

Beyond Flat Classifiers: Practical methodologies for regionally accurate and relevant land use and land cover classification using Landsat and Sentinel data

CHAHAT BANSAL, Amarnath and shashi khosla school of information technology, IIT Delhi, New Delhi, India

ANANJAN NANDI, Computer science and engineering, IIT Delhi, New Delhi, India

BALAKUMARAN RAMACHANDRAN, Water, Environment, Land and Livelihoods (WELL) Labs, Bengaluru, India

CRAIG DSOUZA, Water, Environment, Land and Livelihoods (WELL) Labs, Bengaluru, India

DIVYANSH MITTAL, Computer science and engineering, IIT Delhi, New Delhi, India

EASHAN BAJAJ, Computer science and engineering, IIT Delhi, New Delhi, India

RAMAN KUMAR, Amarnath and shashi khosla school of information technology, IIT Delhi, New Delhi, India

SACHIT SACHDEVA, Computer science and engineering, IIT Delhi, New Delhi, India

SAKETH VISHNUBHATLA, Amarnath and shashi khosla school of information technology, IIT Delhi, New Delhi, India

VATSAL AGARWAL, Computer science and engineering, IIT Delhi, New Delhi, India

AADITESHWAR SETH, Computer Science and Engineering, IIT Delhi, New Delhi, India

India's diverse agricultural landscapes demand a single Land Use and Land Cover (LULC) product integrating both intra-annually static (built-up, tree cover, barren land) and dynamic (water seasonality, cropping intensity) classes for sustainable Natural Resource Management (NRM). Existing LULC products suffer from limited thematic coverage of dynamic processes, imprecise delineation of fragmented smallholder features, limited reproducibility, and poor performance of monolithic classifiers on spectrally similar categories. We introduce a hierarchical decision-tree framework that breaks complex classification tasks into targeted sub-tasks, offering methodological improvements in detecting monsoon water, delineating tree-croplands, and classifying cropping intensity. Class-wise evaluations demonstrate superior performance: SAR water detection achieves an NRMSE of 0.33 (vs. baselines 0.53–0.75), tree-cropland macro-average of 0.94, and cropping intensity macro-average of 0.88 (vs. baseline 0.77). Crucially, this study ensures transparency and reproducibility in LULC mapping. We publicly release four curated datasets alongside the entire classification pipeline implemented on the Google Earth Engine commodity platform.

Authors' Contact Information: Chahat Bansal, Amarnath and shashi khosla school of information technology, IIT Delhi, New Delhi, Delhi, India; e-mail: chahat.bansal@sit.iitd.ac.in; Ananjan Nandi, Computer science and engineering, IIT Delhi, New Delhi, Delhi, India; e-mail: tgk.ananjan@gmail.com; Balakumaran Ramachandran, Water, Environment, Land and Livelihoods (WELL) Labs, Bengaluru, Karnataka, India; e-mail: balakumaran.ramachandran@ifmr.ac.in; Craig Dsouza, Water, Environment, Land and Livelihoods (WELL) Labs, Bengaluru, Karnataka, India; e-mail: craig.dsouza@ifmr.ac.in; Divyansh Mittal, Computer science and engineering, IIT Delhi, New Delhi, Delhi, India; e-mail: divyanshmittal37@gmail.com; Eashan Bajaj, Computer science and engineering, IIT Delhi, New Delhi, Delhi, India; e-mail: eashanb13@gmail.com; Raman Kumar, Amarnath and shashi khosla school of information technology, IIT Delhi, New Delhi, Delhi, India; e-mail: ramank.1137@gmail.com; Sachit Sachdeva, Computer science and engineering, IIT Delhi, New Delhi, Delhi, India; e-mail: sachitsachdeva4@gmail.com; Saketh Vishnubhatla, Amarnath and shashi khosla school of information technology, IIT Delhi, New Delhi, Delhi, India; e-mail: saketh181199@gmail.com; Vatsal Agarwal, Computer science and engineering, IIT Delhi, New Delhi, Delhi, India; e-mail: vatsalagarwal38@gmail.com; Aaditeswar Seth, Computer Science and Engineering, IIT Delhi, New Delhi, Delhi, India; e-mail: aseth@cse.iitd.ac.in.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2834-5533/2026/4-ART

<https://doi.org/10.1145/3806393>

The resulting pan-India output maps at 10m resolution are hosted on CoRE-Stack (digital public good) for easy accessibility and analysis [24], already powering real-world applications in water-security planning [61] and agricultural studies [50].

CCS Concepts: • **Information Systems** → **Geographic information systems**; • **Applied Computing** → *Environmental sciences; Agriculture*; • **Computing Methodologies** → *Machine learning*.

Additional Key Words and Phrases: Land use land cover (LULC) mapping, Hierarchical classification framework, cropping intensity and water seasonality, Optical-SAR fusion, Google Earth Engine (GEE), Open-source pipeline, Data-scarce regions (India)

1 Introduction

In India's diverse agricultural landscape, which spans from arid, rainfed single-cropping systems to intensively irrigated multi-cropping belts, accurate Land Use and Land Cover (LULC) classification is a critical enabler for sustainable development. As agriculture forms the economic backbone for a significant portion of the population, robust LULC mapping is essential not just for monitoring, but for informing strategic Natural Resource Management (NRM). These maps provide the evidence base for vital policy decisions: for instance, identifying single-cropping areas with under-utilized water potential allows for the prioritization of farm ponds under schemes like MGNREGA [63]. Similarly, tracking transitions from traditional crops to plantations elucidates shifts in agricultural resilience. Ultimately, by bridging the gap between observation and action, targeted LULC classification facilitates climate-resilient planning that ensures water security and resource equity in smallholder-dominated landscapes.

The rapid expansion of open-access satellite imagery, combined with advances in cloud computing and machine learning, has unlocked unprecedented opportunities to generate Land Use and Land Cover (LULC) maps at regional to global scales with high temporal frequency. Several LULC products are available for India [13, 49, 72, 117, 123], yet each suffers from different limitations that constrain their practical applicability: (i) their thematic frameworks often exclude intra-annual dynamic classes crucial for representing features such as surface water availability and cropping intensity; (ii) mapping smallholder landscapes demands higher spatial resolution and precision to capture fine-scale features like farm ponds and fragmented croplands [6, 40]; (iii) limited transparency in annotated datasets, output layers, or methodological details restricts validation, reproducibility, and local adaptation [30]; and (iv) their reliance on flat multi-class classification architectures, which require extensive labeled samples across all categories, results in inconsistent performance for spectrally similar classes [38, 115, 118] and exhibit poor extensibility to new application-relevant classes.

This paper presents a LULC classification framework designed for both practical implementation and agricultural relevance, particularly in data-scarce regions such as India, while maintaining adaptability for global applications.

The framework addresses a critical challenge in existing mapping approaches, the limited representation of intra-annual dynamics, by explicitly integrating both static and dynamic classes. These include phenomena such as cropping intensity and surface water variability, optimized to capture patterns across diverse agro-ecological landscapes. It leverages Sentinel-1 and Sentinel-2 datasets to achieve 10 m spatial resolution, the highest publicly available through open-source products, enabling the detection of fine-scale, smallholder features such as farm ponds and fragmented croplands. Full transparency is ensured through the exclusive use of open-access datasets like cropland-tree annotations, multi-scale water body inventories, and cropping intensity labels collected from the Raichur district, and reproducible workflows implemented on cloud platforms like Google Earth Engine.

The primary contribution of this research is the shift from rigid, "flat" multi-class systems to a **hierarchical decision-tree framework**. This approach deconstructs complex taxonomies into targeted sub-problems, treating them as distinct binary or multi-class classification tasks. Each sub-task is solved using a dedicated feature vector and a method tailored to that specific challenge. These components are then orchestrated into a coherent mapping

solution based on thematic hierarchy, per-class accuracy requirements, and computational efficiency. This allows for the production of annual maps at a national scale without exceeding the resource constraints of standard commodity platforms.

This hierarchical design offers several distinct advantages:

- (1) It accommodates labeled datasets of varying sizes, temporal spans, and spatial origins, ensuring broad applicability.
- (2) It supports task-specific feature selection, mitigating the “curse of dimensionality”, where an excessive number of features prohibits consistent performance across all classes in a flat model.
- (3) It provides methodological flexibility by allowing different classification algorithms to be applied as per dataset characteristics and leveraging existing products through targeted refinements to avoid redundant reclassification.
- (4) It is designed to be extensible, allowing for the addition of new, application-specific classes without requiring a full retraining of the model or access to high-performance computing infrastructure.

Regarding specific classification contributions, we enhance the detection of rainfed surface water bodies by utilizing weather-independent Synthetic Aperture Radar (SAR) data, successfully identifying features often missed by products dependent on coarser or optical sensors. Furthermore, we produce finer tree-cropland delineation through the analysis of SAR time-series data. A key methodological contribution of this work is the advancement of **unsupervised classification for determining cropping intensity**, a critical innovation given the absolute scarcity of labeled training data for this metric. We provide a comparative evaluation of these individual classifiers against current baselines to quantify performance gains.

The subsequent sections of this paper review the existing literature on LULC mapping to contextualize our proposed hierarchical principles. We then detail our comprehensive methodology, encompassing taxonomy definition, dataset curation, and the principles for identifying and isolating smaller classification tasks. We provide a detailed breakdown of the task-wise classification strategies and their corresponding results, as well as the decision factors involved in integrating these classifiers into a single, coherent pipeline to deliver scalable LULC maps. Additionally, we present a **temporal correction framework** applied to multi-year outputs; this post-processing step ensures logical coherence across time, resulting in robust, multi-year mapping. The paper concludes with a discussion of our findings, a generic step-wise framework for replicating similar scalable approaches, and an outline of future research directions.

2 Background and Related Work

This section reviews the main building blocks of contemporary LULC mapping pipelines: existing global and national LULC products, the architectures used to classify land cover, and the strategies for curating labeled datasets and designing feature vectors. It highlights how choices regarding map taxonomy, spatial resolution, transparency, labeling workflows, and feature engineering collectively shape the suitability of LULC products for India’s fragmented and dynamically cultivated landscapes.

2.1 Review of Existing LULC Products

Advances in remote sensing have produced several open-access global LULC products with differing resolutions, update frequencies, and thematic detail. Notable examples include Dynamic World (10 m, 2–5 day updates, 9 classes) [13], ESA WorldCover (10 m, annual since 2020, 11 classes) [117], Esri Land Cover (10 m, annual since 2017, 10 classes) [29, 49], GLC_FCS30D (30 m, annual 1985–2022, 35 classes) [123], GlobeLand30 (30 m, 10 classes) [22], NRSC’s Land Use/Land Cover Atlas of India (56–2.5 m, annual since 2005, 18–54 classes) [72], Copernicus Global Land Service (100 m, annual 2015–2019, ~12 classes) [14], and MODIS Land Cover (MCD12Q1, 500 m, since 2001, 17 classes) [35]. These products provide unprecedented global coverage but must be critically assessed

for India's heterogeneous and dynamic agricultural landscapes. The discussion below examines their thematic scope, spatial properties, and classification architectures to identify gaps that motivate a custom framework for India.

2.1.1 Thematic Scope: Annually Static vs. Intra-Annually Dynamic Classes. Most global products emphasize annually static classes like built-up areas, perennial water, barren land, or forests, because of their stability and relative ease of annotation. This taxonomy supports consistent global mapping but tends to obscure intra-annual variability that is crucial for agricultural and hydrological applications [110]. Even Dynamic World, despite its high temporal frequency, focuses on instantaneous generic states (e.g., “crops”, “water”) rather than explicitly modeling functional dynamics. As a result, key variables such as cropping intensity (single vs. double cropping) and water seasonality (intermittent vs. perennial inundation) remain underrepresented [45, 78]. Specialized products like the JRC Global Surface Water Explorer [78], GFSAD [102], and Global Cropping Intensity (GCI-30) [57] address parts of this gap but operate as thematic silos that provide standalone layers rather than integrated, multi-class LULC maps.

The NRSC LULC Atlas for India illustrates the value of integrating dynamic agricultural classes by explicitly mapping seasonal cropping patterns (Kharif, Rabi, Zaid¹) within a broader LULC framework at 56 m [72]. This integration underscores the need to move beyond purely annual, static taxonomies and to incorporate intra-annual classes, potentially through hierarchical or progressive classification schemes [26], when developing application-oriented LULC products for sectors such as water management and agriculture [87, 112].

2.1.2 Spatial Resolution and Methodological Transparency. Spatial resolution is critical in fragmented smallholder landscapes to prevent mixed pixels from sub-pixel features like fields, ponds, and narrow channels. Products such as MODIS Land Cover (500m) and Copernicus Global Land Service (100m) frequently miss or generalize these features, resulting in mixed pixels and loss of detail. Similarly, the NRSC Atlas, although thematically rich, is constrained by IRS-AWiFS imagery at 56m, limiting field-scale analyses. In contrast, recent global products leveraging Landsat and Sentinel achieve 30m (GLC_FCS30D, GlobeLand30) and 10m (Dynamic World, Esri Land Cover, ESA WorldCover) resolution, which now represent practical benchmarks for open LULC mapping. To improve these baselines, custom frameworks must utilize the highest open-source resolution while also expand the thematic detail.

Equally important is transparency in data, outputs, and methodology. Open access to map products, training data, and code enables rigorous validation, reproducibility, and adaptation for local applications. NRSC Atlas rasters require formal requests via the Bhuvan portal, constraining seamless integration at scale. By contrast, Dynamic World, ESA WorldCover, and Esri Land Cover are directly available on platforms like Google Earth Engine (GEE), but their generation depends on proprietary, largely opaque pipelines, limiting reproducibility and custom class extensions. Products such as GCI-30 and the JRC Global Surface Water Explorer demonstrate a more transparent model by publishing both methodology and code, allowing the community to scrutinize and build upon existing frameworks. For India-focused mapping, an ideal approach would combine high-resolution open imagery with fully open-source pipelines and freely accessible outputs [32].

