ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume X-5/W4-2025
Philippine Geomatics Symposium (PhilGEOS) 2025 "Enhancing Human Quality of Life through Geospatial Technologies", 24–25 November 2025, Quezon City, Philippines

Semi-Supervised Graph Transformer Neural Network for Illegal Dumping Detection in CALABARZON Using Remote Sensing and Multivariate Time-Series Data

Mark James J. Broqueza [1], Nathaniel V. Navoa [1], Bienvenido G. Carcellar III [1, 2]

[1] Department of Geodetic Engineering, University of the Philippines Diliman, Quezon City, Philippines - (mjbroqueza, nvnavoa, bgcarcellar)@up.edu.ph
[2] New Space Intelligence Inc. - 329-22 Nishikiwa, Ube, Yamaguchi 755-0151 Japan - bien-carcellar2022@newspaceint.com

Keywords: Waste Detection, Semi-supervised Learning, Spatiotemporal, SAR, Sentinel, Environmental Monitoring.

Abstract

Illegal waste dumping poses growing environmental risks as rapid urbanization overwhelms traditional monitoring, prompting the need for scalable, data-driven anomaly detection using remote sensing and machine learning. This study presents a semi-supervised graph transformer neural network for detecting illegal dumping in CALABARZON, a region increasingly burdened by unregulated waste due to rapid urbanization. The model leverages multivariate time-series data from Sentinel-1 SAR and Sentinel-2 MSI imagery (2020–2021) to capture complex spatiotemporal patterns in environmental indicators. A harmonized graph-based representation integrates spectral indices, radar features, and spatial relationships across time, enabling semi-supervised training under extreme class imbalance and limited ground-truth labels. The model exhibited strong generalization, with PR AUC stabilizing between 0.80 and 0.85, and consistent reductions in Huber and total loss values, indicating reliable modeling of typical site dynamics. However, performance on the minority “waste” class was constrained by data sparsity, achieving high accuracy (0.88) and precision (0.98) for the “non-waste” class but limited recall (0.50) and low precision (0.17) for “waste,” resulting in a high false positive rate. Sensitivity analysis highlighted shortwave infrared bands (B11, B12), vegetation-sensitive bands (B6, B9), indices such as MNDWI, NDWI, NDI45, NDTI, and RVI, and radar-derived features (e.g., VH backscatter variability, terrain angle) as critical for detection. These findings underscore the value of integrating optical and radar data for monitoring heterogeneous dumping behavior. While the model shows strong potential for early detection of illegal dumping, further refinement is needed to improve class precision and ensure actionable deployment in environmental monitoring systems.

1. Introduction

1.1 Background of the Study

Illegal dumping, the unauthorized disposal of waste in non-designated locations, remains a pervasive environmental threat that causes pollution, habitat degradation, and public health risks (United Nations Environment Programme & International Solid Waste Association, 2024). Globally, municipal solid waste generation is projected to increase sharply from 2.1 billion tonnes in 2023 to 3.8 billion tonnes by 2050, accompanied by an exponential rise in management costs (United Nations Environment Programme & International Solid Waste Association, 2024). In the CALABARZON region of the Philippines, rapid urbanization and population growth have overwhelmed existing waste management infrastructure, exacerbating illegal dumping incidences and their environmental consequences.

Conventional monitoring approaches depend heavily on manual field inspections and community reporting, which are labor-intensive, resource-constrained, and geographically limited. These methods often result in delayed detection, allowing illegal dumpsites to persist and worsen. Remote sensing technologies, including multispectral satellite imagery and Synthetic Aperture Radar (SAR), offer continuous, large-scale environmental monitoring, producing rich spatial and temporal datasets. However, analyzing such multivariate time-series data to detect illegal dumping remains challenging due to complex spatial-temporal dependencies and the scarcity of labeled data needed for fully supervised learning models.

This study addresses these challenges by proposing a semi-supervised graph transformer neural network that captures spatial relationships among geographical locations and temporal trends in environmental variables. Notably, few studies have applied advanced deep learning techniques within the Philippine context. By leveraging both labeled and unlabeled multivariate remote sensing data, this model aims to improve the accuracy and timeliness of illegal dumping detection in CALABARZON.

1.2 Research Objectives

This study aims to develop a semi-supervised graph transformer neural network for detecting illegal dumping in CALABARZON using multivariate time-series data from remote sensing. Specifically, it aims to (1) develop a graph-based anomaly model that captures spatial-temporal patterns in environmental indicators; (2) validate its detection performance against baseline methods; and (3) assess the feature importance using gradient-based sensitivity analysis.

