maize-main/tc-springcereals season. (c) Fraction of irrigated land in the tc-maize-second season.

4.4 SEASONAL ACTIVE CROPLAND MAPS

WorldCereal active cropland products indicate whether a pixel identified as temporary crop has been actively cultivated during each of the mapped growing seasons. In order for a pixel to be labeled as "active" during a particular growing season, a full crop growth cycle (sowing, growing, senescence and harvesting) needs to take place within the designated time period. Note that this active marker is not crop-type specific and will capture other crop types aside from cereals and maize, as long as they show a similar growing season. In practice, this also means that any crop grown (slightly) outside the predefined growing seasons will not be flagged as active cropland in any of the seasons covered by the system.

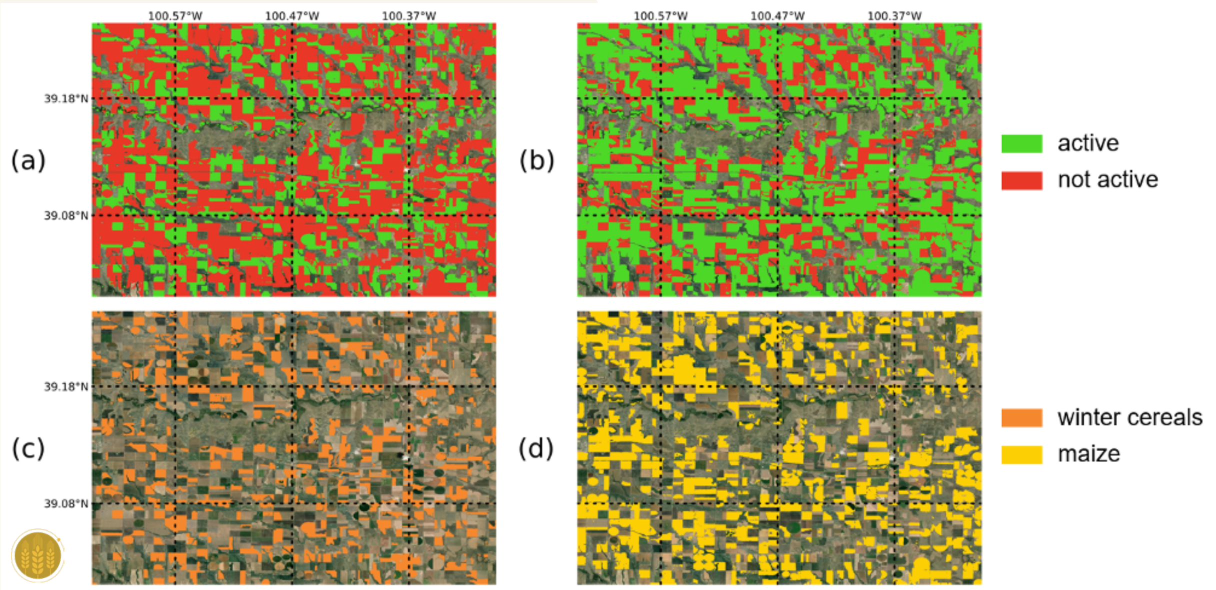

Figure 11: Example of seasonal active cropland maps near Grainfield, Kansas, USA. Active cropland for (a) tc-wintercereals and (b) tc-main-maize seasons show different crop seasonality at parcel level. The (c) winter cereals and (d) maize maps overlap with active cropland for their respective season. Parcels showing up as active cropland but outside winter cereals and maize masks indicate other crops that follow the seasonality for which the respective map was created.

5. WORLDCEREAL PRODUCT ACCURACY

5.1 TEMPORARY CROPLAND

Table 7 summarizes the results of the validation of the annual temporary crop map at global level and by continents. It includes overall accuracy and user's and producer's accuracies for cropland, as well as 95% confidence intervals calculated by applying bootstrapping with replacement. The most informative are user's and producer's accuracies, which are 88.5% and 92.1% respectively, for the globe. Overall, the accuracy numbers are very high for all the continents, while being a bit lower in Asia and Africa. Those two regions are characterised by a significant coverage of very small fields and a very fragmented landscapes. The largest commission errors as well as omission errors are detected in Africa.