2.1.3 Classification Architecture: Flat vs. Hierarchical Models. The architecture of LULC classifiers strongly influences both accuracy and extensibility. Most global products use “flat” multi-class classifiers in which all classes are predicted simultaneously. This design is computationally straightforward but presents three main limitations. First, flat supervised models demand a large, unified training dataset that becomes increasingly difficult to curate as class taxonomies grow. Second, they often struggle with spectrally similar classes (e.g.,

¹Kharif, Rabi, and Zaid refer to India's three main agricultural seasons: Kharif (June/July – September/October) aligns with the southwest monsoon; Rabi (November/December – February/March) follows the monsoon and utilizes residual moisture or irrigation; and Zaid (April–May) is the short summer season, primarily in irrigated areas.

cropland vs. tree cover), especially when using large, undifferentiated feature stacks that exacerbate the “curse of dimensionality” and reduce class separability [55, 67, 74, 77, 91, 109]. Third, flat architectures are poorly extensible: adding a new class usually requires full model retraining, with potential degradation in existing classes and high computational cost.

Hierarchical classification frameworks, such as the FAO Land Cover Classification System (LCCS) [26] and the modular approach of MapBiomas [34], offer a more flexible alternative. By organizing classes into a tree that moves from broad categories (e.g., vegetation vs. non-vegetation) to finer subtypes (cropland, trees, grass, etc.), each node can use its own optimized classifier and feature subset. This structure: (1) reduces dimensionality and improves accuracy by tailoring features to specific class groups—studies report 3–9% gains from such curated hierarchical setups over flat models [112]; (2) allows training on disjoint, class-specific datasets rather than a single monolithic set; and (3) enables incremental extension by appending new branches without disrupting the baseline map. Consequently, hierarchical architectures provide a practical route to thematically rich and extensible LULC products.

2.2 Methods of Curating Labeled Datasets

Labeled data underpin supervised LULC mapping and accuracy assessment [23]. Many global products depend on proprietary reference datasets such as ESA WorldCover’s PlanetScope polygons [117], Esri’s World Imagery labels [49], or Copernicus Land Cover’s licensed survey points [14]. For India, open annotated datasets like Dynamic World expert labels [13], LandCoverNet [1], and CGLS provide important resources but face three key limitations: (1) fragmented spatial coverage [68], (2) labels that are typically generalized or object-based are well suited for segmentation-style mapping but too coarse for precise per-pixel sampling, and (3) lack of classes that are difficult to annotate from imagery alone, particularly intra-annually dynamic variables such as water seasonality and cropping intensity [106].

Existing studies broadly distinguish two complementary data curation strategies—

- (1) **Desk-based annotation** uses high- to very-high-resolution imagery (e.g., Google Earth Pro, GEE, Collect Earth Online) and auxiliary sources such as OSM or government records to digitize polygons or points without field visits [8, 11, 91]. Active-learning strategies and crowdsourced platforms further reduce effort and expand coverage [2, 105]. However, polygon labels in heterogeneous landscapes often suffer from mixed-pixel impurity, and while pixel-wise labeling is most precise for pixel-level classifiers, it is difficult to scale to large areas [73, 116]. Delineating polygons over predominantly homogeneous regions offers a pragmatic compromise: it reduces mixed pixels compared with sampling from coarse segmentation tiles but still requires careful quality control and temporal alignment with satellite acquisitions. Additional challenges include annotator bias, variable skill levels, and the need for systematic cleaning.
- (2) **Field-based collection with ground verification**, typically using GPS-enabled devices, mobile apps (e.g., ODK Collect, QField), and standardized protocols, remains the most reliable way to capture complex or dynamic classes that are difficult to annotate from limited-date images [31, 41, 76, 85]. Yet field campaigns are labor- and cost-intensive, sensitive to positional errors and enumerator variability, and often require repeated visits to capture intra-annual dynamics. These constraints highlight the importance of well-designed field protocols, real-time or near-real-time validation, and rigorous post-processing to avoid costly re-surveys [33, 46, 124].

Combining desk-based and field-based approaches, using desk annotation for stable, visually distinct classes and targeted field campaigns for dynamic classes, has emerged as a practical strategy in recent LULC studies [69, 75].

2.3 Feature Selection for Classification

Feature design is central to LULC classification, and the literature emphasizes the complementary strengths of optical and SAR data. Optical reflectance in visible and infrared bands supports vegetation indices, seasonal trajectories, and material-specific signatures for crops, forests, water, bare soil, built-up areas, and wetlands, but its utility is reduced under persistent cloud cover. SAR, unaffected by clouds and solar illumination, is especially sensitive to structure and moisture, enhancing discrimination of tree cover, flooded and vegetated wetlands, soil moisture, and built-up areas. Effective LULC mapping therefore requires aligning feature choice with class characteristics and seasonal data availability.

Remote-sensing features are usually derived as static composites or temporal profiles. Seasonal or annual composites (e.g., median or maximum NDVI, SAR backscatter) suppress noise and highlight persistent spectral or structural traits, supporting mapping of relatively invariant classes such as built-up areas, bare soils, and evergreen forests. Time-series metrics from optical and SAR data capture intra-annual dynamics such as phenological transitions, cropping cycles, flooding, and wetland moisture, enabling separation of temporally dynamic classes like cropping intensity and seasonal water bodies. Multi-sensor fusion of optical and SAR, implemented as stacked composites or synchronized time series, has repeatedly been shown to improve class separability and boost accuracy for crop and land-cover mapping [104].

Beyond optical and SAR, studies incorporate topographic variables (elevation, slope), texture measures (e.g., GLCM metrics), thermal bands, and climate or hydrometeorological data to better characterize mountainous, heterogeneous, or urban environments and seasonal water dynamics [3, 66, 90]. However, large undifferentiated feature sets are prone to redundancy and overfitting, illustrating the “curse of dimensionality.” Feature-selection and dimensionality-reduction approaches, particularly when applied in a class-wise or hierarchical manner, have been shown to improve performance compared with using all available features, motivating shorter, class-specific feature subsets tuned to each LULC class’s spectral, structural, and temporal properties [67, 77, 99, 112].

2.4 LULC Classification Methodology

LULC classification methods in the literature mainly fall into pixel-based and segmentation-based (object-based) categories. Pixel-based approaches classify each pixel independently using models such as Random Forests, SVMs, and other machine-learning or deep-learning classifiers, and are effective at capturing fine-scale heterogeneity in spectrally mixed classes like cropland. Segmentation-based approaches first delineate homogeneous objects and then classify these segments, reducing salt-and-pepper noise and performing well for spatially coherent classes such as built-up areas, barren land, and scrub. Both strategies can achieve competitive accuracies, but segmentation may underrepresent small or rare classes that are merged into larger objects, whereas pixel-based methods preserve spatial detail and support post-classification refinement of minority classes [5, 10, 27, 28, 47, 70].

Across both paradigms, supervised, semi-supervised, and unsupervised techniques are used depending on label availability. Supervised classifiers generally deliver the highest accuracies where representative training data exist for all classes [10]. Semi-supervised approaches exploit abundant unlabeled data when some classes are poorly sampled [98, 100], and unsupervised clustering (e.g., k-means, hierarchical clustering) is employed when ground truth is scarce or reserved for validation [59, 111]. Hybrid workflows of combining pixel- and object-based methods or mixing supervised and unsupervised steps, are increasingly explored to balance accuracy, robustness, and computational efficiency [15, 119]. Classic machine-learning models remain widely adopted for their robustness and scalability, while deep-learning architectures (e.g., U-Net, SegFormer, transformer-based models) provide additional gains where high-quality training data and compute resources are available [88, 96, 107].

3 Methodology and Technical Validation

3.1 LULC taxonomy

Our taxonomy comprises 12 hierarchically structured LULC classes tailored for India’s agricultural landscapes, blending stable categories (built-up, barren land, trees, shrubs & scrubs) with dynamic ones (water seasonality, cropping intensity). See Figure 1 and Table A1 for details. This design, informed by NRSC [72], Dynamic World [13], FAO LCCS [26], ESRI GLC [49], and ESA WorldCover [117], enables flexible class aggregation/disaggregation for diverse applications [4].




Fig. (1) Land Use and Land Cover Taxonomy

3.2 Dataset Curation

3.2.1 Remote-sensing datasets. This study uses open-access optical and SAR products, accessed via GEE. Optical spectral information is drawn from Landsat-7 ETM+ (USGS Landsat 7 Collection 2 TOA reflectance; LANDSAT/LE07/C02/T1_TOA), Landsat-8 OLI (USGS Landsat 8 Collection 2 TOA reflectance; LANDSAT/LC08/C02/T1_TOA), and Sentinel-2 MSI (Level-1C Harmonized TOA ; (COPERNICUS/S2_HARMONIZED), and MODIS Terra MOD13Q1 NDVI product (MODIS/061/MOD13Q1) [16–19]. Synthetic Aperture Radar (SAR) data is obtained from Sentinel-1 (C-band Interferometric Wide swath & Ground Range Detected ; COPERNICUS/S1_GRD). [20]

Top of Atmosphere (TOA) vs Surface Reflectance (SR) collections: We opted for TOA over SR products because TOA collections typically retain a larger number of usable scenes, particularly in earlier years and in persistently cloudy regions such as monsoon-affected India, where SR retrievals can fail additional quality checks or lack complete coverage [93]. Moreover, TOA products also avoid mixing different SR processing versions across sensors and epochs, simplifying harmonization at the feature level [39, 43, 97].

Cloud Correction: For optical sensors (Landsat-7/8, Sentinel-2, and MODIS), clouds are handled by masking pixels via native QA bands: Sentinel-2 QA60 (masking clouds and cirrus using bits 10–11), Landsat QA_PIXEL (masking medium–high confidence clouds, cirrus, and cloud shadows), and MODIS Summary_QA to mask low quality cloud pixels. Per-pixel cloud masking maximizes temporal density while excluding contaminated observations. More aggressive strategies (e.g., CloudScore+, Fmask confidence thresholds) could reduce commission errors but cause substantial data loss, compromising detection of intra-annual dynamics such as cropping intensity and seasonal water presence [86, 125, 126].

3.2.2 Annotated Datasets Across Different LULC Classes. A prior research study released broad 4-class dataset covering greenery, water, built-up, and barrenland across India's 15 agro-climatic zones [7]. This dataset enables training and testing across these broad categories and provides a starting point for sampling polygons toward disaggregated subclasses. With this paper, we additionally curate and release *four new datasets* targeting fine-grained classes: tree vs. cropland, differently sized water bodies, seasonal vs. perennial water instances, and cropping intensity.

(1) Cropland & Tree Dataset

Starting from the greenery polygons in the broad 4-class dataset [7], we sampled cropland and tree/forest polygons, followed by filling coverage gaps across all 15 agro-climatic zones through targeted manual annotation using high-resolution Google Earth Pro imagery. Tree annotations were supplemented using the forest clearance dataset [62] as a spatial guide. Polygon sizes were maintained comparable to avoid oversampling bias, with a 10m inward buffer applied to minimize edge mixed pixels. All markings were cross-verified by two independent annotators for inter-coder reliability.




Fig. (2) Distribution of cropland and tree annotations over India's agro-climatic zones.

Class	#Polygons	#Pixels (10m resolution)
Cropland	132	4,00,196
Tree	174	4,11,471

Table (1) Cropland-Tree labeled dataset size



Fig. (3) Visual examples of cropland & tree annotations: Subfigures (a,b) illustrate polygons covering croplands, while (c,d) show tree-covered polygons.

This yielded 132 cropland polygons and 174 tree polygons, producing approximately 400k pixels per class at 10m resolution. Figure 2 illustrates annotation distribution across India, Figure 3 shows example polygons, and Table 1 covers the details of the dataset size, confirming comprehensive spatial coverage suitable for supervised classification of this thematic pair.

Note on limitations of polygon annotation- While careful delineation over homogeneous regions minimizes mixed pixels relative to segmentation-derived labels over broad image tiles, India's heterogeneous agricultural landscapes challenge complete pixel purity at scale. For example, some within-farm tree clusters,

tree plantation plots can be part of cropland annotations. Quantitative assessments should therefore be complemented by qualitative validation over broader regions for more nuanced understanding of model's performance.

(2) Differently Sized Water & Non-Water Dataset

Small freshwater bodies, such as farm ponds, are critical for agricultural sustainability but remain challenging to detect in medium-resolution imagery. To stress-test surface water classifiers and assess the minimum water-body size that can be reliably identified from medium-resolution spectral data, a water and non-water dataset of varying sizes was curated across different regions of India. From 12.7 lakh MGNREGA farm pond records [63, 64], two annotators manually verified and delineated 149 accurate water polygons across 110 districts, complemented by 139 similarly sized non-water controls (Fig. 4). Water polygons from the broad 4-class dataset were also incorporated, enriching the collection to 292 water polygons spanning multiple size classes (Fig. 5 and Table 2).




Fig. (4) Varying positional accuracy of MGNREGA farm pond locations used as annotation guide to mark differently sized water bodies.




Fig. (5) Size distribution of water bodies from MGNREGA farm pond annotations.

Size Range (sqm)	#Water Bodies
0–100	98
100–500	93
500–1000	51
1000–5000	32
>5000	18

Table (2) Size distribution of water bodies in test data.

(3) Seasonal & Perennial Water Body Dataset

We defined three water seasonality classes: (i) water present in Kharif, (ii) water present in Kharif and Rabi, and (iii) water persisting through Kharif, Rabi, and Zaid. Desk-based annotation of these classes is challenging, as it requires high-resolution imagery for each cropping season for which data is not

consistently available from Google Earth Pro. Furthermore, to the best of author's knowledge, no known auxiliary datasets currently support reliable delineation of season-specific surface water.

