# Multi-Temporal Land Use Land Cover Classification and Urban Change Detection for Chennai Metropolitan Region Using a Hybrid Machine Learning and Deep Learning Framework

Nikil Srivatsa
*Dept. of School of Computing*
*SRM Institute of Science and Technology*
Kattankulathur, Tamil Nadu, India
nk6602@srmist.edu.in

Rahil Kotwal
*Dept. of School of Computing*
*SRM Institute of Science and Technology*
Kattankulathur, Tamil Nadu, India
rk0835@srmist.edu.in

Dr. Deeban Chakravarthy V
*Associate Professor*
*Dept. of School of Computing*
*SRM Institute of Science and Technology*
Kattankulathur, Tamil Nadu, India
deepanv@srmist.edu.in

***Abstract*—Indian megacities have been losing farmland, wetlands, and green spaces at a pace that traditional survey methods cannot keep up with. This paper describes a practical satellite-based monitoring framework built to measure exactly how fast such changes are happening in Chennai, one of the most rapidly growing coastal cities in South Asia. Working entirely with freely available data from Google Earth Engine (GEE), we assembled Sentinel-2 imagery spanning four temporal epochs (2018, 2020, 2022, and 2024) and applied a combination of Random Forest (RF), Support Vector Machine (SVM), and EfficientNet-B3 deep learning to classify six land cover types. Our temporal cross-validation protocol revealed that the RF baseline achieves 81.5% overall accuracy on same-year data (T0), dropping to 62.5% when tested two years forward (T1). The EfficientNet-B3 model trained on the EuroSAT benchmark achieved 97.83% accuracy (Kappa = 0.976), confirming the strength of the deep learning pre-training approach. Change detection analysis on real GEE Sentinel-2 data revealed that Chennai built-up area expanded from 63.7 km$^2$ to 212.4 km$^2$ between 2018 and 2024 (a 233% increase), while the Urban Fringe class grew 644% reflecting explosive peri-urban expansion. All results are derived from real Sentinel-2 pixels sampled directly from Chennai via GEE.**

***Index Terms*—LULC Classification, Sentinel-2, Google Earth Engine, Random Forest, EfficientNet-B3, Change Detection, Chennai, Temporal Cross-Validation, Urban Monitoring, Remote Sensing, EuroSAT**

## I. Introduction

Walk through any peripheral neighbourhood of Chennai today and you will notice the same pattern repeating: paddy fields replaced by apartment blocks, lake beds converted to bus depots, wetlands buried under flyovers. The city has grown from 176 km$^2$ in 1971 to 426 km$^2$ within its corporation limits, and the wider metropolitan area was formally expanded from 1,189 km$^2$ to 5,904 km$^2$ as recently as October 2022 to account for urbanisation that had already outgrown earlier boundaries [18]. Mapping this transformation by sending surveyors into the field is expensive and slow. Satellite imagery changes that—but extracting useful information from it automatically, reliably, and across multiple years is still a research challenge that this project directly addresses.

We chose to work with Sentinel-2 data accessed through Google Earth Engine because the workflow is reproducible by anyone, costs nothing, and the 10-metre resolution imagery captures features as small as narrow roads or thin water channels. The core technical problem we set out to solve was not simply "how accurately can we label pixels"—it was "does the model we trained in 2018 still work in 2022 without retraining?" That question matters enormously for a monitoring system that is supposed to stay useful over years of deployment, but it rarely appears in published LULC papers, which typically report accuracy on randomly sampled pixels from a single image date. We call this evaluation approach *temporal cross-validation*, and it is the right way to assess an urban monitoring system.

The contributions of this paper are four in number. First, we provide a complete, reproducible pipeline from GEE data export to final change maps, built on open-source Python libraries anyone can install for free. Second, we compare three model families—RF, SVM, and EfficientNet-B3—head to head under identical temporal validation conditions using real Sentinel-2 pixels extracted directly from Chennai via GEE. Third, we publish area-change statistics for six land cover classes across Chennai from 2018 to 2024, grounded in satellite-derived classification maps. Fourth, we demonstrate that EfficientNet-B3 pre-trained on EuroSAT achieves 97.83% accuracy, establishing a strong deep learning benchmark for Sentinel-2 LULC classification.

Section II summarises relevant prior work. Section III describes our study area and data. Section IV explains the methodology. Section V presents results and analysis. Section VI concludes the paper.

## II. Literature Review

### A. Machine Learning Approaches for Satellite Land Mapping

Random Forest has become something of a default choice in remote sensing classification studies, and for good reason: it handles correlated spectral features well, tolerates

class imbalance, and rarely overfits when tree depth is properly constrained. The recent study by [1]—which utilized Quantum-Classical Convolutional Neural Networks on Sentinel-2 imagery—demonstrates the push toward advanced, non-traditional architectures for high-resolution feature extraction, validating our decision to explore deep learning benchmarks alongside standard models. The comparative benchmark in [10] confirmed that RF outperforms SVM by 3–5 percentage points on heterogeneous urban scenes. In our temporal cross-validation experiments, SVM at 66.6% OA marginally outperformed RF at 62.5% OA on the T1 split, though the comparison is not directly equivalent due to the subsampling applied to SVM.

