TETRA TECH

Enabling Crop Analytics At Scale

ENABLING CROP ANALYTICS AT SCALE (ECAAS)

Next Generation Crop Production Analytics

Dynamic Area Sampling Frames for Improved Crop Analytics

Contents

List of Figures

Supporting Documents

1

Background

Background

1.1 Project Overview

Generating accurate and season-specific crop analytics such as statistics on crop area, yield, and production is essential to enable decision-making by various actors, ranging from governments and private companies to smallholder farmers.

However, collecting in-field agricultural statistical data is expensive, slow, and subject to error due to methods that insufficiently recognize spatial-temporal patterns regarding the 'where' and 'when' specific cropping occurs. These inaccuracies make scaling up site-to-area or site-to-landscape estimates problematic because a sampled site must represent a known population, i.e., all map units or strata put to a specific agricultural cropping system. Currently, it is challenging to find landscape level cropping system maps that clearly indicate the classification and terminology of the system used in reporting agricultural statistics. This remains one of the key challenges to achieving crop analytics at scale.

To address this gap, we propose a method to improve effective sampling strategies at a landscape level. This research strongly conveys the message that collecting site-specific data must represent known, relatively homogeneous strata at a landscape level. Information gathered can then be scalable to known strata and beyond. The approach presented here:

- Delineates homogeneous land use and land cover strata to support stratified sampling required to generate crop area estimates
- Generates dynamic strata that present relevant spatial differences in season-specific crop system performance and that function to extrapolate season and site-specific yield measurements to zonal and regional crop production estimates.

By strategically locating fields where and when the value of ground-truthing data is the largest, this method will provide major accuracy gains in generating crop area and crop type production statistics with less fieldwork than traditional methods. Further, combined with the 3D imaging-based yield estimation method developed in the Work Stream #2 of this project, we anticipate that these two complementary approaches can synergistically contribute to the scaling of crop analytics.

Our method focuses on processing remotely-sensed imagery in order to produce a homogenized strata map. We integrate existing secondary data sources on crop-area statistics (by administrative areas) and information on crop calendars into the maps. The focus of this analysis is on rice systems in Odisha, India. The resulting analysis represents groundbreaking improvements in accuracy for landscape-level crop area and productions statistics, building on the Area Frame Sampling (AFS) method (Box 1).

BOX 1:
The sample was unbalanced and unrepresentative

The main finding was that the small number of samples for rice and other crops relative to maize, and the clumpy geographic distribution of the sample, made it hard to develop a robust mapping model, and complicated the ability to objectively assess map accuracy. The class imbalance and lack of geographic representativeness were due in large part to logistical reasons described in the two subsequent findings.

Recommendation: Recent country-scale work found that 2,000-3,000 samples per crop class are sufficient for achieving close to the maximum possible model performance (Azzari et al, 2021). Our mapped region was <10% of the area in that study, thus the size of the maize sample in our analysis may be sufficient, but future efforts should focus on increasing samples from the other crops so that they are closer in size to that of maize. Increasing the spread of the sample across the entire mapping region is also critical. We, therefore, propose an updated sampling design that combines drone-based sampling with ongoing field efforts (see scaling-up plan in D2.5), which can increase both class balance and geographic coverage.

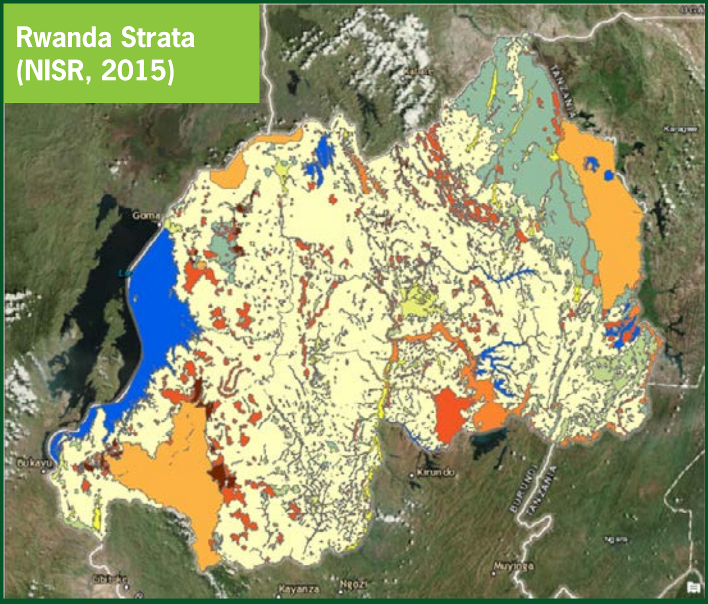

Intensive cropland (50-100% cultivation)
→ 61% of the country
Extensive cropland (15-50% cultivation)

1.2 Situating the Analytical Approach

When creating sampling schemes, it is important to use a benchmark map showing all relevant and different "statistical populations" that are pertinent to the generation of agricultural area and crop production statistics. Crop area or production-related statistics are typically missing when following traditional sampling schemes organized purely on an administrative area basis (often through list-frames). Here we advocate for the use of the Area Frame Sampling method[1] to fill this gap. The approach proposes developing sampling agricultural areas within administrative areas, but does not fully integrate remote sensing-based options to create crop production system zones (CPSZs)[2] and land use and land cover (LULC)[3]-zones that represent the populations surveyed.

Each sampled administrative area consists of a unique mix of CPSZs. Creating more homogenous strata will reduce inaccuracies when scaling up site-level data to area-level statistics. Currently, however, most surveys attribute equal weight to each sample in the respective administrative area, thereby missing differences at the landscape level within the surveyed administrative area. As an illustrative example, Box 2 shows the extent of spatial heterogeniuty within strata in Rwanda, analyzed by our team in a previous study. Our analysis showed an incredible variation within the 50-100% cultivation mask which is otherwise considered homogeneous by the National Institute of Statistics Rwanda.

Once site-specific agricultural area statistics are collected, all sampled sites representing a specific Enumeration Area (LULC-zone) can be used to make season-specific legends on the EA-map (e.g., by updating the area fractions cropped to a specific cropping system within an EA). At any time, the data in that map can contribute to tables that contain summarized statistics by different analytical units such as blocks, administrative areas, or regions (using area and fraction weights based on the extent of each LULC-zone present).

Sampling data must also account for changing weather patterns. Agroecological zone-(AEZ[4]) and CPSZ-based zoning account for climate but are relatively static and do not capture season-specific performance differences due to weather-specific anomalies frequently occur within and across zones. The Long Term Normal (LTN), is another typical weather measurement that presents seasonal differences in rainfall at regional levels (at ± $7km^2$ grids). The weather anomalies occur mostly following patterns of larger weather systems such as El Niño Southern Oscillation (ENSO). Differences in landforms and terrain only marginally influence such patterns (**Figure 1**).

While crop performance can be affected by the severity of large-scale weather anomalies, the local aspects of terrain, soil, and land management are far more significant (**Figure 2**). Therefore, performance indicators must include impacts of terrain, soil, and land management. One such indicator is NDVI, a widely accepted land cover greenness metric representing the performance of cropping systems and other land cover classes. For this work, we focus on anomalies in the response of systems rather than anomalies in inputs; the latter is, however, highly relevant regarding the production of timely performance predictions (**see Annex 3**).

NDVI anomalies can be overlaid on static CPSZ-maps (representing current land use) to create a Dynamic Sample Frame (a season-specific dynamic area frame, DAF). That DAF represents an ideal solution to scale up site-specific yield data from various sources such as Crop Cutting Estimate surveys (CCE surveys) to crop production estimates by area or region. Typically, a DAF does not guide sample schemes but instead creates a layer that presents a season-specific stratification on crop performance.

[1] FAO recommends the use of remote sensing data to improve the Area Frame Sampling method. Here we propose a method to integrate Earth Observation products to enhance area frame sampling.
[2] Crop Production System Zone (CPSZ): A zonation defined by actual clustering of long-duration NDVI profiles. A CPSZ at landscape level is characterized by its climate, landform, terrain, soil, land-cover and land-use. An NDVI-profile reflects the performance of vegetation as influenced by these factors. Vegetation greenness over time captures the integrated effect of the above on the temporal active chlorophyll density, species-mix (floristics), phenology, structure, growing season, cropping calendar, crop-management, etc.
[3] LULC-zones = Land Use Land Cover (or Land use systems) zones.
[4] Agro Ecological Zone (AEZ): A zonation defined by climate, terrain and soils.

BOX 2:
Improving existing Strata (Enumeration Areas)

Within the intensive cropland area (50–100% cultivation) as used for area sampling by the National Institute of Statistics Rwanda (NISR), we identified many different NDVI-strata. We surveyed four of these in detail (Classes 24, 54, 70, and 82). Results showed that the 4 classes (strata) had considerably different crop mixes and cropping calendars. Accordingly, the NISR-mask requires further stratification to create EAs that have relatively homogeneous "cropping system" populations.

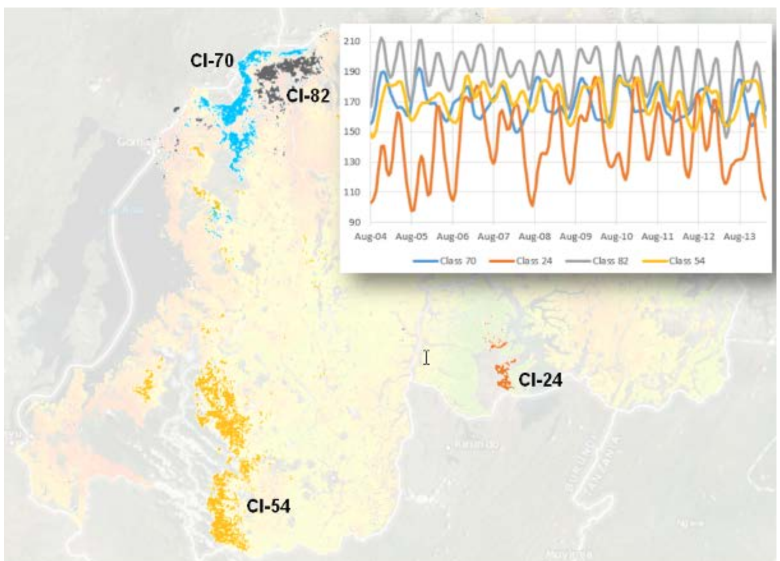

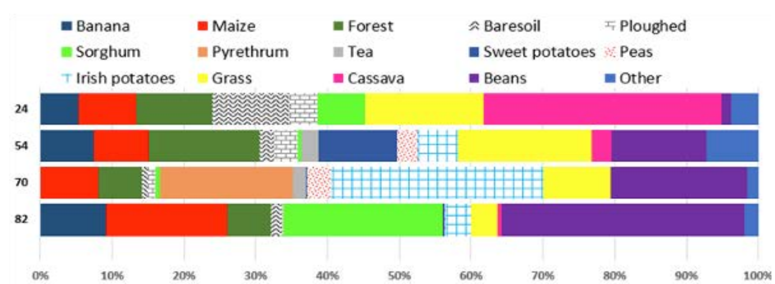

Differences in land-cover during Sep.'15

(Averages of 12 surveyed segments by strata; Modified from Mugabowindekwe, 2016)

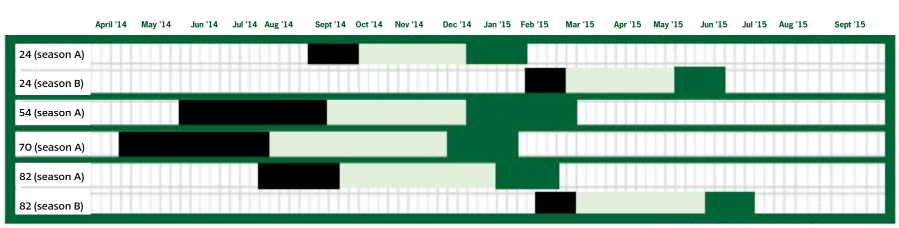

Differences in practiced maize crop calendars 2014/15

(Median data of 12 surveyed segments by strata; Muyizere, 2016)

Planting Period
Growing Period
Harvesting Period

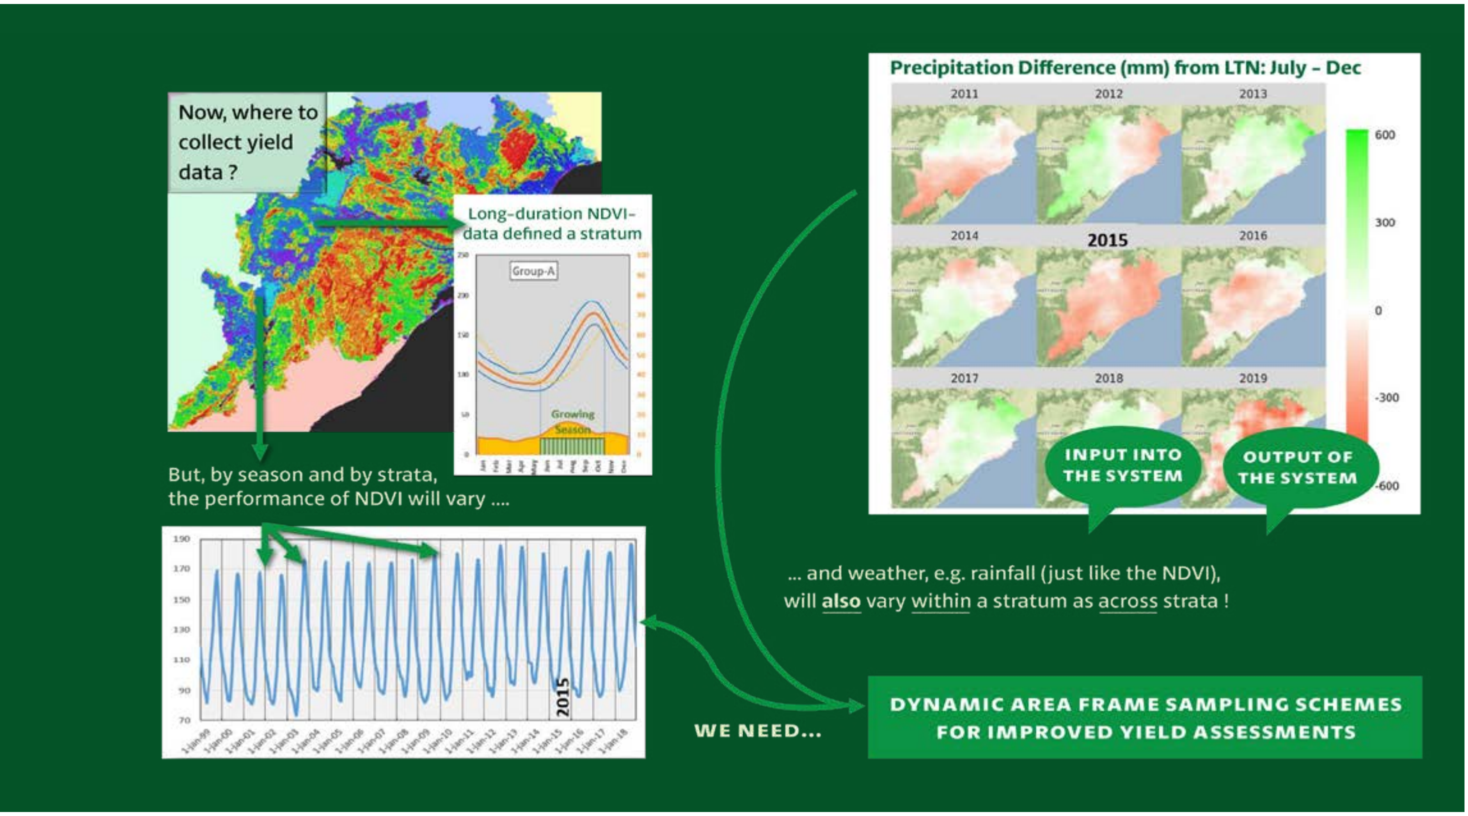

Figure 1:
Dynamic frames can be generated using rainfall or NDVI data. Rainfall data have greater predictive capabilities at larger geographic scales, while NDVI is more useful at sub-national scales.

