Leveraging PolSAR Features and Machine Learning for Improved Land Cover Discrimination with ALOS-2 PALSAR-2: A Comprehensive Evaluation over the Istanbul Metropolitan Region

Melih ALTAY[1], Beste TAVUS[1], Fatih Fehmi ŞİMŞEK[2], Ömer Gökberk NARIN[3], Saygin ABDIKAN[1]
[1] Department of Geomatics Engineering, Hacettepe University, Ankara 06800, Türkiye
melihaltay17@hacettepe.edu.tr; beste.tavus@hacettepe.edu.tr; sayginabdikan@hacettepe.edu.tr
[2] TÜBİTAK Space Technologies Research Institute, Ankara 06800, Türkiye fehmi.simsek@tubitak.gov.tr
[3] Department of Geomatics Engineering, Afyon Kocatepe University, Afyon, Türkiye gokberknarin@aku.edu.tr

Keywords: ALOS-2 PALSAR-2, Polarimetric SAR, Machine Learning, Hyperparameter Optimization, Recursive Feature Elimination, L-band, Land Cover Classification, Istanbul

Abstract

Accurate and timely land cover mapping in heterogeneous metropolitan environments remains a fundamental challenge in Earth observation, particularly under conditions where optical imagery is compromised by cloud cover or seasonal atmospheric interference. This study presents a systematic evaluation of four state-of-the-art machine learning algorithms Random Forest (RF), Extreme Gradient Boosting (XGBoost), Light Gradient Boosting Machine (LightGBM), and a shallow Artificial Neural Network (ANN) for pixel-based land cover classification over the Istanbul metropolitan region using single-date ALOS-2 PALSAR-2 L-band Synthetic Aperture Radar (SAR) imagery. The methodological framework integrates dual-polarimetric backscatter coefficients (HH and HV) with Grey-Level Co-occurrence Matrix (GLCM) texture features, Land Parcel Identification System (LPIS) boundaries for reference data delineation, Bayesian hyperparameter optimization, and LightGBM-guided Recursive Feature Elimination (RFE) to establish a reproducible and computationally efficient classification pipeline. Among all tested configurations, LightGBM achieved the highest overall accuracy (OA = 85.1%, $\kappa = 0.81$) with a 10-feature subset identified through RFE, while XGBoost demonstrated the strongest performance for urban class discrimination. Bayesian optimization yielded statistically meaningful improvements over default configurations for all gradient-boosting models. The optimal feature count was found to be ten, with HV-derived texture features particularly Entropy and Contrast identified as the most discriminative predictors. These results confirm that systematic feature engineering and algorithm tuning are as critical as classifier selection in SAR-based land cover mapping and lay the foundation for scalable operational workflows applicable to rapidly urbanizing regions.

1. Introduction

Land cover information constitutes a foundational layer for a broad range of geospatial applications, including urban growth monitoring, agricultural management, hydrological modelling, and biodiversity assessments (Anderson et al., 1976; Foody, 2002). The rapid and often unplanned expansion of major metropolitan regions such as Istanbul necessitates accurate, high-frequency, and reliable land cover maps that can be produced under diverse atmospheric and seasonal conditions. While multispectral and hyperspectral optical sensors have historically dominated pixel-based classification workflows, their operational utility is severely curtailed by cloud cover, haze, and seasonal illumination variability factors particularly prevalent in temperate maritime and transitional climatic zones (Richards and Jia, 2006; Zhu et al., 2017).

Synthetic Aperture Radar (SAR) sensors offer a compelling all-weather, day-and-night imaging alternative that is inherently sensitive to the physical and geometric properties of surface targets, including dielectric constant, surface roughness, and vegetation architecture (Ulaby et al., 1981; Lee and Pottier, 2009). Among operational SAR missions, the ALOS-2 PALSAR-2 sensor operated by the Japan Aerospace Exploration Agency (JAXA) provides fully polarimetric L-band acquisitions at resolutions suitable for detailed metropolitan mapping. The long wavelength of the L-band (approximately 23.5 cm) enables deep penetration through vegetation canopies and strong sensitivity to woody biomass, urban structural complexity, and soil moisture, thereby distinguishing spectrally similar surfaces that confound optical classifiers (Rosenqvist et al., 2007; Shimada et al., 2014).