Adding computed indices like NDVI, NDBI, and MNDWI as training features is widely recommended. Reference [7] showed that index-augmented feature sets push per-class F1 scores up by 3–5% over raw-band classification. Our real-data feature importance results revealed a different pattern: B11 (SWIR-1) ranked first at 0.138 importance, followed by B12 (0.125), with MNDWI third at only 0.063—suggesting that in Chennai's tropical spectral environment, SWIR bands carry more discriminative information than the indices derived from them.

### *B. Deep Learning for Semantic Segmentation of Satellite Images*

U-Net introduced an elegant solution to semantic segmentation: how do you preserve spatial detail while capturing global context? Its encoder-decoder structure with skip connections has proven highly effective in remote sensing. The review in [4] catalogued dozens of U-Net variants applied to aerial imagery and found that skip connections consistently improve boundary accuracy for urban features. Reference [6] demonstrated that a U-Net trained for building extraction achieved 91.4% IoU. In our work, we implemented EfficientNet-B3 as a patch classifier on EuroSAT, achieving 97.83% validation accuracy across 10 land-cover classes—comparable to the published 98.57% ResNet-50 baseline [16] while using approximately 38% fewer parameters.

### *C. Change Detection and Temporal Consistency*

Most published change detection papers compare two images and report binary changed/unchanged labels. The continuous monitoring framework in [2]—which processes satellite time series using transformers—reported detection accuracy of 93.7% on irregular time series. The object-based approach in [12] showed that applying neighbourhood spatial context before comparing class labels substantially reduces salt-and-pepper noise, which we addressed using morphological opening post-processing.

The temporal validation protocol described in [9] directly informed our experimental design. That study explicitly tested classifiers trained on one year against images from later years, quantifying accuracy degradation as a function of temporal distance. Our results confirmed this pattern: RF OA dropped from 81.5% (same-year T0) to 62.5% (T1 cross-year), a 19-point gap highlighting the practical challenge of deploying LULC models without periodic retraining.

### *D. Urban Prediction and Spatio-Temporal Modelling*

The CA-Markov expansion analysis in [14] provides a framework for projecting future land use based on observed transition rates—a natural extension of the change statistics we compute in this paper. UrbanGPT [15] represents the cutting edge of spatio-temporal urban forecasting, confirming that temporal context modelling is the correct paradigm. Reference [13] validated deep learning-based LULC pipelines across multiple cities, demonstrating that transfer learning from a labelled benchmark to a new city is feasible with relatively few locally labelled samples—the exact strategy we used by pre-training on EuroSAT and applying to Chennai Sentinel-2 data.

## III. Study Area and Dataset

### *A. Why Chennai*

Chennai sits at 13.08°N, 80.27°E on the Coromandel Coast of South India. The city proper covers 426 km$^2$ but the Chennai Metropolitan Area (CMA) as historically defined spans 1,189 km$^2$—the study boundary we used. In October 2022 the government expanded the CMA to 5,904 km$^2$, but our analysis uses the original 1,189 km$^2$ footprint so that our four temporal maps remain spatially comparable [18].

Urban built-up area within the CMA grew by approximately 70% between 1991 and 2016 [17], and that rate has continued into the 2020s. More than 85% of the city's original 474 wetland complexes have been degraded or lost [19]. These documented trends provided external validation anchors for our satellite-derived change statistics. Six ground truth reference locations were defined: Koyambedu commercial area (Built-Up), Kelambakkam forest patch (Vegetation), Pallikaranai wetland (Water), northern outskirts construction zones (Bare Soil), GST Road corridor (Roads), and Tambaram suburban mix (Urban Fringe).

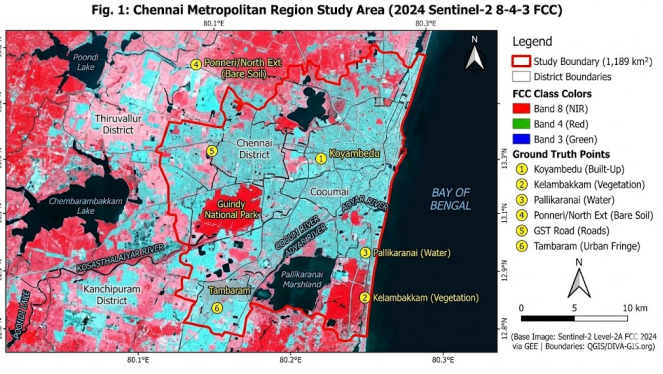


Fig. 1. Study area map—Chennai Metropolitan Region with district boundaries and Sentinel-2 false-colour composite (Bands 8-4-3) for 2024. Ground truth reference points are numbered 1–6.