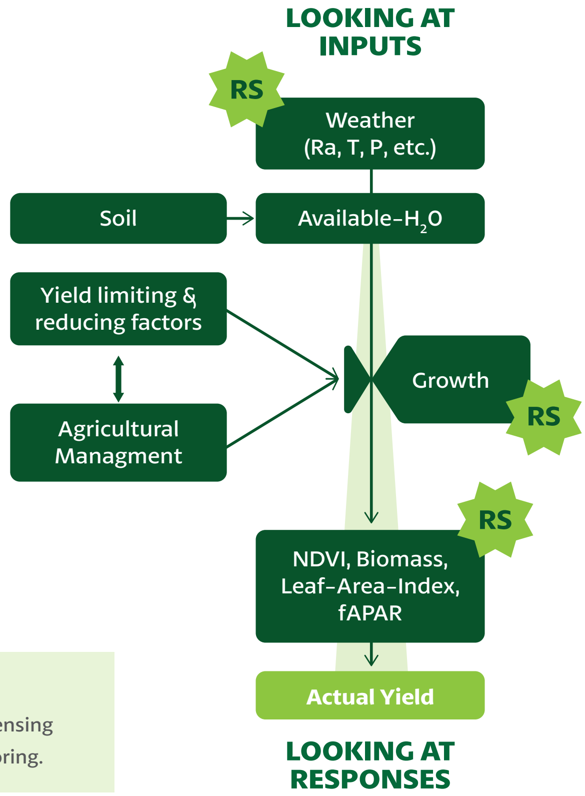

Figure 2:
Logic of general options for remote sensing (RS-) based crop performance monitoring.

2

Method of Analysis

This section outlines key steps performed for our development of dynamic area frames (a graphic summary is provided (**Figure 3 and 4**). The full process is not standardized through fixed code because national-level published land use statistics differ. Analysis must thus remain flexible, following the framework described below (steps 1-5).

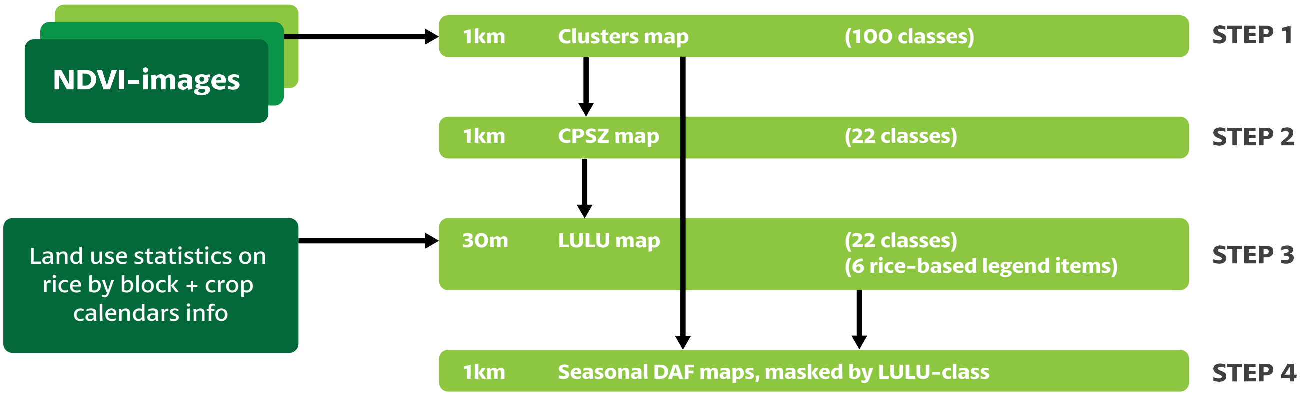

Figure 3:
General overview of processing steps.

Step 1: Classify 1km^2 resolution NDVI time-series into NDVI-clusters

The first step in this method consisted of developing a zonation of the entire state of Odisha through clustering pixel-specific long-duration (20-years) hyper-temporal NDVI profiles (greenness-profiles) obtained from Spot-VGT and Proba-V imagery (2000-2019). This process was achieved using cleaned (quality flags) and upper-envelop processed[5] sequences of dekad-specific NDVI-imagery. From that stack of imagery, we derived many NDVI-classes using the ISODATA[6] unsupervised clustering procedure as available in Erdas-Imagine. This process resulted in the simplification of a 3D NDVI data-cube into a 2D map that included a temporary legend of class-specific NDVI profiles. These profiles feed into subsequent analyses. We assume that clustering the 3-dimensional data cube did not create loss of any relevant information.

[5] This related to upwards NDVI-adjustments using an iterative Savitsky-Golay smoothing process.
[6] ERDAS Field Guide, 2005 Leica Geosystems Geospatial Imaging, LLC.

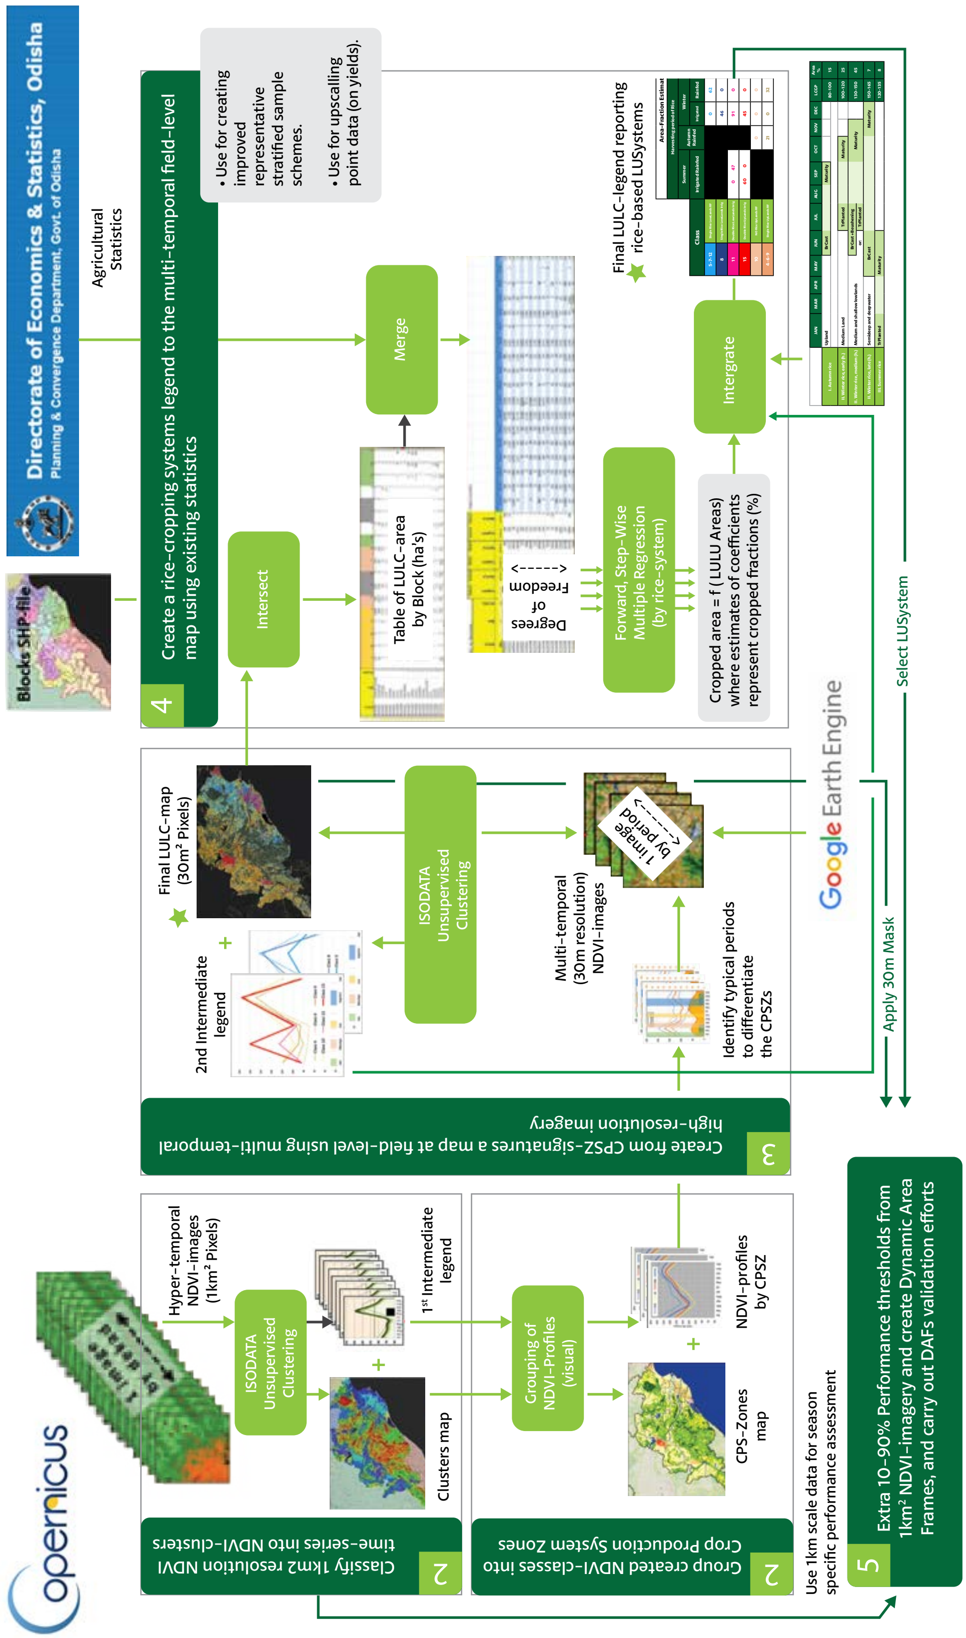

Figure 4:

Overview of the key processing steps performed.

Step 2: Group NDVI-classes into Crop Production System Zones (CPSZs)[7]

Using a combination of statistical analysis and visual grouping, we further reduced the data-cube to a map containing only 22 Crop Production System Zones (CPSZs). Automatically creating multiple clusters, and then reducing them into groups that show clear differences (as related to cropping calendars practiced), has proven to be the easiest and most efficient method for grouping zones. These zones consist of 1km^2 pixels with relatively similar NDVI profiles and are assumed to consist of homogeneous land cover and land use patterns. Across CPSZs, these patterns will differ substantially (**Figure 5**). Each set of NDVI profiles in the CPSZs reflects the average temporal performance of vegetation from January to December, as recorded through the amount of active chlorophyll present. The performance of vegetation is defined by the standing land cover species-mix, its phenology, density and structure, aspects of the growing season, cropping calendar, as well as various aspects of crop and land management.

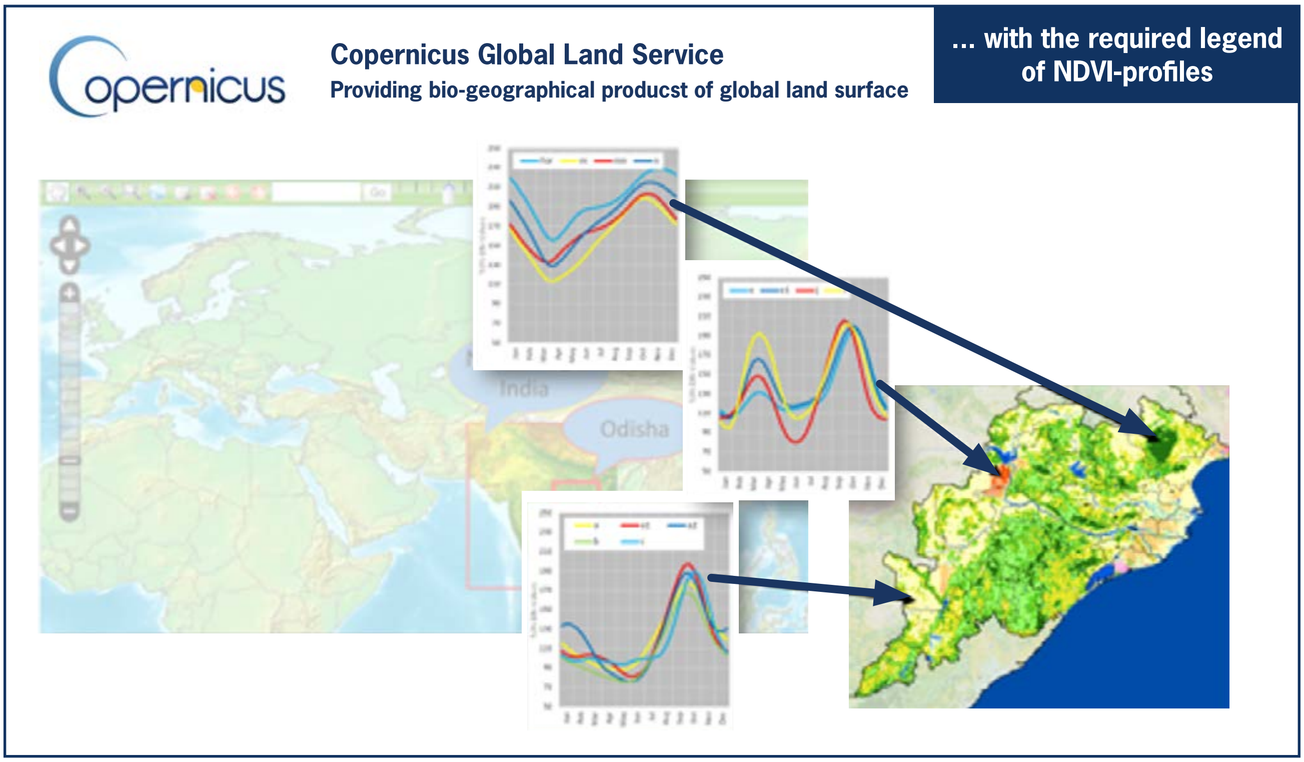

Figure 5:
Grouping of NDVI profiles created clear Crop Production System Zones

[7] Refer to Annex 1: Creating Crop Production System Zones Technical Note– for additional information

Step 3: Use CPSZ profiles to create a map at field-level using multi-temporal high-resolution imagery[8]

Based on NDVI profiles of the 1km² NDVI-based strata (CPSZs), we identified four periods during a typical year that provide sufficient information to allow differentiation between the created CPSZs (**Figure 6**). For instance, some CPSZs have one distinct period of higher greenness at the end of the rainy season, while other zones have two peaks in greenness during the year. For these four periods, we extracted median NDVI data using TM8-imagery available for each period over the past five years (see annex-3 for the complete GEE code). The resulting NDVI images were then stacked and classified into LULC-clusters[9]. This exercise resulted in a higher-resolution map (30m) where fewer pixels consist of mosaics of land cover types, meaning most pixels relate better to one specific agricultural land use or land cover. As with the coarser 1km product, the higher-resolution map included 22 LULC-clusters (**Figure 7**).