Field surveys offer an alternative approach, but would require revisiting the same sites across all three seasons with fine-scale delineation of dried-up areas, which is logistically demanding. Therefore, we focused on collecting a limited yet reliable set of examples for both seasonal and perennial water bodies. During 2022, field visits to 6 districts in central and southern India (Jamui, Koderma, Ranchi, Hazaribagh, Vadagar, and Giridih) supported this effort. Pre-monsoon surveys helped distinguish dried (seasonal) and water-filled (perennial) sites, and geo-tagged photos were used to delineate reference polygons on Google Earth Pro (Figure 6). In total, 16 seasonal and 13 perennial ponds were annotated which were insufficient for model training but valuable for evaluating classification accuracy.




Fig. (6) **Examples of seasonal and perennial water body annotations.** Field photos, satellite views, and polygon markings of perennial and seasonal water bodies from semi-arid regions of India.

(4) Cropping Intensity Dataset

Just like water seasonality, desk-based labeling of single Kharif, single non-Kharif, double, or TAP cropping would require season-specific, high-resolution imagery for each crop cycle which was largely unavailable on Google Earth Pro. To overcome this limitation, we conducted a field data collection campaign in the Raichur district of Karnataka, in collaboration with Gram Vaani, WELL Labs, and the grassroots organization Prarambha. A QField-based mobile application (Figure 7a) enabled enumerators to delineate field boundaries via GPS, record cropping information (like crop name, sowing and harvest periods, irrigation source, etc.) through structured survey forms, and synchronize data with QFieldCloud.

Ten trained enumerators (Figure 7b) surveyed diverse sites across Raichur during April–May 2023, collecting data for the 2021 and 2022 agricultural years. A total of 287 field records were initially submitted (Figure 7). These records underwent a three-step quality control (QC) process to ensure spatial, temporal, and spectral accuracy:

- **Spatial QC:** This initial filter focused on geometric validity. After removing duplicate, overlapping, and incorrectly delineated polygons, 152 unique cropland entries remained.




Fig. (7) On-field data collection of cropping intensity labels: Survey snapshot, training, and labeled data distribution.

- **Temporal QC:** This step checked consistency in the sowing and harvest intervals reported with the associated crop name to validate correct data entry. With this step, 63 field entries were eliminated.
- **Spectral QC:** The final validation step involved NDVI time-series validation at random points within each field to eliminate entries that lacked the expected number of vegetation peaks corresponding to the reported crop cycle (e.g., detecting two peaks in a field annotated as a single crop). Following this step, we were left with a final set of 88 confirmed field annotations- 29 representing single Kharif cropping and 59 representing double cropping.

The final clean dataset did not include any single non-Kharif or TAP-cropping instances, exhibiting class imbalance. This exercise highlighted the complexity of annotated data generation for dynamic agricultural classes. Beyond field mobilization, high-quality data collection depends on carefully designed survey questions, balanced spatial sampling, and real-time moderation to reduce annotation errors. Without these measures, extensive field efforts may yield only a small proportion of usable annotations. Accordingly, the final dataset was used exclusively for evaluating cropping-frequency classification performance rather than model training.

3.3 LULC Class Grouping and Task Decomposition

Constructing a unified, monolithic supervised classifier for the entire target taxonomy is infeasible, primarily due to the heterogeneity of label availability and the scarcity of annotations for dynamic classes. To address these real-world constraints, we adopted a hierarchical decomposition strategy governed by four key factors: (1) label sufficiency, (2) inherent class hierarchy, (3) coherence of feature representations, and (4) empirical performance validation.

The taxonomy partitions naturally based on data availability- Annually stable classes (tree, cropland, perennial water, built-up, barrenland) possess sufficient ground truth for supervised learning. In contrast, dynamic classes (water seasonality, cropping intensity) and spectrally ambiguous categories (shrubs/scrubs) lack adequate annotations, necessitating unsupervised or rule-based alternatives.

Furthermore, distinct feature requirements dictate specific groupings. Structurally sensitive SAR time-series can identify vegetation sub-types, such as distinguishing trees from cropland [9, 82, 89], while optical NDVI time-series are better suited for capturing phenological cycles to identify cropping intensity [58, 83]. Similarly, SAR backscatter provides the unique capability to detect water during the cloud-intensive Kharif season [12, 60, 103]. Finally, a broad mix of features can be tested to separate high-level categories—greenery, water, built-up, and

barrenland, potentially serving as a foundational masking layer that allows specialized downstream classifiers to operate on relevant subsets of the landscape [54, 84].

Performance assessments of a multi-class classifier targeting the primary group (greenery, water, built-up, and barrenland) revealed critical limitations (Figure A2 in Appendix B). Specifically, we observed significant spectral confusion between built-up areas and barren land [7], as well as the inability of a generic water classifier to capture dynamic seasonality.

These empirical findings led to a further decomposition into targeted, specialized modules to maximize separability:

- **Binary Classifiers:** Independent classifiers for *Built-up vs. Non-Built-up*, *Water vs. Non-Water* by identifying seasonal and perennial water using season-specific composites and SAR features, and *Barrenland vs. Greenery (vegetation)*.
- **Vegetation Sub-classification:** A dedicated *Tree vs. Cropland* discriminator operating within a vegetative mask.
- **Unsupervised Modules:** Clustering-based mapping for *cropping intensity* followed by identification of *Shrubs/Scrubs*.

Through this hierarchical decomposition, each classifier operates on a compatible feature space and an appropriately curated label set. Subsequent sections provide detailed descriptions and evaluations of each classifier, focusing particularly on Kharif water detection, tree vs. cropland discrimination, and cropping intensity classification- culminating in a robust pixel-level LULC mapping product.

3.4 Targeted Integration of Dynamic World in the Classification Framework

Dynamic World (DW) is used in the hierarchical pipeline to provide reliable information for relatively stable land-cover types and to support seasonal mapping of surface water. Its segmentation-based 10m prediction maps help delineate persistent features such as built-up areas, non-vegetative barren land, and shrub/scrub, while its dense time series enables aggregation of monthly water classifications into agriculturally relevant seasons. This section describes how DW is integrated for these roles, and how its outputs are refined through per-pixel corrections and SAR-based methods before feeding into the final LULC pipeline.

3.4.1 Built-up Classification. The built-up class represents an annually static category, consistently included in nearly all state-of-the-art global LULC products. Consequently, we first benchmark these global products against a locally trained baseline to determine the optimal approach for this classification.

To perform this comparative evaluation, we utilized the broad 4-class polygon dataset (greenery, water, built-up, and barrenland) established in previous work [7]. The dataset was split into an 80:20 ratio. The 80% training split was used to develop a custom supervised baseline: a Random Forest classifier (100 trees) trained on pixels sampled at 10m resolution, following the classification strategy using Sentinel-2 bands described in [7]. The remaining 20% constituted a held-out test set, comprising approximately 12,000 built-up and 370,000 non-built-up pixels at 10m resolution.

On this held-out test set, as shown in Table A3, the custom 4-class classifier achieves the highest recall but suffers from very low precision mirroring the performance of GLC_FCS30D. This is primarily due to barrenland being frequently misclassified as built-up owing to their spectral similarities that pose a fundamental challenge for pixel-based classification approaches. Among the other global products evaluated, ESA WorldCover demonstrates the best overall performance; however, its limited annual timestamps preclude continuous monitoring. Dynamic World provides the next-best balance of precision and recall, supported by near-real-time Sentinel-2 outputs since 2017 and the inherent advantages of segmentation-based models in delineating cohesive urban clusters [13, 109]. However, its outputs often over-generalize built-up predictions in peri-urban areas (misclassifying urban

greens and nearby croplands), and exhibit confusion with barren land near water bodies and rocky exposures [115].




Fig. (8) NDVI-based correction of Dynamic World built-up predictions in Adilabad. Threshold of 0.5 resolves false positives over urban greens and nearby croplands.

To mitigate these limitations, we apply per-pixel error-correction layers using NDVI (threshold 0.5) and NDWI (threshold 0.25), with thresholds optimized via ROC analysis (Figure A2). These corrections effectively remove vegetation- and water-adjacent false positives, as can be seen in Figure 8 and Figure A4-. The corrected Dynamic World outputs outperform all baselines on the held-out test set (Table A3), achieving superior built-up precision suitable for use as an upstream mask in the full LULC pipeline. It is noted that DW nonetheless struggles with fine-scale features such as small constructions and linear infrastructure like roads, indicating potential for future improvement. For a comprehensive description of the validation methodology, ROC curve analysis, and supplementary performance tables, refer to Appendix C.

3.4.2 Non-Vegetative Barrenland Classification. Like the built-up class, non-vegetative barrenland also represents another annually static category exhibiting minimal temporal variation in a year. So, on the lines of preceding subsection, we evaluated the barrenland vs. non-barrenland performance of the same supervised classifier against the "bare" class of Dynamic World, ESA WorldCover, Esri Land Cover, and GLC_FCS30D. The held-out test set comprised approximately 75,000 non-vegetative barrenland and 310,000 non-barrenland pixels sampled at 10m resolution.

As shown in Table A4, ESA WorldCover again demonstrates superior performance, but its limited temporal coverage renders it unsuitable for continuous monitoring. DW serves as the next-best alternative with its near-real-time outputs.

Unlike the built-up class, no direct corrections were applied to DW's bare predictions, as non-vegetative barrenlands can spectrally overlap with fallow croplands, a risk we avoided eliminating given the agriculture-specific focus of this work. Instead, we allow these predictions to propagate through the hierarchical LULC pipeline, where subsequent classifiers re-assign confident labels to the non-vegetative barrenland class, improving precision to 0.94 (methodology detailed in Section 3.5).

Notably, low recall across all products, including the pipeline-enhanced DW (F1 0.60), highlights the challenge of comprehensively delineating barren lands, representing a key area for future improvement. A detailed breakdown of barrenland classification performance, including class-wise metrics and examples of typical misclassifications, is provided in Appendix D.

3.4.3 Surface Water Body Classification. The water seasonality classifier detects water presence per pixel and attributes it to agricultural seasons, with emphasis on identifying Kharif-only water bodies often missed by optical datasets. Pixels are classified into three classes: (1) water in Kharif only, (2) water in Kharif and Rabi, and

(3) water in Kharif, Rabi, and Zaid. This design assumes that water detected in drier seasons is also present in wetter seasons, so isolated “Rabi-only”, “Zaid-only”, and “Rabi & Zaid-only” cases are reassigned to Kharif & Rabi or Kharif, Rabi & Zaid, respectively.

To select a suitable baseline for water detection, existing LULC and water products—Dynamic World (DW), ESA WorldCover, Esri Land Cover, GLCFC30D, and Global Surface Water (GSW) were first evaluated against water annotations from the broad 4-class dataset, split into roughly 180,000 water and 200,000 non-water pixels at 10m resolution. Figure A6 shows that all products achieve high (> 0.95) precision, recall, and F1-score for these large-scale labels, but a stratified analysis by water body size reveals important differences.




Fig. (9) Performance of existing LULC products on water classification across different water-body size categories, showing precision, recall, and F1-score for the water class.

As highlighted by Figure 9, across water bodies smaller than 5000 sqm, all products show high precision but very low recall, indicating that most small water bodies remain undetected even when correctly classified when found. DW is the only product that begins to capture a non-trivial fraction of 100–500 sqm water bodies and, although it achieves higher recall than other products in the 500–1,000 sqm range, its recall remains modest, reinforcing the need for improved methods for these agriculturally important features. Given this relative advantage and its near-real-time outputs, DW is better suited as the operational baseline for subsequent seasonal water analysis [13, 109].

SAR-based Kharif Water Detection. Both GSW and DW provide monthly water vs. non-water classifications that can be aggregated using a majority rule to generate seasonal water maps [52]. While these approaches produced reliable coverage during the Rabi and Zaid seasons, they failed during Kharif, as shown in Figure A7 due to persistent monsoon cloud cover, which created extensive data gaps and omission errors [56, 95].

To overcome this Kharif-specific limitation, the framework incorporated Sentinel-1 SAR, which provides weather independent observations suitable for monsoon water mapping [56, 79, 95]. Backscatter from VV and VH polarizations was analyzed using the 4-class ground truth revealing strong separability between water and non-water classes [44], with VV showing greater sensitivity to open water surfaces than VH (Figure A9) [48, 53, 81]. A VV threshold of -16dB was selected via ROC analysis (Figure A10) as a robust decision boundary, and monthly VV composites below this threshold were converted into binary Kharif water masks [37, 103]. To reduce false positives arising from transient soil moisture or short-lived inundation, a pixel was treated as Kharif water only if it was detected as water in more than two of the four Kharif months [121], and this SAR-based Kharif layer was combined with DW-derived Rabi and Zaid layers to form a unified seasonal water presence map.

We used DW instead of GSW for Rabi and Zaid because our analysis, supported by published evaluations [42, 78], showed that GSW systematically underestimated water extent and missed many small water bodies [42, 78].

Comparative Evaluation: Perennial water body polygons provided a pragmatic proxy for evaluating Kharif performance, since pixels within these polygons are expected to contain water throughout all four Kharif months. For each classifier, a monthly water-count (0–4) was computed per pixel, and performance was evaluated via Normalized Root Mean Squared Error (NRMSE) that compares observed month counts to the ideal value of four, thereby penalizing larger temporal omissions more strongly than minor gaps. On nearly 5000 perennial water pixels, the SAR-based Kharif classifier achieved a markedly lower NRMSE (0.33) than both DW (0.75) and GSW (0.53), indicating substantially more consistent detection of water during the monsoon. Figure A11 illustrates the distribution of perennial water pixels classified as water across Kharif months by different methods. Notably, Dynamic World identifies water in zero Kharif months for nearly 50% of these pixels, while GSW detects water in only 2 out of 4 Kharif months for the majority of these pixels, highlighting the severe data limitations imposed by monsoon cloud cover on both approaches.

Combined SAR-in-Kharif plus DW-in-Rabi/Zaid configuration not only maintains strong performance on large, perennial water bodies (Figure A6) but also substantially improves detection of smaller, often Kharif-only structures that are underrepresented in existing products (Figure 9).