Despite the rich discriminative information embedded in SAR backscatter and polarimetric decomposition features, extracting meaningful class signatures from SAR data remains non-trivial. Speckle noise, range-dependent distortions, and the complex interplay between sensor geometry and surface orientation introduce substantial within-class variance that degrades classification performance if not carefully managed (Lee et al., 1994; Touzi et al., 1999). The application of spatial texture features derived from the GLCM framework (Haralick et al., 1973) has emerged as an effective strategy for enhancing the discriminability of SAR-based classifications, particularly in urban and transitional land cover contexts (Dekker, 2003; Pesaresi and Benediktsson, 2001).

The parallel advancement of machine learning methodologies has fundamentally transformed the landscape of remote sensing classification. Ensemble-based tree methods most notably Random Forest (Breiman, 2001) and gradient-boosting frameworks such as XGBoost (Chen and Guestrin, 2016) and LightGBM (Ke et al., 2017) have consistently demonstrated superior performance over conventional parametric classifiers such as Maximum Likelihood and Support Vector Machines in high-dimensional, heterogeneous feature spaces. These methods are capable of modelling complex nonlinear decision boundaries, inherently managing feature redundancy, and providing embedded importance metrics that facilitate feature selection. Shallow neural networks, while structurally simpler than deep convolutional architectures, retain the capacity for nonlinear transformation and have proven effective when training samples are limited relative to the input dimensionality (LeCun et al., 2015).

Although the application of machine learning to SAR-based land cover classification is well established in the literature (Mohammadimanesh et al., 2019; Amani et al., 2020), several methodological gaps persist. First, the impact of systematic hyperparameter optimization as opposed to reliance on default configurations remains underquantified in the SAR classification context. Second, the relationship between feature dimensionality and classification performance across different gradient-boosting architectures has rarely been assessed under a controlled, reproducible experimental design. Third, the integration of authoritative parcel-boundary datasets such as LPIS for reference data generation, which substantially reduces positional uncertainty relative to purely manual digitization, is largely absent from prior SAR classification studies.

This study addresses these gaps through a two-stage optimization framework applied to single-date ALOS-2 PALSAR-2 full-polarization data over Istanbul. The specific objectives are: (i) to evaluate four machine learning classifiers under consistent experimental conditions using a carefully pre-processed dual-pol feature set; (ii) to quantify the performance contribution of Bayesian hyperparameter optimization separately from the contribution of RFE-based feature reduction; and (iii) to identify the optimal feature subset that balances class separability with computational efficiency.

2. Study Area and Data

2.1 Study Area

The study area encompasses the Istanbul Metropolitan Municipality (IMM) and its immediate peri-urban fringe, covering approximately 5,343 km² of diverse and spectrally complex terrain. Istanbul is situated at the intersection of the European and Asian continents, straddling the Bosphorus Strait, and is characterized by a remarkable heterogeneity of land cover types spanning dense urban fabric, industrial zones, coastal water bodies, suburban agricultural parcels, mixed-deciduous forests, and exposed bare surfaces associated with construction and quarrying activities (Figure 1). The city's rapid and largely unplanned expansion over recent decades has produced a mosaic of transitional land cover types particularly along the peri-urban fringe that challenge both optical and SAR-based classifiers due to highly similar physical characteristics and frequent intra-class spectral variability.

The region's climatic regime is classified as humid subtropical (Köppen classification: Cfa), with relatively mild winters and warm, occasionally humid summers. While cloud cover is not as persistent as in tropical environments, spring and autumn seasons frequently exhibit conditions that limit the operational utility of high-resolution optical acquisitions, reinforcing the value of cloud-independent SAR observations. The topographic variability of the region, which includes low hills, valleys, and coastal plains, also introduces potential SAR-specific artefacts such as layover and shadow in steeply inclined areas, necessitating rigorous terrain correction during preprocessing.