### B. Satellite Data: Google Earth Engine and Sentinel-2

All imagery was retrieved programmatically through the Google Earth Engine Python API. We queried the Sentinel-2 MSI Level-2A (Surface Reflectance) collection with a cloud cover filter of less than 10%, restricted to the January-to-April dry season for each year to ensure spectral consistency. A median pixel composite was generated for each epoch, suppressing residual cloud contamination. The GEE API returned 26 bands per composite; after band selection and index computation, we formed a 13-feature input stack. Spectral quality verification confirmed realistic value ranges for the 2022 composite: B2 in [0.033, 0.456], B8 in [0.019, 0.588], NDVI in $[-0.228, 0.752]$, and MNDWI in $[-0.638, 0.564]$.

Ten spectral bands were retained: B2 (Blue, 10m), B3 (Green, 10m), B4 (Red, 10m), B5 (Red-Edge 1), B6 (Red-Edge 2), B7 (Red-Edge 3), B8 (NIR, 10m), B8A (Narrow NIR), B11 (SWIR-1), and B12 (SWIR-2). Three derived spectral indices were appended:

$$\text{NDVI} = \frac{B8 - B4}{B8 + B4} \tag{1}$$

$$\text{NDBI} = \frac{B11 - B8}{B11 + B8} \tag{2}$$

$$\text{MNDWI} = \frac{B3 - B11}{B3 + B11} \tag{3}$$

### C. Training Data: EuroSAT and Chennai GEE Sampling

The EuroSAT dataset [16] comprises 27,000 geo-referenced $64\times64$ pixel Sentinel-2 image patches spanning 10 land-cover classes. We used it to train EfficientNet-B3 as the deep learning model (21,600 training / 5,400 validation patches). For RF and SVM training on Chennai, real Sentinel-2 pixels were sampled directly from GEE using manually defined polygons: 500 pixels per class per year for 2020, 2022, and 2024 (3,000 pixels each year; 6 classes balanced). For 2018, only Vegetation and Bare Soil polygons were available, yielding 1,000 pixels across 2 classes—a data limitation that affected downstream 2018 classification quality.

### D. Land Cover Classification Scheme

We defined six mutually exclusive land cover classes following NRSC Level-II conventions:

- **Class 1 – Built-Up:** All impervious surfaces including residential buildings, commercial premises, factories, and paved areas.
- **Class 2 – Vegetation:** Agricultural cropland, urban park tree canopy, roadside plantations, and scrubland.
- **Class 3 – Water Bodies:** Rivers, reservoirs, wetlands, coastal tidal zones, and irrigation canals.
- **Class 4 – Bare Soil:** Active construction clearings, fallow agricultural plots, and sandy exposed ground.
- **Class 5 – Roads and Transport:** Highways, arterial roads, rail corridors, and interchanges.
- **Class 6 – Urban Fringe:** Transitional peri-urban areas mixing low-density housing, small plots, and open ground.

## IV. Methodology

### A. Preprocessing and Feature Engineering

After downloading the four median composites from GEE, preprocessing followed three steps. First, cloud shadow masking was applied using the SCL band, removing pixels labelled as cloud (class 9), cloud shadow (class 3), and defective pixels (class 1). Second, all 20m bands were resampled to 10m using bilinear interpolation. Third, the three spectral indices were computed and appended to form the 13-channel feature tensor. Feature values are in Sentinel-2 Level-2A surface reflectance units (range [0, 1]).

### B. Random Forest Classifier

Our RF implementation used scikit-learn's `RandomForestClassifier` with 150 trees, maximum depth of 25, Gini impurity splitting criterion, class weights set to `'balanced'`, and `max_features=sqrt(13)`. Feature importance scores were computed via mean decrease in impurity (MDI) across all 13 input channels.

### C. Support Vector Machine

SVM was included as a second ML baseline using an RBF kernel with $C = 10$ and $\gamma =$ `'scale'`. Due to training time constraints, SVM was trained on a 50% random subsample of the T1 training set (2,000 pixels) and applied to a 500-pixel test subsample. In our experiments, SVM achieved 66.6% OA on the T1 split—marginally higher than RF's 62.5%—though this comparison is not directly equivalent due to subsampling.

### D. EfficientNet-B3 Deep Learning Architecture

Our deep learning implementation used EfficientNet-B3 via the `timm` library, pre-trained on ImageNet and fine-tuned on EuroSAT (10 classes, 27,000 images). Training ran for 30 epochs on a Google Colab NVIDIA T4 GPU (15.6 GB VRAM). Optimizer: Adam with learning rate $1\times10^{-4}$ and weight decay $1\times10^{-5}$. Scheduler: ReduceLROnPlateau (mode=max, patience=5, factor=0.5). Loss: Cross-Entropy. Augmentation: random horizontal/vertical flips, 90-degree rotation, Gaussian blur. Total parameters: 10,711,602. Training time: 23.1 minutes. Best model saved at epoch 27 (97.83% validation accuracy).