Continents	Overall accuracy	95% confidence intervals for overall accuracy		User's accuracy for cropland	95% confidence intervals for user's accuracy		Producer's accuracy for cropland	95% confidence intervals for producer's accuracy	
		lower bound	higher bound		lower bound	higher bound		lower bound	higher bound
Global	97,8%	97,8%	97,9%	88,5%	88,0%	89,0%	92,1%	91,7%	92,5%
Africa	97,2%	97,0%	97,4%	76,7%	75,0%	78,3%	85,9%	84,5%	87,5%
Asia	97,3%	97,2%	97,5%	85,3%	84,4%	86,1%	93,9%	93,3%	94,5%
Australia and Oceania	99,0%	98,8%	99,2%	91,1%	89,0%	93,6%	96,1%	94,8%	98,0%
Europe	97,8%	97,6%	98,1%	96,6%	96,0%	97,2%	92,9%	92,0%	93,9%
North America	98,7%	98,5%	98,8%	95,6%	94,7%	96,5%	93,3%	92,3%	94,3%
South America	98,9%	98,8%	99,1%	95,7%	94,8%	96,8%	90,4%	89,1%	91,8%

Table 7: Summary accuracy metrics with confidence intervals calculated using bootstrapping with replacement

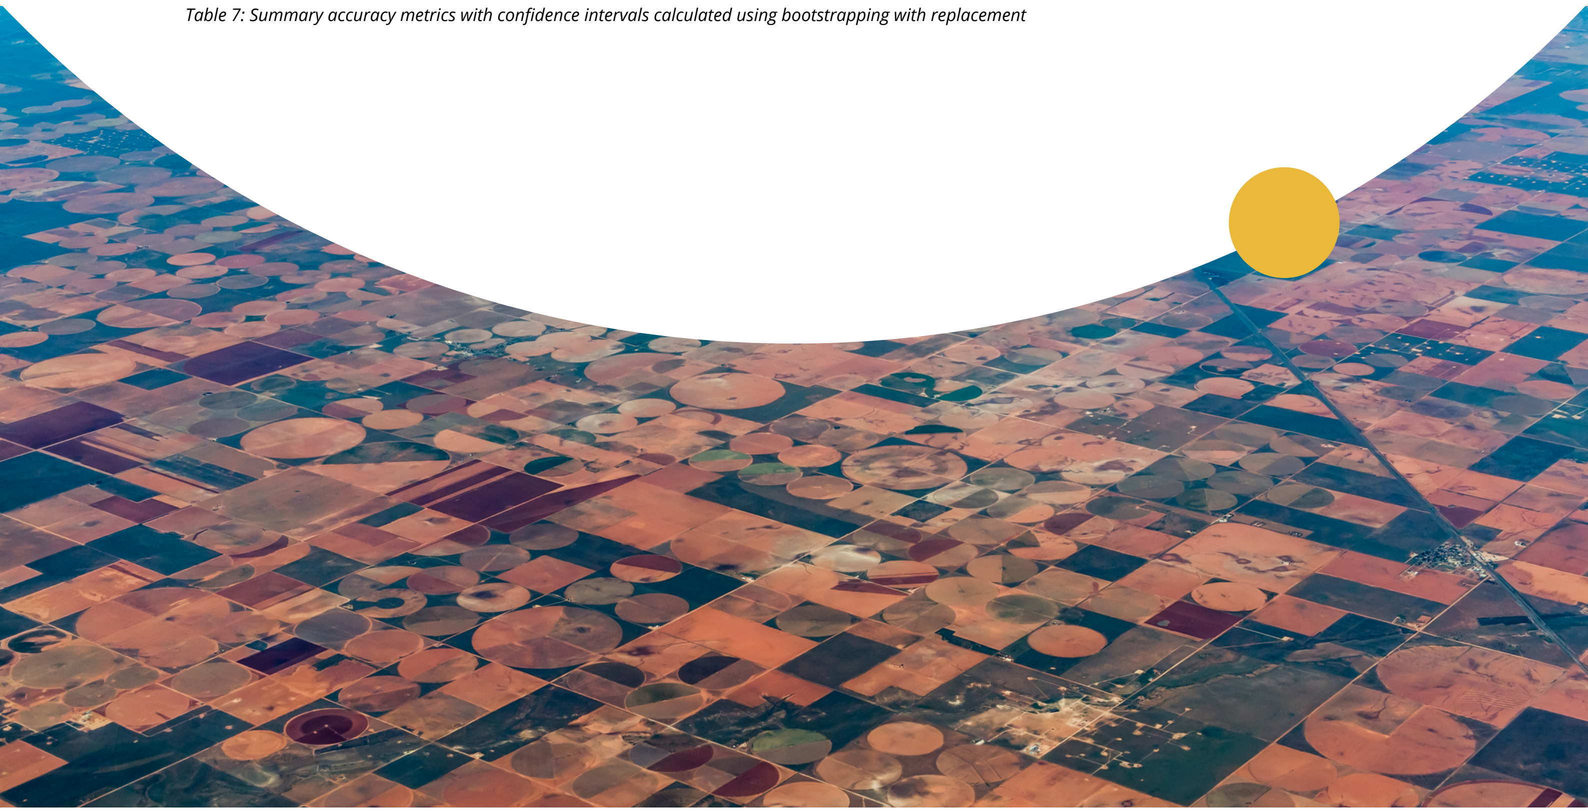

5.2 CROP TYPE

Table 8 shows the results of crop type validation at global level by using the new crop type validation data set collected in the Street Imagery Validation Tool. To calculate the confusion matrix, maize in season 1 and season 2 was combined into one single class "maize", and spring cereals and winter cereals were also combined into one single class "cereals".

