

SICKLE: A Multi-Sensor Satellite Imagery Dataset Annotated with Multiple Key Cropping Parameters

Depanshu Sani¹, Sandeep Mahato², Sourabh Saini¹, Harsh Kumar Agarwal¹,
 Charu Chandra Devshali², Saket Anand¹, Gaurav Arora¹, Thiagarajan Jayaraman²

¹Indraprastha Institute of Information Technology, Delhi, India

²MS Swaminathan Research Foundation, Chennai, India

<https://sites.google.com/iiitd.ac.in/sickle/home>

Abstract

The availability of well-curated datasets has driven the success of Machine Learning (ML) models. Despite greater access to earth observation data in agriculture, there is a scarcity of curated and labelled datasets, which limits the potential of its use in training ML models for remote sensing (RS) in agriculture. To this end, we introduce a first-of-its-kind dataset called **SICKLE**, which constitutes a time-series of multi-resolution imagery from 3 distinct satellites: Landsat-8, Sentinel-1 and Sentinel-2. Our dataset constitutes multi-spectral, thermal and microwave sensors during January 2018 – March 2021 period. We construct each temporal sequence by considering the cropping practices followed by farmers primarily engaged in paddy cultivation in the Cauvery Delta region of Tamil Nadu, India; and annotate the corresponding imagery with key cropping parameters at multiple resolutions (i.e. 3m, 10m and 30m). Our dataset comprises 2,370 season-wise samples from 388 unique plots, having an average size of 0.38 acres, for classifying 21 crop types across 4 districts in the Delta, which amounts to approximately 209,000 satellite images. Out of the 2,370 samples, 351 paddy samples from 145 plots are annotated with multiple crop parameters; such as the variety of paddy, its growing season and productivity in terms of per-acre yields. Ours is also one among the first studies that consider the growing season activities pertinent to crop phenology (spans sowing, transplanting and harvesting dates) as parameters of interest. We benchmark **SICKLE** on three tasks: crop type, crop phenology (sowing, transplanting, harvesting), and yield prediction.

1. Introduction

Satellite imagery has emerged as a widely applicable RS tool for researchers across domains like wildlife conservation, climate science, and economics among others; due to

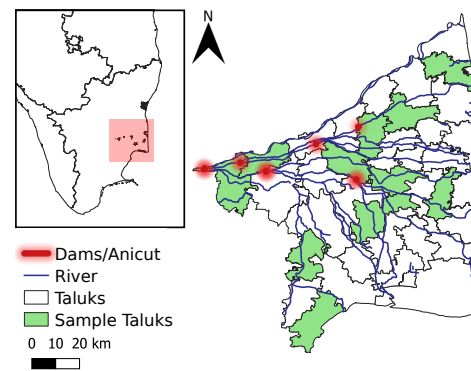


Figure 1. Schematic diagram of the complete study region in Tamil Nadu, India. Field data was collected from the highlighted blocks.

its ability to generate policy-relevant data for large and inaccessible areas in a cost-effective manner. Despite the abundance of satellite data, there is a lack of curated datasets that are annotated with field data, thereby limiting its potential applicability for training and evaluating ML models. Agriculture is also one such domain where collecting field data involves conducting ground-based surveys, which is expensive and laborious. There exist publicly available datasets that target specific tasks, such as crop identification [5, 26, 31], land cover classification [5, 31] and yield estimation [31] among others. These datasets provide value when researchers are interested in singular line of inquiry that requires one dataset at a time. For instance, a crop yield prediction model can be trained using SUSTAINBENCH’s crop yield dataset [31] and a transfer learning approach can then be adopted. However, the dataset only has county-level annotations, which precludes a high-resolution, plot-level analysis. Furthermore, most available datasets have different characteristics and cover different geographies; for instance, the PASTIS [19] has multi-spectral information for crop type mapping within the French metropolitan territory; whereas Agriculture-Vision [2] has only Red, Green, Blue

and NIR spectral bands for crop anomaly detection across the US; and SUSTAINBENCH’s field delineation task has only visible spectrum bands in France. Therefore, it is not possible to train and evaluate a single model for different, yet closely related tasks, e.g., harvesting dates and yield estimates. Moreover, the time-series samples in these datasets are produced using an arbitrary sequence length, which might affect the overall performance. For example, each sequence in PASTIS-R [20] contains observations taken between Sep’18 – Nov’19 and has no details regarding the crop phenology dates. It is likely that multiple crops have been cultivated in this period and hence the quality of input is degraded, ultimately affecting the performance. All the aforementioned datasets lack labels of key cropping parameters, which, if predicted accurately, could lead to better-informed downstream decisions, whereas [9] provides the key cropping parameters but lacks the access to satellite imagery data, which adds an overhead to collect and process the data from multiple satellites making the training process cumbersome. To the best of our knowledge, there exists no prior dataset that contains multi-sensor satellite data annotated with multiple cropping parameters for the same set of plots. To tackle the mentioned challenges, we make the following contributions that can significantly impact the quality of deployment and hence promote applicability of RS and ML in agriculture:

1. We introduce a first-of-its-kind dataset, **SICKLE** (Satellite Imagery for Cropping annotated with Key-parameter LabEls), containing images from multiple satellites (Landsat-8, Sentinel-1 and Sentinel-2) having a variety of sensors (optical, thermal and radar) with annotations of multiple agricultural cropping parameters for each plot. These annotations are created at 3 scales, i.e., 30m, 10m and 3m spatial resolution.
2. We organize the dataset in a way that can be readily used by researchers from multiple domains, including agronomy, RS and ML. Multiple cropping parameters for the same set of plots allow multi-task learning, which can help learn better feature representation. The availability of multi-scale annotations enables the generation of high-resolution (HR) inference maps from low-resolution (LR) images. Moreover, the individual bands at actual resolutions and the metadata are provided, which can help develop a better band interpolation strategy. The presence of time-series image sequences from different satellites of the same geographical location can help evolve cross-satellite fusion, synthetic band generation and forecasting techniques.
3. We include the plot-level crop phenology dates, i.e., sowing, transplanting and harvesting dates, in the dataset. In the absence of these dates, the estimation of cropping parameters using time-series data can suffer

due to interference from other seasonal crops planted in cases of heterogeneous farming.

4. We present a novel strategy for preparing time-series data over a seasonal temporal window that is consistent with the *regional cropping standards* that are typically followed by the farmers for crop production, which improves the robustness and correctness of a solution in a real-world deployment.

2. Related Work

2.1. Crop Type Mapping

With the increased application of multi-temporal datasets for crop type mapping, the classification methods have transitioned from applying Maximum Likelihood based approach [32] to modern algorithms such as Support Vector Machines [10, 14, 27], Decision Trees [4, 30], Random Forest [8, 12] and neural networks [1, 33]. However, the availability of such datasets is limited. Often, there is an imbalance with respect to the availability of labelled datasets across geographical regions. For instance, Europe has an abundance of large, densely annotated datasets, whereas, for regions like South Asia and Africa, labelled datasets are sparsely available. Kehs et al. [9] provide a field dataset collected from May 6, 2019 – June 9, 2019 having plot boundaries, crop type labels, irrigation status, density of green leaf area and other key parameters for 474 plots in Northern Busia county, Kenya. CV4A Kenya Crop Type Competition [6] dataset with segmented sentinel-2 tiles is also available for Kenya. For Ghana and Sudan, SUSTAINBENCH provides data inputs as growing season time series of imagery from three satellites: Sentinel-1, Sentinel-2, and PlanetScope in 2016 and 2017 for different crop types [31]. Some other datasets published particularly focusing on crop type target class for the European regions are [18, 20, 24]. At the global level CropHarvest collection is available with crop/non-crop and agricultural class labels, collected from publicly available datasets and covering 343 labels [26]. Except for [9] all the datasets listed here are labelled with crop type information only, but on the other hand, [9] has no satellite imagery associated with it.

2.2. Crop Phenology Dates

Estimating start and end of the season requires phenology metrics to be derived from satellite images. Methods such as shape model fitting of time-series indexes such as NDVI and EVI are widely used for generating phenology metrics of a crop season [11, 23]. For paddy crops, SAR-based time-series are also applied to derive these season dates [29]. Phenology dates are often used for rule-based crop type classification [23]. The training data provided by [9] has labels on the planting and harvest dates for Maize at plot level. However, this data has not been benchmarked

yet, possibly due to unavailability of any satellite images with the dataset.

2.3. Crop Yield

Current methods for yield estimation use empirical as well as process-based models. ML-based empirical models can produce reliable estimates of paddy yield [15]. Random Forest models with Sentinel-1 data have been used to predict paddy yield with high accuracy [3]. [22] used CNN-LSTM with Sentinel-2 images for tehsil (block) level wheat predictions and achieved 50% better performance over other models. The Process-based models mentioned here use dynamic crop models with remote sensing inputs along with other parameters to simulate crop yield. SUSTAINBENCH crop yield datasets provides county-level yields for 857 counties in the United States, 135 in Argentina and 32 in Brazil for the years 2005 – 2016 [28].

2.4. Lack of Multiple Downstream Tasks

There is a lack of publicly available datasets that can be used for predicting multiple key cropping parameters. One of the largest datasets available for multiple tasks, such as crop type mapping and field delineation, is [31]. However, the geographical coverage of these datasets varies depending on the specific tasks involved. In contrast, both semantic and panoptic segmentation masks are supplied by [19, 20] for a specific area; a feature that is lacking in [31]. Addressing anomaly patterns, plot boundaries along with nine distinct types of these patterns are provided by [2]. Despite the fact that few available datasets provide multiple cropping patterns, neither of them can be used to train a multi-tasking end-to-end trainable network for reasons outlined above.

3. Dataset Description

The Cauvery Delta is a major rice cultivation region in Tamil Nadu, India, where farmers cultivate one or two crops of paddy, depending upon water availability. However, due to significant shifts in the its agrarian landscape, an Indian Council for Agricultural Research (ICAR) study in 2013 reclassified four districts of the Cauvery Delta from dry semi-humid to semi-arid conditions [17]¹. Given its centrality to rice production in India, the inferences via well-designed ML methods from this dataset would enable analysis of the impact of aridification on paddy yields and other cropping patterns in the region. Moreover, we anticipate that the challenges of the *SICKLE* dataset (discussed hereafter in section 4) will also drive novel vision techniques. We highlight the key features of *SICKLE*, mentioned in section 1, and show a detailed comparison with the recent literature in Tables 1 and 2.

¹The supplementary material has more details about the study region, data sampling and dataset statistics.