1.3 Research Significance

This study directly contributes to several United Nations Sustainable Development Goals (SDGs), particularly SDG 11 (Sustainable Cities and Communities), and SDG 13 (Climate Action. Specifically, it addresses Target 11.6, which aims to reduce the adverse environmental impact of cities by improving municipal waste management, and Target 13.3, which seeks to enhance institutional and human capacity on climate-related monitoring and impact reduction. By developing such a model, the research enhances early detection and response mechanisms for environmental hazards, supporting sustainable urban development and improved resilience against waste-related pollution.

1.4 Related Literature

Effective detection and monitoring of illegal dumping require methodologies capable of addressing both spatial complexity and temporal variability inherent in waste disposal activities. Advances in remote sensing (RS) technologies have enhanced the capability to detect environmental anomalies associated with landfills and illegal dumpsites by providing improved spatial resolution and increased revisit frequency (Cicala et al., 2024; Fraternali et al., 2024). Sensors such as Landsat, Sentinel-2, and MODIS facilitate the identification of waste sites through thermal and spectral analyses, utilizing indices like the Normalized Difference Vegetation Index (NDVI) and thermal anomaly detection to distinguish waste-related signatures (Al-Harbi et al., 2024). The fusion of optical data with Synthetic Aperture Radar (SAR) enhances monitoring capabilities by overcoming challenges posed by cloud cover and varying illumination conditions and enables structural change detection through techniques such as Differential Interferometric SAR (DInSAR) (Agrawal & Khairnar, 2019; Mohammed et al., 2014). This multisensor integration proves critical for robust spatial detection and characterization of illegal dumping sites. While remote sensing effectively captures spatial features, the temporal dynamics of illegal dumping necessitate analytic approaches that model complex, interdependent environmental variables over time.

Multivariate time series (MVTS) anomaly detection provides this temporal dimension by leveraging correlations among multiple environmental indicators, thus improving detection sensitivity and specificity beyond univariate or static analyses (Bhumika & Das, 2022; Cheng et al., 2009). Recent developments in deep learning, including stacked autoencoders, Long Short-Term Memory (LSTM) networks, and Graph Neural Networks (GNNs), enable the extraction of latent spatial-temporal patterns from high-dimensional MVTS data, offering enhanced capabilities to identify anomalous events (Deng & Hooi, 2021; Kong et al., 2023). Integration of these advanced MVTS techniques with remote sensing data sources supports a comprehensive framework for early detection and real-time monitoring of illegal dumping activities (Shahab & Anjum, 2022; Sharma & Sood, 2024). Nonetheless, environmental factors such as cloud cover, seasonal variability, and sensor noise present ongoing challenges that require sophisticated preprocessing and data fusion strategies to maintain detection accuracy (Cheng et al., 2009). Collectively, these developments underscore the value of combining multisensor remote sensing with MVTS anomaly detection for effective surveillance in complex, rapidly changing environments.

2. Methodology

2.1 Study Area

CALABARZON, comprising the provinces of Cavite, Laguna, Batangas, Rizal, and Quezon, with over 16 million residents (PSA, 2020), is a highly urbanized and diverse region generating over 8,200 tons of solid waste daily, with a 14.33% increase in 2023 (EMB CALABARZON, 2024). Its varied land cover and limited monitoring capacity presents a complex environment for illegal dumping detection using remote sensing. Studying the entire region allows the model to learn from this environmental diversity and varied dumping behaviors, improving its ability to generalize across different contexts and enhancing detection performance at a broader, more actionable scale (Figure 1).

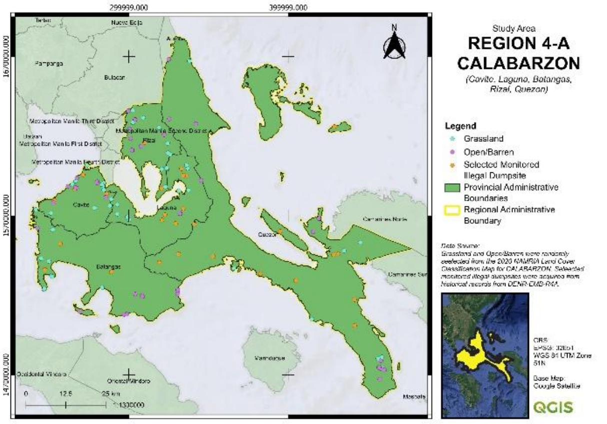

Figure 1. Study Area showing the Provinces in CALABARZON with the training points.