Overall, omission errors are larger than commission errors for both crop types, maize and cereals. This could be explained by a lack of training data. It is important to note that the presented results are biased towards the areas covered by the validation data set, e.g. large gaps in Africa.

WorldCereal products/ SV validation data set	Other crops	Maize	Cereals	Agreement by classes (User's accuracy)
Other crops	1010	167	161	75,5%
Maize	90	544	12	84,2%
Cereals	31	10	604	93,6%
Agreement by classes (Producer's accuracy)	89,3%	75,5%	77,7%	
Overall agreement				82,1%

Table 8: Confusion matrix for crop types at global level

6. USE OF THE WORLDCEREAL PRODUCTS

At the start of WorldCereal, three Champion Users (GEOGLAM, AMIS and FAO) were selected, and they helped guide the WorldCereal work and requirements. These Champion Users were selected from across the agricultural sector for their dominant and far-reaching roles and mandates in global agricultural monitoring, representing a diverse set of interests and communities. Additional users were also included to capture regional and national requirements. At the end of the project, products are made available, and these users and the larger community are picking them up for implementation.

6.1 AVAILABILITY OF THE WORLDCEREAL PRODUCTS

WorldCereal products, maps, in-situ data and code are made available to the agricultural monitoring community.

6.1.1 maps

From discussions with the users, it is clear that the WorldCereal maps are the ones that attract the most attention, among all the three freely available products. The maps can be viewed at:
https://vdm.esa-worldcereal.org/

Figure 12: Welcome screen of the WorldCereal Visualization Module

As can be seen in Figure 12, there are different ways to explore the products. It is possible to explore in detail the different production zones (AEZs), by zooming into a certain area. This will give information on the available products for that region (i.e. the different time stamps of the products) and the confidence layer of the product can be displayed.

The detailed exploration viewer is shown in Figure 13. A maximum of 4 windows can be opened for the same area, with each window displaying a different product.
The list of products can be found in the bottom section of the window, where the time stamp of the products is also shown.

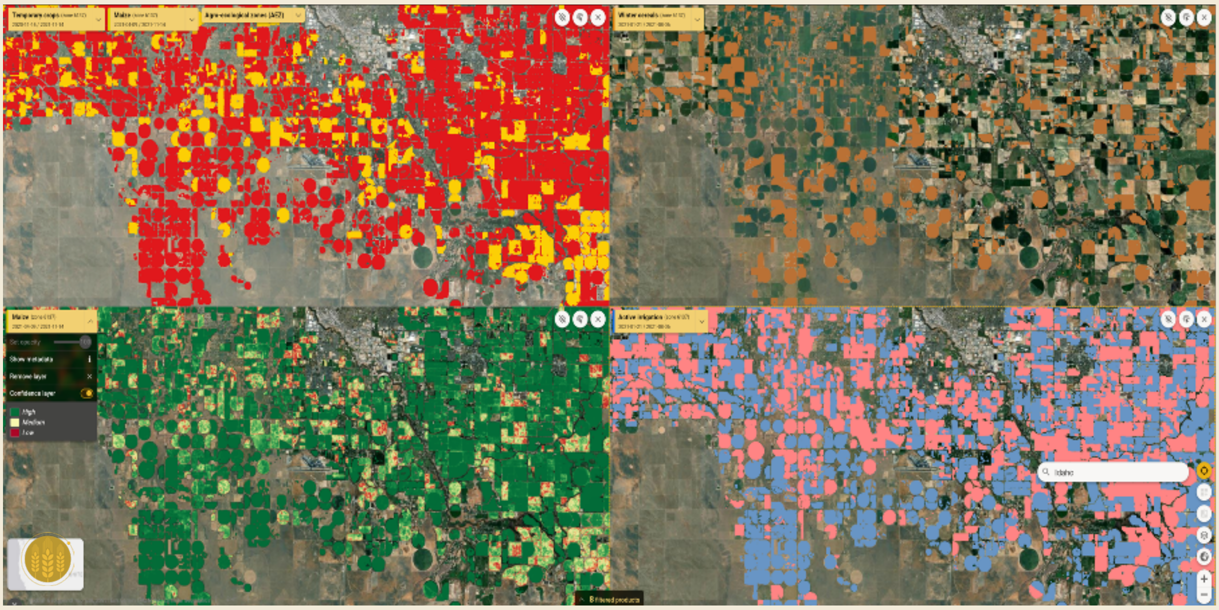

Figure 13: Overview of the WorldCereal products in the detailed exploration mode. The screen can be split in up to four windows. In this example, the top left window has red fields which are the fields classified as temporary crops for Idaho in the US. The yellow field are the maize fields cultivated in the tc-maize-main season. The top right window is for the same area in Idaho, showing the winter cereals detected during tc-wintercereals. The bottom left window displays the confidence of the maize map (shown in the top left). And the bottom right window is the irrigation map for the main cereal season in Idaho.