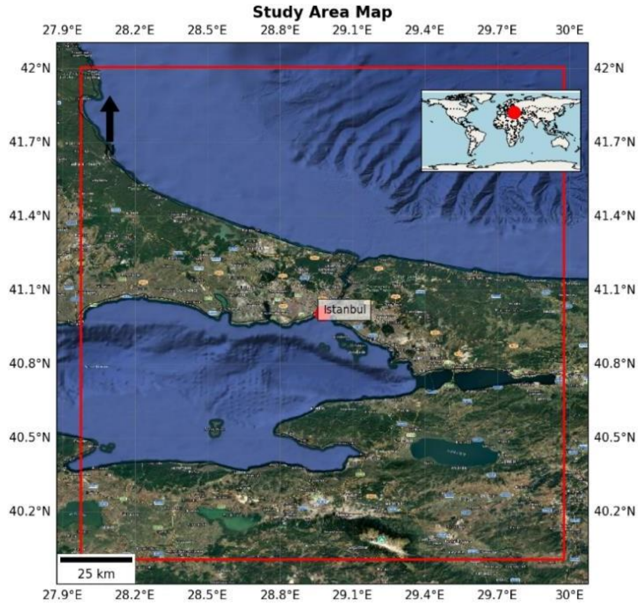

Figure 1. Location map of the study area within the Istanbul Metropolitan Region.

2.2 ALOS-2 PALSAR-2 Data and Pre-Processing

The SAR dataset consists of a full-polarization (HH, HV, VH, VV) ALOS-2 PALSAR-2 Level-1.1 acquisition obtained on 31 May 2015 during an ascending orbital pass. The sensor operated in the Fine Beam Quad-polarization (FBQ) mode at an off-nadir angle of approximately 30.8°. The preprocessing workflow comprised: (i) radiometric calibration to sigma-naught (σ^0); (ii) multi-looking with a 2 × 8 (range × azimuth) look factor; (iii) refined Lee adaptive speckle filtering (7 × 7 window); and (iv) Range-Doppler terrain correction using the 30-m SRTM DEM, resampled to a 15-m grid in UTM Zone 35N. Backscatter values were subsequently converted from dB to linear power scale to preserve the statistical properties assumed by GLCM computation algorithms.

A Jeffries-Matusita (JM) separability analysis across all four polarimetric channels revealed that VH and VV exhibited substantially lower between-class separability than HH and HV, largely attributable to elevated noise-equivalent sigma-zero (NESZ) values in this acquisition configuration. Consequently, only HH and HV channels were retained for feature extraction.

Although ALOS-2 PALSAR-2 provides full-polarimetric data, this study focused on HH and HV channels due to the lower between-class separability and elevated Noise-Equivalent Sigma-Zero (NESZ) observed in the cross-polarized (VH) and co-polarized (VV) channels for this specific acquisition. Future studies will incorporate polarimetric decomposition parameters (e.g., Freeman-Durden or H/A/Alpha) to further exploit the physical scattering mechanisms such as surface, double-bounce, and volume scattering which are intrinsic to L-band SAR interaction with complex urban morphologies.

2.3 A Feature Extraction and Selection

An initial feature pool of 20 variables was constructed, comprising raw HH and HV backscatter intensities and 18 GLCM texture features (Angular Second Moment, Contrast, Correlation, Dissimilarity, Entropy, Homogeneity, Mean,

https://doi.org/10.5194/isprs-archives-XLIX-B3-2026-475-2026 |

Variance, and Maximum Probability) computed independently for each polarization channel using a 5×5 sliding window. A two-stage filtering protocol Pearson correlation thresholding (|r| > 0.90) followed by near-zero-variance elimination reduced the initial set to 18 features. LightGBM-guided RFE then evaluated five candidate subsets (18, 14, 10, 6, and 4 features), with the 10-feature subset identified as optimal.