NDWI-based post-classification error correction. To further reduce false positives, especially where non-water classes such as water-intensive croplands or bright built-up surfaces mimic water in optical or SAR signatures, an NDWI-based error-correction layer was applied [36]. Hard-negative mining over misclassified crop/vegetation and barrenland areas (29,050 pixels), combined with correctly labeled water regions from multiple districts (31,141 pixels), yielded balanced water/non-water pixel samples. ROC analysis on this combined dataset identifies an NDWI threshold of 0.15 as effective for distinguishing true water from spectrally similar non-water surfaces (Figures A12a, A12b). Pixels whose maximum NDWI over the year never exceeds 0.15 are reclassified as non-water, under the assumption that genuine surface water bodies exhibit at least one high-NDWI observation annually. This correction substantially increases precision from 0.91 to 0.99 (while maintaining high recall) on the hard-mined set, indicating a meaningful reduction in false positives for seasonal water mapping.

For a comprehensive breakdown of SAR backscatter distributions, ROC analyses for threshold selection, and stratified performance metrics across different water body sizes, refer to Appendix E.

3.5 Cropland vs Tree Classification

This section describes the development of a SAR-based tree–cropland classifier tailored to heterogeneous Indian landscapes, where trees occur as plantations, field-boundary rows, and small clusters embedded within croplands. Existing global LULC products provide cropland–tree labels but often overgeneralize in such settings, either due to segmentation-based smoothing or spectral confusion among vegetated subclasses in multi-class flat classifiers [120]. To obtain finer delineation of tree cover within agricultural mosaics, a custom classifier based on structure-sensitive Sentinel-1 SAR time series was developed.

Feature engineering. Time-series features can be highly effective for delineating tree cover from croplands because trees remain relatively static within a given year, whereas croplands exhibit pronounced temporal variability due to cropping cycles. To determine which sensor modality better captures these distinctions, we constructed synchronized 16-day time series of optical indices (NDVI, EVI) from Sentinel-2 and backscatter intensity (VV, VH) from Sentinel-1.

For the optical time series, we utilized cloud-masked Sentinel-2 imagery, aggregating pixels into 16-day median composites to minimize atmospheric artifacts. For the SAR time series, Sentinel-1 VV and VH bands from both ascending and descending orbits were combined to maximize temporal coverage. Notably, no speckle filtering

was applied to the SAR data to preserve the high-frequency spatial detail required for detecting individual trees and narrow tree lines [94]. Missing observations in both datasets were filled via linear interpolation to ensure continuity.

We first assessed the discriminative power of these features using a Random Forest classifier. Feature importance analysis, measured by Mean Decrease in Impurity (MDI), indicated that NDVI contributed most significantly to the model, reflecting its high sensitivity to vegetation phenology relative to SAR's structural backscatter.




Fig. (10) Inter-annual performance of SAR- and Optical-data based models for tree–cropland delineation, showing Macro F1 scores for years 2018, 2019, and 2020.

However, a multi-year robustness assessment revealed a critical limitation in the optical approach. When models trained on 2019 data were evaluated on held-out polygons from 2018, 2019, and 2020 (assuming stable land cover), NDVI-based predictions exhibited high inter-annual variance (as shown in Figure 10). This instability likely stems from variable clear-sky availability and shifts in phenological timing (e.g., peak greenness dependence on rainfall). In contrast, the SAR-based model demonstrated remarkable temporal consistency. As shown in Figure 11, the joint VV–VH trajectories effectively capture the structural stability of trees (high, consistent volume scattering) versus the lower, variable backscatter of croplands. Because these structural signals are cloud-independent and less sensitive to short-term climatic shifts, SAR backscatter was selected as the discriminative feature vector for this study.

Classification method. Tree–cropland polygons were split into an 80:20 ratio for training and testing. From the training polygons, 25,000 pixels per class were randomly sampled at 10m resolution. A Random Forest classifier with 100 trees, trained on the Sentinel-1 time-series features, was implemented on Google Earth Engine for efficient large-area deployment.

Post-classification error correction. Post-processing corrections were applied to reduce systematic false positives and enforce landscape constraints (Appendix F, Figure A14). First, a slope-based correction used SRTM-derived slope to reassign cropland predictions on slopes greater than 20° to tree cover, reflecting that Indian agriculture predominantly occurs on gentle slopes, with steep-slope cropping largely limited to rare terrace systems in the Himalayas [65, 71, 127]. Boxplots of slope values for labeled tree and cropland points (Appendix F, Figure A15) confirm this pattern, and example maps (Figure A16) illustrate cropland pixels on steep terrain correctly reclassified as tree.

A second stage applied NDVI-based corrections. Separate NDVI thresholds were derived for cropland and tree classes using a combination of NDVI distributions, ROC curve analysis, and spatial inspection of corrected regions.




Fig. (11) VH backscatter boxplots (16-day intervals, 2019 hydrological year): Trees exhibit higher, stable values; croplands show greater variability. (VV time series: Appendix F, Figure A13.)

For cropland, pixels whose maximum annual NDVI remained below 0.15, typically barren or highly degraded areas, were reassigned to the barren class. For tree cover, pixels with maximum NDVI below 0.3, characteristic of sparse vegetation or low, bushy cover, were also reclassified as barren. The supporting quantitative and visual analyses are presented as panels in Figure A17 (cropland) and Figure A19 (tree) in Appendix F, ensuring that the threshold choices balance statistical robustness with practical mapping accuracy.

Performance evaluation. The SAR-based classifier was evaluated against four global LULC products—Dynamic World, ESA WorldCover, Esri Land Cover, and GLC_FCS30D—on a held-out 10 m test set for 2019 (101,481 cropland and 56,962 tree pixels). Table 3 reports per-class precision, recall, and F1-score, along with macro averages to account for class imbalance. Quantitatively, ESA WorldCover and Esri Land Cover achieve the highest aggregate scores on this polygon-based test set, while the SAR-based model performs competitively, particularly for tree recall.

Method	Cropland (101,481)			Tree (56,962)			Macro Average
	Precision	Recall	F1-score	Precision	Recall	F1-score	
Dynamic World	1.00	0.88	0.94	0.83	1.00	0.91	0.92
ESA Worldcover	1.00	0.98	0.94	0.99	0.93	0.96	0.98
Esri Land Cover	1.00	0.99	1.00	1.00	1.94	0.97	0.98
GLC_FCS30D	0.95	0.91	0.93	0.89	0.86	0.87	0.90
SAR timeseries based	0.99	0.93	0.96	0.88	0.98	0.93	0.94

Table (3) **Performance Metrics for Cropland vs Tree Classification.** Comparative evaluation of existing LULC products against SAR-based classifier on held-out test dataset at 10m resolution.

Visual inspection of model predictions over the held-out test polygons (Figure 12) shows that the dominant misclassification type is “tree-in-cropland”: many pixels predicted as tree but labeled as cropland in the annotated dataset lie along field boundaries, within plantations, or in small clusters inside cropland polygons (as noted in Section 3.2.2). Although these cases reduce the apparent accuracy of Dynamic World and the SAR-based classifier




Fig. (12) Reference satellite image (2019) and corresponding cropland–tree predictions from Dynamic World, ESA WorldCover, Esri Land Cover, GLC_FCS30D, and the SAR-based classifier over a test polygon, illustrating fine-scale tree-in-cropland patterns (notably in plantation areas) captured by DW and the SAR model but largely missed by other products.

Legend: ■ polygon boundary ■ cropland ■ tree

in per-pixel metrics, they correspond to genuine tree cover embedded in agricultural fields and are therefore desirable from a mapping perspective.

To understand how this behavior manifests beyond individual polygons, predictions are examined over larger heterogeneous landscapes, such as Adilabad (Figure 13). Dynamic World and GLC_FCS30D tend to underpredict cropland in such fragmented agricultural mosaics and overgeneralize tree cover around existing woody patches and along forest margins. ESA WorldCover and Esri Land Cover, in contrast, frequently overgeneralize cropland, assigning large areas almost entirely to cropland with very limited tree representation. Taken together, these patterns indicate that, among the available products, the SAR-based method provides the most detailed and structurally consistent delineation of the cropland–tree mosaic.

3.6 Cropping Intensity Classification

Due to the scarcity of labeled annotations, we developed an unsupervised hierarchical K-means clustering method for cropping intensity classification. Applied to 30m NDVI time series from cropland pixels, this approach distinguishes four categories: single Kharif, single non-Kharif, double, and triple/annual/perennial (TAP) cropping. Methodological details follow, with extended procedures in Appendix G.

Data Preparation. To capture the full diversity of cropping patterns across India, we sampled pixels, at 30m resolution, from the annotated cropland polygons described in Section 3.2.2. This sampling strategy spanned all agro-climatic zones, ensuring our dataset represented the unique phenological signatures driven by India’s varied temperature, rainfall, and soil conditions.

For each sampled pixel, we constructed an NDVI time series to serve as the feature vector for classification. This vector captures the distinct phenological fluctuations associated with sowing, peak growth, senescence, and harvest. Adopting the workflow of Zhang et al. [122], we generated continuous 16-day composite time series over a hydrological year by harmonizing Landsat-7/8 TOA reflectance with Sentinel-2 data [21], filling in the gaps with MODIS and applying linear interpolation.




Fig. (13) Tree-cropland predictions over a heterogeneous region in Adilabad (reference image centered at [19.35327°N, 78.7826°E]), comparing different LULC products and highlighting the finer tree-cropland delineation achieved by the SAR-based classifier; non-tree-cropland pixels for each product are shown as masked.

Legend: ■ cropland ■ tree

We selected a 16-day interval as it offers an optimal trade-off: it provides sufficient temporal density to distinguish the Kharif, Rabi, and Zaid seasons while maintaining computational feasibility for large-scale analysis [80, 108]. Crucially, while our 16-day median compositing inherently minimizes high-frequency noise, we deviated from conventional approaches by omitting additional temporal smoothing algorithms such as the Whittaker filter. Although such filtering effectively reduces residual noise, it often attenuates the sharp, transient signal peaks characteristic of intensive multi-cropping systems (Figure 14).




Fig. (14) **Whittaker Smoothing of NDVI Time Series in Raichur, India.** Application of Whittaker smoothing to NDVI data could reduce peak intensities, which may obscure detection of actual cropping intensity pattern.

Hierarchical Time Series Clustering and Labeling. To classify cropping intensity, we employed hierarchical K-means clustering on the NDVI time series of cropland pixels, grouping similar phenological patterns for subsequent annotation. Guided by elbow and silhouette analyses (Figures A23-A24, Appendix G), we initialized

the model with $k=18$ clusters. To reduce intra-class variability and generate tighter clusters, we hierarchically refined these clusters by recursively splitting them based on a distortion threshold of 0.23. This threshold was determined to balance the minimization of intra-cluster variance with a practical limit on the number of clusters required for manual labeling. This process yielded 33 distinct clusters (Figure A25, Appendix G), which served as the basis for subsequent annotation and refinement.

Cluster Labeling. Prior to annotating the 33 derived clusters, we established a reference baseline by examining NDVI time series from known single, double, and TAP cropping locations across India (Figure A26). This qualitative validation allowed us to characterize the distinct phenological signatures, such as peak timing and separation, associated with each cropping intensity, serving as a guide for the subsequent cluster labeling process.

Cluster centroids (mean NDVI profiles of member pixels) and data spreads represent distinct classes (NDVI > 0.3) [92, 114] with roughly 40–50 day separations (5–6 time points) to distinguish cropping frequencies while avoiding intra-seasonal confusion. Standard seasonal windows were applied with flexibility for regional variations, such as early or delayed Kharif onset, offering advantages over rigid automated peak detection.

Ambiguous clusters lacking clear seasonal alignment were visually flagged and recursively subdivided ($k = 2$) to isolate consistent sub-patterns. Clusters remaining ambiguous after subdivision were discarded (9%), ensuring only robust phenological signatures contributed to the final classification (Figure 15). Detailed examples of this refinement process are provided in Appendix G (Figure A27).

Final Pan-India Clusters after re-clustering ambiguous clusters




Fig. (15) Final labeled clusters for cropping intensity classification. Centroids and data spreads represent distinct classes: single Kharif, single non-Kharif, double, and TAP cropping.

Pixel-Wise Mapping and Distance Metric Selection. Following the definition of cluster centroids, cropland pixels were assigned to the nearest labeled centroid using Euclidean distance on their NDVI time series, generating pan-India cropping intensity maps. Euclidean distance was selected over cosine similarity (which ignores

amplitude differences critical for cropping intensity) and dynamic time warping (infeasible due to memory limits in GEE). This approach balances shape and magnitude sensitivity for scalable pixel-wise classification.

While this unsupervised method effectively captures diverse cropping patterns without assuming fixed phenologies, manual cluster labeling introduces some subjectivity. Scalable alternatives for flexible distance metrics that handle temporal shifts and incorporation of auxiliary datasets like precipitation could further enhance classifier performance in future work.

Performance Evaluation. We compare the performance of our unsupervised classifier against GCI-30 as a baseline. For this evaluation, we use 88 final annotated field samples, described in Section 3.2.2, that pass all three Quality Control (QC) checks.

Method	Single Kharif (31,164)			Double (5027)			Macro Average
	Precision	Recall	F1-score	Precision	Recall	F1-score	
GCI-30	0.92	1.00	0.95	0.95	0.43	0.59	0.77
Unsupervised Classifier	0.97	0.97	0.97	0.81	0.79	0.80	0.88

Table (4) Pixel-level comparative evaluation of GCI-30 against our unsupervised classifier for cropping intensity classification on Raichur dataset.




Fig. (16) Pixel-level confusion matrices comparing GCI-30 against the proposed Unsupervised Classifier. The matrices reveal GCI-30's systematic underestimation of cropping intensity by mis-classifying Double cropping as Single.

We conducted evaluation at both **pixel** and **field** levels. From the final clean 88 field boundaries, we sampled 10m pixels, yielding 31,164 single Kharif cropping pixels and 5,027 double cropping pixels for model assessment. The results are summarized in Table 4. The pixel-level confusion matrices (Figure 16) reveal that GCI-30 exhibits limited performance, particularly for the Double class, achieving a macro average score of 0.77. Its tendency to underestimate cropping intensity, frequently misclassifying Double as Single Kharif, likely arises from smoothing artifacts or rigid thresholding that obscure multiple cropping cycles. In contrast, our unsupervised classifier, trained on pan-India cropland clusters, achieves a superior macro average of 0.88, demonstrating robust discrimination between Single Kharif and Double classes.