The global view can display the same products, however, in there, products are merged to fit a global scale and therefore users lose the information on the time domain (i.e. different production dates for the different AEZs). Looking at Figure 14, it is clear that product selection between the detailed viewer (Figure 13) and the global viewer (Figure 14) is different. In the global viewer, the products can only be selected, and they are immediately displayed as global products. Of course, a user can still zoom into specific areas but the information on production date is lost. The main goal of this viewer is to produce fast global products for all seasons together.

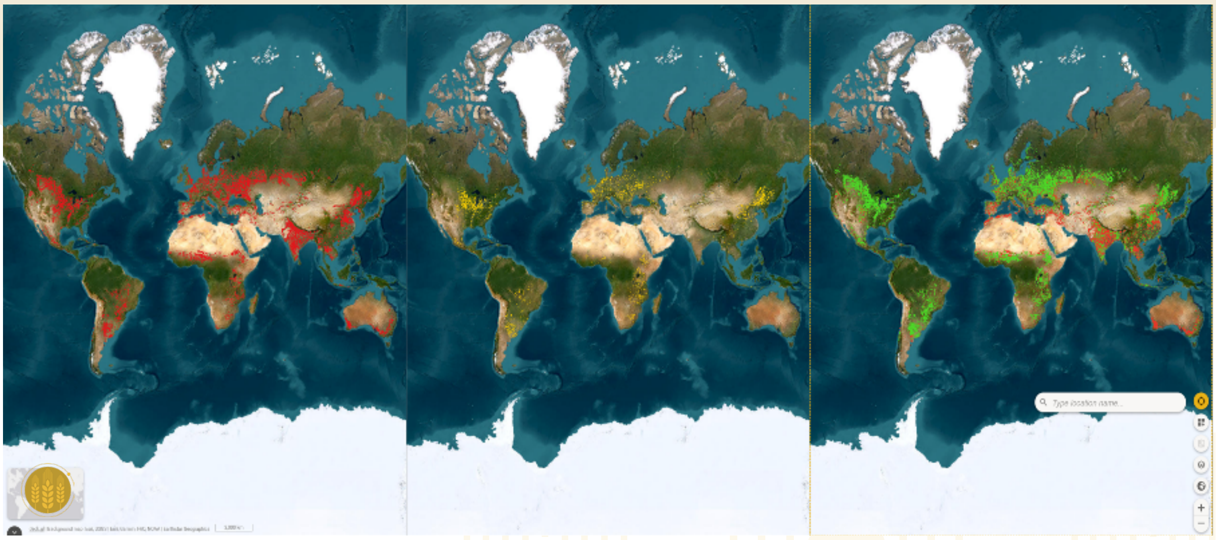

Figure 14: Overview of the WorldCereal products in the global view. The screen can be split in up to four windows (in this example only three are shown). On the left display, the red fields are the fields classified as temporary crops. In the middle one, the yellow field are the maize. In the right display active cropland for the tc-maize-main is shown.

Finally, in the statistics view, the user can extract statistics per country of the different layers. However, one has to keep in mind that these are not official statistics and the information provided here in the viewer is based on pixel counting. It is possible to make a comparison per product and per time period, either at a global level or at country level. At global level, countries are ranked from high to low. At country level, the nuts Level 2 regions are ranked from high to low. An example of global ranking for temporary cropland is shown in Figure 15.