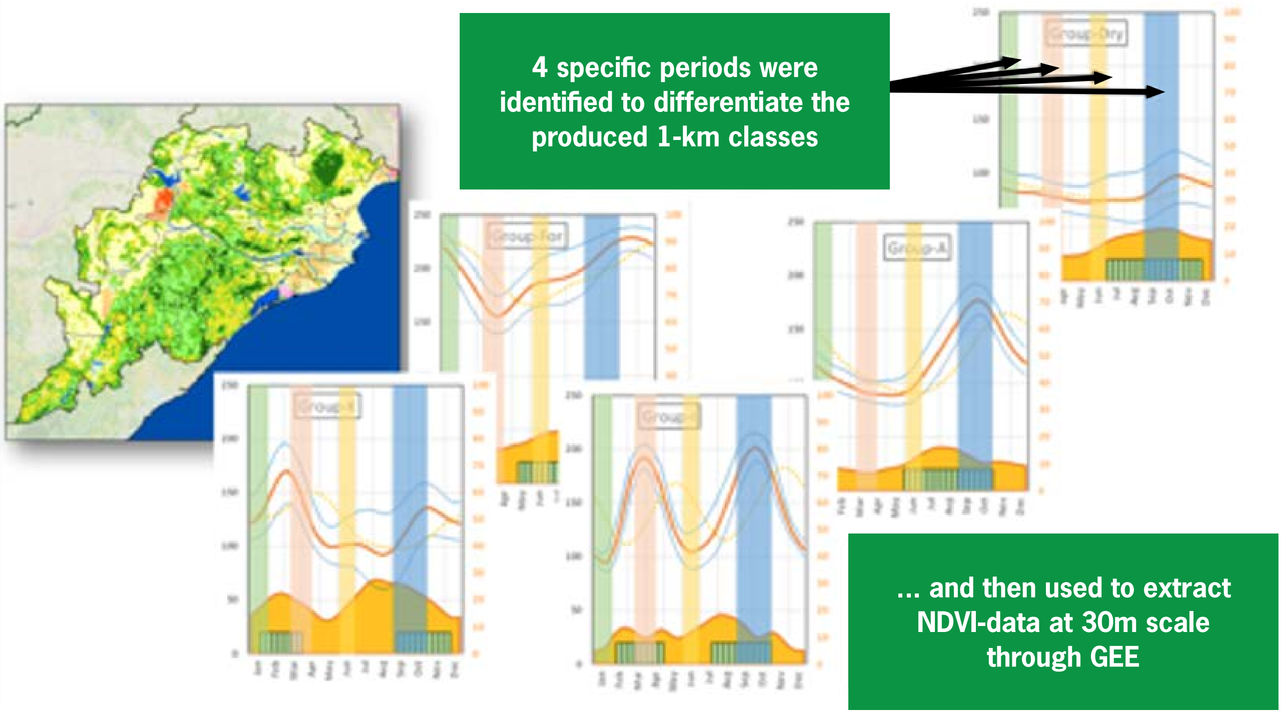

Figure 6:
Visually identifying the periods that allow differentiating identified CPSZs.

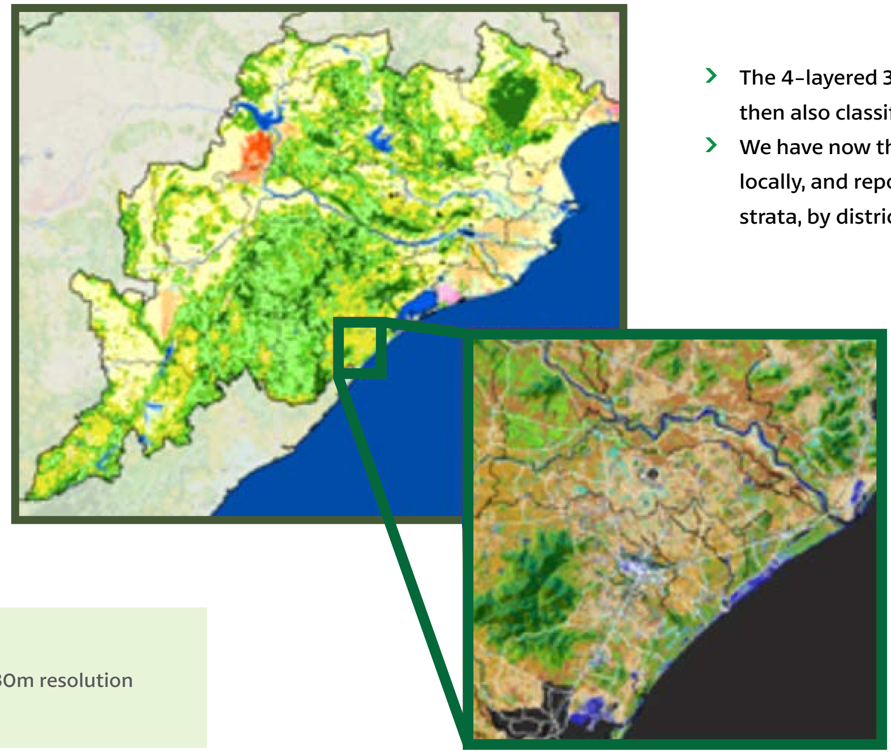

- The 4-layered 30m NDVI-image was then also classified into strata
- We have now the means to sample locally, and report regionally: by strata, by district, by country.

Figure 7:
Creating the 30m resolution LULC-map.

[8] Refer to Annex 2: Create a Field-Level CPSZs Map Technical Note- for additional information
[9] The ISODATA unsupervised classification routine is available in ERDAS. Different versions of ISODATA code are freely available (via Github) and in use by all major GIS software packages.

Step 4: Create a rice-cropping system legend to the multi-temporal field-level map using existing statistics

Out of the 22 30m-resolution NDVI-classes, only 10 related to annual cropping. (**Figures 8 and 9**). We converted the intermediate legend into a final legend by replacing all NDVI profiles with factual cropping system information. This was achieved by establishing spatial correlations between the mapped classes and reported agricultural crop area statistics.

We present the final legend in **Figure 10**. We created the final legend by relating the extent and location of the 22 classes to official rice-area statistics by Block[10]. The legend reports the average fraction cropped by pixel to each class mapped (for 3 seasons and for rainfed versus irrigated). However, retrieving the required polygon file on Blocks as used in EARAS was challenging because India uses various block/boundary definitions interchangeably (**Box 3**).

We further added data on rice-based crop calendars, specific for Odisha, to the legend (IRRI, 2012; **Figure 11**). The final synthesis of the required 30m resolution benchmark map shows all details on rice-based cropping systems of Odisha (**Figure 9**). Note: The top-left inset shows the initial 1km² resolution NDVI-map. The produced 30m-resolution map now provides the means to sample rice locally (representative for a map-unit), and to scale up those data to report rice-based cropped area statistics regionally by strata, district, and/or state.

[10] Obtaining the surveyed "Block" polygons proved a practical challenge and very problematic (confidentiality issues). We finally used multiple sources to synthesize Block-specifics and created our own 'guesstimated' polygon file. The official rice-based data at block-level were retrieved from 2 TXT-files (Source: EARAS):
http://www.desorissa.nic.in/pdf/tables_1617.txt
http://www.desorissa.nic.in/pdf/tables_1718.txt

BOX 3:
Enumeration Areas (EAs) of Odisha.

India uses interchangeably Tehsil and Block definitions. However, boundary differences between them frequently occur because their basic function differs, and their definition maintenance takes place by different organizations. EARAS manages Block specifications. They also specify within a Block, which GPs (Gram Panchayats or villages) contribute to its Enumeration Area. We present below, the GPs of the actual 2020 EAs of 8 Blocks of Angul district in Odisha. Together they cover all agricultural areas of Angul. Only for several districts proper Block boundaries were available; for all remaining districts, we adjusted provided Tehsil boundaries and guesstimated actual Block boundaries4

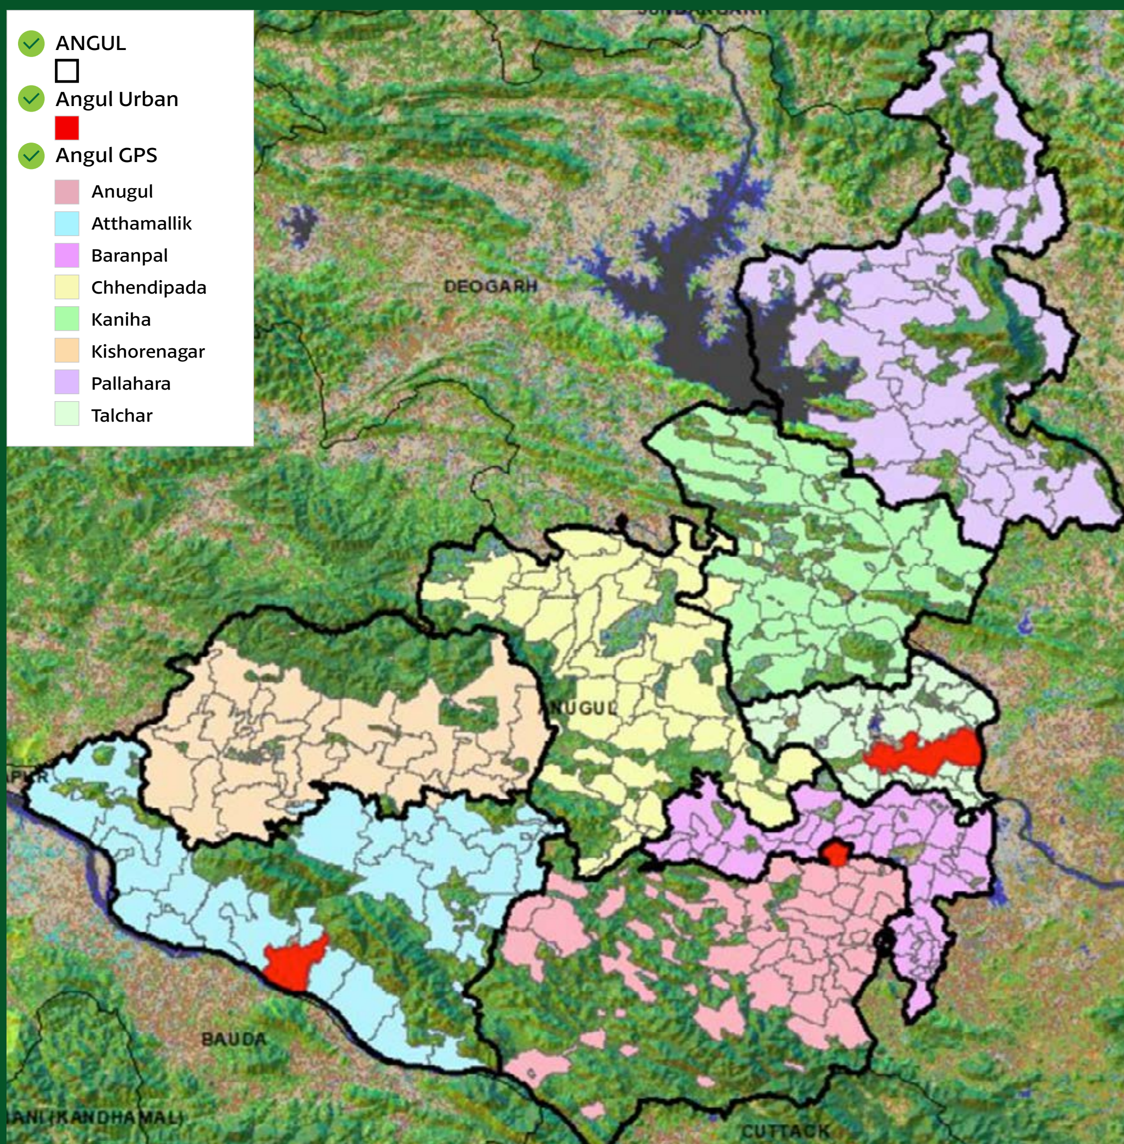

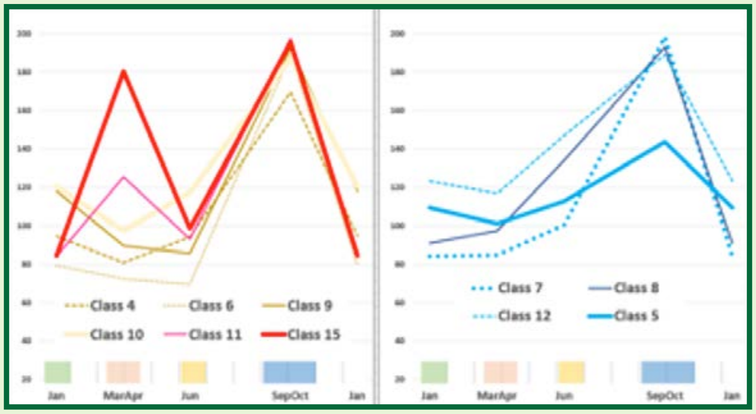

Figure 8:

The 2nd intermediate legend to the 30m resolution NDVI-map (NDVI profiles that relate to annual cropping).

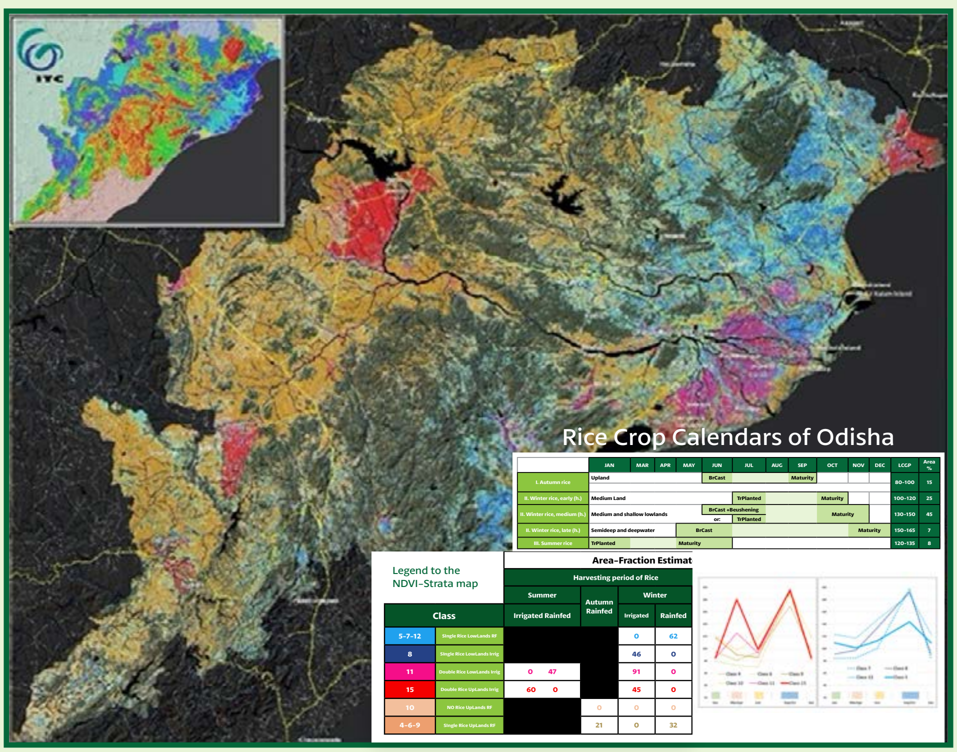

Figure 9:

The derived 30m resolution, rice-based, cropping systems map of Odisha.