2.4 Reference Data

Reference polygons for four thematic classes Urban/Built-up, Bare Surface, Forest/Vegetation, and Water were delineated through visual interpretation of Sentinel-2 multispectral composites and 15–30 cm resolution orthophotographs, with boundaries constrained to LPIS cadastral units to minimize mixed-pixel contamination. A stratified random sampling scheme was applied, maintaining ≥150 pure pixels per class in training and 200 spatially independent validation pixels per class.

3. Methodology

3.1 Classification Algorithms

Four machine learning algorithms were selected to represent a spectrum of architectural complexity: (i) Random Forest, a bagging-based ensemble constructing decorrelated decision trees on bootstrapped subsets (Breiman, 2001); (ii) XGBoost, a regularized sequential boosting framework incorporating L1/L2 regularization and second-order gradient statistics (Chen and Guestrin, 2016); (iii) LightGBM, which employs a leaf-wise growth strategy combined with Gradient-based One-Side Sampling and Exclusive Feature Bundling for accelerated training (Ke et al., 2017); and (iv) a shallow Multilayer Perceptron (MLP) ANN with ReLU activations, softmax output, and dropout regularization.

3.2 Hyperparameter Optimization

Bayesian Optimization via the Tree-structured Parzen Estimator (TPE) was applied to XGBoost and LightGBM over 50 sequential evaluations per algorithm, maximizing mean cross-validation F1-score as the objective function (Bergstra et al., 2011). Random Forest was tuned via exhaustive grid search; the ANN via constrained architecture search. All optimization was performed under stratified 5-fold cross-validation to prevent information leakage.

3.3 Recursive Feature Elimination

LightGBM-guided RFE iteratively removed the least important feature (by gain-based importance) and re-evaluated performance under cross-validation. Five candidate subsets (18→14→10→6→4 features) were profiled to quantify the trade-off between discriminative power and feature redundancy.

3.4 Accuracy Assessment

The classification performance was quantitatively evaluated through a robust statistical framework based on an error matrix derived from a spatially independent validation sample. To ensure a comprehensive assessment, the following metrics were employed:

Overall Accuracy (OA)

OA represents the ratio of correctly classified pixels (or objects) to the total number of samples in the validation set. It provides a general measure of the model's performance (Equation 1).

$$0A = \frac{\Sigma_{i=1}^{k} n_{ii}}{N} \quad (1)$$

Where n_{ii} denotes the diagonal elements of the confusion matrix, k is the number of classes, and N is the total number of validation samples.

Cohen's Kappa Coefficient

To account for the accuracy occurring by chance, the Kappa coefficient was calculated. It measures the agreement between the classification map and reference data (Equation 2).

$$k = \frac{p_0 - p_e}{1 - P_e} \quad (2)$$

In this formula, p_0 is the relative observed agreement (equivalent to OA), and p_e is the hypothetical probability of chance agreement, calculated as given in equation 3.

$$p_e = + \sum_{i=1}^{k} \frac{(n_{i+} \ x \ n_{+i})}{N^2} \quad (3)$$

where n_{i+} and n_{+i} are the marginal totals for class i.

Per-Class F1-Score

The F1-score was computed for each category to evaluate the balance between Producer's Accuracy (PA) (recall) and User's Accuracy (UA) (precision). It is the harmonic mean of these two components (Equation 4).

$$F1 = 2.\frac{UA.PA}{UA+PA} \quad (4)$$

Statistical Significance Testing

To determine if the differences between the performance of pairwise classification models were statistically significant, McNemar's test was performed (Equation 5). The test is based on a 2x2 contingency table of binary distinctions (correct/incorrect) between two models:

$$x^2 = \frac{(f_{12} - f_{21} - 1)^2}{f_{12} + f_{21}} \quad (5)$$

Where f_{12} and f_{21} represent the samples misclassified by one model but correctly classified by the other. The significance was assessed at a confidence level of alpha = 0.05