Fig. (17) Macro average scores for GCI-30 and Unsupervised Classifier across the 3-stage Quality Control pipeline. Upward trends show the impact of spatial, temporal, and spectral validation on test label refinement.

Performance metrics improved systematically with successive data quality control (QC) steps, as shown in Figure 17. This progressive enhancement across both methods confirms that our cleaning steps enhanced data reliability without introducing method-specific bias.

While pixel-level analysis is valuable, assigning a single cropping intensity label to every pixel within a farm boundary can be confounding due to inter-cropping and intra-field variability. We therefore conducted **field-level analysis** by aggregating per-pixel predictions via majority voting within each field boundary, enabling direct comparison with field-assigned labels. The field-level results, summarized in Table A5, Figure A28, and Figure A29, reinforce the trends observed at the pixel level but with greater polarity. The unsupervised classifier maintained robust performance with a macro average of 0.85. However, GCI-30 performance degraded significantly to 0.52 at the field level. This stark contrast highlights that while GCI-30 may capture some pixel-level heterogeneity, it fails to consistently characterize the dominant cropping pattern of smallholder fields, providing strong evidence that our unsupervised approach is better suited for agricultural monitoring in this landscape.

It should be noted that due to the bias in test data towards Single Kharif and Double cropping classes, a comprehensive evaluation of classifier performance remains lacking. This limitation underscores the need for future studies and data collection initiatives to assemble a more balanced and representative dataset covering all cropping intensity classes.

Selection of geographic scope for phenological clustering: To determine the optimal spatial extent for training the unsupervised classifier, we evaluated clustering performance at three geographic scales: district-level (Raichur), agro-ecological zone (AEZ-level), and Pan-India. As detailed in Table A6, the district-level approach failed to capture sufficient phenological diversity, resulting in high labeling ambiguity (50%). Conversely, while AEZ-specific clustering offered marginal accuracy gains (94% vs. 91% on Raichur dataset) by fitting local patterns more closely, it failed to capture the full spectrum of cropping intensity classes, specifically missing Single Non-Kharif (SNK) patterns, and posed significant scalability challenges regarding labeling consistency across 20 distinct AEZs. The Pan-India strategy ultimately proved to be the most robust; it achieved 90% unambiguous labeling, captured all four target intensity classes (Single Kharif, SNK, Double, and Triple), and maintained high class diversity (Shannon Entropy = 1.18). Although localized refinements can yield specific improvements, as demonstrated by a 7.8% accuracy gain in a separate case study on Anantapur (see Appendix G), the Pan-India

approach provides the necessary balance of broad class coverage and consistent performance required for scalable nationwide mapping.

3.7 Decision-Tree Based Sequencing of Classifiers for Hierarchical LULC Mapping

The final LULC classification pipeline was implemented through a hierarchical decision tree structure integrating multiple classifiers developed from distinct input datasets and temporal contexts. The hierarchy was designed to balance thematic consistency, classifier performance, and computational efficiency, as shown in Figure 18. Accordingly, three guiding principles shaped the decision order: (i) broad thematic separations precede finer subclass distinctions; (ii) high-precision and scalable classifiers are placed early to capture “high-confidence” regions; and (iii) computationally inexpensive classifiers are prioritized in the initial layers to facilitate national-scale mapping.




Fig. (18) LULC classification pipeline as a sequence of binary classifiers with pixel-level error correction layers, leveraging annual composites for static classes and time series for dynamic classes.

At the top of the hierarchy, the built-up versus non-built-up distinction was derived from the Dynamic World product. This decision was motivated by the relatively high accuracy and processing speed associated with the annual DW built-up layer. Among the DW-based classifiers, the built-up class exhibits both strong positional precision and straightforward implementation using a single annual composite, making it ideal for the first node. This early extraction efficiently masks anthropogenic surfaces, minimizing confusion in subsequent classifiers.

For the non-built-up areas, the next branch targets water seasonality, which leverages multi-temporal information to identify hydrologically dynamic regions. Water classes were sequentially characterized using a threshold-based SAR classification for Kharif season water bodies, supplemented by DW-derived water layers

for Rabi and Zaid seasons. Although DW-generated water masks are reliable, the necessity of generating multi-season composites renders them computationally more intensive than the built-up classification, justifying their placement below the first node. The use of SAR backscatter thresholds during kharif enhances sensitivity to monsoon dynamics and ensures accurate seasonal water delineation.

Pixels classified as non-built-up and non-water proceed to a barrenland versus greenery classifier derived from DW output. This step delineates vegetated and non-vegetated surfaces to establish a broad thematic separation. Although this classifier is computationally light due to its DW dependency, its relatively low recall for barren regions limits its use at higher hierarchy levels. The retained barrenland class captures stable non-vegetated areas, whereas the greenery outputs become the input mask for the next set of vegetation sub-classifiers.

Within the greenery class, further subclassification was undertaken through a SAR time-series-based tree versus cropland classifier. The inclusion of SAR data at this stage is justified by its ability to discriminate perennial woody vegetation from actively cultivated fields, especially under varying cloud conditions and phenological phases. The use of temporal backscatter metrics allows a robust separation consistent with thematic hierarchy and minimizes vegetation-type confusion.

The areas identified as cropland were further categorized through an NDVI-based cropping intensity classifier that distinguishes single kharif, single non-kharif, double, and triple/annual/perennial (TAP) cropping regimes. This step was placed deeper in the hierarchy as it represents a computationally intensive analysis requiring intra-annual time-series processing. The logical sequence of broad agricultural detection first, followed by cropping intensity mapping ensures scalability and minimizes redundant computation.

Finally, shrub/scrub regions were extracted from single kharif cropland classes using DW composites. This final module reassigns areas exhibiting transitional or sparse vegetation phenology typically misclassified as cropland into a more appropriate class, thereby enhancing thematic correctness in semi-arid environments.

This sequential structure ensures that the classification proceeds from rapid, high-confidence separations to more specialized and data-intensive analyses to produce annual LULC maps. Each classifier operates only on a restricted spatial subset defined by preceding masks, optimizing computational load and enabling controlled error propagation across thematic levels. Class-specific accuracy assessments corresponding to these classifier stages are presented separately. The complete pipeline could not be evaluated against existing LULC products because no available product directly matches the thematic detail and class definitions adopted in this study. Furthermore, labeled reference datasets are not uniformly available across all final classes, particularly for themes such as water seasonality and cropping intensity, making a single comprehensive validation infeasible.

3.8 Temporal Correction

Temporal inconsistencies are a common challenge in yearly LULC predictions, often stemming from missing data, salt-and-pepper noise [51], or limited generalizability of model across years [8]. For example, in Raichur district, our LULC predictions from 2017–2020 (Figures 19a–19d) show noticeable inconsistencies, particularly a pronounced over-prediction of shrubs and barren land in 2018. Such artifacts hinder reliable temporal change detection and necessitate a robust temporal correction framework [113].

3.8.1 Inconsistent Pattern Identification and Rule Derivation. To detect temporal inconsistencies, we analyzed pixel-wise land cover transitions within a three-year window which is the minimum span needed to capture temporal patterns. With 11 LULC classes, possible transitions are numerous and often rare in real data. We thus analyzed representative subsets from regions with diverse land dynamics and strong field knowledge: Raichur, Adilabad, and upper/lower Tungabhadra basins. LULC maps were generated for these regions over six hydrological years (2017–2022), and approximately 400,000 pixels were randomly sampled per class. We analyzed transition patterns in two 3-year windows: 2017–2019 and 2020–2022. Frequency distribution of such sequences in Adilabad are shown in Figure A31–A32.




Fig. (19) Snapshots of Temporal LULC Maps of Raichur District Without Temporal Correction. (a)–(d) show district LULC maps for HY 2017–2020, where year 2018 shows reveal sudden conversion of tree cover into barren area and (e) shows the LULC legend.

Each high-frequency inconsistent pattern was identified (Refer to Figure A33 for examples) and verified using Google Earth Pro historical imagery. These observations informed targeted correction rules to improve consistency within each three-year sequence. The complete rule tables are provided in Appendix H. We divided these correction rules into following categories-

- **Missing Data Interpolation:** Gaps from cloud cover or sensor errors are filled using adjacent-year predictions (Table A7). For a missing observation at year t , if classifications at $(t - 1)$ and $(t + 1)$ are identical, the class is filled via temporal persistence. In case of different adjacent classes, the prediction of year closer to the 2019 training baseline is used to preserve model integrity.
- **Intermediate-Year Inconsistency:** These address abrupt reversals where a middle year differs from its neighbors. While many are simple "flips" (e.g., *Tree* $\rightarrow$ *Shrub* $\rightarrow$ *Tree* $\Rightarrow$ *Tree*), some require domain-specific logic. For instance, *Water* $\rightarrow$ *Cropland* $\rightarrow$ *Water* transitions often indicate sedimented zones and are corrected to Barren, while *Builtup* $\rightarrow$ *Tree* $\rightarrow$ *Builtup* patterns are resolved using an extended five-year window to account for persistent urban structures. Refer Table A8 for the complete set of these handcrafted rules.
- **Edge-Year Inconsistency:** We also identified edge-year inconsistencies (Table A9), particularly *builtup* $\rightarrow$ *cropland* $\rightarrow$ *cropland* and *builtup* $\rightarrow$ *tree* $\rightarrow$ *tree* patterns, which are unlikely as built-up areas rarely revert to vegetative cover in such short window of years. These sequences were corrected by updating the first year to match the prediction of subsequent years.

3.8.2 Temporal Correction Algorithm. A multi-step algorithm (Pseudocode 1, Appendix H) ensures temporal consistency in multi-year LULC predictions through iterative refinement. It begins with a single-pass interpolation of missing classifications using adjacent-year predictions. A three-year moving window then identifies candidate pixels by counting intermediate-year inconsistencies per pixel. District-wide analysis (Figure A34, Appendix H) showed that most correctable pixels exhibit low transition frequencies, indicating isolated artifacts rather than systemic classification failures. Pixels with inconsistencies in more than 50% of years are excluded due to high temporal instability. Stable pixels receive a second pass of correction rules for intermediate-year consistency. Training-year predictions remain unmodified to preserve data integrity. Finally, specialized rules are applied to address inconsistencies in the edge years of the time series.

The proposed algorithm systematically improves the temporal coherence of LULC maps by filling missing values, enforcing logical land cover transitions, and minimizing classification errors. These improvements are shown in the corrected maps for the Raichur district (HY 2017–2020), where the previous over-prediction of




Fig. (20) Corrected temporal LULC maps for Raichur district (2017–2020HY). (a–d) annual LULC maps; (e) legend.

barren land in 2018 is significantly reduced, resulting in a more stable temporal sequence (Figure 20). Refer to Figure A35 for further localized examples in the cropland regions of Adilabad that undergo temporal correction. Ultimately, this framework ensures that the resulting multi-year LULC products are sufficiently consistent to support robust, high-confidence temporal change detection and long-term land-use trend analysis.

For comprehensive details regarding the specific correction rules, the multi-pass algorithm, and additional validation examples, refer to Appendix H.

4 Discussion and Conclusions

This research establishes a practical, hierarchical LULC classification framework tailored to the heterogeneity of India's landscapes, demonstrating that a class-specific approach significantly outperforms traditional flat multi-class models. Since a collective comparison against global products was precluded by thematic inconsistencies and a lack of coherent datasets, we validated our results through rigorous class-wise evaluations against appropriate baselines. Our SAR-based classification for cloud-obscured Kharif water bodies achieved superior consistency with a markedly lower NRMSE (0.33) compared to Dynamic World (0.75) and Global Surface Water (0.53), while our SAR time-series metrics distinguished trees with a peak recall of 0.98 and a cropland F1-score of 0.96. Similarly, our unsupervised pipeline for cropping intensity outperformed the GCI-30 baseline, securing a pixel-level macro average of 0.88 (vs. 0.70) with a substantial gain at the field level (0.85 vs. 0.52). To support these methods and the broader research community, we curated and released four high-quality annotated datasets validated through rigorous desk-based and field data collection. Furthermore, the framework's implementation on GEE ensures high computational efficiency suitable for national-scale deployment; for instance, generating a 10m resolution map for Agro-Ecological Zone-11 (138,311 sq.km) requires approximately 7 hours and yields a lightweight asset (385 MB), confirming that this open-source, cloud-native approach is both technically feasible and accessible for widespread adoption.

4.1 Generalizable Framework for Hierarchical LULC Mapping

Developing regional LULC products with arbitrary thematic detail can follow a modular, five-stage hierarchical design framework (Figure 21):

- (1) **Class Grouping and Task Decomposition.** Partition the target taxonomy into thematically coherent groups separable by shared feature sets (e.g., NDVI phenology for cropping intensity subclasses, or SAR backscatter for tree–crop structure). Decompose complex multi-class problems into simpler binary or few-class tasks where joint separability is challenging.




Fig. (21) Hierarchical classification framework for regional LULC mapping.

- (2) **Task-Specific Classifier Development.** For each decomposed task, apply standard ML workflows: curate labels (field, desk-based, or product-derived); engineer class-relevant features (composites, time series from optical/SAR/ancillary data); train models with regularization and tuning; validate on held-out data for error patterns; apply post-processing corrections; and evaluate against baselines. The framework accommodates any validated algorithm.
- (3) **Hierarchical Decision-Tree Assembly.** Sequence classifiers respecting: (i) thematic hierarchy (broad classes before subclasses); (ii) performance (high-precision modules early to mask "easy" regions); and (iii) scalability (cheap, broad classifiers first to constrain compute-intensive steps). This yields a pixel-wise decision cascade superior to flat multi-class models.
- (4) **Annual Mapping.** Execute the pipeline year-by-year, aligning temporal windows to regional calendars (agricultural/hydrological seasons) at target resolution.
- (5) **Multi-Year Temporal Smoothing.** Post-process annual outputs with temporal filters, transition rules, or probabilistic smoothing to enforce plausibility (e.g., stable built-up extents, gradual cropland shifts), enhancing longitudinal coherence.