3.1. Data Acquisition

Ground-based surveys were conducted throughout the study region to collect field data for the time span of January 2018 – March 2021 for 388 individual plots. We interacted with the farmers of these plots to collect information on the type and variety of the crop grown in each agricultural season, its growing season duration and productivity in terms of yield. Paddy being the primary focus of this study, the phenology dates and the crop yield were only gathered for the seasons in which paddy was cultivated. Moreover, the surveys were conducted from June 2021 - February 2022 and hence the collected field data is entirely dependent on the ability of the farmer to recall the details of the crops grown in the past. As a result, out of 2,370 samples, we have phenology dates and crop yield estimates for 351 paddy samples only. For the remaining 656 paddy samples, the farmers could only recall rough estimates of the season duration but not the crop yields. We also collected the GPS coordinates of the centroid of the plot, the block and the district it lies in, the relative location to the Cauvery Delta, and area of the plot in acres. The specific coordinates of each plot are withheld for privacy reasons. It is to be noted that ground-based surveys are expensive, inefficient and difficult to scale. Therefore, we surveyed only a subset of all the plots from multiple blocks at random.

For the entire study region, we acquired the tile images from 3 publicly available satellites, Landsat-8, Sentinel-2² and Sentinel-1, using the products ‘LANDSAT/LC08/C02/T1_L2’, ‘COPERNICUS/S2_SR’ and ‘COPERNICUS/S1_GRD’ from the Google Earth Engine platform [7]. Unlike previous works, we downloaded all the available bands, including the Quality Assessment and the derived bands, from all these satellites at their original resolutions band-wise, along with metadata. These bands and metadata files can be used as a prior for various tasks. For instance, Sentinel-2 provides a derived Scene Classification (SCL) band, which classifies each pixel as bare soil, vegetation, water, cloud, etc., that can provide more information about the image. Now, if the crop cover map suggests that the water pixel (based on the prior) is a paddy crop, we can conclude that the model is not accurate.

3.2. Data Annotation

For the annotation task, we created vector files for each of the 388 surveyed plots using QGIS [16] by manually drawing the polygon boundaries around the GPS coordinates of each plot by visualizing the high-resolution image-based geographical maps available with the product. Since the broader objective of *SICKLE* is to be able to gener-

²The surface reflectance (SR) data for Sentinel-2 is not available before Dec’18. Moreover, the top-of-atmosphere (TOA) data is also archived and is not accessible via Google Earth Engine. We requested the TOA tiles from Copernicus and converted them to SR images using Sen2Cor [13].

Tasks	[31]	[5]	[2]	[21]	[20]	[26]	<i>SICKLE</i>
Crop Type Semantic Segmentation	✓	✓		✓	✓	✓	✓
Crop Type Panoptic Segmentation					✓		✓
Cropland Segmentation	✓	✓					
Field Delineation	✓		✓				✓
Phenology Date Prediction							✓
Crop Yield Prediction	✓						✓
Crop Anomaly Detection			✓				
All above tasks for same set of plots	✗	✗	✓		✓		✓
Multi-Task Learning			✓				✓
Multi Image Super Resolution							✓
Cross-Satellite & Cross-Sensor Fusion					✓		✓
Synthetic Band Generation					✓		✓
HR prediction using LR images							✓

Table 1. A comparison of *SICKLE* with related datasets (SUSTAINBENCH [31], Radiant ML Hub [5], Agriculture-Vision [2], Pixel Set [21], PASTIS-R [20] and Crop Harvest [26]) based on the tasks that can be performed using them. The bottom 4 tasks are not only related to the agricultural domain but are also applicable for remote sensing community.

Characteristics	[31]	[5]	[20]	[2]	<i>SICKLE</i>
Time series data	✓	✓	✓		✓
Multiple annotations for same plots				✓	✓
Annotations at multiple resolutions					✓
Multi-sensor data for all tasks			✓		✓
Consistent with the regional cropping practises					✓
Number of time-series samples	[1, 966 – 10, 332]	Variable	2, 433	NA	2, 370

Table 2. A comparison of *SICKLE* with the related datasets that can be used for multiple tasks (SUSTAINBENCH [31], Radiant ML Hub [5], PASTIS-R [20] and Agriculture-Vision [2]) based on their characteristics.

ate inferences for the entire Cauvery region rather than a plot-wise analysis, we are interested in a region-wise analysis where the individual plots are unknown. Therefore, we create a rectangular buffer area of $320m \times 320m$ around each plot. As many of our surveyed plots were from the same local region, the buffer area for such plots was highly overlapping. Thus, to regularize the relative locations of all the plots within these extended boundaries, we translated the centroid of the polygon buffers randomly in X and Y directions. These polygon buffer vectors were used for clipping out image patches from the satellite tile images. We identify all the plots within a buffer using the actual polygon vectors and then annotate all of them with *plot_id*, *crop_type*, *sowing_date*, *transplanting_date*, *harvesting_date* and *crop_yield* labels at 30m, 10m and 3m resolutions based on the preparation strategy mentioned in section 3.3. For crop yield, we assume that cropping distribution was uniform within a plot and therefore distribute the amount of yield such that the sum of pixel-wise yield in a plot is equal to the total yield provided by the farmer.