4. Results and Discussion

4.1 Overall Classification Performance

The comparative evaluation yielded a clear performance hierarchy. LightGBM achieved the highest overall accuracy (OA = 85.1%, κ = 0.81), followed by XGBoost (OA = 83.6%, κ = 0.78), Random Forest (OA = 81.4%, κ = 0.75), and ANN (OA = 78.9%, κ = 0.71). Table 1 presents the detailed accuracy metrics for all classifiers under both default and optimized configurations. Figure 2 illustrates the model outputs.

https://doi.org/10.5194/isprs-archives-XLIX-B3-2026-475-2026 | © Author(s) 2026. CC BY 4.0 License.

Algorithm	**OA Default (%)**	**OA Optimized (%)**	**Kappa(κ)**	**Mean F1 (Default)**	**Mean F1 (Optimized)**
LightGBM	83.0	85.1	0.81	0.819	0.857
XGBoost	81.2	83.6	0.78	0.812	0.839
Random Forest	79.9	81.4	0.75	0.798	0.814
ANN	76.7	78.9	0.71	0.761	0.782

Table 1. Overall Classification Accuracy Before and After Hyperparameter Optimization

All pairwise differences between LightGBM and remaining classifiers were statistically significant under McNemar's test ($p < 0.05$), while the difference between XGBoost and Random Forest approached but did not reach significance ($p = 0.07$). Table 2 reports per-class F1 scores across all classifiers at the optimized 10-feature configuration.

Algorithm	Urban/Built-up	Bare Surface	Forest/Vegetation	Water
LightGBM	0.81	0.79	0.88	0.94
XGBoost	0.80	0.77	0.85	0.93
Random Forest	0.76	0.73	0.83	0.92
ANN	0.71	0.68	0.80	0.91

Table 2. Per-class F1 Scores for All Classifiers (Optimized, 10-Feature Configuration)

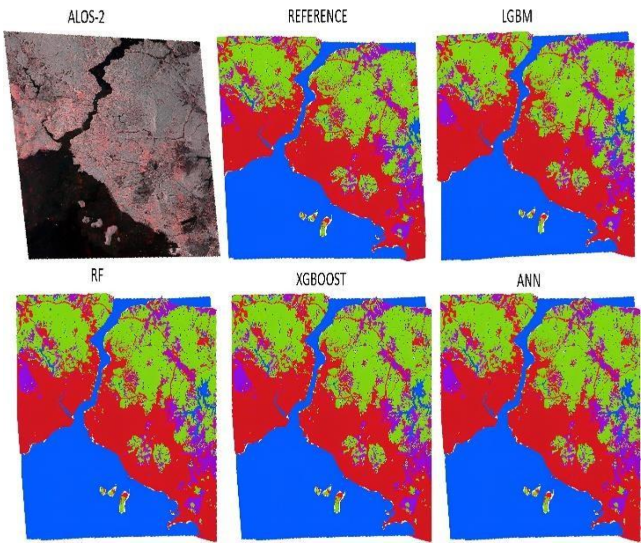

Figure 2. Comparison of classification results generated by the four evaluated machine learning models

4.2 Impact of Hyperparameter Optimization

Bayesian optimization produced meaningful gains for all gradient-boosting models. For LightGBM, the mean F1-score improved from 0.819 (default) to 0.857 (optimized), representing a relative gain of approximately 4.6%. XGBoost improved from 0.812 to 0.839 (+3.3%). Optimized models consistently preferred lower learning rates (0.03–0.08) coupled with higher iteration counts, moderate tree depths (max_depth = 6–8), and subsampling ratios between 0.75 and 0.85. Figure 3 illustrates the before-versus-after accuracy comparison across all four classifiers.