Class		Area-Fraction Estimates (%): Harvesting period of Rice – Summer: Irrigated	Summer: Rainfed	Autumn Rainfed	Winter: Irrigated	Winter: Rainfed	Rice Intensity	Net Area Sown (to any arable crop): Irrigated	Net Area Sown (to any arable crop): Rainfed	Hectares
5-7-12	Single Rice LowLands RF				0	62	62	0	41	2,325,472
8	Single Rice LowLands Irrig				46	0	46	48	19	767,675
11	Double Rice LowLands Irrig	0	47		91	0	138	93	11	396,670
15	Double Rice UpLands Irrig	60	0		45	0	105	58	0	369,938
10	NO Rice UpLands RF			0	0	0	0	0	12	1,498,681
4-6-9	Single Rice UpLands RF			21	0	32	53	0	47	3,277,574

Figure 10:
The final LULC-legend to the 30m resolution NDVI-map containing fractions cropped by pixel to rice-based LUSystems, season and water management system (irrigated versus rainfed).

	JAN	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	LCGP	Area %
I. Autumn rice	Upland				BrCast			Maturity				80-100	15
II. Winter rice, early (h.)	Medium Land					TrPlanted			Maturity			100-120	25
II. Winter rice, medium (h.)	Medium and shallow lowlands				BrCast +Beushening / or:	TrPlanted			Maturity			130-150	45
II. Winter rice, late (h.)	Semideep and deepwater			BrCast						Maturity		150-165	7
III. Summer rice	TrPlanted			Maturity								120-135	8

Figure 11:
Retrieved information on season-specific rice-based crop-calendars as practiced in Odisha.

Step 5: Extract 10-90% Performance thresholds from 1km^2 NDVI-imagery and create Dynamic Area Frames[11]

Using the pre-processed 1km^2 NDVI-imagery archive, we extracted the 10 and 90 percentile values for every cluster (100) – dekad (36) combination. The NDVI-data of all pixels in a cluster and all dekad-repeats (20 years) generated one specific percentile value. We required the percentile values to prepare season-specific Dynamic Area Frames (DAFs). The 10 to 90% percentiles represent the lower and upper-thresholds that we assume to represent 0 to 100% crop-performance for a specific cluster-dekad combination (**Figure 12**).

The thresholds are not derived on a pixel-by-pixel basis, but on a cluster-by-cluster basis. For all pixels by cluster (100) and dekad (36; see step-1), we extracted pixel-specific values across years and retrieved the required percentile values from their frequency-histograms.

We consider that extracted 10 and 90% threshold data remain valid for several years as the thresholds (benchmarks) are based on long-duration population performance dynamics. Using the threshold is, however, pixel and season specific. This is equivalent to comparing climate versus weather. Therefore, a repeat rate of 10-years for re-calibration of the created zonation and parameterization would suffice and could, for instance, follow the agricultural census survey regime for ease of performance.

Figure 12:
Use of the 10 and 90% thresholds.
Likely, the expected Crop Performance is not linearly related to NDVI.

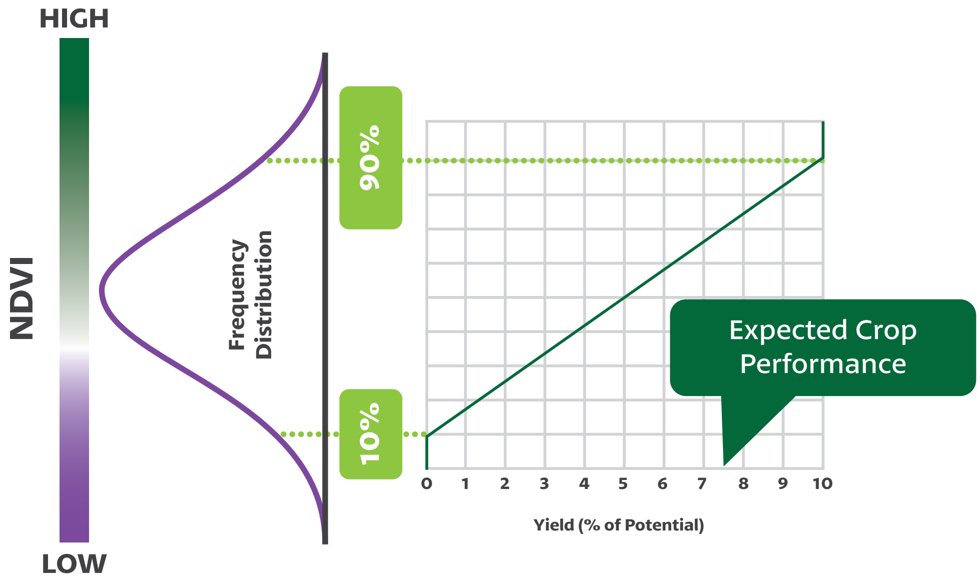

Next, we used cluster-specific NDVI profiles to specify which dekads are part of a growing season[12]. For these dekads we generated performance assessments. Subsequently, we averaged the dekad-specific performance data to a pixel-specific seasonal performance assessment (see Annex-3).

[11] Refer to Annex 3: Creating the DAFs – for additional information
[12] In contrast to 'growing season' specifics as extracted from NDVI profiles, a 'crop-growing period' is specified by actual planting and maturity/harvesting dates. Such data were not used. See annex-5.

Regarding spatial scale, any specific $1km^2$ level assessment will apply to every 30m pixel within that larger pixel. Note that field-to-field performance variability, caused by site-specific farm management practices, can change assessment accuracy significantly. Despite this challenge, our approach presents a method to capture season-specific spatial performance variability across a given region.

The Dynamic Area Frames provide the means to extrapolate series of field and season-specific yield measurements to whole regions cropped to that specific LULC-class. (**see Figure 13**). The creation of DAFs is achieved using 1km pixels and use of the original map with 100 NDVI-clusters. We superimposed selected 30m maps as masks to provide further spatial detail.

We anticipate that the pixel-specific performance for a specific rice-based cropping system and season can significantly follow this function:

Yield (kg/ha)[13] *= a . Reference NDVI*[14] *(DN-value) + b . DAF value*[15] *(%)*

The data examples show that impacts on the performance of crops are different every year due to differences in weather patterns. Theoretically, it is incorrect to simply average all field measurements in an area without considering spatial variability. Thus, DAFs provide an optional approach to extrapolate available plot-specific CCE-measurements. We do not anticipate that a DAF can be used as a tool to guide a crop-production survey because DAFs can only be generated at the end of a growing season (GS). Thus, season-specific DAFs cannot be produced in time to support a survey scheduled during a harvesting period. The DAF will, however, help to improve the spatial generalization of collected field data. The 20-year series of DAFs shows patterns of season-to-season and local-to-regional differences in crop performance (**Figure 14**). The map further supports a practical need to use DAFs to scale up local point data (CCE-data) to regional crop production statistics.

DAFs show impacts of seasonal weather differences as deviations from the "reference" performance at a specific location. Accordingly, we must make the long-duration, seasonal-NDVI (i.e., as measured during a growing season) available as a map. It will function as the "reference" performance of the crop production systems present. Two seasonal reference-performance maps (showing DN NDVI-values) are presented in **Figure 15**. The figure shows the spatial difference in average[16] crop performance caused by spatial differences in soil, terrain, or other permanent land characteristics.

[13] Location specific data from CCE-measurements (Crop Cutting Estimate surveys).

[14] The Reference NDVI varies between 75 and 190 (DN-values). **See Figure 15**.

[15] The DAF value varies between 0 and 100 (%). **See Figure 13**.

[16] Averages of all dekad-specific “NDVI-values” that cover the cluster-specific growing season. “NDVI-values” are the 50-percentile (50%) values of the 20-years repeats for that dekad and pixel.

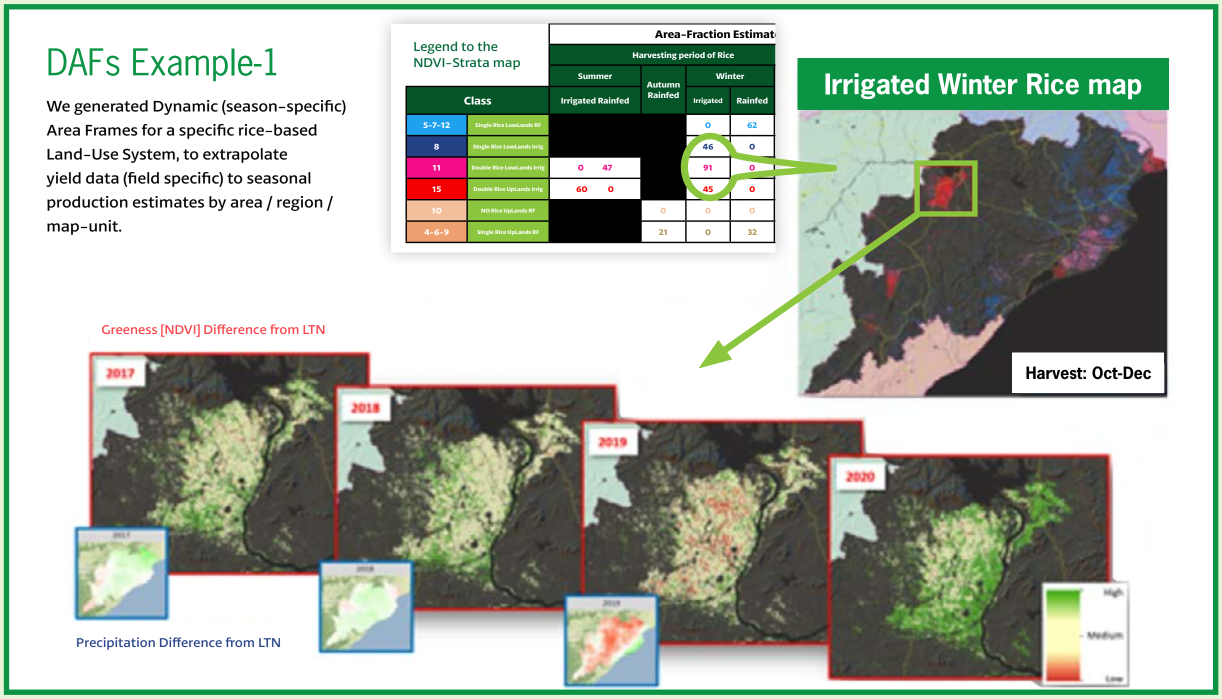

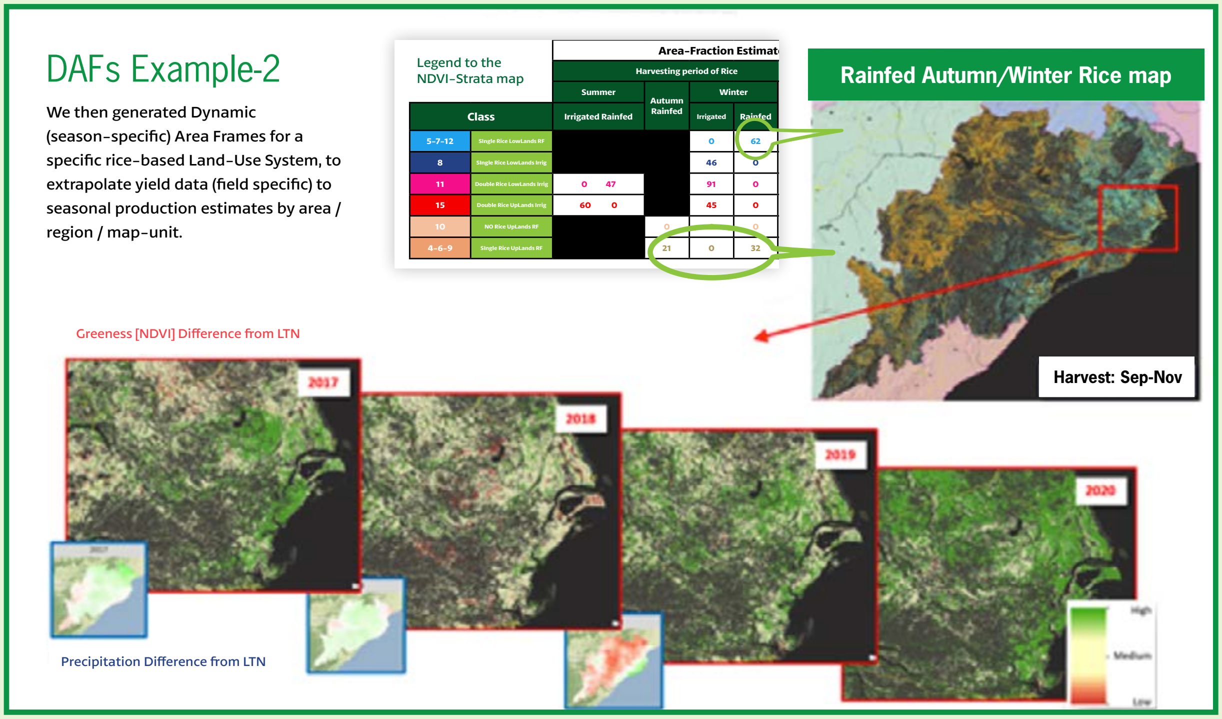

Figure 13:

Two DAF-examples, each covers 4 specific seasons for a specific Land Use System.