3.3. Dataset Preparation

Regional Standard Growing Season: As discussed in section 1, using time-series data of an arbitrary sequence length will likely deteriorate the quality of cropping parameter prediction due to interference from crops planted in other seasons. If the time-series data spans an arbitrary duration (t_a) that is substantially smaller than the actual growing season of paddy (t_p), then the data might not be able to capture the complete phenological structure and hence will lack the information necessary for accurate prediction. On the other hand, if $t_a \gg t_p$, then the data will likely span multiple seasons and may witness interference due to a multi-crop cultivation leading to poor predictions. Therefore, we argue that the duration of the observations to create time-series input data should depend on the downstream task. Moreover, the crop phenology dates gathered from the farmers might differ from that region’s standard paddy season duration. For instance, while the standard growing season starts in September, a farmer might sow the seeds in August or October, depending on the water availability and




Figure 2. Geographical extent of the surveyed region. Each marker denotes the presence of a surveyed plot. The blue color represents plots in the train set, the green color represents the validation set and the red color represents the test set.

other factors. On the other hand, there might be crops, such as coconut and banana, that can span an entire year for its cultivation. Thus, instead of using arbitrary sequences as in prior work [2, 5, 20, 21, 26, 31], we propose a novel method for preparing time-series data where we consider the length of the regional standard growing season as the duration for creating the time-series data. This strategy leverages the domain knowledge of the crop’s growing season in the region, which is available (as a *default*) even in the absence of information on the crop’s *actual* growing season in a particular plot, which is often the case in a real-world setting. For this purpose, we use Table 2 (supplementary material) as a regional standard obtained from Tamil Nadu Agricultural University (TNAU) [25]. We then curate a data set based on this regional standard for paddy crop cultivation. Specifically, for each regional standard season, we consider all the satellite images of a particular plot between that duration as a single sample. From an application standpoint, we argue that leveraging domain knowledge is crucial for success.

Dataset Creation: As mentioned in section 3.2, instead of clipping satellite images at a plot level, we clipped a patch of $320m \times 320m$ around each plot, which might also contain other plots from the region. Following the technique we proposed in the previous section, we downloaded and clipped the images available in the regional standard growing season only.³

Creating Train/Val/Test Splits: Randomly dividing the plots into training, validation and test sets created uneven distributions across these splits. We observed that the distribution is highly skewed in terms of almost all the key parameters when using random splitting. Some of the possible reasons for this could be the sensitivity of crop parameters on geographical location, uneven distribution of size and count of the plots across different patches, and unavailability of all the parameters for some samples. Thus, we created plot-level parameter histograms using pixel-wise annotations of all the samples associated with each plot; and then concatenated all the parameter histograms to create a

³The dataset directory structure is explained in supplementary material.

1-dimensional histogram for each plot in the dataset. We then iteratively split the plots using a stratified split with different random seeds. For each iteration, we compute the Wasserstein metric (Earth Mover’s Distance) between the training, validation and test plot histograms. We executed this method for over 200,000 iterations and chose the split with the least Mean Wasserstein Distance (Fig 2).

4. Challenges

Along with the applicability of *SICKLE* for developing better algorithms to predict multiple key cropping parameters, we also aim to bridge the gap between the agriculture, RS and ML communities. The proposed dataset has different modalities of raw data from distinct satellites that would allow ML engineers to leverage data-driven methods as well as domain-specific knowledge, which could be exploited to develop practically usable systems.

1. Currently, the publicly available datasets ignore optical images having clouds based on a certain threshold. But paddy is also grown during rainy seasons, thus will contain heavy cloud covers for such samples. When the images corresponding to these samples are ignored during the development phase of a solution, the inferences generated after the deployment will be unanticipated. It is also important to know that including such images will not contribute much towards the solution’s performance because of the limitations of optical sensors. However, *SICKLE* is also encapsulated with radar observations, which are not sensitive to clouds and other atmospheric interference. Thus sophisticated methods for fusing multiple sensory observations is a necessity.
2. The annotations masks are based on the ground surveys conducted during January 2018 - March 2018. Compared to the crop-cutting based data collection, conducting ground-based surveys is economically more feasible, but the ground truths are much noisier. Thus, considering these annotations as ground truths would be logically incorrect. Alternatively, one can assume these annotations to be weak signals about the ground truth information and thus develop algorithms based on weak supervision.
3. Majority of the plots in this study region are small farms, with more than 95% plots having an area of less than or equal to 1 acre and an average size of 0.38 acres. Thus low-resolution crop parameter masks tend to smoothen the edges of individual plots (Fig 3). Moreover, the masks at 30m resolution are unavailable for very small plots (Fig 4). Sometimes, small plots adjacent to each other are merged due to the resolution (Fig 5). The availability of high-resolution




Figure 3. Masks at LR tend to smoothen the plot boundaries.




Figure 4. Masks are not available for smallholding farms at LR.




Figure 5. Masks for small adjacent plots are merged at LR.

masks allows us to develop methods for generating high-resolution inferences. Figures 3, 4 and 5 depict the boundary and area of individual surveyed plots.

4. Despite conducting an extensive ground survey, most of the area in the study region remains unsurveyed. Thus the number of data samples available over a season for predicting crop parameters is small. This restriction inspires small-sample and data-efficient learning methods.
5. Although the sampling strategy allowed us to gather data from all over the study region, there are several other factors that influence the distribution of crop parameters. Factors such as distance from the irrigation source, soil conditions, etc. can vary across different plots and may also influence the prediction accuracy. The predictive models would be required to factor in these changes or be designed to be robust to them in order to be deployed.