This paper instantiated the framework for India's 12-class taxonomy, blending global products with custom classifiers.

4.2 Operational Deployment and Real-World Applications

The framework's outputs, 10m pan-India LULC maps, have been deployed via **CoRE Stack's GEE application**, an open-source digital public infrastructure promoting bottom-up landscape planning for climate resilience and sustainable NRM [24]. These maps, particularly the intra-annual dynamic classes (water seasonality, cropping intensity), power several decision-support tools within CoRE Stack.

Key applications include: (i) **Commons Connect**, which currently supports more than 15 civil society organizations in over 500 villages to demand water assets under MGNREGA; our LULC layers help model crop water demand, drought sensitivity, and agricultural intensification risks, powering hydrological reports and visualizations for equitable resource allocation [61]; (ii) **Know Your Landscape (KYL) dashboard**, delivering micro-watershed and village-level insights on agricultural intensity and seasonal water availability for policy-makers and researchers [25]; and (iii) **pan-India NRM impact assessments**, as demonstrated by Kaur et al., where water and cropping intensity layers identify farm ponds/check dams and quantify their effects on local agricultural productivity [50].

This deployment bridges research and practice, enabling community-led NRM planning across hundreds of Indian villages while demonstrating the framework's scalability and actionability.

4.2.1 Field Validation and Community Impact: Insights from Murgathal village in Hathiyapathar Gram Panchayat (Jharkhand).

Murgathal is a village in the Masalia Block of Dumka district. With field insights

from the deployment of Commons Connect app in this village, we demonstrate how the water seasonality and cropping intensity layers produced by our LULC framework aids in micro-watershed development in this village.

Implementation Framework. The planning process, led by facilitator Rupdhan Hansda (SUPPORT CSO [101]), utilized a hybrid methodology combining traditional Participatory Rural Appraisal (PRA) with our satellite-derived LULC maps. The workflow involved four distinct stages: (i) community consultations (Figure 22a) and Gram Sabhas, (ii) digital vulnerability mapping, (iii) ground-truthing via transect walks (Figure 22b- 22c), and (iv) the finalization of a Detailed Project Report (DPR).



Fig. (22) Field visit snapshots from Murgathal village, Masalia block, Jharkhand, conducted by Rupdhan Hansda (SUPPORT).

Precision Planning via LULC Analytics. The integration of temporal LULC data provided three critical advantages over traditional manual planning methods-

- **Vulnerability Assessment:** By utilizing the cropping intensity layers, the team performed a longitudinal analysis to identify "vulnerability hotspots"—regions exhibiting a steady decline in agricultural productivity over a five-year horizon. This allowed for the creation of a Social-Economic-Ecological Profile that prioritized land parcels belonging to the most vulnerable household beneficiaries.
- **Micro-Watershed Characterization:** The LULC maps facilitated a granular audit of all micro-watersheds intersecting the settlement. By overlaying temporal land-use changes with existing Natural Resource Management (NRM) assets, the team identified specific spatial gaps where existing water structures (wells and ponds) were no longer sufficient to support current cropping demands.
- **Hydrological Profiling:** The water seasonality layers were used to plot the inter-annual variability of surface water within the watershed. This informed the community about shifting water availability patterns, which had previously been noted only through anecdotal observations.

Outcomes and Policy Implications. The data-driven approach culminated in the Plan Murgathal DPR, presented to the Gram Sabha on March 26, 2025. The precision afforded by the LULC layers led to the proposal of targeted NRM interventions (Figure 23)-

- (1) **Recharge Optimization:** The identification of declining cropping intensity led to a community-backed demand for **three new percolation tanks** located at strategic recharge points identified by the slope and LULC intersection.
- (2) **Sustainability Targets:** The mapping of seasonal water shifts allowed for the design of **irrigation tasks** tailored to stabilize the non-monsoon crop, directly addressing the identified "low-intensity" zones.




Fig. (23) Spatial distribution of NRM works proposed by the community in Murgathal village via Commons Connect.

This case study demonstrates that high-resolution LULC maps do not merely serve as classification outputs; they act as a "spatial bridge" between satellite-based observations and community-led climate resilience strategies. This bottom up demand generation becomes all the more significant in areas like Jharkhand, which is predominantly inhabited by a marginalized caste community of Scheduled Tribes. As of October 2025, SUPPORT organization [101] has submitted 30 DPRs powered by the Commons Connect App.

4.3 Limitations And Future Directions

While the hierarchical framework demonstrates strong performance for agriculturally relevant classes, several limitations highlight opportunities for refinement-

Thematic and Classification Gaps. Barrenland and shrub/scrub modules were deprioritized due to the study's agricultural focus, relying on Dynamic World proxies that exhibit known confusions with croplands, grasslands, and trees. Future work can prioritize these classes through dedicated SAR/texture-based classifiers or ancillary datasets (e.g., soil maps, elevation) to enhance overall thematic completeness.

Temporal Windows and Regional Variability. Fixed Kharif, Rabi, and Zaid definitions that are anchored to southwest monsoon patterns, do not fully capture regional onset shifts. cropping intensity clustering proved robust due to pan-India sampling and flexible labeling, but water seasonality remains sensitive to these variations. Dynamic seasonal boundaries derived from monsoon onset data (e.g., IMD rainfall records) or precipitation time series could yield more adaptive outputs.

Temporal Coverage Constraints. Pipeline dependency on Sentinel-1 SAR (post-2017) and Dynamic World (post-2015) limits maps to 2017–2024. Backcasting SAR/DW-reliant classes requires auxiliary classifiers using harmonized historical products (Landsat 5/7/8, IRS LISS-III). Forward extension via Landsat-9 and emerging sensors (e.g., PRSS-1) will enable continuous monitoring.

Validation Data Limitations. cropping intensity assessment suffered from skewed, limited test samples (primarily Kharif/double cropping), while water seasonality classes lacked precise reference datasets. Comprehensive evaluation requires broad-scale, carefully designed field collection programs incorporating real-time

or near-real-time validation, with precise class delineation through repetitive seasonal visits combined with desk-based annotations for available timestamps.

The framework's modularity facilitates these enhancements, allowing integration of new classifiers, temporal extensions, and validation datasets while preserving scalability, transferability, and open accessibility. These directions will strengthen validation, expand thematic resolution, and deepen temporal coverage for robust, long-term LULC analysis.

Acknowledgments

We gratefully acknowledge funding and on-ground data collection support from WELL LABS, Prarambh, NYAS, and Gram Vaani. Special thanks to Gram Vaani for their pivotal role in deploying the LULC framework at pan-India scale, ensuring accessibility for end users. We also thank GIZ and Tower Research for their engagement, which enhanced the real-world impact of our work. Additionally, we appreciate the Information and Communication Technologies for Development (ICTD) lab at IIT-Delhi for developing the QField App, which streamlined field data collection, and the research groups for integrating our LULC framework into their ongoing projects and providing valuable feedback that informed key methodological improvements.

References

- [1] Hamed Alemohammad and Kevin Booth. 2020. LandCoverNet: A global benchmark land cover classification training dataset. *arXiv preprint arXiv:2012.03111* (2020).
- [2] Usman Ali, Travis J Esau, Aitazaz A Farooque, Qamar U Zaman, Farhat Abbas, and Mathieu F Bilodeau. 2022. Limiting the collection of ground truth data for land use and land cover maps with machine learning algorithms. *ISPRS international journal of geo-information* 11, 6 (2022), 333.
- [3] Jesús A Anaya, René R Colditz, and Germán M Valencia. 2015. Land cover mapping of a tropical region by integrating multi-year data into an annual time series. *Remote Sensing* 7, 12 (2015), 16274–16292.
- [4] James Richard Anderson. 1976. *A land use and land cover classification system for use with remote sensor data*. Vol. 964. US Government Printing Office.
- [5] Akanksha Balha, Javed Mallick, Suneel Pandey, Sandeep Gupta, and Chander Kumar Singh. 2021. A comparative analysis of different pixel and object-based classification algorithms using multi-source high spatial resolution satellite data for LULC mapping. *Earth Science Informatics* 14, 4 (2021), 2231–2247.
- [6] Subhajit Bandopadhyay, Sourav Dey, Latika Grover, Subhasis Ghosh, and Barnali Das. 2024. Sensing the dynamic nature of fragmented croplands in India through earth observation: a comprehensive review. (2024).
- [7] Chahat Bansal, Hari Om Ahlawat, Mayank Jain, Om Prakash, Shivani A Mehta, Deepanshu Singh, Harshavardhansushil Baheti, Suyash Singh, and Aaditeshwar Seth. 2021. IndiaSat: A Pixel-Level Dataset for Land-Cover Classification on Three Satellite Systems-Landsat-7, Landsat-8, and Sentinel-2. In *Proceedings of the 4th ACM SIGCAS Conference on Computing and Sustainable Societies*. 147–155.
- [8] Chahat Bansal, Arpit Jain, Phaneesh Barwaria, Anuj Choudhary, Anupam Singh, Ayush Gupta, and Aaditeshwar Seth. 2020. Temporal prediction of socio-economic indicators using satellite imagery. In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD*. 73–81.
- [9] Xin Bao, Rui Zhang, Jichao Lv, Renzhe Wu, Hongsheng Zhang, Jie Chen, Bo Zhang, Xiaoying Ouyang, and Guoxiang Liu. 2023. Vegetation descriptors from Sentinel-1 SAR data for crop growth monitoring. *ISPRS Journal of Photogrammetry and Remote Sensing* 203 (2023), 86–114.
- [10] Mariana Belgiu and Lucian Drăguț. 2016. Random forest in remote sensing: A review of applications and future directions. *ISPRS journal of photogrammetry and remote sensing* 114 (2016), 24–31.
- [11] Adia Bey, Alfonso Sánchez-Paus Díaz, Danae Maniatis, Giulio Marchi, Danilo Mollicone, Stefano Ricci, Jean-François Bastin, Rebecca Moore, Sandro Federici, Marcelo Rezende, et al. 2016. Collect earth: Land use and land cover assessment through augmented visual interpretation. *Remote Sensing* 8, 10 (2016), 807.
- [12] Filsa Bioresita, Anne Puissant, André Stumpf, and Jean-Philippe Malet. 2018. A method for automatic and rapid mapping of water surfaces from Sentinel-1 imagery. *Remote Sensing* 10, 2 (2018), 217.
- [13] Christopher F Brown, Steven P Brumby, Brookie Guzder-Williams, Tanya Birch, Samantha Brooks Hyde, Joseph Mazzariello, Wanda Czerwinski, Valerie J Pasquarella, Robert Haertel, Simon Ilyushchenko, et al. 2022. Dynamic World, Near real-time global 10 m land use land cover mapping. *Scientific Data* 9, 1 (2022), 251.
- [14] Marcel Buchhorn, Myroslava Lesiv, Nandin-Erdene Tsendbazar, Martin Herold, Luc Bertels, and Bruno Smets. 2020. Copernicus global land cover layers—collection 2. *Remote Sensing* 12, 6 (2020), 1044.