2.2 Data Acquisition

This study utilized multi-source geospatial datasets to train and evaluate a semi-supervised graph transformer model for illegal dumping detection. Sentinel-1 Ground Range Detected (GRD) Synthetic Aperture Radar (SAR) imagery was acquired in Interferometric Wide (IW) mode, descending orbit, with dual-polarization (VV and VH), a central frequency of 5.405 GHz (C-band), and a spatial resolution of 10 meters. The radar imagery enabled consistent detection of surface changes indicative of illegal dumping, such as mounds, pits, and other deformations, regardless of cloud cover or time of day.

Complementing SAR, Sentinel-2 Multispectral Instrument (MSI) Level 2 imagery was also acquired from GEE, covering the same time frame. These data offer 13 spectral bands, with four bands (Red, Green, Blue, and Near-Infrared) at 10-meter resolution, and others (e.g., SWIR, Red Edge) at 20 meters. The atmospherically corrected surface reflectance allowed for precise discrimination of land cover changes and vegetation health, which support anomaly detection associated with illegal waste accumulation.

Data from 2020 to 2021 were obtained from Google Earth Engine and divided into training (January–December 2020), validation (January–June 2021), and test (July–December 2021) sets. This timeframe was chosen due to the availability of key historical data, including digitized records from the Department of Environment and Natural Resources and consistent Sentinel-1B satellite imagery before its retirement in 2022. Notably, this period coincides with extended COVID-19 lockdowns in CALABARZON, which likely impacted illegal dumping behavior through altered human activity and reduced enforcement capacity. Hence, the data may reflect atypical dumping behaviors, which should be considered when interpreting the model's applicability to other periods.

To define non-dumping baseline areas, the 2020 Land Cover Map of CALABARZON from the National Mapping and Resource Information Authority (NAMRIA) was used to identify grasslands and barren or open lands, assumed to be naturally undisturbed. Random points within these land cover classes served as negative training examples for the model. Furthermore, OpenStreetMap (OSM) was utilized to refine training point selection by identifying areas tagged as waste areas to create buffers and strengthen the assumption with non-waste points.

2.3 Data Preprocessing

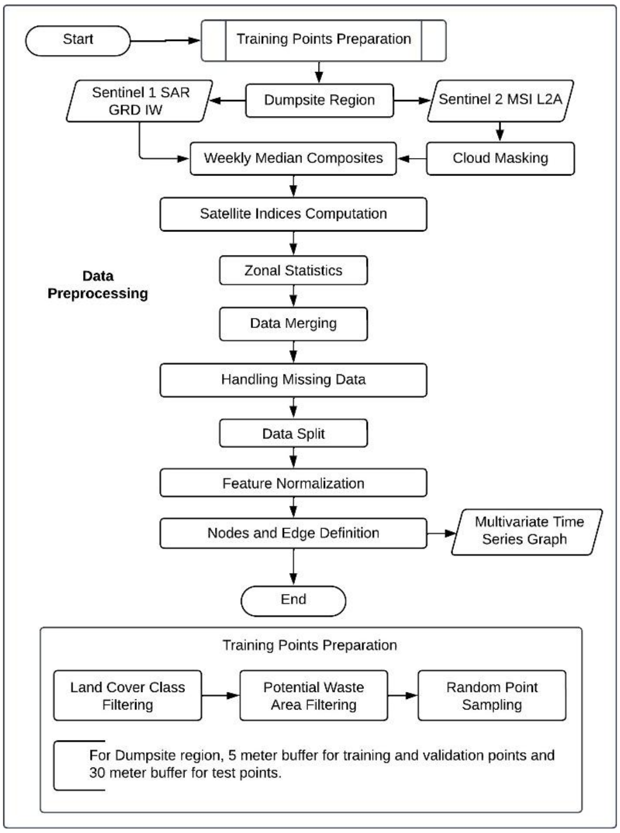

Figure 2. Data Preprocessing Methodology Flowchart.

2.3.1 Training Points Definition: To define training points for the semi-supervised anomaly detection model, two categories of locations were established: known waste and non-waste areas. A total of 28 waste points were manually identified and georeferenced based on previous site records, proximity to industrial and urban zones, and expert consultations. For non-waste locations, 100 points were generated using a stratified random sampling approach within areas classified as “Grassland” and “Open/Barren” in the 2020 Land Cover Map from NAMRIA. To minimize potential contamination from undocumented waste activity, exclusion buffers were applied around known dumpsites and solid waste facilities. OpenStreetMap (OSM) data were used to identify and exclude areas near anthropogenic features such as landfills, transfer stations, or material recovery facilities (MRFs). All spatial processing was conducted in UTM Zone 51N (EPSG:32651), and results were reprojected to geographic coordinates (EPSG:4326).