5. Methodology

Although *SICKLE* can be used to perform various tasks, such as synthetic band generation, image super-resolution,

multi-task learning, etc., we benchmark it on 3 tasks, i.e. crop segmentation, phenology date and crop yield prediction. Phenology date prediction is further composed of sowing, transplanting and harvesting date predictions. We benchmark the dataset using the code repository of [19], which includes the implementation of U-Net 3D, ConvLSTM and U-TAE among others. We slightly modified the backbone architecture of all the models to output a pixel-wise embedding, instead of a class map, and add a convolution output layer. This modification allows us to perform cross-satellite fusion.

5.1. Crop Segmentation

We pose this as a binary semantic segmentation problem of classifying each pixel as paddy or non-paddy. But one can also pose this as a multi-class semantic segmentation problem as finer labels for all the non-paddy crops are also provided. We use Pytorch’s implementation of the Cross-Entropy Loss. Therefore we don’t need to use any activation function explicitly on the logits. This makes the architecture reusable for regression tasks, discussed in the next sections. Along with the time-series crop segmentation, we also demonstrate the performance of U-Net 2D and DeepLabV3+ using a single image as the input.

5.2. Phenology Dates Prediction

Phenology date prediction can be posed as a segmentation as well as a regression problem. We know the duration of regional standard season a priori; hence, we can classify one of the season days as the sowing, transplanting or harvesting day. A common approach in such a problem setting is to use a cross-entropy loss, which is not ideal in this scenario because it treats all the misclassifications equally; hence a date that is misclassified by ± 1 day will be the same as a misclassification of ± 100 days. Therefore, we pose this problem as a regression problem and use the root mean squared error (RMSE) loss. A regressor also makes it possible to account for outliers, as mentioned in section 4. Because the architecture discussed in the previous section does not contain any activation function at the output layer, it is directly reusable for regression tasks.

5.3. Crop Yield Prediction

Similar to 5.2, we pose yield prediction as a regression task and use the RMSE loss. We emphasize that even after incorporating the regional standards, it might be challenging to estimate the yield if multiple vegetation signals are available for a pixel. Let’s consider a case where the start of the regional standard season contains the harvesting period of a crop *A* grown in the previous season and the actual growing season for the crop *B* starts just after the harvesting period of *A*. Because there are multiple vegetation signals within the standard season, the estimates will be ambigu-

ous. If B turns out to be a damaged crop but A has high productivity, the yield estimates will be unanticipated and incorrect. Therefore, we propose a novel strategy of using the phenology dates to estimate the *actual growing season*, which is then used for yield prediction. This follows two implementations. One can argue that time-series data is necessary to monitor productivity throughout the season. One can also claim that a single image of the crop can estimate crop yield during the harvest season. Therefore, we benchmark the results of crop yield prediction on both these implementations.

5.4. Cross-Satellite Fusion

The data from different satellites can be fused in various ways. The prominent fusion techniques are early fusion: the inputs are concatenated and then processed, and late fusion: the embeddings from each satellite data are concatenated and further processed. Because the satellites have different revisiting frequencies and hence acquire images on different dates, making it inappropriate to do an early fusion. Therefore, we adopt the late fusion technique as shown in Fig 6, wherein we use the architectures mentioned in the previous sections to generate a feature embedding instead of a class map. We concatenate the embeddings of all the participating satellites channel-wise and then use a convolution output layer.

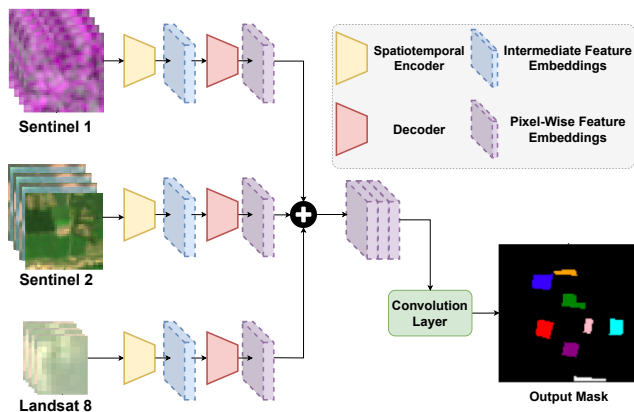


Figure 6. We adopted this architecture to benchmark the dataset on all the tasks. The time-series sequence from each participating satellite is processed using a spatiotemporal encoder to generate an intermediate feature embedding. The decoder accepts these embeddings and generates a pixel-wise embedding. All the pixel-wise embeddings from all the participating satellites are channel-wise concatenated and fed to a convolution output layer. If only one satellite is participating in fusion, it is equivalent to a single satellite prediction task.