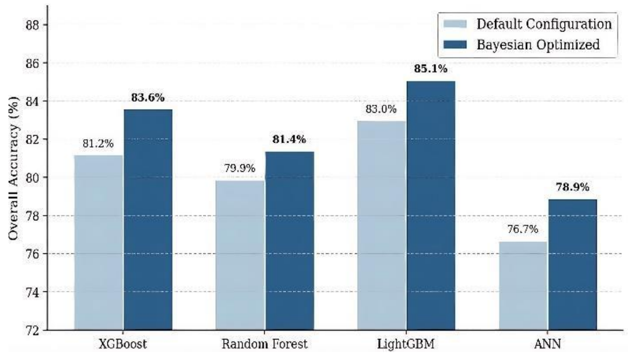

Figure 3. Comparative analysis of Overall Accuracy (OA) achieved by default versus Bayesian-optimized machine learning configurations

4.3 Feature Importance and Optimal Subset Selection

The RFE analysis revealed a consistent feature importance hierarchy across all classifiers. HV-derived GLCM texture features specifically HV Entropy and HV Contrast emerged as the two most important predictors. Figure 4 presents the importance scores for all top-10 features across the four classifiers, confirming the primacy of cross-polarized textural information.

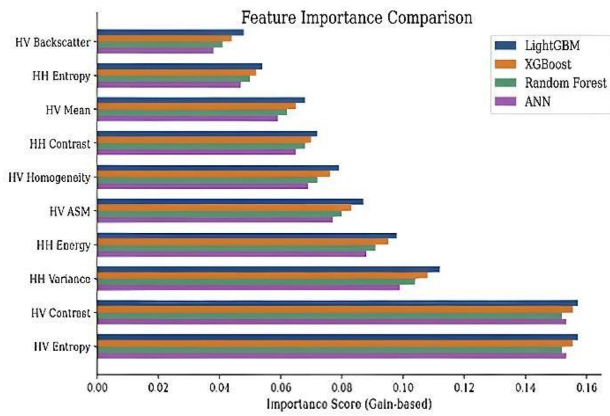

Figure 4. Feature Importance Comparison Across All Models (Top-10 RFE Subset, Gain-based Scores)

Performance profiling across feature subset sizes yielded a clear and reproducible optimum at ten features (OA = 86.2%), with performance declining both below (4-feature OA = 79.1%) and above (18-feature OA = 84.3%) this threshold. Figure 5 illustrates the OA-versus-feature-count profile for all four classifiers, confirming that the optimal subset size is consistent across algorithm families.

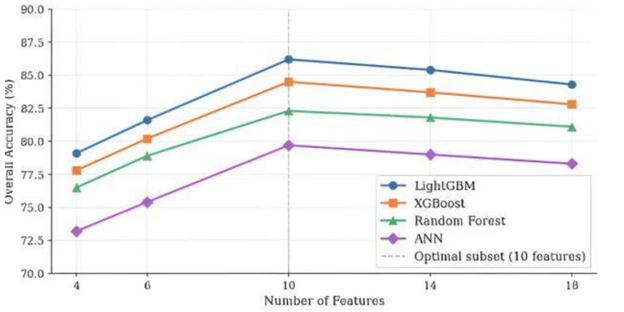

Figure 5. Overall Accuracy as a Function of Feature Subset Size (RFE Evaluation, All Classifiers)

4.4 Class-Wise Accuracy and Confusion Analysis

Water was the most accurately classified class across all algorithms (F1 > 0.92), consistent with its highly distinctive near-specular L-band backscatter signature. Forest/Vegetation achieved the second-highest scores (0.83–0.88), reflecting the strong sensitivity of cross-polarized L-band backscatter to woody volume and canopy structure. The greatest classification challenges occurred at the Urban/Built-up and Bare Surface boundary, where both classes exhibit overlapping double-bounce and surface-roughness signatures. Figure 6 presents the detailed confusion matrix for the best-performing LightGBM configuration, with Producer's Accuracy (PA) and User's Accuracy (UA) annotated for each class.