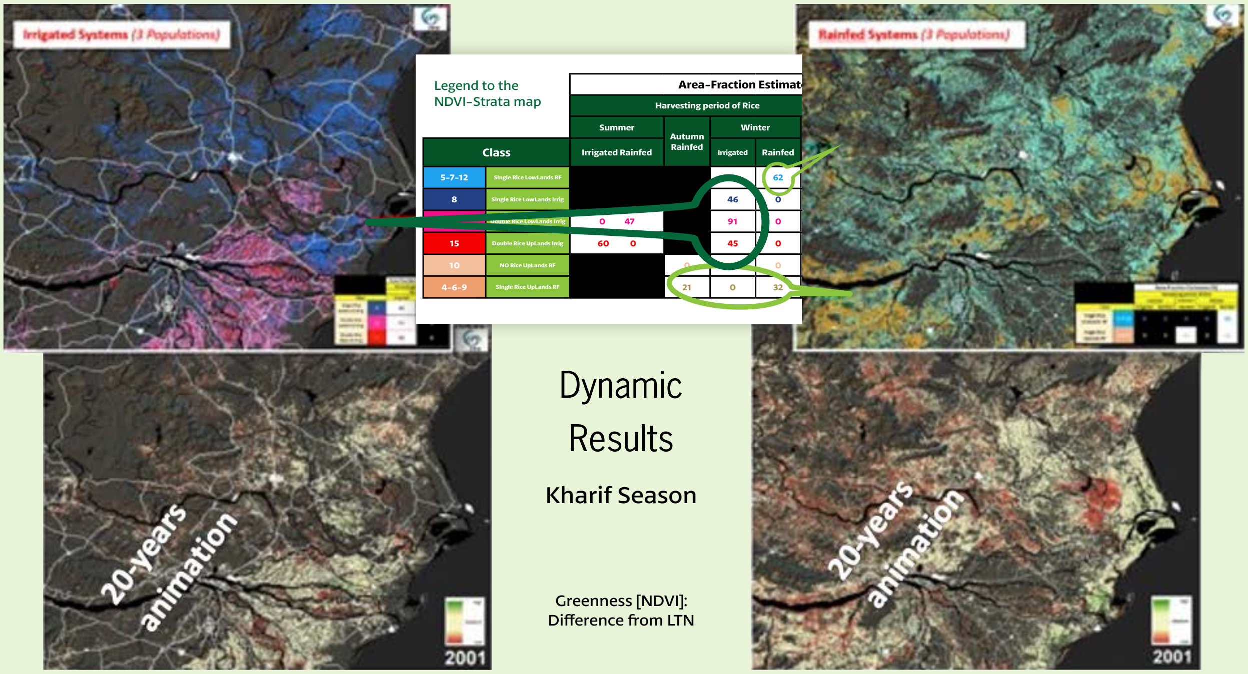

Figure 14:

Two 20-year animations, where each region covers three unique rice-based Land Use Systems.

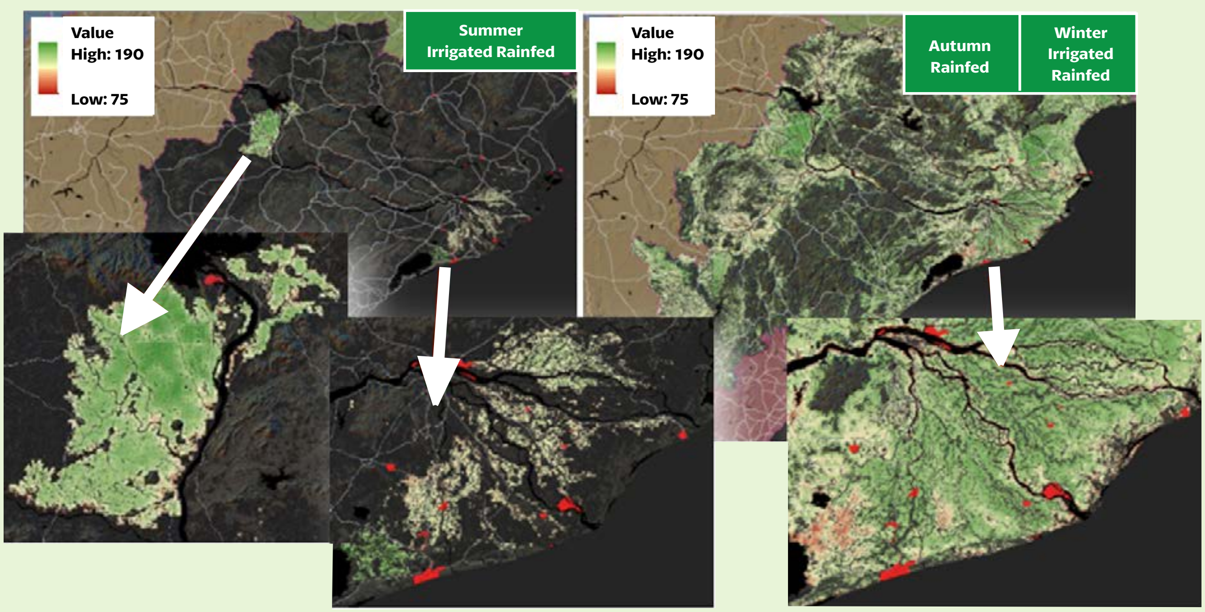

Figure 15:

Reference performance maps of the summer- and autumn/winter-season rice systems.

3

Outcomes

3.1 Key Results

This research aims to delineate homogeneous land use and land cover strata to support the stratified sampling required to generate crop area estimates that present relevant spatial differences in season-specific crop systems focusing on the Odisha State in India. The presented approaches enable improved crop analytics at scale.

Project Contributions:

- Using hyper-temporal NDVI-imagery, we created a 100-class NDVI-cluster map. The map was further generalized by grouping the obtained 100 NDVI-profiles into 22 Crop Production Systems Zones (CPSZs). The CPSZs clearly differentiate four distinct periods in a year. Using those periods, we extracted average NDVI-values from recent ETM-imagery, which in turn, we classified into a 30m LULC-map (**Figure 9**).
- Using 2 years of official (EARAS) agricultural area statistics by administrative units (Blocks), we established that 10 out of the mapped 22 LULC-classes related to arable cropping. We reported data for three seasons of rice harvests and rainfed versus irrigated rice. Through cross-correlation techniques, we estimated the area fractions by legend entry used for rice-cropping. Non-rice crop information is not part of the produced benchmark map because we could not obtain relevant non-rice crop area statistics by Block from EARAS.
- Using the 100-class NDVI-cluster map, we extracted the required growing season periods by class and established, by zone, their long-duration—10 and 90 percentile values (using data from all dekad repeats across years and data from all pixels belonging to a class).
- We applied the thresholds to season-specific, pixel-based, NDVI-values to translate them into Dynamic Area Frames (DAFs) that capture seasonal-specific and spatial performance differences of a specific LULC-class. We used the LULC-map to mask the specific application areas and created a series of example-DAFs for Odisha to illustrate their relevance.
- Each DAF is comprised of two 1km scale maps: (i) the year, and season-specific anomaly estimates that are assumed to relate to the performance of the cropping system monitored, and (ii) the normally expected season-specific cropping system performance (median NDVI-values).

Innovations:

- The generated rice map for Odisha reports essential details on practiced crop calendars and cropping systems. It fully adheres to terminology and agricultural statistical specifications in use by EARAS. It also fully defines the "populations" that matter when creating a sample scheme for crop-area surveys. Sample schemes must evolve from using highly generalized maps depicting strata with different cropping intensities, to more detailed maps showing the location of populations of relevant LULC-defined cropping systems. On a seasonal basis, a crop-area survey should only update the cropped-area fractions by LULC-class to further produce tables on crop-area statistics.
- We generated all required logic to produce Dynamic Area Frames (DAFs) that show spatially (at 1km^2 resolution) the impacts of varying weather conditions on the performance of the in-situ cropping systems. We assumed, similar to most early warning systems, that the use of NDVI-imagery is the best way to capture seasonal crop performance. The way of interpreting the NDVI-data through thresholds, however, is a novel approach. We further shared the logic to superimpose the created LULC-based benchmark map as a mask on top of any produced DAF. Validation of DAFs was not conducted due to the lack of primary CCE-data.

Notes for Future Research:

- It is generally assumed that benchmark maps, like the maps produced for rice systems in Odisha with growing details and location information, are readily available anywhere. Unfortunately, this is seldom true. However, if benchmark maps can be obtained, the process developed in this project can be scaled and applied to other geographies without fail using the most up-to-date NDVI-imagery from the Copernicus Global Land Service system (CGLS) in tandem with Google Earth Engine script that creates multi-temporal 30m resolution NDVI layers. Past time series of Spot-VGT and Proba-V NDVI imagery can be seamlessly continued from Sentinal-3 data from CGLS. This approach applies to the BRDF-adjusted 1km^2 catalog of 1999 to 2020.
- Many future users will face a data confidentiality (or quality) issue as this project experienced during the data-mining exercise in Step-4. When confronted with data confidentiality or quality issues, a detailed legend of the 30m resolution LULC-map should be created to define the 'populations' under review. Besides that, using data that are already too generalized or carrying out correlation studies with too few statistical degrees of freedom will hamper the creation of the required results. Frequently, LULC-map must be used to start surveying a specific region from scratch (Box 1 & 2).
- Surprisingly, the relationship between yields obtained and the DAF-maps is linear and static across years. Therefore simple re-calibration between the DAF-values and CCE-measurements will likely remain a seasonal task. Furthermore, DAFs have no function in creating sample schemes. They were designed to scale up point-data to area estimates. For these reasons, field-specific crop yields measurements must represent specific random locations and randomly cover the full extent of the map unit under survey.

3.2 Lessons Learned

The innovative approach presented above (i) applies a proven concept to map in detail all rice-based cropping systems of Odisha (**Figure 10**), and (ii) presents a novel approach to create Dynamic Area Frames (DAF) to improve scaling up collected point-based yield data to area-specific yield-maps and production estimates (**Figure 13**). The first item functions as a benchmark map to improve sampling schemes that produce crop area statistics and as a mask to specify cropping specifics plus the area extent of a created DAF.

Our analysis indicates that using publicly accessible satellite remote sensing data, combined with secondary data sources and local expertise, can support and significantly improve agricultural survey schemes as well as the quality of crop area and crop production statistics. We highlight three key lessons relevant for selecting a reliable source of satellite remote sensing and ground-truthing data:

- **Future-proofing satellite data:** Technologies enabling Earth observation advance rapidly. The analytical workflow incorporating satellite remote sensing data should not rely on a unique dataset — especially given that sensors fail, they become obsolete, or data structures change. Instead, a key recommendation is to design a research framework that is flexible and reproducible with data from different sources and formats. During the project period, we learned that the ESA Copernicus program upgraded the entire catalog of NDVI to ensure compatibility with future Sentinel-3 satellites. Unfortunately, the upgrade process caused a temporal data inconsistency in the construction of CPSZs (e.g., extracting the NDVI thresholds for defining dynamic area frames from historical data). Despite a setback in the timeline, we repeated the whole process, starting from step 1 using the newly extracted datasets. This reanalysis process allowed the workflow to be fully compatible with the ESA's upgraded data products and made our workflow future-proof.
- **Prepare for changes in data sharing policies:** The utility of seasonally collected georeferenced crop production data extends beyond crop analytics research. Such data are critical for food security and policymakers. During the project planning phase, ICRISAT agreed to provide CCE-data collected from Odisha to validate our benchmark map and to fine-tune the use of created DAFs. However, in late 2020, a new government-mandated policy was imposed to restrict the sharing of such data. This restriction also partially applied to the detailed, block-level crop area statistics data. As a result, only data on rice-based cropped area statistics were available. We relied on secondary data sources to fill some data gaps (e.g., rice-based cropping calendars of Odisha as described by IRRI).
- **Improvements for Applied Practice:** Present crop-area statistics poorly indicate which crops are grown where and when. Natural resource managers would benefit from maps showing crop-area statistics (Steps 1-4). Remote sensing analysts also require landscape-based benchmark maps to give insight into what they are monitoring and when. With these maps, agricultural service organizations and Ministries/Departments of Agriculture could target their work better and improve future survey campaigns.

Enabling Crop
Analytics At Scale

Supporting Documents

Annexes 1-3

Next Generation Crop Production Analytics

Dynamic Area Sampling Frames for Improved Crop Analytics

The Enabling Satellite-based Crop Analytics at Scale (ECAAS) Initiative is a multi-phase project that aims to catalyze the development, availability, and uptake of agricultural remote-sensing data and subsequent applications in smallholder farming systems. The initiative is funded by The Bill & Melinda Gates Foundation and implemented by Tetra Tech.

info.ecaas@tetratech.com
cropanalytics.net

Final report
September 2021

Author
© Prepared by Kees de Bie and Andy Nelson, University of Twente/ITC

1

Annex

Creating Crop Production System Zones Technical Note

1. Processes Summary

In step 2, we grouped NDVI-classes into Crop Production System Zones (CPSZs) (**Figure 1**).

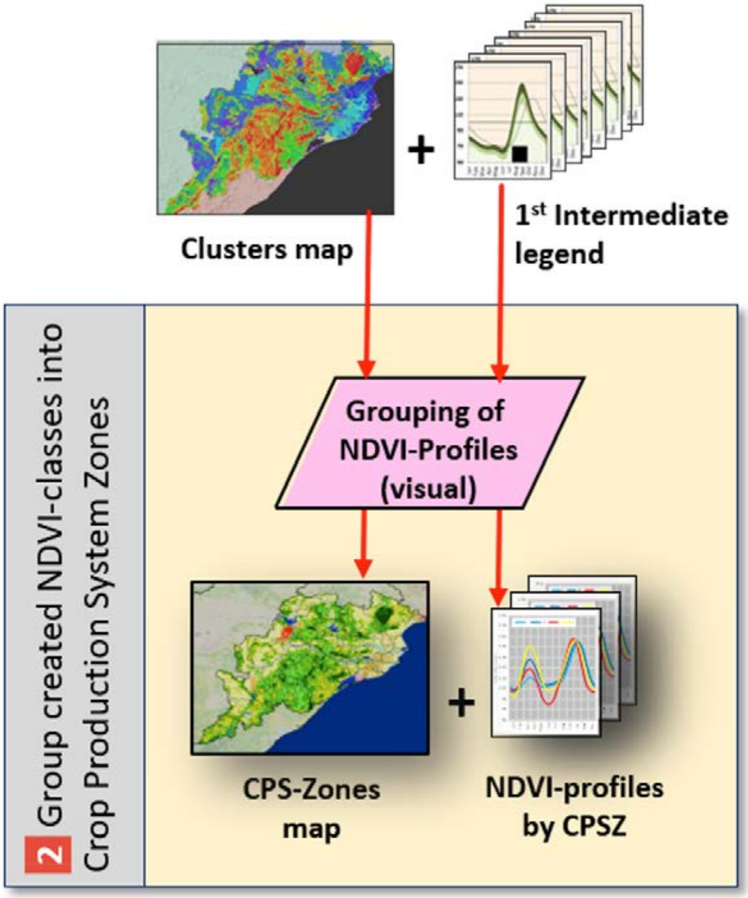

Sub-steps:

Extract all 100 NDVI-profiles to a spreadsheet
Create from it a Dendrogram

Use that and the Profiles to visually complete grouping.

Annotate the groups (CPSZs).

Figure 1:
The process of taking NDVI-profiles and creating CPSZs

2. Detailed Description

The legends of the 100 clusters, consisting of 144 averaged DN-values across all pixels classified to a specific cluster (class), were subsequently subjected to further statistical grouping. The shown dendrogram materialized.

The dendrogram's NDVI-profiles of each suggested group were depicted in a spreadsheet, and further adjustments were made by hand, based on visual comparison of the profiles.