These points were cross validated through high-resolution satellite imagery on Google Earth Pro to confirm the presence of waste materials or dumping activity during the observation period. Although there were minimal images available, these were used as labels for semi-supervised learning. This balanced dataset of verified waste and non-waste points provided the foundation for model training, enabling the graph-based architecture to learn both anomalous and typical spatial-temporal patterns associated with illegal dumping.

2.3.2 Satellite Data Preprocessing: Preprocessing of Sentinel-1 SAR data utilized the COPERNICUS/S1_GRD ImageCollection accessed via Google Earth Engine, which provides analysis-ready, calibrated, and terrain-corrected products. To enhance the consistency of backscatter measurements across varying terrain, linear normalization was applied to the vertical–vertical (VV) and vertical–horizontal (VH) polarization backscatter coefficients using the local incidence angle (Equation 1).

$$\sigma_{normalized} = \sigma_{raw} + k(\theta - \theta_{ref}) \tag{1}$$

where σ_{raw} = backscatter coefficient (in dB), VV or VH
k = scaling factor (in dB per degree), 12°/dB
θ = local incidence angle (in degrees)
θ_{ref} = reference incidence angle, 35°

Additionally, the Radar Vegetation Index (RVI) was computed from the dual-polarized bands to characterize vegetation structure, aiding in distinguishing natural versus disturbed surfaces common in dumping sites (Equation 2).

$$RVI = \frac{4VV}{VV + VH} \tag{2}$$

where VV = normalized vertical–vertical backscatter
VH = normalized vertical–horizontal backscatter

Sentinel-2 Level-2A surface reflectance images, atmospherically corrected via the European Space Agency’s (ESA) Sen2Cor processor and accessed through Google Earth Engine, underwent cloud masking using the Scene Classification Layer (SCL) to exclude clouds and shadows. Multiple spectral indices were calculated to capture various land surface characteristics relevant to waste dumping, including the Normalized Difference Vegetation Index (NDVI) and Modified Soil-Adjusted Vegetation Index (MSAVI) for vegetation health; the Normalized Difference Built-up Index (NDBI) for urban areas; the Normalized Difference Water Index (NDWI) and Modified Normalized Difference Water Index (MNDWI) for water bodies and moisture content; the Normalized Difference Bareness Index (NDBaI) for exposed soil; the Normalized Difference Tillage Index (NDTI) for soil disturbance; and the Normalized Difference Index using red-edge and red bands (NDI45) for detecting subtle vegetation stress and soil condition variations (Equation 3-10).

$$NDVI = \frac{NIR - Red}{NIR + Red} \tag{3}$$

$$MSAVI = \frac{2NIR + 1 - \sqrt{(2NIR + 1)^2 - 8(NIR - Red)}}{2} \tag{4}$$

$$NDBI = \frac{SWIR - NIR}{SWIR + NIR} \tag{5}$$

$$NDWI = \frac{Green - NIR}{Green + NIR} \tag{6}$$

$$MNDWI = \frac{Green - SWIR}{Green + SWIR} \tag{7}$$

$$NDBaI = \frac{SWIR - Red}{SWIR + Red} \tag{8}$$

https://doi.org/10.5194/isprs-annals-X-5-W4-2025-105-2026 | © Author(s) 2026. CC BY 4.0 License.

$$NDBTI = \frac{Red - Green}{Red + Green} \tag{9}$$

$$NDI45 = \frac{Red\ Edge\ 1 - Red}{Red\ Edge\ 1 + Red} \tag{10}$$

where *NIR* = near-infrared band reflectance, Band 8
Red = red band reflectance, Band 4
Green = green band reflectance, Band 3
SWIR = shortwave infrared band reflectance, Band 11
Red Edge 1 = red-edge band 1 reflectance, Band 5

For both datasets, zonal statistics (mean, minimum, maximum, and standard deviation) of these indices were computed within 5-meter buffers around dumpsite locations to provide detailed spatial characterization of surface conditions. These data were integral to modeling anomaly detection by highlighting subtle changes.

2.3.3 Graph-based Input Preparation: To integrate Sentinel-1 SAR and Sentinel-2 MSI datasets for anomaly detection, a harmonized spatiotemporal structure was developed. Weekly compositing was employed using week definitions from ISO 8601 (starting every Monday), enabling standardized alignment across sensors regardless of revisit inconsistencies. All available images within each week were merged, improving temporal density and minimizing the impact of cloud contamination and SAR acquisition gaps. To address missing data, a temporal nearest-neighbor imputation was applied. The algorithm identified the closest valid observation in time, preceding or succeeding, and filled missing values, accordingly, preserving temporal coherence without introducing artificial fluctuations.