6. Experiments

6.1. Experimental Setup

We use all the 2,370 samples for benchmarking crop segmentation. Although rough estimates of the phenology dates are available for all the 1,007 paddy samples, we benchmark phenology date and yield estimation on 351 paddy samples for which all the other crop parameters are also available. We divided the samples into 80%, 10% and 10% splits to create training, validation and test sets using the technique mentioned in section 3.3. Although, while making the splits we ensured that no plot from the validation and testing sets is exposed during the training, it was observed that plots in validation and test sets were still being exposed during training because some of these plots are adjacent to the plots in the training set. Therefore, by masking them out, we ensure that such pixels are ignored while training. We use a 50% dropout strategy to train all the models for 100 epochs with a cosine anneal scheduler having starting and minimum learning rate of 0.1 and $1e-4$, respectively, which is kept constant after 75 epochs. We set the starting learning rate for all the single image tasks to be $1e-3$. We reduced the number of parameters by reducing the number of convolution blocks because of the small size of the dataset as compared to [19]. We also adopted data augmentation techniques, which included random horizontal and vertical flip along with random brightness adjustment and Gaussian blur. While the flipping augmentation helps in better spatial regularization of the plots, brightness and blur augmentations help by synthesizing the effect of climatic and atmospheric interference. Moreover, because there is an imbalance of crop type distribution with paddy being less, we use a weighted cross-entropy loss with weights 0.62013 and 0.37987 for paddy and non-paddy crops, respectively. For a fair comparison between the reported results, we resize the Landsat 8 images to generate predictions at $10m$ resolution.

6.2. Metric

We report the mean and standard deviation of the results obtained with different seed values (ranging from 0 to 4) on the test set. For segmentation tasks, we report the results using pixel-wise accuracy, f1-score and intersection-over-union (IoU). The results include macro-averaged as well class-wise metrics. For regression tasks, we report the root mean squared error (RMSE), mean absolute error (MAE) and mean absolute percentage error (MAPE). Because the phenology date prediction has a specified range of dates, we normalize the MAE by the maximum possible date, i.e. 183, instead of the ground truth value to compute MAPE, because otherwise, its value overshoots due to the minimal range of sowing date (≈ 0).

Task	Metric	L8	S2	S1	Fusion
Crop Type (SI)	IoU (%)	47.73% \pm 1.77%	54.87% \pm 3.08%	64.35% \pm 4.82%	-
Crop Type	IoU (%)	56.04% \pm 5.84%	78.12% \pm 3.48%	81.77% \pm 6.60%	81.07% \pm 5.77%
Sow Date	MAE (days)	2.66 \pm 0.961	2.30 \pm 0.611	3.61 \pm 0.898	2.33 \pm 0.639
Transplant Date	MAE (days)	6.20 \pm 1.030	6.36 \pm 2.164	7.23 \pm 0.779	6.16 \pm 1.770
Harvest Date	MAE (days)	9.86 \pm 0.736	8.83 \pm 1.520	10.08 \pm 0.561	10.75 \pm 3.389
Crop Yield (SI)	MAPE (%)	46.74% \pm 3.82%	60.44% \pm 14.50%	48.35% \pm 7.64%	-
Crop Yield (RS)	MAPE (%)	54.00% \pm 9.67%	72.38% \pm 8.74%	71.81% \pm 17.27%	70.35% \pm 13.75%
Crop Yield (AS)	MAPE (%)	59.38% \pm 14.75%	73.59% \pm 9.81%	65.66% \pm 16.24%	64.56% \pm 13.77%

Table 3. Results for the benchmarking tasks. Single-image experiments are denoted with SI in parenthesis. The results are reported using the same benchmarking model (U-Net 3D for time-series and U-Net 2D for single image) for a fair comparison. RS denotes the experiment when using Regional Standards to create the time-series input, whereas AS denotes the one using Actual Season. The supplementary material includes a list of exhaustive experiments and their results.

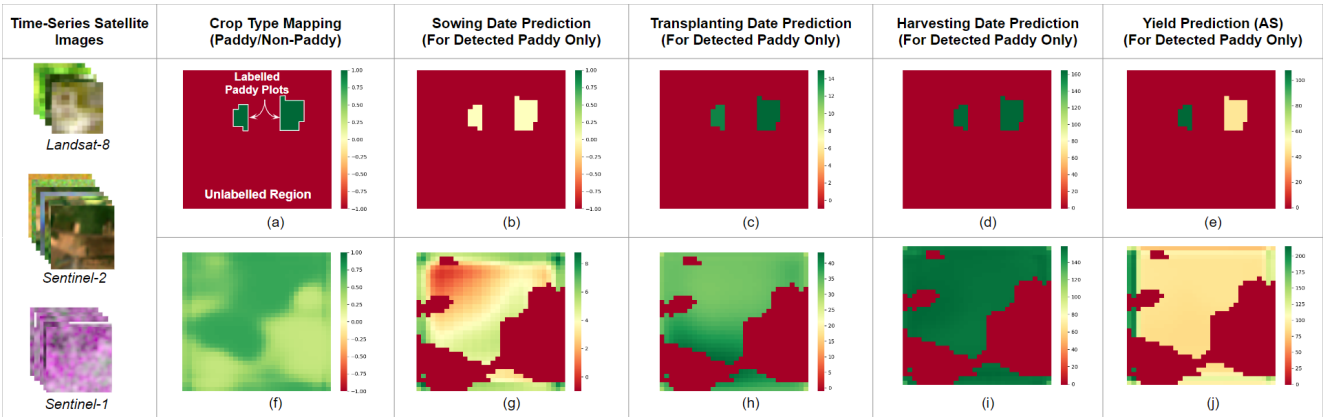


Figure 7. An illustration of the outputs obtained from U-Net 3D (Fusion) on the test set. Figures (a) - (e) are the ground truth annotations and figures (f) - (j) are the output maps generated. All the pixels that were predicted as non-paddy in (f) were masked out in (g) - (j).