Ultimately, just 22 groups (based on grouping the 100 ISO clusters) were considered sufficiently different from one another.

Through this process, we knowingly ignored the specific years "when" deviations in patterns did occur. The approach is similar to comparing 'weather versus climate' or 'actual versus normal.'

The next pages show graphical details of the created 22 groups.

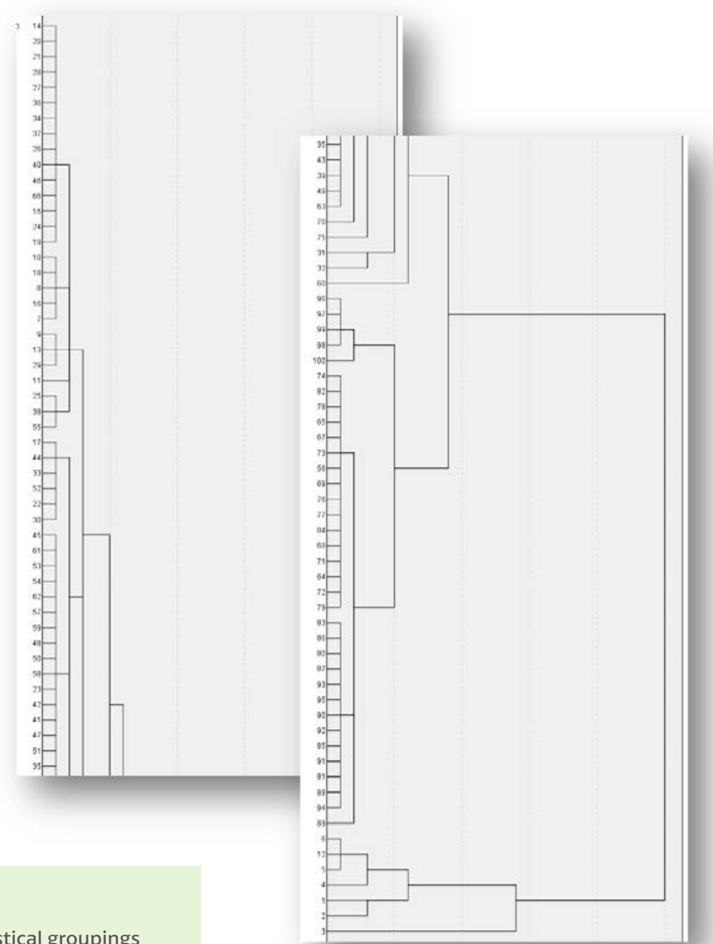

Figure 2:
Dendrogram depicting statistical groupings

Of the 22 created NDVI-groups, some primary sequences can be detected representing 'gradients':

1. from HIGH to LOW (10 groups): FOR →N →MN→ M → F→ F1→ D → O → DRY →H2O
2. from a LONG to an SHORT clear season: [F, F1, D] → [A, A1, A2, B, C]
3. ... to TWO clear SHORT seasons: → [E, E1, I, J]
4. ... and to mixed-up patterns: → [G, H, K]

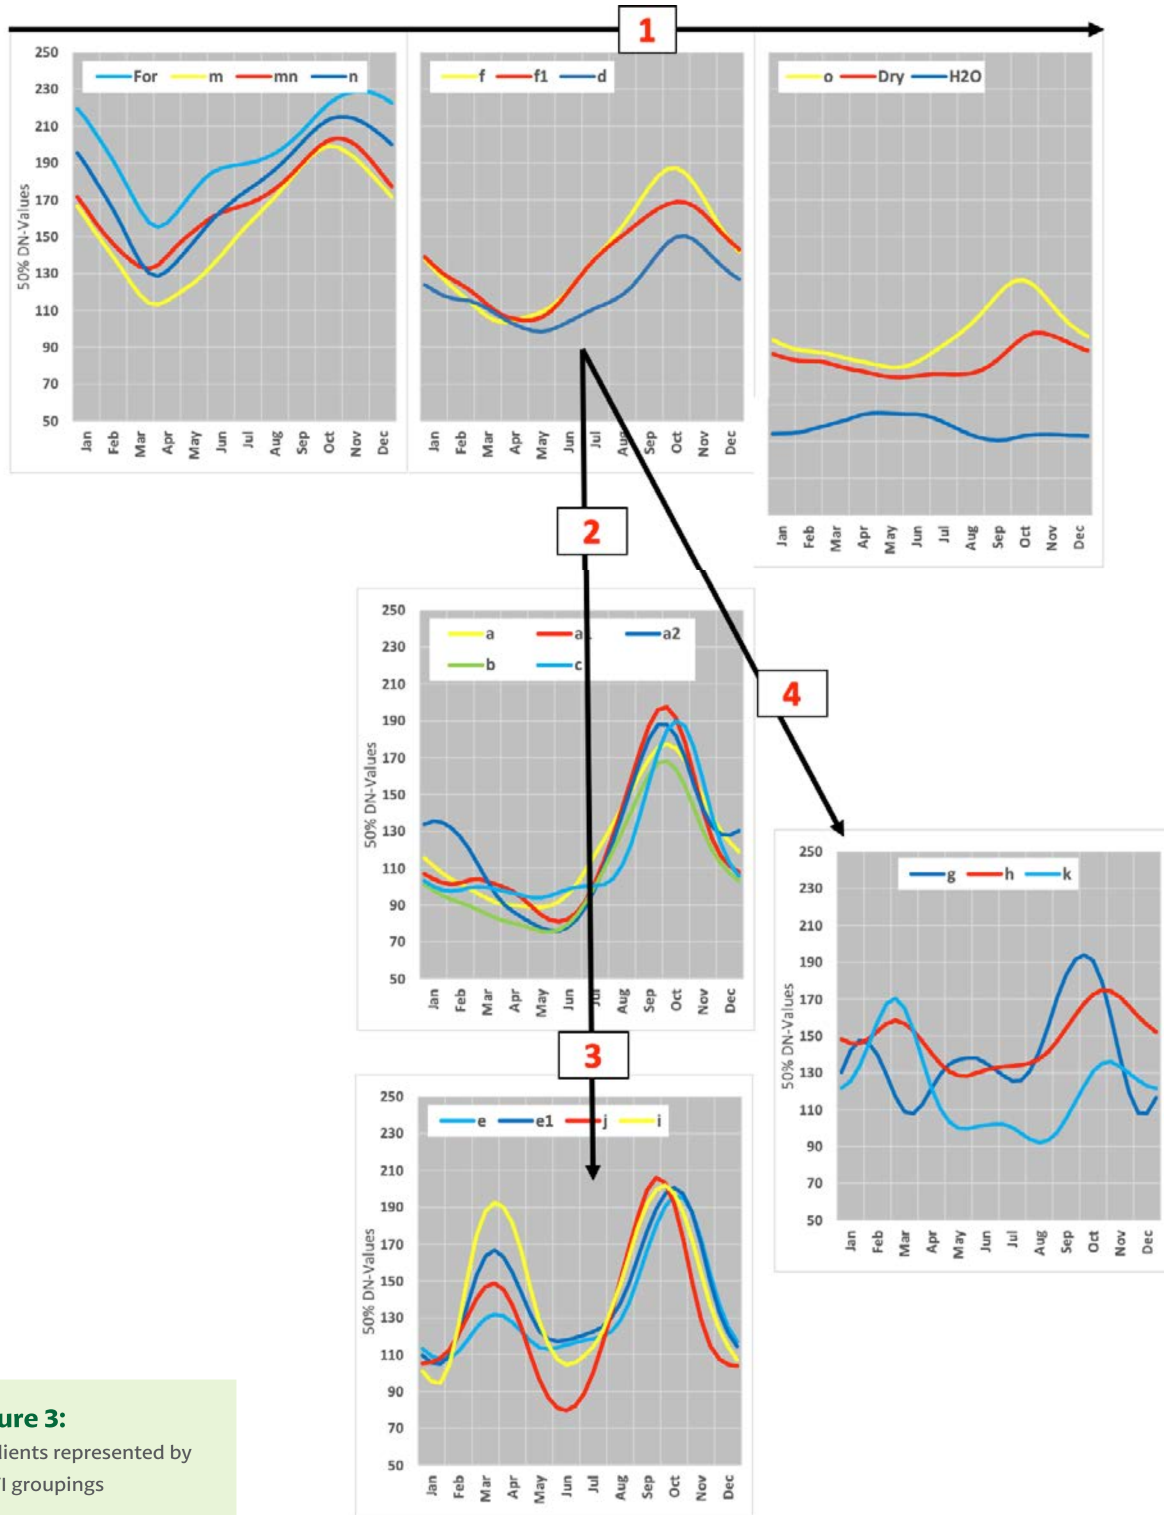

Figure 3:
Gradients represented by NDVI groupings

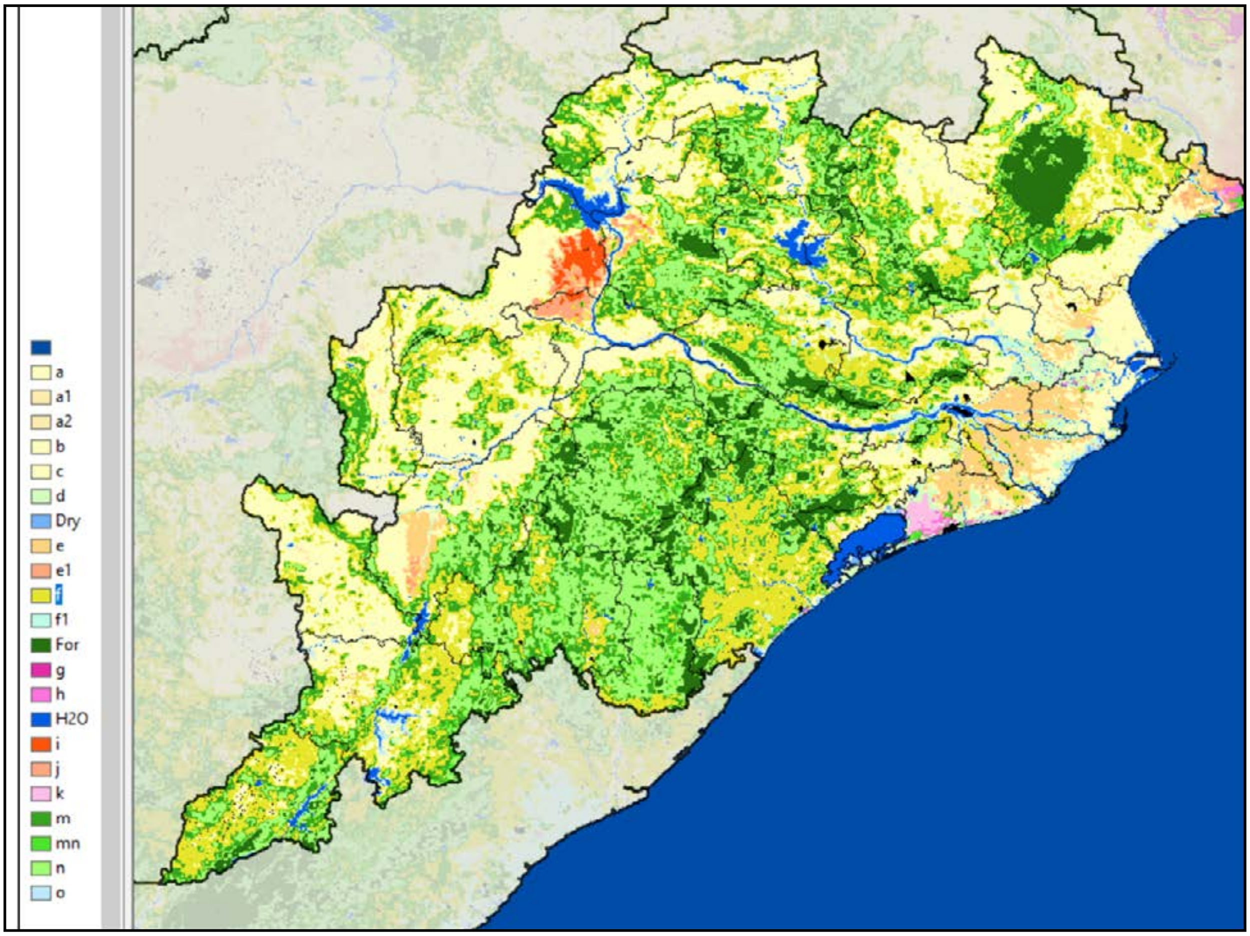

Figure 4:
Gradients represented by NDVI groupings

2

Annex

Generating Field-Level LULC-Maps Technical Note

1. Processes Summary

In step 3, we created a 30m resolution Land Use Land Cover (LULC) map at field-level, after interpreting 1km resolution CPSZ-signatures, and creating a multi-temporal 30m-resolution image (**Figure 1**).

Figure 1:
Summary process of creating CPSZ-signatures.

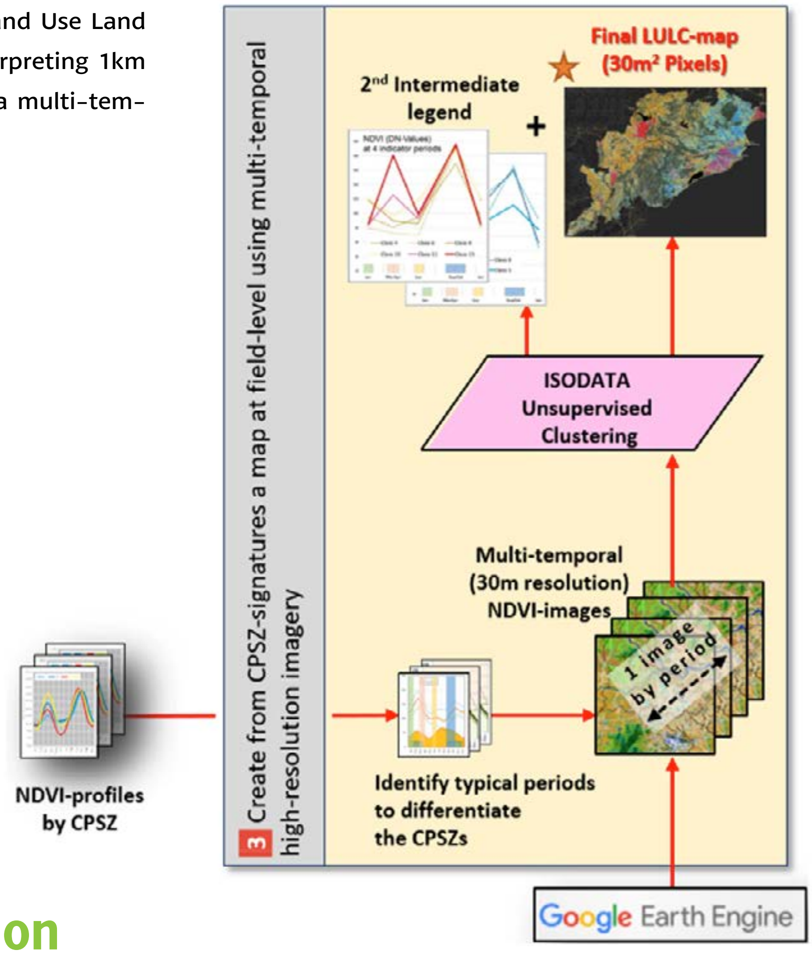

2. Detailed Description

We experimented with a very straightforward approach to refine the 1km map to a 30m map. For that, we required the periods during which NDVI was instrumental in differentiating the 22 classes and the averaged NDVI-data during these periods from all available 30m imagery (TM8) as available for all recent years.