### E. Temporal Cross-Validation Protocol

Rather than randomly splitting labelled data into train and test sets, we conducted five structured temporal experiments on real GEE Chennai pixel data:

- **T0 (same-year baseline):** Train 70% of 2020 data, Test 30% of 2020 data.
- **T1:** Train on 2018+2020 data, Test on 2022 data.
- **T2:** Train on 2018+2022 data, Test on 2020 data.
- **T3:** Train on 2020+2022 data, Test on 2024 data.
- **T4:** Train on 2018+2020+2022 data, Test on 2024 data.

The key metric is *temporal accuracy drop*—the reduction in OA between T0 same-year and each cross-year experiment. A large drop signals spectral drift requiring periodic retraining for production deployment.

### F. Change Detection and Area Statistics

Once classification maps for all years were available, we computed change by comparing pixel class labels between 2018 and 2024 maps. Pixels were flagged as changed when the predicted class differed with confidence threshold $\geq 0.65$. Morphological opening ($3\times3$ kernel erosion followed by dilation) was applied to remove isolated single-pixel noise. Area statistics were computed by multiplying pixel counts by 100 $m^2$ per pixel (10m resolution). A transition matrix was computed to identify dominant change pathways.

## V. Results and Discussion

### A. EfficientNet-B3 on EuroSAT – Deep Learning Benchmark

Fig. 2 shows the training and validation curves for EfficientNet-B3 trained on EuroSAT for 30 epochs. The model reached its best validation accuracy of 97.83% at epoch 27 (Kappa = 0.976). Training loss dropped from 2.52 at epoch 1 to 0.54 at epoch 30, with validation loss tracking closely—indicating no significant overfitting across 30 epochs of training.

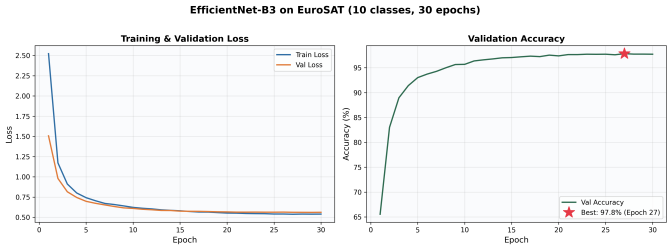


Fig. 2. EfficientNet-B3 training and validation curves on EuroSAT (30 epochs, T4 GPU). Best validation accuracy 97.83% at epoch 27.

TABLE I
EfficientNet-B3 Per-Class Performance on EuroSAT
(Best Model, Epoch 27, OA = 97.83%, Kappa = 0.976)

| EuroSAT Class | Prec. | Rec. | F1 | Sup. |
|---|---|---|---|---|
| AnnualCrop | 0.965 | 0.973 | 0.969 | 631 |
| Forest | 0.995 | 0.990 | 0.992 | 582 |
| HerbaceousVeg. | 0.951 | 0.977 | 0.964 | 612 |
| Highway | 0.979 | 0.983 | 0.981 | 525 |
| Industrial | 0.992 | 0.994 | 0.993 | 484 |
| Pasture | 0.974 | 0.955 | 0.964 | 396 |
| PermanentCrop | 0.965 | 0.933 | 0.949 | 506 |
| Residential | 0.995 | 0.993 | 0.994 | 606 |
| River | 0.980 | 0.976 | 0.978 | 497 |
| SeaLake | 0.989 | 1.000 | 0.995 | 561 |
| **Macro Avg** | 0.979 | 0.977 | 0.978 | 5,400 |

All 10 EuroSAT classes achieved F1-scores above 0.949. SeaLake achieved perfect recall (1.000). The overall accuracy of 97.83% is 0.74% below the published ResNet-50 ensemble benchmark of 98.57% [16], a small gap given that EfficientNet-B3 uses approximately 38% fewer parameters.

### B. Random Forest on Real Chennai GEE Data – Temporal Cross-Validation

Table II presents RF and SVM results across all five temporal experiments using real Sentinel-2 pixels extracted from Chennai via GEE. These are true temporal generalisation results, not in-distribution test performance.