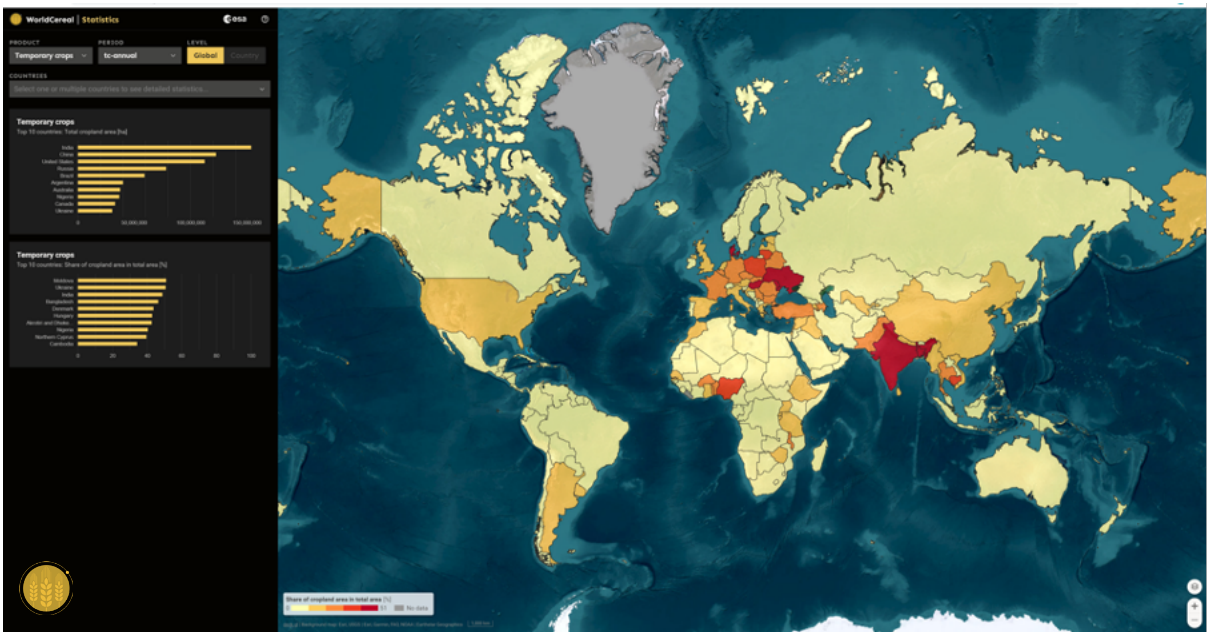

Figure 15: Overview of the WorldCereal statistics viewer. On the left side of the screen, the selection of product and period can be made, and it is possible to choose an analysis at global level or at country level analysis. The right panel displays the selection: in this case, a global analysis of temporary crops where the percentage of temporary cropland vs total area is displayed. The dark red areas correspond to a percentage of around 50%, and the lighter the color, the lower the percentage.

At the moment, products can be downloaded via an FTP site (details below), and products are stored according to their AEZ.

The ftp site login and password are:
cvbftp-02.vgt.vito.be
account : WorldCereal_RO
pwd : Grasland!Waterzooi

6.1.2 In-situ

As mentioned in section 2.2, collecting, gathering and harmonizing in-situ data for training and validation of the algorithms was a huge and important task. The work will be continued in the GEOGLAM in-situ data working group where the WorldCereal in-situ data reference will form the basis on which the rest will be built. The in-situ data is already freely available for the agricultural monitoring community. The data can be discovered in several ways.

The first way to find the data is using the Geo-Wiki-hosted reference data module, available at https://worldcereal-rdm.geo-wiki.org, where users can explore through the different collections visible on a global map. All data and metadata can also be downloaded from the website.

A second way to access the data is by entering the WorldCereal community in the Zenodo data repository, available at https://zenodo.org/communities/worldcereal-rdm/ . The repository shows the harmonized data divided in three parts, each one having its own license and based on the licence of the original data sets. Furthermore, the protocol to harmonize the reference data is also shown. The Zenodo repository also has API access.

And finally, it is also possible to contribute to the in-situ data through the link below. A login and password are needed, to follow up the contribution and data harmonisation.
https://auth.cloud.esa-worldcereal.org/realms/worldcereal/protocol/openid-connect/auth?redirect_uri=https%3A%2F%2Frdm.cloud.esa-worldcereal.org%2Fcb&scope=openid&response_type=code&state=81f9624e73f5a0950d078e75fd93329f&client_id=rdm&nonce=b6b589029cc8c7367854ebd12a6c5e16

The methodology and datasets are in the process of being published as a scientific article (Boogaard et al., in review)

6.1.3 Code

Besides the maps and in-situ data, the code is also made available to the global agricultural monitoring community. This code can be found on the following link:
https://github.com/WorldCereal/worldcereal-classification

6.2 INTEGRATION OF WORLDCEREAL PRODUCTS

A couple of examples of potential usage are presented below.

6.2.1 Integration of the WorldCereal products into the GEOGLAM Crop Monitor Maps & Pie Charts.

The GEOGLAM Crop Monitor is a global initiative that provides timely, easy to understand information on global crop conditions for major commodity and food security crops across the globe. It is based on a collaborative effort between over 40 national and international organizations and Ministries of Agriculture. The initiative uses a variety of Earth Observation data sources, such as satellite imagery, meteorological data, and ground-based observations, to provide accurate and up-to-date information on crop conditions.

The GEOGLAM Crop Monitor produces 3 global monthly reports on crop conditions reports, namely the Crop Monitor for AMIS (launched in 2013), The Crop Monitor for Early Warning (CM4EW) (launched in 2016) and the Global Crop Monitor (launched in 2022). The reports are based on a variety of indicators, including vegetation indices, weather data, and other relevant factors that affect crop growth and development. These are all filtered by cropland and crop type masks, as well as by growing season calendars. The reports also provide analysis and interpretation of the data, highlighting trends and potential risks to crop production.

The GEOGLAM Crop Monitor is an important tool for decision-makers in the agricultural sector because it provides accurate and timely information on crop conditions and yield forecasts, and therefore helps ensuring that food markets remain stable and transparent.