We could identify four relevant periods that differentiated the 1km resolution classes. We visually identified the periods from the created NDVI-profiles of the 1km-scale classes. The four periods are defined below with justification for the various classes provided in figure 2:

OFF	1 Jan – 31 Jan	[2014-2019]
QOFF	10 Mar – 20 Apr	[2014-2019]
DIP	1 Jun – 30 Jun	[2014-2019]
ON	1 Sep – 31 Oct	[2014-2019]

Classes with one long green season with a short-dry spring: Groups FOR, N, MN, M

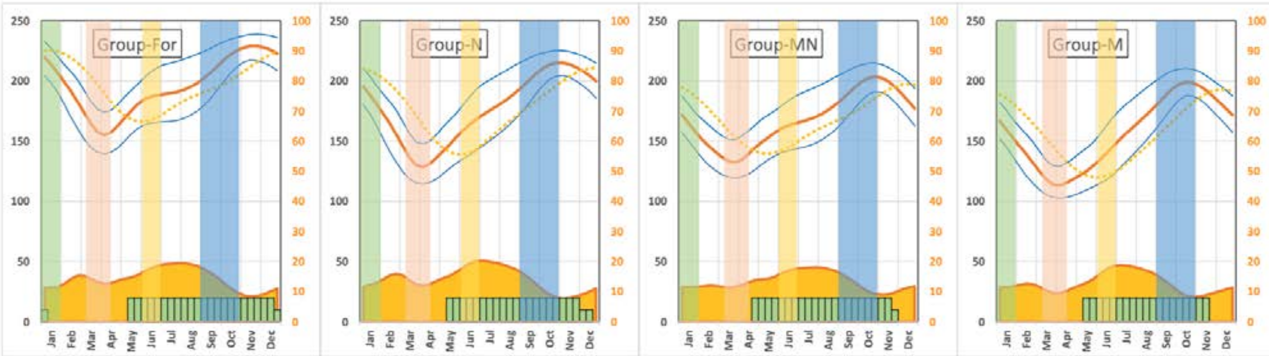

Classes with a gradually disappearing crop growing period (during monsoon): Groups F, F1, D, O, Dry, H2O

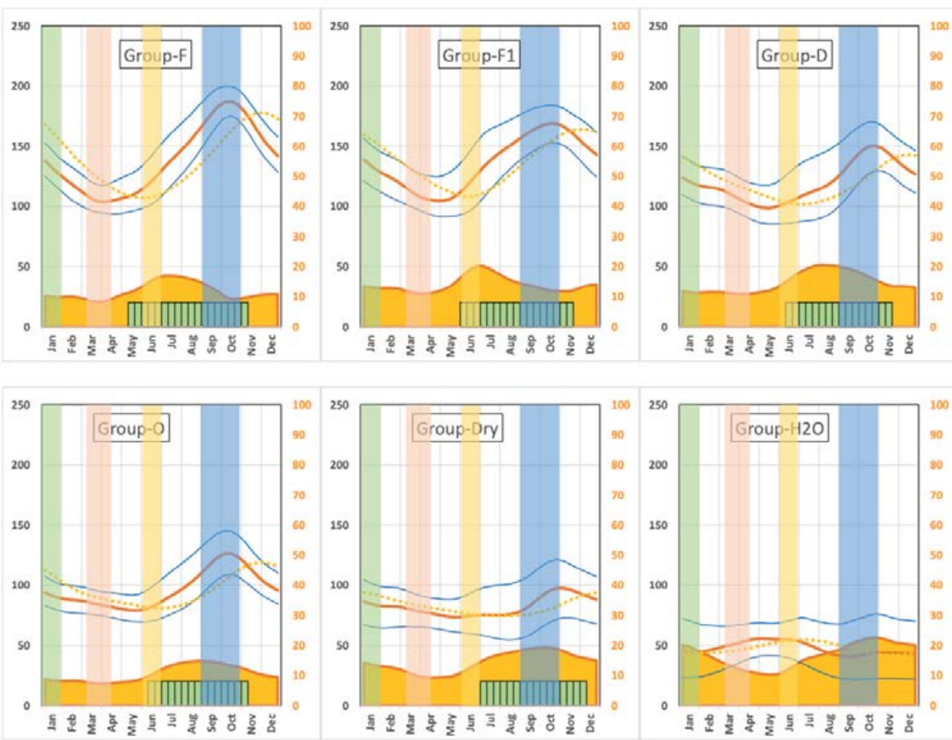

Classes with a very clear (distinct) crop growing period (during monsoon): Groups A, B, A1

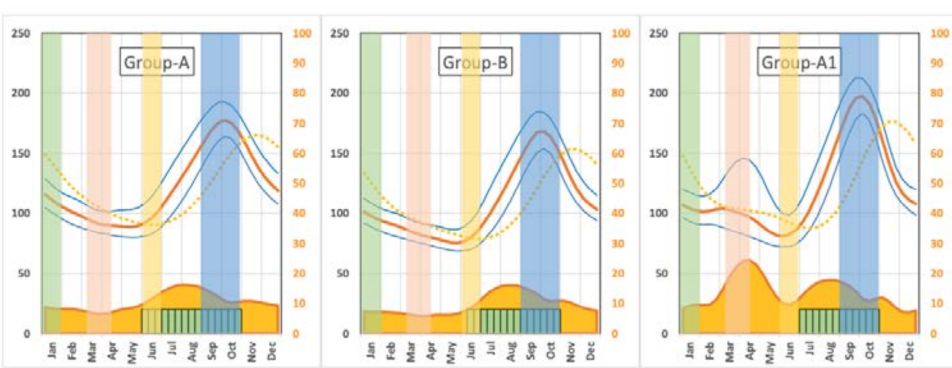

With and without a mid-winter season: Group A2, C

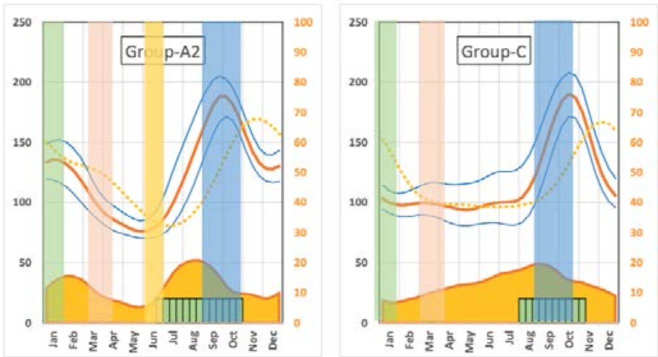

Classes with gradually an emerging earlier growing season: Groups E, E1, I, J

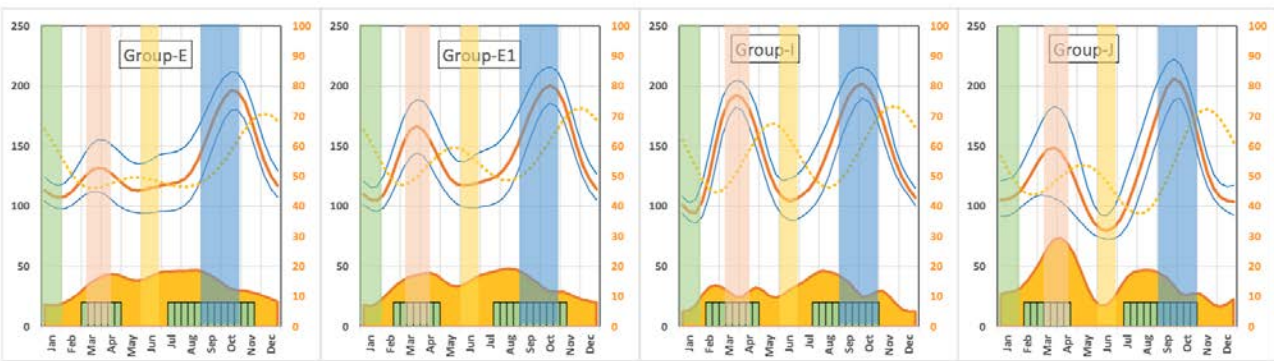

The three remaining smaller groups: Groups G, H, K

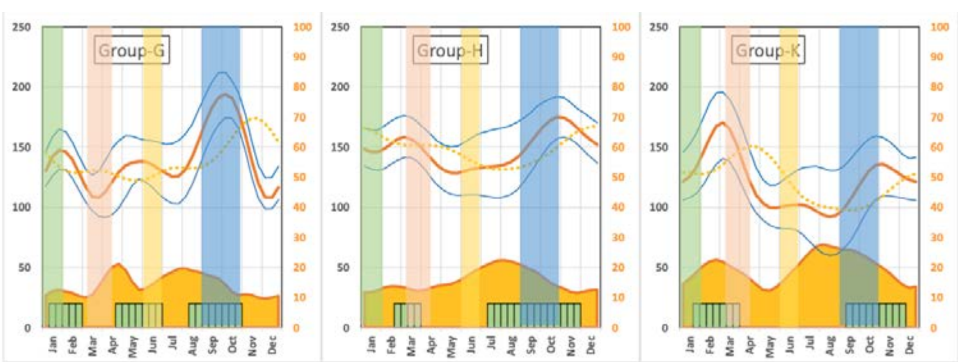

Figure 2:
Classes with grouping by period.

A Google Earth Engine (GEE) code was used to extract the median NDVI-values for the 4-periods, as required to differentiate 22 classes at 30m scale is available upon request. We used the NDVI-medians of all cloud-free pixel data obtained from TM8 imagery (30m) covering 2014-2019.

Once the 4 layers were extracted and stacked, we used the ISODATA unsupervised classification once more to convert those 30m NDVI-layers into a classified map containing 22 classes (including the background). Enormous spatial gain was achieved by using the temporal info extracted from the 1km Hyper-Temporal imagery, as shown in **Figure 3**.

Please note that the twenty-two 1km-classes are NOT similar to the twenty-two 30m-classes (see "intermediate legend" as depicted in **Figure 4**).

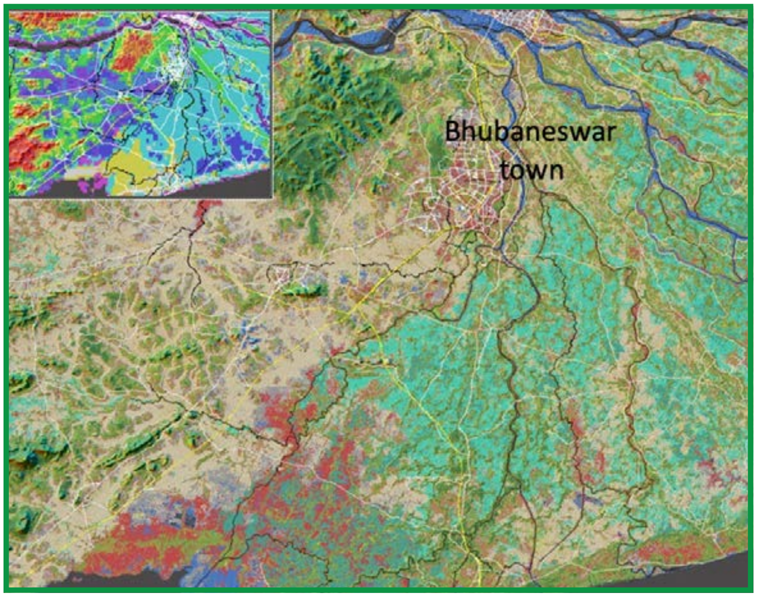

Figure 3:
Final LULC Map (30m).

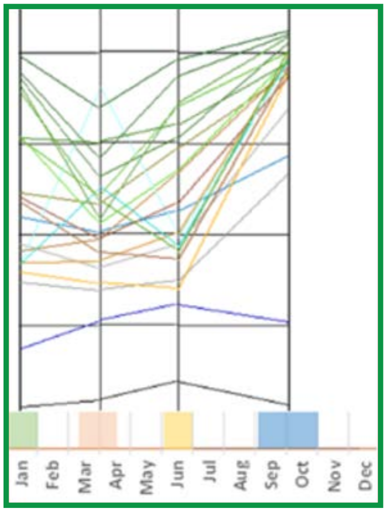

A significant gain is that the 30m classes are relatively more homogeneous than the 1km classes. The presence of mosaics/complexes/mixtures of land cover types is thus seriously reduced. We can now assume that we can link the present NDVI-classes directly to specific land use and cover types and not to mixtures of notably different LULC's.

The original four NDVI-layers at 30m scale covering the Bhubaneswar that generated were used to generate the Final LULC Map are shown in Figure 5 and Figure 6. These images depict the median NDVI-values of all cloud-free pixel-level data covering the provided period and the years 2014-2019. The images also show considerable fragmentation of the landscape, with scattered "clusters of" small fields.

Figure 4:
Final LULC Map (30m).

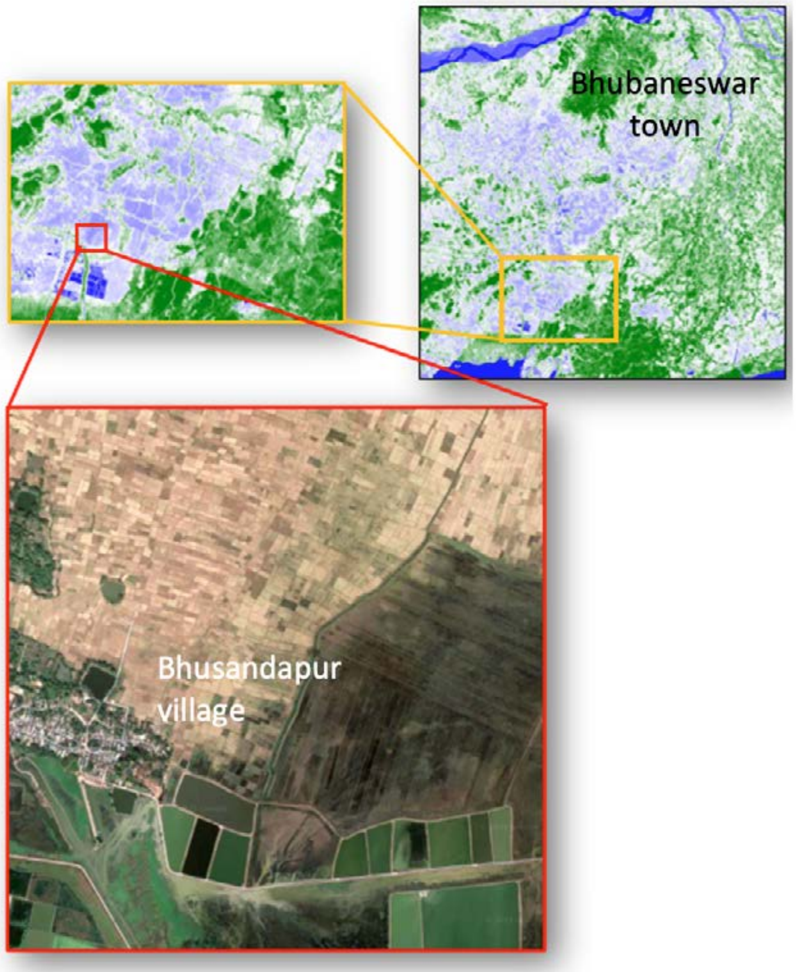

Figure 5:
NDVI Layers at 30m.