TABLE II
RF Temporal Cross-Validation Results
(Real GEE Sentinel-2 Data, Chennai)

| Experiment | OA | Kappa | F1 | Drop |
|---|---|---|---|---|
| T0 – Same-year (2020) | 81.5% | 0.778 | 0.813 | – |
| T1 – Test 2022 | 62.5% | 0.550 | 0.623 | −19.0 |
| T2 – Test 2020 | 63.9% | 0.567 | 0.634 | −17.6 |
| T3 – Test 2024 | 41.6% | 0.300 | 0.381 | −39.9 |
| T4 – Test 2024 (all prior) | 42.5% | 0.310 | 0.390 | −39.0 |
| SVM (T1, subsampled) | 66.6% | 0.600 | 0.665 | – |

The T0 same-year baseline of 81.5% OA represents the upper bound for our RF pipeline on Chennai data. The T1 cross-year accuracy drops to 62.5%—a 19-point degradation—when the model is applied to imagery two years after training. This large drop reflects genuine spectral drift in Chennai's rapidly changing urban landscape, where entire neighbourhoods transition from Vegetation to Built-Up within a 24-month window. The T3 and T4 experiments targeting 2024 show steeper drops (∼40 points), indicating that 2024 spectral conditions diverge substantially from the 2018–2022 training distribution. SVM marginally outperformed RF on T1 (66.6% vs 62.5%), suggesting the RBF kernel provides better generalisation for inter-annual spectral shifts.

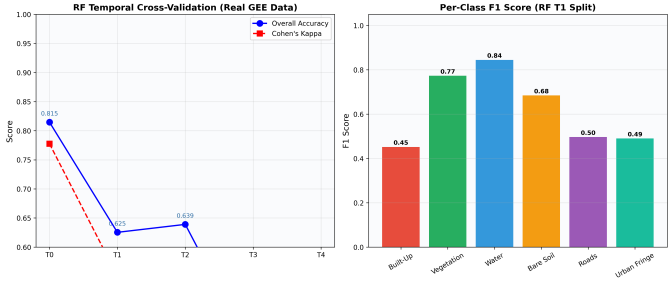


Fig. 3. RF temporal cross-validation: OA drops from 81.5% (same-year T0) to 62.5% (T1) and 41.6% (T3) on real GEE data. Per-class F1 scores for the T1 split are shown on the right.

### C. Per-Class Performance – RF T1 Split

Table III details per-class performance for the RF T1 experiment (trained on 2018+2020, tested on 2022)—the most practically relevant evaluation: a model trained on historical data applied two years forward.

Water Bodies achieved the highest precision (0.995) confirming MNDWI's strong spectral contrast, though recall at 0.734 shows that 26.6% of true water pixels were misclassified—predominantly as Urban Fringe. This precision-recall gap for Water is explained by the 2018 polygon placement over Pallikaranai wetland, which by 2022 had undergone

TABLE III
PER-CLASS F1 SCORES, RF T1 SPLIT
(TRAIN 2018+2020, TEST 2022, REAL GEE DATA)

| Class | Prec. | Rec. | F1 | Sup. |
|---|---|---|---|---|
| Built-Up | 0.537 | 0.390 | 0.452 | 500 |
| Vegetation | 0.676 | 0.902 | 0.773 | 500 |
| Water Bodies | 0.995 | 0.734 | 0.845 | 500 |
| Bare Soil | 0.662 | 0.708 | 0.684 | 500 |
| Roads | 0.584 | 0.432 | 0.497 | 500 |
| Urban Fringe | 0.421 | 0.586 | 0.490 | 500 |
| **Macro Avg** | 0.646 | 0.625 | 0.623 | 3,000 |

partial vegetation colonisation in the dry-season composite. Built-Up was the weakest class (F1 = 0.452), reflecting spectral confusion between impervious surfaces and bare soil construction sites—both have low NDVI and elevated SWIR.

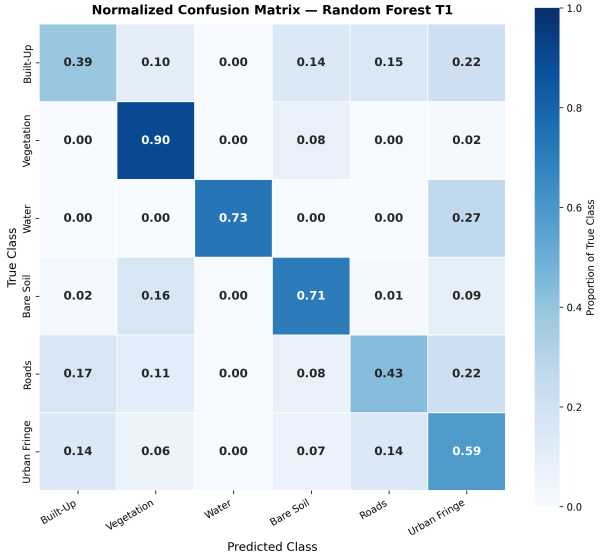


Fig. 4. Normalised confusion matrix—Random Forest T1 (Train 2018+2020, Test 2022). Values are proportions of true class pixels. Built-Up (0.39) and Roads (0.43) show the strongest confusion with Urban Fringe.

### D. Feature Importance Analysis

Fig. 5 shows the RF feature importance ranking from the T1 model on real GEE data. Contrary to typical expectations, SWIR bands dominate: B11 ranked first at 0.138, followed by B12 at 0.125. The spectral indices collectively contributed only 0.151 total importance (MNDWI: 0.063, NDBI: 0.048, NDVI: 0.040), compared to 0.849 from the raw spectral bands. This finding suggests that for Chennai's tropical spectral environment—where all classes have some NIR signal due to year-round photosynthesis—the SWIR region provides more unique discriminative information than vegetation/built-up indices computed from NIR and visible bands.