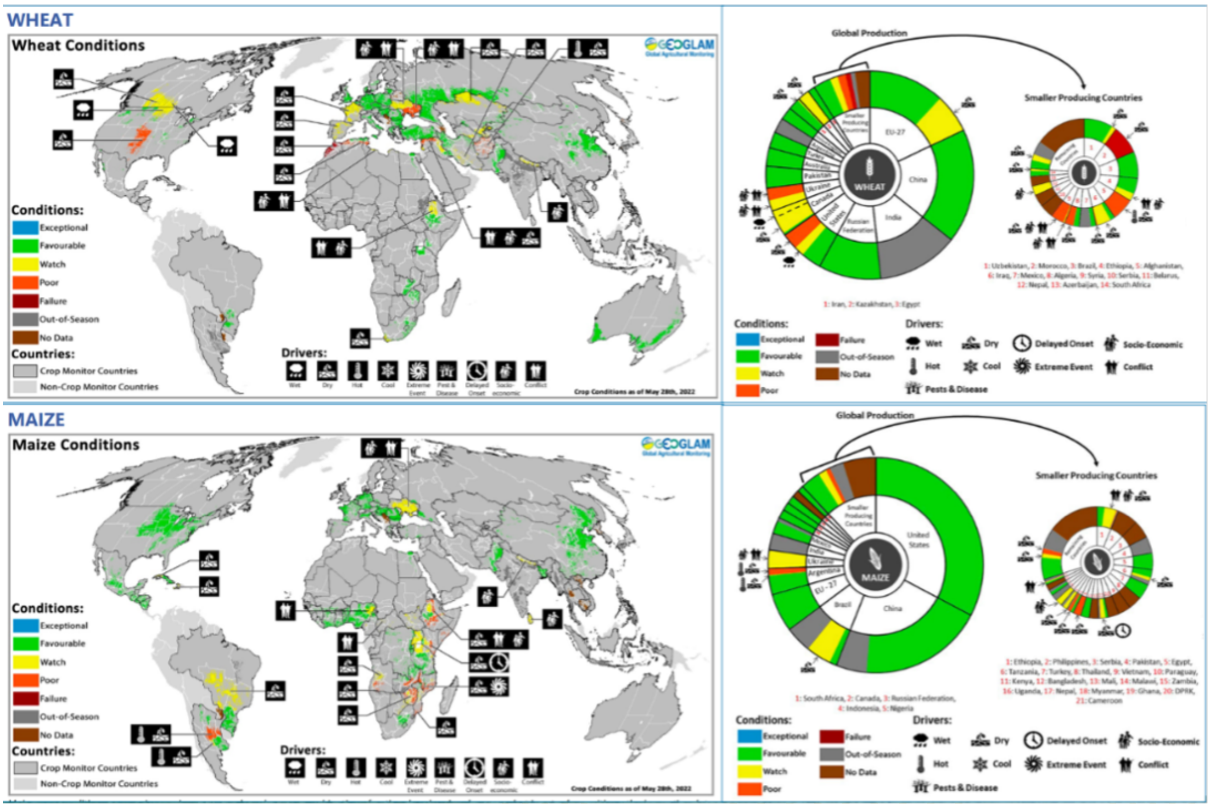

Figure 16: Example of GEOGLAM Maps and Pie Charts from the Global Crop Monitor

One of the key datasets that underpins its systems, analysis and reports are cropland and crop type maps. As such, a key use case is the integration and evaluation of the WorldCereal temporary crop land, cereals and maize masks into the Crop Monitors. More specifically, this means integration as a base layer into:

- The 18 maps that are produced operationally on a monthly basis and included in the various Crop Monitor reports.
- Recomputing the areas under the various segments of 18 pie charts that are produced operationally every month, which reflect the proportion of croplands under the various crop condition categories, within each country covered by the Crop Monitors.

This is expected to have a major impact, particularly in countries where the crop masks are out of date or of low quality. Specifically, in the CM4EW, this will be a huge improvement as currently it utilizes a general cropland mask rather than crop type masks. Maize and wheat are both key crops in many of the countries that the CM4EW covers.

The Crop Monitors where World Cereal data will be integrated are:

1. Global Crop Monitor
2. Crop Monitor for AMIS
3. Crop Monitor for Early Warning
4. Regional East Africa Crop Monitor

6.2.2 Integration of Worldcereal products into FAO platforms and services.

The FAO GIEWS team is revamping their in-house crop monitoring analysis capabilities and systems. As part of this the WorldCereal Cropland map will be integrated as a base layer into the Agricultural Stress Index System (ASIS). This system is routinely used by the GIEWS Crop Analysts and is designed to allow easy identification of areas of crop land with a high likelihood of water stress. Analysts use these indicators as input for their crop condition assessments and for their crop production forecasts across the globe.