6.3. Results

We performed experiments with multiple state-of-the-art architectures for time series and single-image predictions wherever applicable. Table 3 includes the benchmarking results for each task. It is important to note that these results only have the experiments done on the same backbone architecture for a fair comparison. A complete list of experiments is available in the supplementary material. The best results are highlighted in **green** color and the competing (second best) are represented using a **bold font**.

7. Conclusion

This paper introduces **SICKLE**, a first-of-its-kind dataset consisting of satellite image time-series data from multiple satellites with multiple key cropping parameters annotated at multiple resolutions. The dataset enables researchers to develop an end-to-end pipeline wherein one can predict the key cropping parameters at a plot-level given the satellite images of the complete region. Such an end-to-

end pipeline is necessary for real-world deployments where gathering the individual plot-level insights is a crucial task.

Besides the cropping parameter annotations, the dataset also serves as the first step towards building a research ecosystem for researchers across agriculture, remote sensing and machine learning domains by emphasizing on various fundamental challenges that could potentially benefit the overall development of the community.

Acknowledgement

Depanshu Sani was supported by Google’s AI for Social Good “Impact Scholars” program, 2021. Saket Anand gratefully acknowledges for the partial support from the Infosys Center for Artificial Intelligence at IIIT-Delhi. We also appreciate Parichya Sirohi’s contributions in the early stages of the project. Additionally, we wish to express our gratitude to Dr. Gopinath R. and Dr. Rajakumar R. from Ecotechnology, MS Swaminathan Research Foundation, Chennai, for their valuable inputs concerning the study area and assistance with field data collection.

References

- [1] Hugo Crisóstomo de Castro Filho, Osmar Abílio de Carvalho Júnior, Osmar Luiz Ferreira de Carvalho, Pablo Pozzobon de Bem, Rebeca dos Santos de Moura, Anesmar Olino de Albuquerque, Cristiano Rosa Silva, Pedro Henrique Guimarães Ferreira, Renato Fontes Guimarães, and Roberto Arnaldo Trancoso Gomes. Rice crop detection using lstm, bi-lstm, and machine learning models from sentinel-1 time series. *Remote Sensing*, 2020. **2**
- [2] Mang Tik Chiu, Xingqian Xu, Yunchao Wei, Zilong Huang, Alexander G. Schwing, Robert Brunner, Hrant Khachatryan, Hovnatan Karapetyan, Ivan Dozier, Greg Rose, David Wilson, Adrian Tudor, Naira Hovakimyan, Thomas S. Huang, and Honghui Shi. Agriculture-vision: A large aerial image database for agricultural pattern analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. **1, 3, 4, 5**
- [3] Kersten Clauss, Marco Ottinger, Patrick Leinenkugel, and Claudia Kuenzer. Estimating rice production in the mekong delta, vietnam, utilizing time series of sentinel-1 sar data. *International journal of applied earth observation and geoinformation*, 73:574–585, 2018. **3**
- [4] B Deschamps, H McNairn, J Shang, and X Jiao. Towards operational radar-only crop type classification: comparison of a traditional decision tree with a random forest classifier. *Canadian Journal of Remote Sensing*, 38(1):60–68, 2012. **2**
- [5] Radiant Earth Foundation. ML hub. Open Library for Earth Observations Machine Learning. **1, 4, 5**
- [6] Radiant Earth Foundation. Cv4a competition kenya crop type dataset. <https://doi.org/10.34911/RDNT.DW605X>, 2020. **2**
- [7] Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment*, 2017. **3**
- [8] Siddharth Hariharan, Dipankar Mandal, Siddhesh Tirodkar, Vineet Kumar, Avik Bhattacharya, and Juan Manuel Lopez-Sanchez. A novel phenology based feature subset selection technique using random forest for multitemporal polsar crop classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(11):4244–4258, 2018. **2**
- [9] Annalyse Kehs, Peter McCloskey, and David Hughes. Busia county, kenya agricultural fields may 2019. *Frontiers in Sustainable Food Systems*, 2021. **2**
- [10] Pradeep Kumar, Dileep Kumar Gupta, Varun Narayan Mishra, and Rajendra Prasad. Comparison of support vector machine, artificial neural network, and spectral angle mapper algorithms for crop classification using liss iv data. *International Journal of Remote Sensing*, 36(6):1604–1617, 2015. **2**
- [11] Licong Liu, Ruyin Cao, Jin Chen, Miaogen Shen, Shuai Wang, Ji Zhou, and Binbin He. Detecting crop phenology from vegetation index time-series data by improved shape model fitting in each phenological stage. *Remote Sensing of Environment*, 277:113060, 2022. **2**
- [12] John A Long, Rick L Lawrence, Mark C Greenwood, Lucy Marshall, and Perry R Miller. Object-oriented crop classification using multitemporal etm+ slc-off imagery and random forest. *GIScience & Remote Sensing*, 50(4):418–436, 2013. **2**
- [13] Magdalena Main-Knorn, Bringfried Pflug, Jerome Louis, Vincent Debaecker, Uwe Müller-Wilm, and Ferran Gascon. Sen2Cor for Sentinel-2. In Lorenzo Bruzzone, editor, *Image and Signal Processing for Remote Sensing XXIII*, volume 10427, page 1042704. International Society for Optics and Photonics, SPIE, 2017. **3**
- [14] Ajay Mathur and Giles M Foody. Crop classification by support vector machine with intelligently selected training data for an operational application. *International Journal of Remote Sensing*, 29(8):2227–2240, 2008. **2**
- [15] Kodimalar Palanivel and Chellammal Surianarayanan. An approach for prediction of crop yield using machine learning and big data techniques. *International Journal of Computer Engineering and Technology*, 10(3):110–118, 2019. **3**
- [16] QGIS Development Team. *QGIS Geographic Information System*. Open Source Geospatial Foundation, 2009. **3**
- [17] B. M. K. Raju, K. V. Rao, B. Venkateswarlu, A. V. M. S. Rao, C. A. Rama Rao, V. U. M. Rao, B. Bapuji Rao, N. Ravi Kumar, R. Dhakar, N. Swapna, and P. Latha. Revisiting climatic classification in india: a district-level analysis. *Current Science*, 105(4):492–495, 2013. **3**
- [18] Marc Rußwurm, Charlotte Pelletier, Maximilian Zollner, Sébastien Lefèvre, and Marco Körner. Breizhcrops: A time series dataset for crop type mapping. *arXiv preprint arXiv:1905.11893*, 2019. **2**
- [19] Vivien Sainte Fare Garnot and Loic Landrieu. Panoptic segmentation of satellite image time series with convolutional temporal attention networks. *ICCV*, 2021. **1, 3, 6, 7**
- [20] Vivien Sainte Fare Garnot, Loic Landrieu, and Nesrine Chehata. Multi-modal temporal attention models for crop mapping from satellite time series. *ISPRS Journal of Photogrammetry and Remote Sensing*, 2022. **2, 3, 4, 5**
- [21] Vivien Sainte Fare Garnot, Loic Landrieu, Sebastien Giordano, and Nesrine Chehata. Satellite image time series classification with pixel-set encoders and temporal self-attention. *CVPR*, 2020. **4, 5**
- [22] Sagarika Sharma, Sujit Rai, and Narayanan C Krishnan. Wheat crop yield prediction using deep lstm model. *arXiv preprint arXiv:2011.01498*, 2020. **3**
- [23] Mrinal Singha, Bingfang Wu, and Miao Zhang. An object-based paddy rice classification using multi-spectral data and crop phenology in assam, northeast india. *Remote Sensing*, 8(6):479, 2016. **2**
- [24] Dimitrios Sykas, Maria Sdraka, Dimitrios Zografakis, and Ioannis Papoutsis. A sentinel-2 multiyear, multicountry benchmark dataset for crop classification and segmentation with deep learning. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 15:3323–3339, 2022. **2**
- [25] Season and varieties :: Rice, 2013. Tamil Nadu Agricultural University. **5**