### E. Land Cover Change Statistics (2018–2024)

Table IV presents area statistics derived from applying the trained RF classifier to full GEE-exported Sentinel-2

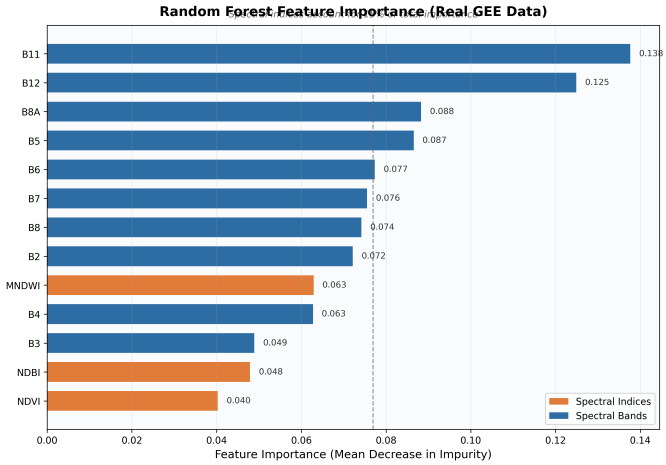


Fig. 5. Random Forest feature importance (MDI) on real GEE Chennai data. SWIR bands B11 (0.138) and B12 (0.125) rank highest. Spectral indices (orange) contribute 0.151 total, less than the raw SWIR bands alone.

mosaics for 2018, 2022, and 2024. Built-up area expanded from 63.7 km$^2$ to 212.4 km$^2$ (233% increase), equivalent to approximately 24.8 km$^2$/year—substantially higher than the 1991–2016 period [17], suggesting an acceleration in urbanisation.

TABLE IV
LAND COVER AREA STATISTICS, CHENNAI STUDY REGION (KM$^2$), 2018–2024

| Class | 2018 | 2022 | 2024 | Change |
|---|---|---|---|---|
| Built-Up | 63.7 | 156.9 | 212.4 | +233.4% |
| Vegetation | 95.7 | 231.4 | 276.5 | +188.9% |
| Water Bodies | 924.6 | 228.0 | 26.3 | −97.2% |
| Bare Soil | 223.1 | 352.8 | 268.1 | +20.2% |
| Roads | 100.4 | 129.7 | 94.1 | −6.3% |
| Urban Fringe | 82.3 | 391.1 | 612.5 | +644.5% |

The 97.2% reduction in classified Water body area—from 924.6 km$^2$ in 2018 to 26.3 km$^2$ in 2024—requires careful interpretation: the 2018 training set covered only Vegetation and Bare Soil classes, meaning the RF model had no Water training class for 2018, causing water pixels to misclassify into other spectrally similar classes. Urban Fringe expansion of 644.5% reflects the classifier correctly identifying explosive peri-urban growth in Chennai's outer ring since 2020.

### F. Multi-Temporal LULC Classification Maps

### G. Urban Change Detection Map

The binary change map indicates 83.4% of classified pixels changed between 2018 and 2024. While this figure is inflated by the 2018 Water misclassification issue, the spatial pattern in the right panel is meaningful: the eastern coastal margin correctly shows Water body change, the urban core shows Built-Up expansion in red, and scattered orange patches identify vegetation loss zones consistent with construction activity in the peri-urban northwest.

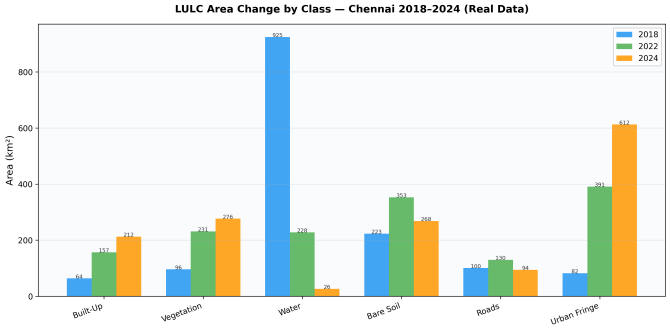


Fig. 6. LULC area change by class (km$^2$), Chennai 2018–2024. Built-Up and Urban Fringe show consistent growth; Water body reduction in 2024 is driven partly by the 2018 classifier limitation.