- [15] Dang Hung Bui and László Mucsi. 2021. From land cover map to land use map: A combined pixel-based and object-based approach using multi-temporal landsat data, a random forest classifier, and decision rules. *Remote Sensing* 13, 9 (2021), 1700.
- [16] Earth Engine Data Catalog. 2024. Harmonized Sentinel-2 MSI: MultiSpectral Instrument, Level-2A. https://developers.google.com/earth-engine/datasets/catalog/COPERNICUS_S2_SR_HARMONIZED
- [17] Earth Engine Data Catalog. 2024. Landsat 7 Collection 1 Tier 1 8-Day TOA Reflectance Composite. https://developers.google.com/earth-engine/datasets/catalog/LANDSAT_LE07_C01_T1_8DAY_TOA
- [18] Earth Engine Data Catalog. 2024. Landsat 8 Collection 1 Tier 1 8-Day TOA Reflectance Composite. https://developers.google.com/earth-engine/datasets/catalog/LANDSAT_LC08_C01_T1_8DAY_TOA
- [19] Earth Engine Data Catalog. 2024. MOD13Q1.061 Terra Vegetation Indices 16-Day Global 250m. https://developers.google.com/earth-engine/datasets/catalog/MODIS_061_MOD13Q1#bands
- [20] Earth Engine Data Catalog. 2024. Sentinel-1 SAR GRD: C-band Synthetic Aperture Radar Ground Range Detected, log scaling. https://developers.google.com/earth-engine/datasets/catalog/COPERNICUS_S1_GRD
- [21] Robert Chastain, Ian Housman, Joshua Goldstein, Mark Finco, and Karis Tenneson. 2019. Empirical cross sensor comparison of Sentinel-2A and 2B MSI, Landsat-8 OLI, and Landsat-7 ETM+ top of atmosphere spectral characteristics over the conterminous United States. *Remote sensing of environment* 221 (2019), 274–285.
- [22] Jun Chen, Jin Chen, Anping Liao, Xin Cao, Lijun Chen, Xuehong Chen, Chaoying He, Gang Han, Shu Peng, Miao Lu, et al. 2015. Global land cover mapping at 30 m resolution: A POK-based operational approach. *ISPRS Journal of Photogrammetry and Remote Sensing* 103 (2015), 7–27.
- [23] Russell G Congalton and Kass Green. 2019. *Assessing the accuracy of remotely sensed data: principles and practices*. CRC press.
- [24] CoRE Stack. 2024. CoRE Stack GEE Application. <https://ee-corestackdev.projects.earthengine.app/view/core-stack-gee-app> Accessed: 2025-10-25.
- [25] CoRE Stack. 2025. CoRE Stack: Commoning for Resilience and Equality. <https://core-stack.org/> Accessed: 2025-05-03.
- [26] Antonio Di Gregorio and Louisa JM Jansen. 1998. Land Cover Classification System (LCCS): classification concepts and user manual. *FAO, Rome* (1998).
- [27] Laura Dingle Robertson and Douglas J King. 2011. Comparison of pixel- and object-based classification in land cover change mapping. *International journal of remote sensing* 32, 6 (2011), 1505–1529.
- [28] Dennis C Duro, Steven E Franklin, and Monique G Dubé. 2012. A comparison of pixel-based and object-based image analysis with selected machine learning algorithms for the classification of agricultural landscapes using SPOT-5 HRG imagery. *Remote sensing of environment* 118 (2012), 259–272.
- [29] esri. 2020. Sentinel-2 10-Meter Land Use/Land Cover. <https://livingatlas.arcgis.com/landcover/>
- [30] Fabian Ewald Fassnacht, Joanne C White, Michael A Wulder, and Erik Næsset. 2024. Remote sensing in forestry: current challenges, considerations and directions. *Forestry: An International Journal of Forest Research* 97, 1 (2024), 11–37.
- [31] Yelena Finegold and Antonia Ortmann. 2016. Map accuracy assessment and area estimation: a practical guide. (2016).
- [32] Food and Agriculture Organization of the United Nations. 2020. Land Cover Classification System (LCCS). <https://www.fao.org/land-water/land/land-governance/land-resources-planning-toolbox/category/details/en/c/1036361/>. Accessed: 5 December 2025.
- [33] Giles M Foody, Mahesh Pal, Duccio Rocchini, Carol X Garzon-Lopez, and Lucy Bastin. 2016. The sensitivity of mapping methods to reference data quality: Training supervised image classifications with imperfect reference data. *ISPRS International Journal of Geo-Information* 5, 11 (2016), 199.
- [34] Washington JS Franca Rocha, Rodrigo N Vasconcelos, Diego P Costa, Soltan Galano Duverger, Jocimara SB Lobão, Deorgia TM Souza, Stefanie M Herrmann, Nerivaldo A Santos, Rafael O Franca Rocha, Jefferson Ferreira-Ferreira, et al. 2024. Towards Uncovering Three Decades of LULC in the Brazilian Drylands: Caatinga Biome Dynamics (1985–2019). *Land* 13, 8 (2024), 1250.
- [35] M. Friedl and D. Sulla-Menashe. 2022. *MODIS/Terra+Aqua Land Cover Type Yearly L3 Global 500m SIN Grid V061*. doi:10.5067/MODIS/MCD12Q1.061
- [36] Bo-Cai Gao. 1996. NDWI—A normalized difference water index for remote sensing of vegetation liquid water from space. *Remote sensing of environment* 58, 3 (1996), 257–266.
- [37] Geohackweek. 2016. SAR Water Body Classification. <https://mbonnema.github.io/GoogleEarthEngine/07-SAR-Water-Classification/>
- [38] Frane Gilić, Mateo Gašparović, and Martina Baučić. 2023. Current state and challenges in producing large-scale land cover maps: review based on recent land cover products. *Geocarto International* 38, 1 (2023), 2242693.
- [39] Google Earth Engine Team. 2025. Landsat Algorithms - Earth Engine. <https://developers.google.com/earth-engine/guides/landsat>. Accessed: 2025-12-09.
- [40] Adam S Green, Aftab Alam, Shruti Bhogal, Sandeep Dixit, Kamal Vatta, and Cameron A Petrie. 2025. Ponds and Their Potential for Agricultural Sustainability in Punjab: Insights from a Century of Surface Water Change in the Granary of India. *The Journal of Development Studies* (2025), 1–21.
- [41] Carl Hartung, Adam Lerer, Yaw Anokwa, Clint Tseng, Waylon Brunette, and Gaetano Borriello. 2010. Open data kit: tools to build information services for developing regions. In *Proceedings of the 4th ACM/IEEE international conference on information and*

- communication technologies and development*. 1–12.
- [42] Kelsey Herndon, Rebekke Muench, Emil Cherrington, and Robert Griffin. 2020. An assessment of surface water detection methods for water resource management in the Nigerien Sahel. *Sensors* 20, 2 (2020), 431.
 - [43] Yunfeng Hu, Yu Dong, et al. 2018. An automatic approach for land-change detection and land updates based on integrated NDVI timing analysis and the CVAPS method with GEE support. *ISPRS journal of photogrammetry and remote sensing* 146 (2018), 347–359.
 - [44] Wenli Huang, Ben DeVries, Chengquan Huang, Megan W Lang, John W Jones, Irena F Creed, and Mark L Carroll. 2018. Automated extraction of surface water extent from Sentinel-1 data. *Remote Sensing* 10, 5 (2018), 797.
 - [45] Meha Jain, Pinki Mondal, Ruth S DeFries, Christopher Small, and Gillian L Galford. 2013. Mapping cropping intensity of smallholder farms: A comparison of methods using multiple sensors. *Remote Sensing of Environment* 134 (2013), 210–223.
 - [46] Meha Jain, Amit K Srivastava, Rajiv K Joon, Andrew McDonald, Keitasha Royal, Madeline C Lisaius, and David B Lobell. 2016. Mapping smallholder wheat yields and sowing dates using micro-satellite data. *Remote sensing* 8, 10 (2016), 860.
 - [47] Amita Jangid, Pankaj Dadheech, and Mukesh Kumar Gupta. 2025. Enhanced Land Cover Classification by Integrating Multispectral Satellite Imagery and Machine Learning Techniques. *Journal of Information Systems Engineering and Management* 10, 1 (2025), 1–15. doi:10.55267/iadt.2587
 - [48] Chaowei Jiang, Hong Zhang, Chao Wang, Ji Ge, and Fan Wu. 2022. Water Surface Mapping from Sentinel-1 Imagery Based on Attention-UNet3+: A Case Study of Poyang Lake Region. *Remote Sensing* 14, 19 (2022), 4708.
 - [49] Krishna Karra, Caitlin Kontgis, Zoe Statman-Weil, Joseph C Mazzariello, Mark Mathis, and Steven P Brumby. 2021. Global land use/land cover with Sentinel 2 and deep learning. In *2021 IEEE international geoscience and remote sensing symposium IGARSS*. IEEE, 4704–4707.
 - [50] Ramneek Kaur, Kshitiz Bansal, Devang Garg, Ramita Sardana, Saketh Vishnubhatla, Sanjali Agrawal, Shruti Kumari, Parag Singla, and Aaditeshwar Seth. 2024. Assessing the impact of farm ponds on agricultural productivity in Northern India. In *Proceedings of the 7th ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies*. 281–293.
 - [51] Hamidreza Keshkar, Winfried Voigt, and Esmail Alizadeh. 2017. Land-cover classification and analysis of change using machine-learning classifiers and multi-temporal remote sensing imagery. *Arabian Journal of Geosciences* 10 (2017), 1–15.
 - [52] Ankush Khandelwal, Anuj Karpatne, Praveen Ravirathinam, Rahul Ghosh, Zhihao Wei, Hilary A Dugan, Paul C Hanson, and Vipin Kumar. 2022. ReaLSAT, a global dataset of reservoir and lake surface area variations. *Scientific data* 9, 1 (2022), 356.
 - [53] NE Kolarik, A Roopsind, A Pickens, and JS Brandt. 2023. A satellite-based monitoring system for quantifying surface water and mesic vegetation dynamics in a semi-arid region. *Ecological Indicators* 147 (2023), 109965.
 - [54] Jiayi Li, Ben Zhang, and Xin Huang. 2022. A hierarchical category structure based convolutional recurrent neural network (HCS-ConvRNN) for Land-Cover classification using dense MODIS Time-Series data. *International Journal of Applied Earth Observation and Geoinformation* 108 (2022), 102744.
 - [55] Peilin Li, Yan Wang, Chisheng Wang, Lin Tian, Meijiao Lin, Siyao Xu, and Chuanhua Zhu. 2025. A Comparison of Recent Global Time-Series Land Cover Products. *Remote Sensing* 17, 8 (2025), 1417.
 - [56] Yang Li and Zhenguo Niu. 2022. Systematic method for mapping fine-resolution water cover types in China based on time series Sentinel-1 and 2 images. *International Journal of Applied Earth Observation and Geoinformation* 106 (2022), 102656.
 - [57] Chang Liu, Qiang Zhang, Yu Tao, Jianguo Qi, Yongqian Ding, and Ngai Weng Chan. 2020. A new framework to map fine resolution cropping intensity across the globe: Algorithm, validation, and implication. *Remote Sensing of Environment* 251 (2020), 112095. doi:10.1016/j.rse.2020.112095
 - [58] Luo Liu, Xiangming Xiao, Yuanwei Qin, Jie Wang, Xinliang Xu, Yueming Hu, and Zhi Qiao. 2020. Mapping cropping intensity in China using time series Landsat and Sentinel-2 images and Google Earth Engine. *Remote Sensing of Environment* 239 (2020), 111624.
 - [59] Lucrécio Silvestre Macarrigue, Édson Luis Bolfe, and Paulo Roberto Mendes Pereira. 2022. Developments in land use and land cover classification techniques in remote sensing: A review. *Journal of Geographic Information System* 14, 1 (2022), 1–28.
 - [60] Timothy Mayer, Ate Poortinga, Biplov Bhandari, Andrea P Nicolau, Kel Markert, Nyein Soe Thwal, Amanda Markert, Arjen Haag, John Kilbride, Farrukh Chishtie, et al. 2021. Deep learning approach for Sentinel-1 surface water mapping leveraging Google Earth Engine. *ISPRS Open Journal of Photogrammetry and Remote Sensing* 2 (2021), 100005.
 - [61] Shivani A Mehta, Aila Dutt, Ajay Tannirkulam, Akshay Pratap Singh, Aman Verma, Anamitra Singha, Ananda Sreenidhi, Ananjan Nandi, Ankit, Ankit Kumar, et al. 2024. Initial Observations from Field Testing of a Digital Participatory Tool to Improve Water Security in Rural India. In *Proceedings of the 13th International Conference on Information & Communication Technologies and Development*. 337–364.
 - [62] Forest Ministry of Environment and Government of India Climate Change. 2024. Forest Clearance Dataset, India. https://forestclearance.nic.in/Online_Status.aspx
 - [63] Government of India Ministry of Rural Development. 2024. Bhuvan MGNREGA. https://bhuvan-app2.nrsc.gov.in/mgnrega/mgnrega_phase2.php
 - [64] Government of India Ministry of Rural Development. 2024. MIS report NREGA. <https://nreganarep.nic.in/netnrega/MISreport4.aspx>
 - [65] Prabuddh Kumar Mishra, Aman Rai, and Suresh Chand Rai. 2020. Indigenous knowledge of terrace management for soil and water conservation in the Sikkim Himalaya, India. (2020).

- [66] Pegah Mohammadpour, Domingos Xavier Viegas, and Carlos Viegas. 2022. Vegetation mapping with random forest using sentinel 2 and GLCM texture feature—A case study for Lousã region, Portugal. *Remote Sensing* 14, 18 (2022), 4585.
- [67] Mohammed Abdulmajeed Moharram and Divya Meena Sundaram. 2023. Dimensionality reduction strategies for land use land cover classification based on airborne hyperspectral imagery: a survey. *Environmental Science and Pollution Research* 30, 3 (2023), 5580–5602.
- [68] Mohammed Abdulmajeed Moharram and Divya Meena Sundaram. 2023. Land use and land cover classification with hyperspectral data: A comprehensive review of methods, challenges and future directions. *Neurocomputing* 536 (2023), 90–113.
- [69] Daniel Moraes, Manuel L Campagnolo, and Mário Caetano. 2024. Training data in satellite image classification for land cover mapping: a review. *European Journal of Remote Sensing* 57, 1 (2024), 2341414.
- [70] Soe W Myint, Patricia Gober, Anthony Brazel, Susanne Grossman-Clarke, and Qihao Weng. 2011. Per-pixel vs. object-based classification of urban land cover extraction using high spatial resolution imagery. *Remote sensing of environment* 115, 5 (2011), 1145–1161.
- [71] LGK Naidu, V Ramamurthy, O Challa, Rajendra Hegde, and P Krishnan. 2006. Manual soil-site suitability criteria for major crops. *NBSS publ* 129 (2006).
- [72] National Remote Sensing Centre. 2020. Land Use / Land Cover Atlas of India. https://www.nrsc.gov.in/sites/default/files/pdf/Announcements/LULC_Atlas_NRSC.pdf Accessed: 2025-04-29.
- [73] CHARLES T NICHOLS. 2012. Land use/land cover classification: Methods to overcome pixel confusion and the effects of tree shadows in very high resolution multispectral imagery. *Maryville: Northwest Missouri State University* (2012).
- [74] Andrea Puzzi Nicolau, Africa Flores-Anderson, Robert Griffin, Kelsey Herndon, and Franz J Meyer. 2021. Assessing SAR C-band data to effectively distinguish modified land uses in a heavily disturbed Amazon forest. *International Journal of Applied Earth Observation and Geoinformation* 94 (2021), 102214.
- [75] H NRSC. 2012. National land use and land cover mapping using multitemporal AWiFS data. *National Remote Sensing Center (NRSC) and North Eastern Space Applications Centre (NESAC), Hyderabad* (2012).
- [76] OPENGIS.ch. 2025. QField - Efficient Field Work Built for QGIS. <https://qfield.org/> Accessed: 2025-05-02.
- [77] Mahesh Pal and Paul M Mather. 2005. Support vector machines for classification in remote sensing. *International journal of remote sensing* 26, 5 (2005), 1007–1011.
- [78] Jean-François Pekel, Andrew Cottam, Noel Gorelick, and Alan S Belward. 2016. High-resolution mapping of global surface water and its long-term changes. *Nature* 540, 7633 (2016), 418–422.
- [79] Binh Pham-Duc, Catherine Prigent, and Filipe Aires. 2017. Surface water monitoring within Cambodia and the Vietnamese Mekong Delta over a year, with Sentinel-1 SAR observations. *Water* 9, 6 (2017), 366.
- [80] Amy Hudson Pickens, Matthew Hansen, Zhen Song, Andrew J Poulson, Andre Lima, Antoine Baggett, Aleksandra Mikus, and Theo Kerr. 2023. DIST-ALERT: Global vegetation monitoring in near real-time. In *AGU Fall Meeting Abstracts*, Vol. 2023. B23A–05.
- [81] Evelyn Marcia Possa and Philippe Maillard. 2018. Precise delineation of small water bodies from Sentinel-1 data using support vector machine classification. *Canadian Journal of Remote Sensing* 44, 3 (2018), 179–190.
- [82] Abdul Qadir and Pinki Mondal. 2020. Synergistic use of radar and optical satellite data for improved monsoon cropland mapping in India. *Remote Sensing* 12, 3 (2020), 522.
- [83] Laura Recuero, Lilian Maila, Víctor Cicuéndez, César Sáenz, Javier Litago, Lucía Tornos, Silvia Merino-de Miguel, and Alicia Palacios-Orueta. 2023. Mapping cropland intensification in Ecuador through spectral analysis of MODIS NDVI time series. *Agronomy* 13, 9 (2023), 2329.
- [84] Mariane S Reis, Maria Isabel S Escada, Luciano V Dutra, Sidnei JS Sant’Anna, and Nathan D Vogt. 2018. Towards a reproducible LULC hierarchical class legend for use in the Southwest of Pará State, Brazil: a comparison with remote sensing data-driven hierarchies. *Land* 7, 2 (2018), 65.
- [85] Andrew D Richardson, Koen Hufkens, Tom Milliman, Donald M Aubrecht, Min Chen, Josh M Gray, Miriam R Johnston, Trevor F Keenan, Stephen T Klosterman, Margaret Kosmala, et al. 2018. Tracking vegetation phenology across diverse North American biomes using PhenoCam imagery. *Scientific data* 5, 1 (2018), 1–24.
- [86] Mohamad Rifai et al. 2024. Integration of Cloud Score+ with Sentinel-2 Harmonized for land use and land cover classification using machine learning algorithms. In *IOP Conference Series: Earth and Environmental Science*, Vol. 1418. IOP Publishing, 012039.
- [87] Chirasmayee Savitha and Reshma Talari. 2023. Mapping cropland extent using sentinel-2 datasets and machine learning algorithms for an agriculture watershed. *Smart Agricultural Technology* 4 (2023), 100193.
- [88] Linus Scheibenreif, Joëlle Hanna, Michael Mommert, and Damian Borth. 2022. Self-supervised vision transformers for land-cover segmentation and classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1422–1431.
- [89] Christian Schulz, Michael Förster, Stenka Valentinova Vulova, Alby Duarte Rocha, and Birgit Kleinschmit. 2024. Spectral-temporal traits in Sentinel-1 C-band SAR and Sentinel-2 multispectral remote sensing time series for 61 tree species in Central Europe. *Remote Sensing of Environment* 307 (2024), 114162.
- [90] MA Shaban and O Dikshit. 2001. Improvement of classification in urban areas by the use of textural features: the case study of Lucknow city, Uttar Pradesh. *International Journal of remote sensing* 22, 4 (2001), 565–593.