- [26] Gabriel Tseng, Ivan Zvonkov, Catherine Lilian Nakalembe, and Hannah Kerner. Cropharvest: A global dataset for crop-type classification. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. [1](#), [2](#), [4](#), [5](#)
- [27] Shyamal S Virnodkar, Vinod K Pachghare, VC Patil, and Sunil Kumar Jha. Application of machine learning on remote sensing data for sugarcane crop classification: a review. *ICT analysis and applications*, pages 539–555, 2020. [2](#)
- [28] Anna X. Wang, Caelin Tran, Nikhil Desai, David Lobell, and Stefano Ermon. Deep transfer learning for crop yield prediction with remote sensing data. In *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies, COMPASS '18*, New York, NY, USA, 2018. Association for Computing Machinery. [3](#)
- [29] Yanyan Wang, Shenghui Fang, Lingli Zhao, Xinxin Huang, and Xueqin Jiang. Parcel-based summer maize mapping and phenology estimation combined using sentinel-2 and time series sentinel-1 data. *International Journal of Applied Earth Observation and Geoinformation*, 108:102720, 2022. [2](#)
- [30] Wei Wei, Dawid Polap, Xiaohua Li, Marcin Woźniak, and Junzhe Liu. Study on remote sensing image vegetation classification method based on decision tree classifier. In *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, pages 2292–2297. IEEE, 2018. [2](#)
- [31] Christopher Yeh, Chenlin Meng, Sherrie Wang, Anne Driscoll, Erik Rozi, Patrick Liu, Jihyeon Lee, Marshall Burke, David Lobell, and Stefano Ermon. Sustainbench: Benchmarks for monitoring the sustainable development goals with machine learning. In *Thirty-fifth Conference on Neural Information Processing Systems, Datasets and Benchmarks Track (Round 2)*, 12 2021. [1](#), [2](#), [3](#), [4](#), [5](#)
- [32] Le Yu, Lu Liang, Jie Wang, Yuanyuan Zhao, Qu Cheng, Luanyun Hu, Shuang Liu, Liang Yu, Xiaoyi Wang, Peng Zhu, et al. Meta-discoveries from a synthesis of satellite-based land-cover mapping research. *International Journal of Remote Sensing*, 35(13):4573–4588, 2014. [2](#)
- [33] Hongwei Zhao, Zhongxin Chen, Hao Jiang, Wenlong Jing, Liang Sun, and Min Feng. Evaluation of three deep learning models for early crop classification using sentinel-1a imagery time series—a case study in zhanjiang, china. *Remote Sensing*, 11(22):2673, 2019. [2](#)