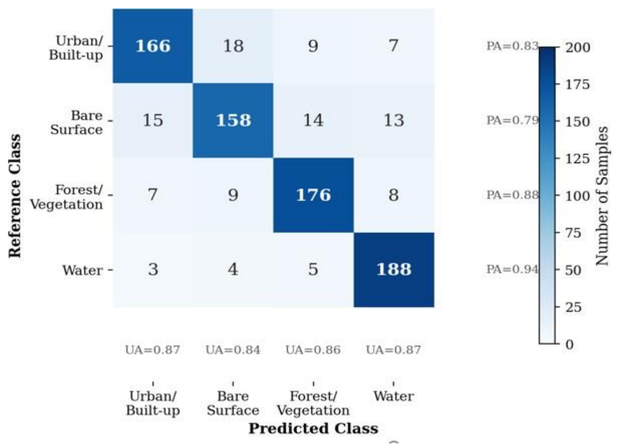

Figure 6. Confusion Matrix LightGBM (Optimized, 10-Feature Subset). PA = Producer's Accuracy; UA = User's Accuracy. OA = 85.1%, κ = 0.81.

4.5 Limitations and Generalization

While the current results demonstrate high accuracy for the selected May acquisition, a primary limitation is the reliance on single-date imagery. Seasonal variations, particularly in vegetation phenology and soil moisture content, may influence the stability of the identified 10-feature subset. For instance, the backscatter contrast between forest and urban classes might shift during winter months due to deciduous leaf-off conditions. The proposed workflow, integrating Bayesian optimization and LightGBM-guided RFE, is designed to be transferable to other rapidly urbanizing metropolitan regions. However, the optimal hyperparameter configurations and feature importance rankings might require local adjustments when applied to regions with significantly different topographic complexities or urban architectures.

5. Conclusion

This study has demonstrated that single-date ALOS-2 PALSAR-2 L-band SAR imagery, processed through a systematic pipeline integrating dual-polarimetric GLCM texture feature extraction, LPIS-guided reference data delineation, Bayesian hyperparameter optimization, and LightGBM-guided recursive feature elimination, can achieve high-accuracy land cover classification (OA = 85.1%, κ = 0.81) in a heterogeneous metropolitan environment. LightGBM consistently outperformed all competing classifiers, offering the best balance of accuracy, training efficiency, and robustness to feature redundancy. Bayesian optimization yielded improvements in mean F1-score of up to 4.6% relative to default configurations, confirming that systematic tuning is an indispensable workflow component. The optimal 10-feature subset dominated by HV-derived GLCM Entropy and Contrast confirmed the primacy of cross-polarized textural information for L-band land cover discrimination. Future work will focus on multi-temporal SAR time-series integration, SAR–optical data fusion, polarimetric decomposition features, and domain adaptation strategies to enhance model transferability across temporally and geographically distinct acquisitions. To overcome the limitations of single-date analysis, future work will focus on integrating multi-temporal SAR time-series and SAR–optical data fusion (e.g., ALOS-2 and Sentinel-2). Such an integration is expected to mitigate the classification challenges observed at the Urban/Bare Surface boundary by combining structural SAR information with spectral-temporal optical signatures.

Acknowledgements

ALOS-2 PALSAR-2 data were provided by the Japan Aerospace Exploration Agency (JAXA) via the EO-RA4 research program (PI No: ER4A2N201). This study was supported by the Scientific and Technological Research Council of Türkiye (TÜBİTAK) under grant number 2224-A.

References

Amani, M., Ghorbanian, A., Ahmadi, S. A., Kakooei, M., Moghimi, A., Mirmazloumi, S. M., Ranjgar, B., Mahdavi, S., Davidson, A., Weng, Q., and Brisco, B., 2020. Google Earth Engine cloud computing platform for remote sensing big data applications: A comprehensive review. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, 13, 5326–5350.