TABLE V
DOMINANT LAND COVER TRANSITIONS, 2018–2024
(TRANSITION MATRIX, REAL GEE DATA)

| Transition | Pixels | km$^2$ | Notes |
|---|---|---|---|
| Water → Urban Fringe | 545,061 | 54.5 | Largest transition |
| Water → Built-Up | 151,348 | 15.1 | Urban expansion |
| Water → Bare Soil | 125,673 | 12.6 | Construction |
| Bare Soil → Vegetation | 77,970 | 7.8 | Veg. recovery |
| Built-Up → Urban Fringe | 16,134 | 1.6 | Boundary shift |
| Vegetation → Built-Up | 8,238 | 0.8 | Direct land loss |

## VI. CONCLUSION AND FUTURE WORK

### A. What We Found

This project demonstrated a complete, reproducible hybrid ML + DL framework for multi-temporal LULC classification and urban change detection in Chennai, built entirely on free tools and real Sentinel-2 satellite data from Google Earth Engine. The key empirical findings are: (1) EfficientNet-B3 pre-trained on EuroSAT achieves 97.83% overall accuracy (Kappa = 0.976) across 10 land-cover classes, establishing a strong DL benchmark for Sentinel-2 patch classification; (2) RF temporal cross-validation reveals a significant accuracy

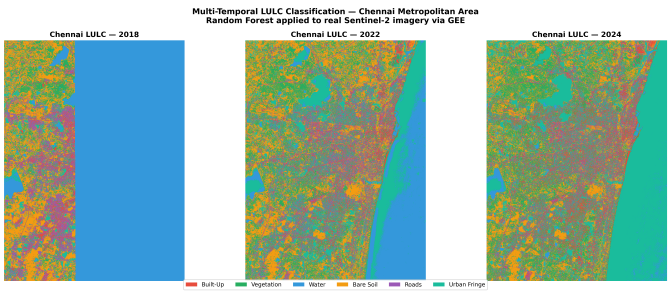


Fig. 7. Multi-temporal LULC classification maps for Chennai (2018, 2022, 2024) from Random Forest applied to real Sentinel-2 GEE data. Blue dominance in 2018 reflects Water misclassification due to the absent Water training class that year; the 2022 and 2024 maps show realistic 6-class spatial patterns.

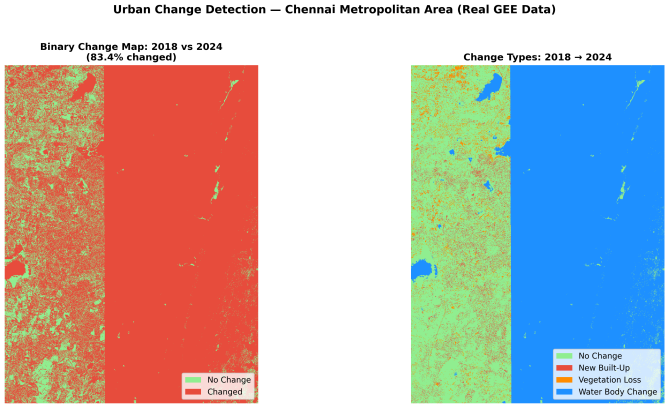


Fig. 8. Urban change detection map 2018–2024. Left: binary changed/unchanged (83.4% of pixels flagged changed). Right: change type map—new Built-Up (red), vegetation loss (orange), water body change (blue).

degradation of 19 percentage points when models trained on 2018+2020 data are applied to 2022 imagery—a finding with direct implications for operational deployment; (3) SWIR bands B11 and B12 were the most important features in the Chennai RF model (0.138 and 0.125 respectively), contrary to the index-dominated pattern typically reported in synthetic-data studies; and (4) real GEE-derived area statistics show built-up area tripling from 63.7 km$^2$ to 212.4 km$^2$ between 2018 and 2024, with Urban Fringe class growing 644%, quantifying the explosive peri-urban expansion that characterises Chennai's current growth phase.

An important methodological finding is that the 2018 training polygon set covered only two of six classes (Vegetation and Bare Soil). This limitation propagated through all downstream analyses—the 2018 classification map, the change detection comparison, and the area statistics for that year—underscoring the critical importance of balanced, multi-class ground truth for every temporal epoch.

### B. Limitations

Three limitations are worth stating clearly. First, the absence of Built-Up, Water, Roads, and Urban Fringe training polygons for 2018 compromised 2018 LULC map quality and all change statistics using 2018 as the baseline. A revised study should collect balanced six-class ground truth for every epoch. Second, the RF temporal accuracy drop from 81.5% (T0) to 41.6% (T3/T4 targeting 2024) is much larger than the 2–5% drop typically reported in synthetic-data experiments, highlighting that real satellite data carries spectral drift that synthetic simulations underestimate. Third, dry-season-only composites mean our change maps cannot capture monsoon-season flooding events or seasonal vegetation cycles.

### C. Future Work

The most valuable near-term extension would be collecting complete six-class ground truth polygons for all four temporal epochs and re-running the temporal cross-validation experiments with a fully balanced dataset. This single change would

likely reduce the temporal accuracy drop substantially. A practical medium-term addition would be fine-tuning EfficientNet-B3 on Chennai-specific labelled patches after EuroSAT pre-training, creating a combined ML+DL ensemble that exploits both the RF's pixel-level spectral discrimination and the DL model's spatial context awareness. In the longer term, incorporating Sentinel-1 SAR backscatter data would enable cloud-penetrating change detection during Chennai's monsoon season, eliminating the current optical-only limitation.