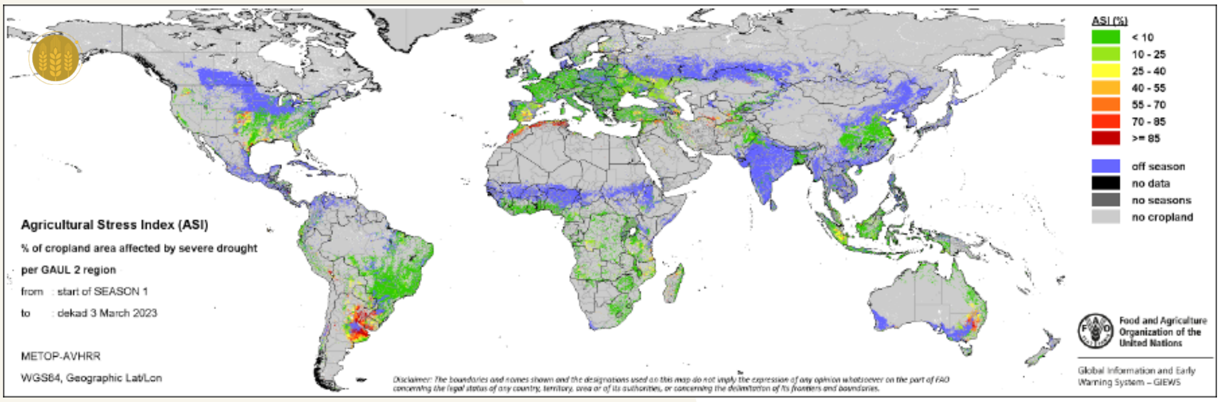

Figure 17. GIEWS ASIS Dashboard

Alongside this, there is also a huge potential to integrate WorldCereal products into the WAPOR database of FAO https://wapor.apps.fao.org/home/WAPOR_2/1.

This database is open and free, and provides information on water consumption and biomass on a near real time basis. Besides, meteorological data, a key layer of this database is the land cover layer, which at the moment uses the Copernicus Global Land Product. This layer is used to extract the agricultural class, and is improved with extra information on irrigated or non-irrigated agriculture based on rainfall deficit information.

It is clear that a seasonally update product (rather than the annual product) would significantly improve the system.

7. ROADMAP

A final task for the WorldCereal project was the development of a roadmap for future improvements of the system. This roadmap is developed based on the discussions that were held with the project's Champion Users during product development and evaluation.

WorldCereal demonstrated the possibility of global seasonally updated crop type classification at field scale. The current system has surpassed the minimum requirements as set by ESA and the user community, but requests for improvement have been collected throughout the process.

A first step in the roadmap would be to make the system even more flexible and light. The WorldCereal system was developed as a modular, open-source system, which allows stakeholders to pick and choose parts to improve, but its main focus was global production. As such, the system as it is currently not easy to install.

One would need a lot of IT background to install it and run it. Furthermore, since it was developed as a global system, the resources to be deployed on a cloud environment are not small, significantly limiting the number of users that could actually implement it. Hence, a first step would be to update the system, so that it becomes lighter and more flexible, and a possible solution for this could be its integration into OpenEO.

A second step would be to set up dedicated trainings for global users as well as for regional, national and sub national users. All of them would be using the same system, and therefore the integration of the products at those scales into a global product would be greatly simplified. This approach could be of significant help to GEOGLAM and the GEOGLAM Crop Monitors.

The integration of new crops of global importance is also part of this roadmap. It could be conceived to include perennial crops (and not only temporary crops). Paddy rice, sugar cane and soy are all possible crops which have a global impact and could be added. Integrating new crops will also have an impact on the production steps and on the production zones of WorldCereal. In close collaboration with international players like NASA Harvest, JRC and FAO, these could be adapted to include new and updated crop calendars.