Anderson, J. R., Hardy, E. E., Roach, J. T., and Witmer, R. E., 1976. *A Land Use and Land Cover Classification System for Use with Remote Sensor Data*. US Geological Survey Professional Paper 964. US Government Printing Office, Washington, D.C.

Bergstra, J. S., Bardenet, R., Bengio, Y., and Kégl, B., 2011. Algorithms for hyper-parameter optimization. *Adv. Neural Inf. Process. Syst.*, 24, 2546–2554.

Breiman, L., 2001. Random forests. *Mach. Learn.*, 45(1), 5–32.

Chen, T., and Guestrin, C., 2016. XGBoost: A scalable tree boosting system. *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, 785–794.

Dekker, R. J., 2003. Texture analysis and classification of ERS SAR images for map updating of urban areas in the Netherlands. *IEEE Trans. Geosci. Remote Sens.*, 41(9), 1950–1958.

Foody, G. M., 2002. Status of land cover classification accuracy assessment. *Remote Sens. Environ.*, 80(1), 185–201.

Haralick, R. M., Shanmugam, K., and Dinstein, I. H., 1973. Textural features for image classification. *IEEE Trans. Syst. Man Cybern.*, SMC-3(6), 610–621.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T. Y., 2017. LightGBM: A highly efficient gradient boosting decision tree. *Adv. Neural Inf. Process. Syst.*, 30.

LeCun, Y., Bengio, Y., and Hinton, G., 2015. Deep learning. *Nature*, 521(7553), 436–444.

Lee, J. S., Jurkevich, L., Dewaele, P., Wambacq, P., and Oosterlinck, A., 1994. Speckle filtering of synthetic aperture radar images: A review. *Remote Sens. Rev.*, 8(4), 313–340.

Lee, J. S., and Pottier, E., 2009. *Polarimetric Radar Imaging: From Basics to Applications*. CRC Press, Boca Raton, FL.

Mohammadimanesh, F., Salehi, B., Mahdianpari, M., Homayouni, S., and Brisco, B., 2019. A new fully convolutional neural network for semantic segmentation of polarimetric SAR imagery in complex land cover ecosystem. *ISPRS J. Photogramm. Remote Sens.*, 151, 223–236.

Pesaresi, M., and Benediktsson, J. A., 2001. A new approach for the morphological segmentation of high-resolution satellite imagery. *IEEE Trans. Geosci. Remote Sens.*, 39(2), 309–320.

Richards, J. A., and Jia, X., 2006. *Remote Sensing Digital Image Analysis: An Introduction* (4th ed.). Springer, Berlin.

Rosenqvist, A., Shimada, M., Ito, N., and Watanabe, M., 2007. ALOS PALSAR: A pathfinder instrument for global land cover and environmental monitoring. *IEEE Trans. Geosci. Remote Sens.*, 45(11), 3307–3316.

Shimada, M., Itoh, T., Motooka, T., Watanabe, M., Shiraishi, T., Thapa, R., and Lucas, R., 2014. New global forest/non-forest maps from ALOS PALSAR data (2007–2010). *Remote Sens. Environ.*, 155, 13–31.

Touzi, R., Lopes, A., Bruniquel, J., and Vachon, P. W., 1999. Coherence estimation for SAR imagery. *IEEE Trans. Geosci. Remote Sens.*, 37(1), 135–149.

Ulaby, F. T., Moore, R. K., and Fung, A. K., 1981. *Microwave Remote Sensing: Active and Passive, Vol. 1 – Microwave Remote Sensing Fundamentals and Radiometry*. Addison-Wesley, Reading, MA.

Zhu, X. X., Tuia, D., Mou, L., Xia, G. S., Zhang, L., Xu, F., and Fraundorfer, F., 2017. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geosci. Remote Sens. Mag.*, 5(4), 8–36.

This contribution has been peer-reviewed.
https://doi.org/10.5194/isprs-archives-XLIX-B3-2026-475-2026