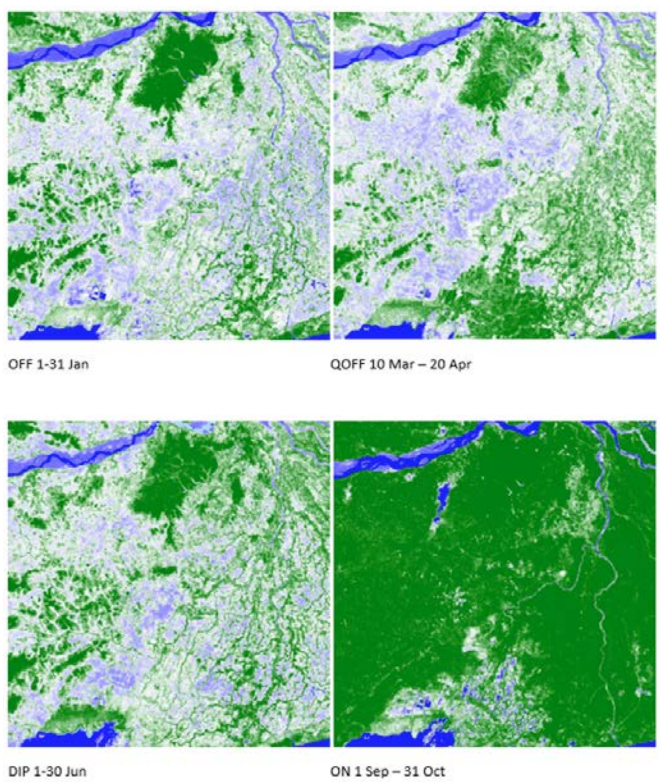

Figure 6:
Four NDVI Layers at 30m depicting periods.

3

Annex

Creating the DAFs Technical Note

1. Processes Summary

In step 5, we extracted 10-90% Performance thresholds from 1km2 NDVI-imagery and create Dynamic Area Frames (**Figure 1**).

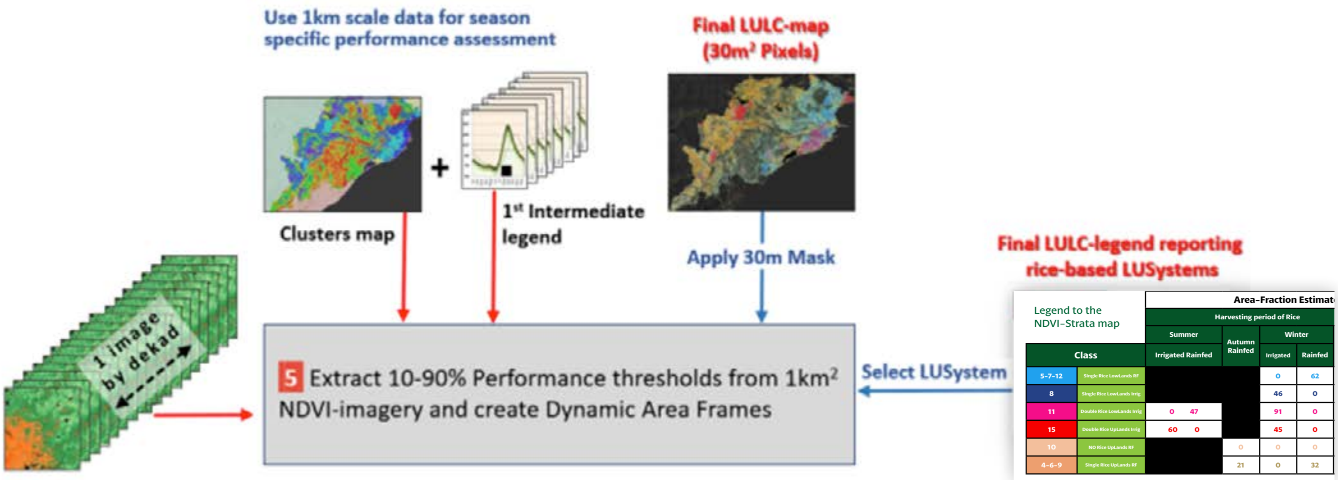

Figure 1:
Summary of DAF creation process.

- The intermediate legend (NDVI-profiles) is used to define the Growing Season(s) by cluster.
- The stack of 20-years (cleaned) NDVI-images is used to extract the 10% and 90% threshold values by cluster.
- The stack of 20-years (cleaned) NDVI-images is used to extract the 50% (median) NDVI-values by pixel.
- The 30m LULC-map (and legend) is used to mask the above for a specific season and rice-based cropping system.

The original 1km clusters map (with 100 classes) was used and not the 1km CPSZ-map (with 22 classes) because the original map was less generalized and therefore more specific regarding performance assessment. Further on, info on LULCs was superimposed through the use of the 30m maps; they provided the required masks.

2. Detailed Description

Determining Growing Seasons for each NDVI-cluster

To create the DAFs, we re-used the cluster map as generated during Step-1 (map at 1km resolution with 100 NDVI-classes) and the cleaned original 20-years' time-series of NDVI imagery. The map created during Step-4 was used as a mask concerning the extend on a DAF of a specific LULC-class.

Construction of a DAF was season-specific and followed growing season (GS) information as interpreted from the 50% NDVI-Profiles of each 100 clusters and further fine-tuned (**Figure 2**). Growing estimates for class-45 (cluster-45) were then mathematically extracted, with grey bars defining the GS (**Figure 3**). The GS start and end are demarcated by the cross-points of the 50% NDVI-line (long duration median NDVI-values by dekad) and its 9th-dekad moving-average (black dotted line). The accuracy of this simple and practical method, in relation to GS-reality, is well studied and accordingly commonly used.

Is a specific dekad$_i$ part of a Growing Season?

	Condition	If Yes, then:	If No, then:
(0)	Is: (9MovAvg Valuei) + 5 > (50%-Valuei) ?	Yes; Goto (1)	No
(1)	Is: (50%-Valuei) > (9MovAvg Valuei) ?	Likely Yes; Goto (2)	Likely No; Goto (3)
(2)	Is: (50%-Valuei) > 80 ?	Yes	No
(3)	Is: (50%-Valuei) / (Max of Annual 50%-Values) > 0.95 ?	Still possible; Goto (4)	No
		Yes	No
(4)	Is: (50%-Valuei) > 80 ?	Yes	No

0.95	User defined setting to identify "Extended Seasons": when a drop in NDVI is only marginal, this can point to the presence of a relatively short drier period that falls within a much longer growing season, or to the presence of still sufficient residual soil moisture causing very gradual decay at the end of the growing season.
5	The difference between 9th moving average and actual must be sufficiently high (5 DN-values). This rule removes vague GS-starts and GS-ends
80	User defined setting to identify "Growth"; below this value crop growth is not deemed likely (the value points to relatively too little cover of active green photosynthetic plant materials).

** The 9MovAvg Value is the Average of the NDVI-values of nine Dekads, i.e. of 8 values that fall immediately prior to the dekad applicable (that dekad provides the 9th value).*

Figure 2:
Defining a Growing Season (GS).

Class	Variable	Count	Band_1	Band_2	Band_3	Band_4	Band_5	Band_6	Band_7	Band_8	Band_9	Band_10	Band_11	Band_12	Band_13	Band_14	Band_15	Band_16	Band_17	Band_18	Band_19	Band_20
45	10%	4574	120	115	110	105	101	99	96	95	94	94	94	94	94	95	97	100	103	108	115	
45	50%	4574	133	127	122	118	114	111	108	106	104	104	106	107	109	111	114	118	124	130	135	
45	50%MovAvg	4574	162	155	148	141	135	129	124	120	116	113	110	109	108	107	108	109	111	114	117	
45	90%	4574	149	142	136	131	128	125	123	121	119	119	122	126	130	134	137	140	145	151	157	
45	Sdev	4574	11	10	10	10	10	10	10	9	9	10	11	12	13	14	14	15	16	16	16	
45	Season	4574															1.0	1.0	1.0	1.0	1.0	

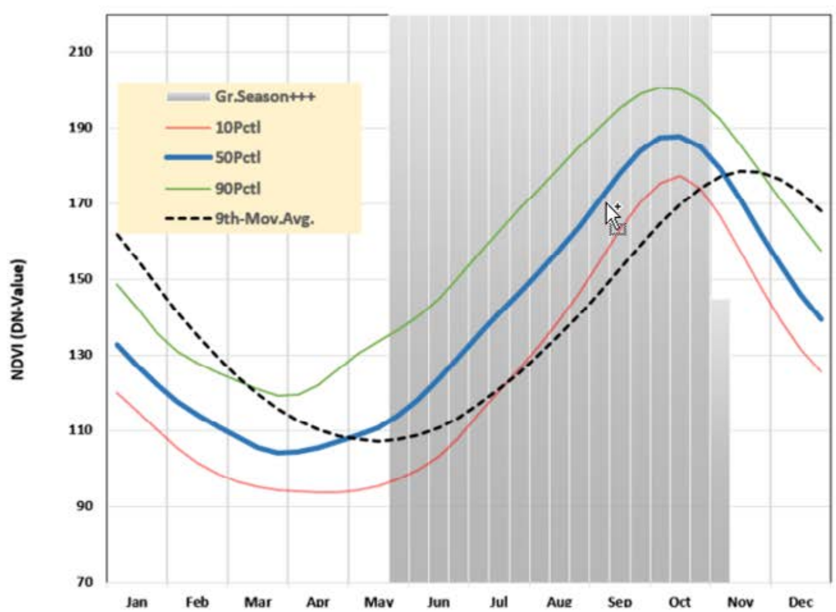

Figure 3:
Growing estimates for class–45.

Determining performance

The growing season is also depicted by dekad 10 and 90% data (red and green lines) as extracted from the original (cleaned) NDVI imagery stack (**Figure 3**). These percentile lines represent data from all years and all 1km pixels of the concerned cluster. The 10 and 90% lines are compared to the actual NDVI-values recorded by pixel, dekad, and season. For each pixel of the cluster and each dekad of a GS, performance was estimated based on a linear interpolation between the 10% value (assumed as 0% performance) and the 90% value (assumed as 100% performance). For each pixel and season, all dekad-specific values are then averaged to provide a season-specific performance estimate. This approach is relatively simple but makes for the best use of all NDVI-data available.

Extra: Relationship between NDVI and Rainfall

Earlier studies looked at the relationship between temporal patterns of average rainfall and NDVI-Profiles. In those studies, 'climatology' was used as a long-duration averaged data and both unimodal as bimodal weather patterns were included (Source: Garcia Velez, 2016). As displayed in prior studies, NDVI data lags in time behind rainfall data (**Figure 4**). This lag varied from 5 to 8 dekads (about 2 months). Temporal patterns of rainfall and NDVI are remarkably similar and that accordingly defining the GS-specifications based on NDVI-data is correctly carried out as also shown in **Figure 4**.

Key differences between rainfall and NDVI remain:

- NDVI integrates the result of growing conditions; rainfall is just one of those (an important one!).
- Averaged Rainfall data show country-regional patterns while NDVI shows regional-local patterns.
- Rainfall is the better predictor (anticipating growth), while NDVI is the better performance assessor (after-season quantified assessment of performance).
- Rainfall excludes influences of local soil and management aspects (etc.), while NDVI includes such impacts.

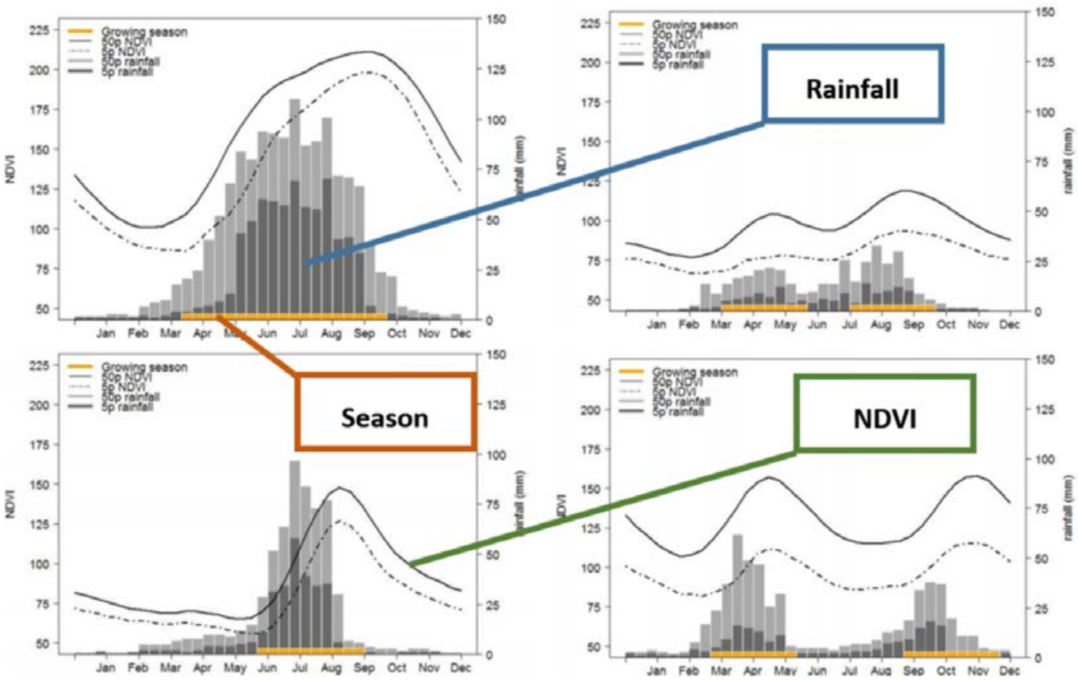

Figure 4:
Rainfall vs NDVI data.

Enabling Crop Analytics At Scale

Dynamic Area Sampling Frames for Improved Crop Analytic

The Enabling Satellite-based Crop Analytics at Scale (ECAAS) Initiative is a multi-phase project that aims to catalyze the development, availability, and uptake of agricultural remote-sensing data and subsequent applications in smallholder farming systems. The initiative is funded by The Bill & Melinda Gates Foundation and implemented by Tetra Tech.

info.ecaas@tetratech.com
cropanalytics.net

Final report

September 2021
IFPRI Project No. 600278.000.001 520-01-01
Final Report for Workstream 1: Dynamic Area Sampling Frame

Authors

Kees de Bie, University of Twente/ITC
Andy Nelson, University of Twente/ITC