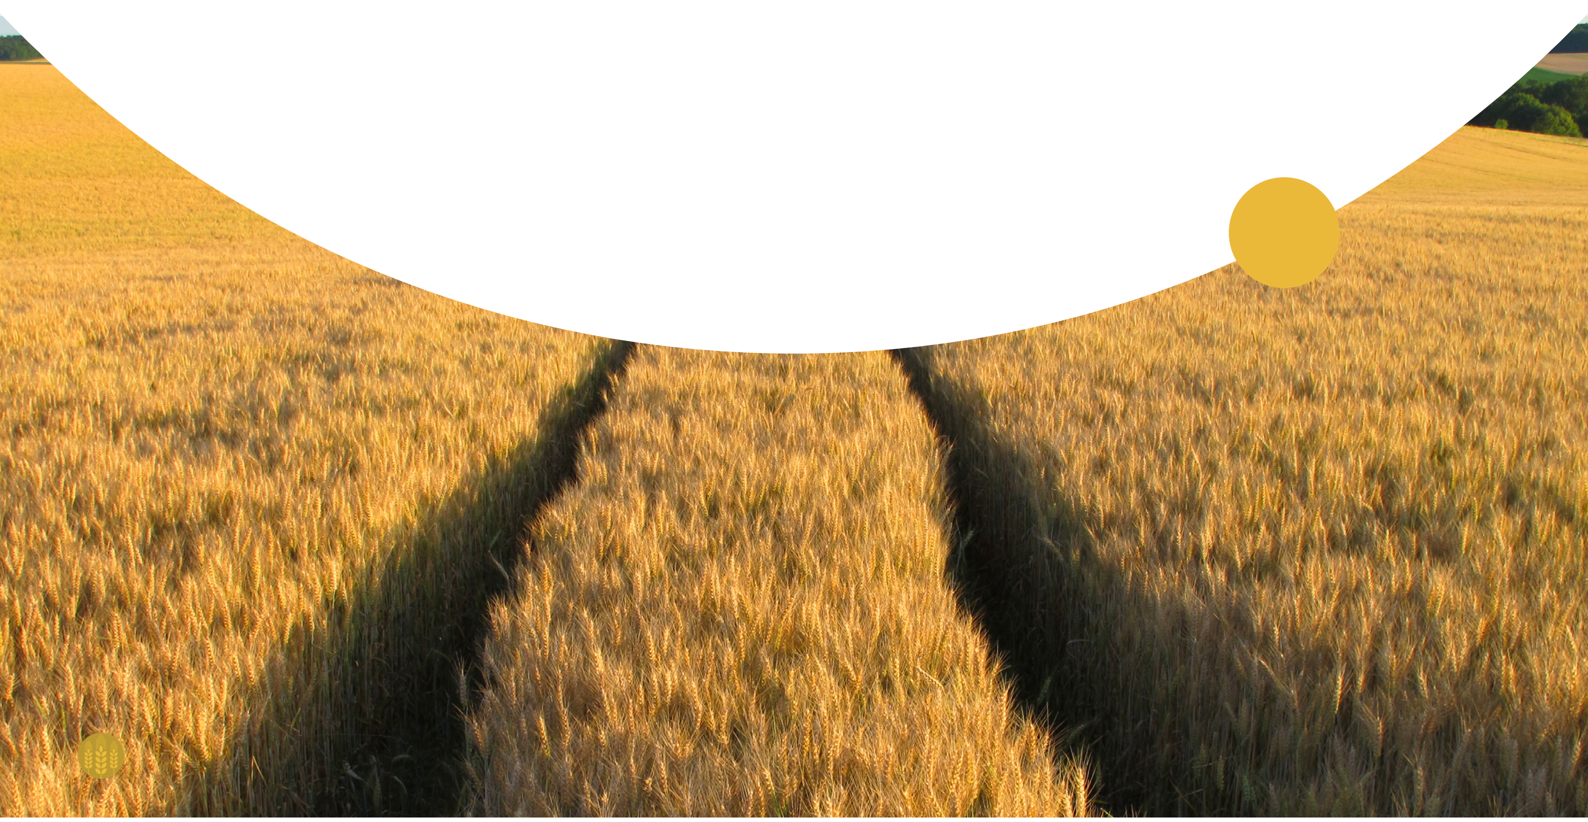

A next step would be to expand the use of the products. The WorldCereal products could potentially support the FAO global statistics. For now, the statistics part in WorldCereal is based on pixel counting. This is not the correct approach for agricultural statistics, and it only allows for a quick comparison between countries and regions.

The consortium was in contact about this with FAO's GIEWS and FAO's Land and Water Division, as well as with FAO's Statistical Department. Increasing such collaborations with FAO could help develop a system that FAO can directly integrate into theirs and is hence an important element in this roadmap.

Improvement of the existing products would also be a step in the roadmap. As mentioned throughout this report, the requirements set out at the beginning of the project are largely met. However, we know from the confidence layers of the products that not all areas in the world are performing at the same level. There can be different reasons for this: field sizes, different agricultural practices, cloud cover, lack of training data, lack of EO data, resolution of the EO data, problem in the defined crop calendars, among many other aspects, as well as combination of several of these reasons.

Until now WorldCereal has focused on the production stages, but in a next phase some more detailed research will need to be done in those areas that show lower confidence levels, to see how they can be improved.
For this we will need to make use of local knowledge and to reach out to local and international experts.
It will also be important to make use of the research networks that have been built within the project, ensuring a smooth integration of national products into the global scale and vice versa.

8. CONCLUSION

There is a lot of excitement in the international agricultural monitoring community for the WorldCereal products. The WorldCereal products are largely under various steps of evaluation and integration into the systems of the various Champion Users, with positive feedback so far.

The main overall feedback is that there is a strong interest for these products to continue to be produced, supported and improved with more onsite training and capacity development. More specifically, there is an interest to focus on improving the products so they can be useful not only at the global level, but also down to the national and sub-national levels. Therefore, the improvements would focus on the integration of ground data (particularly in areas with lower accuracies and where little to no ground data was available in Phase 1 of the WorldCereal project).

The sentiment is that there is an urgent need for these seasonally updated products, and there is an interest from the UN and GEOGLAM community to contribute to enhancing these through further integration of ground data and expertise.