- [91] Hossein Shafizadeh-Moghadam, Morteza Khazaei, Seyed Kazem Alavipanah, and Qihao Weng. 2021. Google Earth Engine for large-scale land use and land cover mapping: An object-based classification approach using spectral, textural and topographical factors. *GIScience & Remote Sensing* 58, 6 (2021), 914–928.
- [92] Sadia Alam Shammie and Qingmin Meng. 2021. Use time series NDVI and EVI to develop dynamic crop growth metrics for yield modeling. *Ecological Indicators* 121 (2021), 107124.
- [93] Andrii Shelestov, Mykola Lavreniuk, Nataliia Kussul, Alexei Novikov, and Sergii Skakun. 2017. Exploring Google Earth Engine platform for big data processing: Classification of multi-temporal satellite imagery for crop mapping. *frontiers in Earth Science* 5 (2017), 232994.
- [94] Jayant Singhal, Tanumi Kumar, Rakesh Fararoda, Prabir Kumar Das, Rakesh Paliwal, Sudhakar Reddy Chintala, and Gopalakrishnan Rajashekar. 2024. Forest characterization using C-band SAR data—Initial results of EOS-04 data. *Journal of the Indian Society of Remote Sensing* 52, 4 (2024), 787–800.
- [95] Patrick Sogno, Igor Klein, and Claudia Kuenzer. 2022. Remote sensing of surface water dynamics in the context of global change—a review. *Remote Sensing* 14, 10 (2022), 2475.
- [96] Jonathan V Solórzano, Jean Francois Mas, Yan Gao, and José Alberto Gallardo-Cruz. 2021. Land use land cover classification with U-net: Advantages of combining sentinel-1 and sentinel-2 imagery. *Remote Sensing* 13, 18 (2021), 3600.
- [97] Conghe Song, Curtis E Woodcock, Karen C Seto, Mary Pax Lenney, and Scott A Macomber. 2001. Classification and change detection using Landsat TM data: when and how to correct atmospheric effects? *Remote sensing of Environment* 75, 2 (2001), 230–244.
- [98] Marjan Stoimchev, Jurica Levatić, Dragi Kocev, and Sašo Džeroski. 2024. Semi-supervised multi-label classification of land use/land cover in remote sensing images with predictive clustering trees and ensembles. *IEEE Transactions on Geoscience and Remote Sensing* (2024).
- [99] Oliver Stromann, Andrea Nascetti, Osama Yousif, and Yifang Ban. 2019. Dimensionality reduction and feature selection for object-based land cover classification based on Sentinel-1 and Sentinel-2 time series using Google Earth Engine. *Remote Sensing* 12, 1 (2019), 76.
- [100] Bo Sun, Yang Zhang, Qiming Zhou, and Xinchang Zhang. 2022. Effectiveness of semi-supervised learning and multi-source data in detailed urban landuse mapping with a few labeled samples. *Remote Sensing* 14, 3 (2022), 648.
- [101] SUPPORT (Society For Upliftment Of People With People's Organization and Rural Technology). 2025. SUPPORT Jharkhand - Pioneering NGO in Jharkhand. <https://www.supportjharkhand.org/>. Accessed: 2025-12-22.
- [102] P. S. Thenkabail, P. Teluguntla, J. Xiong, A. Oliphant, R. Congalton, M. Ozdogan, et al. 2017. *Global Food Security-support Analysis Data (GFSAD) Cropland Extent 2015 Global 30 m V001*. doi:10.5067/MEaSUREs/GFSAD/GFSAD30.001
- [103] Khuong H Tran, Massimo Menenti, and Li Jia. 2022. Surface water mapping and flood monitoring in the Mekong Delta using sentinel-1 SAR time series and Otsu threshold. *Remote Sensing* 14, 22 (2022), 5721.
- [104] Rahat Tufail, Patrizia Tassinari, and Daniele Torreggiani. 2025. Assessing feature extraction, selection, and classification combinations for crop mapping using Sentinel-2 time series: A case study in northern Italy. *Remote Sensing Applications: Society and Environment* 38 (2025), 101525.
- [105] Devis Tuia, Frédéric Ratle, Fabio Pacifici, Mikhail F Kanevski, and William J Emery. 2009. Active learning methods for remote sensing image classification. *IEEE Transactions on Geoscience and Remote Sensing* 47, 7 (2009), 2218–2232.
- [106] Alexandra Tyukavina, SV Stehman, GM Foody, S Bontemps, Linda See, P Olofsson, NE Tsendbazar, J Radoux, Anna Komarova, Bryant M Serre, et al. 2025. Land Cover and Change Map Accuracy Assessment and Area Estimation Good Practices Protocol. (2025).
- [107] Anastasios Tzepkenlis, Konstantinos Marthoglou, and Nikos Grammalidis. 2023. Efficient deep semantic segmentation for land cover classification using sentinel imagery. *Remote Sensing* 15, 8 (2023), 2027.
- [108] Kristof Van Tricht, Jeroen Degerickx, Sven Gilliams, Daniele Zanaga, Marjorie Battude, Alex Grosu, Joost Brombacher, Myroslava Lesiv, Juan Carlos Laso Bayas, Santosh Karanam, et al. 2023. WorldCereal: a dynamic open-source system for global-scale, seasonal, and reproducible crop and irrigation mapping. *Earth System Science Data* 15, 12 (2023), 5491–5515.
- [109] Zander S Venter, David N Barton, Tirthankar Chakraborty, Trond Simensen, and Geethen Singh. 2022. Global 10 m land use land cover datasets: A comparison of dynamic world, world cover and esri land cover. *Remote Sensing* 14, 16 (2022), 4101.
- [110] Gladys Maria Villegas Rugel, Daniel Ochoa, Jose Miguel Menendez, and Frieke Van Coillie. 2023. Evaluating the applicability of global LULC products and an author-generated phenology-based map for regional analysis: a case study in Ecuador's ecoregions. *Land* 12, 5 (2023), 1112.
- [111] Junye Wang, Michael Bretz, M Ali Akber Dewan, and Mojtaba Aghajani Delavar. 2022. Machine learning in modelling land-use and land cover-change (LULCC): Current status, challenges and prospects. *Science of the Total Environment* 822 (2022), 153559.
- [112] Adam Waśniewski, Agata Hościło, and Milena Chmielewska. 2022. Can a hierarchical classification of Sentinel-2 data improve land cover mapping? *Remote Sensing* 14, 4 (2022), 989.
- [113] Barbara Wiatkowska, Janusz Ślodziński, and Aleksandra Stokowska. 2021. Spatial-temporal land use and land cover changes in urban areas using remote sensing images and GIS analysis: The case study of Opole, Poland. *Geosciences* 11, 8 (2021), 312.
- [114] Deji Wuyun, Mengqi Duan, Liang Sun, Luís Guilherme Teixeira Crusiol, Nitu Wu, and Zhongxin Chen. 2024. Pixel-wise parameter assignment in LandTrendr algorithm: Enhancing cropland abandonment monitoring using satellite-based NDVI time-series. *Computers and Electronics in Agriculture* 227 (2024), 109541.

- [115] Panpan Xu, Nandin-Erdene Tsendbazar, Martin Herold, Sytze de Bruin, Myke Koopmans, Tanya Birch, Sarah Carter, Steffen Fritz, Myroslava Lesiv, Elise Mazur, et al. 2024. Comparative validation of recent 10 m-resolution global land cover maps. *Remote Sensing of Environment* 311 (2024), 114316.
- [116] Su Ye, Robert Gilmore Pontius Jr, and Rahul Rakshit. 2018. A review of accuracy assessment for object-based image analysis: From per-pixel to per-polygon approaches. *ISPRS Journal of Photogrammetry and Remote Sensing* 141 (2018), 137–147.
- [117] Daniele Zanaga, Ruben Van De Kerchove, Dirk Daems, Wanda De Keersmaecker, Carsten Brockmann, Grit Kirches, Jan Wevers, Oliver Cartus, Maurizio Santoro, Steffen Fritz, et al. 2022. ESA WorldCover 10 m 2021 v200. (2022).
- [118] Chuanrong Zhang and Xinba Li. 2022. Land use and land cover mapping in the era of big data. *Land* 11, 10 (2022), 1692.
- [119] Ce Zhang, Xin Pan, Huapeng Li, Andy Gardiner, Isabel Sargent, Jonathon Hare, and Peter M Atkinson. 2018. A hybrid MLP-CNN classifier for very fine resolution remotely sensed image classification. *ISPRS Journal of Photogrammetry and Remote Sensing* 140 (2018), 133–144.
- [120] Guanjin Zhang, Siti Nur Aliaa binti Roslan, Ci Wang, and Ling Quan. 2023. Research on land cover classification of multi-source remote sensing data based on improved U-net network. *Scientific Reports* 13, 1 (2023), 16275.
- [121] Meimei Zhang, Fang Chen, Dong Liang, Bangsen Tian, and Aqiang Yang. 2020. Use of Sentinel-1 GRD SAR images to delineate flood extent in Pakistan. *Sustainability* 12, 14 (2020), 5784.
- [122] Miao Zhang, Bingfang Wu, Hongwei Zeng, Guojin He, Chong Liu, Shiqi Tao, Qi Zhang, Mohsen Nabil, Fuyou Tian, José Bofana, et al. 2021. GCI30: a global dataset of 30-m cropping intensity using multisource remote sensing imagery. *Earth System Science Data Discussions* 2021 (2021), 1–22.
- [123] Xiao Zhang, Tingting Zhao, Hong Xu, Wendi Liu, Jinqing Wang, Xidong Chen, and Liangyun Liu. 2023. GLC_FCS30D: The first global 30-m land-cover dynamic monitoring product with a fine classification system from 1985 to 2022 using dense time-series Landsat imagery and continuous change-detection method. *Earth System Science Data Discussions* 2023 (2023), 1–32.
- [124] Shengyu Zhao, Kaiwen Tu, Shutong Ye, Hao Tang, Yaocong Hu, and Chao Xie. 2023. Land use and land cover classification meets deep learning: A review. *Sensors* 23, 21 (2023), 8966.
- [125] Zhe Zhu, Shixiong Wang, and Curtis E Woodcock. 2015. Improvement and expansion of the Fmask algorithm: Cloud, cloud shadow, and snow detection for Landsats 4–7, 8, and Sentinel 2 images. *Remote sensing of Environment* 159 (2015), 269–277.
- [126] Zhe Zhu and Curtis E Woodcock. 2012. Object-based cloud and cloud shadow detection in Landsat imagery. *Remote sensing of environment* 118 (2012), 83–94.
- [127] Rajendra Bhausaheb Zolekar and Vijay Shivaji Bhagat. 2015. Multi-criteria land suitability analysis for agriculture in hilly zone: Remote sensing and GIS approach. *Computers and Electronics in Agriculture* 118 (2015), 300–321.

Received 8 July 2025; revised 22 December 2025; accepted 10 March 2026