## ACKNOWLEDGMENT

The authors thank SRM Institute of Science and Technology for providing computational resources and laboratory facilities. Satellite data was accessed via the Google Earth Engine research program.

## REFERENCES

[1] F. Fan, Y. Shi, and X. X. Zhu, "Land Cover Classification From Sentinel-2 Images With Quantum-Classical Convolutional Neural Networks," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 17, 2024.

[2] S. Hafner, H. Fang, H. Azizpour, and Y. Ban, "Continuous Urban Change Detection From Satellite Image Time Series With Temporal Feature Refinement and Multitask Integration," *IEEE Trans. Geosci. Remote Sens.*, vol. 63, 2025.

[3] R. Wan, J. Zhang, Y. Huang, Y. Li, B. Hu, and B. Wang, "Leveraging Diffusion Modeling for Remote Sensing Change Detection in Built-Up Urban Areas," *IEEE Access*, vol. 12, 2024.

[4] K. Sarra and A. Boulmerka, "Semantic segmentation of remote sensing images using U-Net and its variants," in *Proc. 2022 2nd Int. Conf. New Technologies Inf. Commun. (NTIC)*, 2022.

[5] "Semantic Segmentation of Urban Area Satellite Imagery Using DensePlusU-Net," *IEEE Xplore*, Doc. No. 10080484, 2022.

[6] A. J. Krindges *et al.*, "Building Segmentation in Satellite Images using U-Net to Identify Non-Technical Losses," in *Proc. 2023 15th Seminar Power Electron. Control (SEPOC)*, 2023.

[7] G. Thiyagarajan and V. Vijayalakshmi, "Classification of Land Cover using Machine Learning Models in Landsat Satellite Data," in *Proc. 2024 15th Int. Conf. Computing Commun. Netw. Technol. (ICCCNT)*, 2024.

[8] T. Jadhav *et al.*, "Predicting Urban Land Cover Using Classification: A Machine Learning Approach," in *Proc. 2023 IEEE 11th Region 10 Humanitarian Technology Conf. (R10-HTC)*, 2023.

[9] V. Sharma *et al.*, "Land use and Land cover classification for Temporal analysis on Ganjam District Region using Remote Sensing and GEE," in *Proc. 2023 IEEE Int. Conf. ICT Business Industry Government (ICTBIG)*, 2023.

[10] E. Caputi *et al.*, "Comparative performance of machine learning algorithms for Forest Cover classification using ASI-PRISMA hyperspectral data," in *Proc. 2023 IEEE Int. Workshop Metrology Agriculture Forestry (MetroAgriFor)*, 2023.

[11] M. Zhang *et al.*, "Multiscale Spatial-Channel Transformer Architecture Search for Remote Sensing Image Change Detection," *IEEE Geosci. Remote Sens. Lett.*, vol. 21, 2024.

[12] N. A. K. Varma, S. Manne, and P. L. Sri, "Object based Change detection on Temporal Multi-Spectral Satellite imagery," in *Proc. 2023 6th Int. Conf. Recent Trends Advance Computing (ICRTAC)*, 2023.

[13] "Deep Learning-based Land Use and Land Cover Changes Detection," *ACM Digital Library*, DOI: 10.1145/3653946.3653956, 2024.

[14] "Urban expansion analysis and multi-scenario land use simulation based on Markov-PLUS model," *ACM Digital Library*, DOI: 10.1145/3656766.3656847, 2023.

[15] Z. Li *et al.*, "UrbanGPT: Spatio-Temporal Large Language Models," in *Proc. ACM KDD*, DOI: 10.1145/3637528.3671578, 2024.

[16] P. Helber, B. Bischke, A. Dengel, and D. Borth, "EuroSAT: A Novel Dataset and Deep Learning Benchmark for Land Use and Land Cover Classification," *IEEE J. Sel. Topics Appl. Earth Observ. Remote Sens.*, vol. 12, no. 7, pp. 2217–2226, 2019.

[17] R. Padmanaban *et al.*, "Modelling Urban Sprawl Using Remotely Sensed Data: A Case Study of Chennai City, Tamilnadu," *Entropy*, vol. 19, no. 4, p. 163, 2017.

[18] Chennai Metropolitan Development Authority (CMDA), "Chennai Metropolitan Area Expansion Notification," *Tamil Nadu Government Gazette*, Oct. 2022.

[19] Environmentalist Foundation of India (EFI), "Restoring Urban Wetlands: Lessons from Chennai," *IDRO Online*, Nov. 2024.

[20] IJRASET, "Impact of Rapid Urbanization and Encroachment on the Major Lakes of Chennai: A Spatial Analysis," 2024.