The resulting weekly time series was used to construct a spatiotemporal graph structure. To build the spatiotemporal graph, geospatial records were filtered by date and split into training and validation sets, with high class imbalance addressed by down sampling non-waste observations to match the waste class ratio of 3:1. Each timestamped dataset was cleaned and standardized using a global StandardScaler fitted on the training set, ensuring consistent feature scaling. These normalized features were assigned to nodes representing spatial sites, with centroids used for positioning. A BallTree-based nearest neighbor search, with 500-meter radius, generated bidirectional spatial edges, defining the static spatial structure reused across all time steps.

Each weekly graph instance comprised nodes with statistical attributes, geographic coordinates, and site identifiers. Labels were stored in tensors with masks to support partially observed data. Temporal features such as week of year and month were added to enrich the model's temporal context. The final dataset was prepared as fixed-length temporal windows using a custom PyTorch Dataset, packaging node features, edge indices, time metadata, and labels. A 4-week sliding window (stride = 1 week) was applied to build temporal sequences for each node, improving the model's capacity to detect evolving anomalies. This structure enabled supervised spatiotemporal learning while preserving spatial and temporal continuity.

2.4 Model Architecture

The proposed model captures spatiotemporal patterns in time-series data collected across multiple geographic sites (Figure 3). Each data point is enriched with three types of contextual information: (1) a SiteEmbedding, which encodes the identity or location of each site; (2) a NodeTimeEmbedding, which captures the position of each observation within the time series; and (3) a CalendarEmbedding, which introduces seasonal or periodic context using features like the week of the year and month. These embeddings provide richer input representations that help the model learn location-specific and time-sensitive patterns.

Spatial relationships are modeled through the SpatialBlock, which contains two parallel components. The StaticSpatialLayer applies graph attention over a fixed network of site connections (e.g., based on physical proximity or known administrative boundaries). In contrast, the DynamicSpatialLayer uses attention mechanisms to learn spatial interactions directly from the data at each time step, without relying on predefined connections. These two views are fused using a learnable GatingFusion mechanism, which adaptively weights their contributions for each site and time step.

The TemporalTransformerBlock models the evolution of each site's data over time. It treats each site's sequence of observations as a timeline and applies self-attention to capture both short- and long-term dependencies. Following this, a temporal attention pooling layer highlights the most relevant time steps by assigning attention weights, allowing the model to focus on periods most indicative of anomalous behavior. The processed temporal and spatial features are then combined in an InteractionBlock, which helps the model understand how spatial and temporal signals interact. Finally, a node-level AttentionBlock identifies which locations are most relevant for the prediction task.

In the final stage, the HybridPredictionBlock produces two outputs simultaneously: (1) probabilistic reconstruction of the original input data by predicting Gaussian distribution parameters (mean and variance), and (2) anomaly scores through classification logits. This dual-task setup allows the model to both reconstruct normal patterns and detect unusual activity, such as illegal dumping, within an imbalanced dataset.

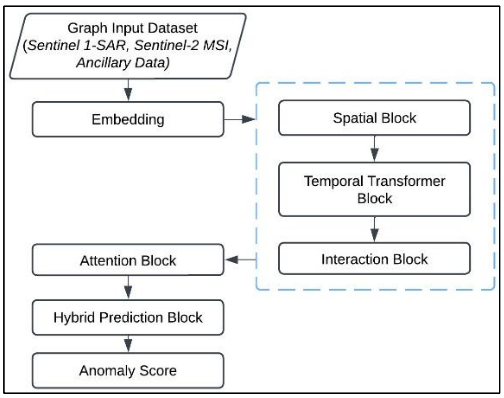

Figure 3. GTNN Model Architecture

2.5 Model Training

The model training process involves optimizing the model to simultaneously perform spatiotemporal reconstruction and anomaly classification on the graph-structured time series data. During each epoch, the model is trained in batches where input sequences, edge indices, site embeddings, and temporal features are fed into the model. The model outputs include reconstructed values of the input features, their associated log variances and logits indicating anomaly likelihoods for each node and timestep. The reconstruction loss is calculated using Huber loss,

$L(y, \hat{y})$, between the predicted and true values for the last time step (Equation 11).

$$L(y,\hat{y}) = \begin{cases} \frac{1}{2}(y-\hat{y})^2, & if\ |y-\hat{y}| \leq \delta \\ \delta \cdot (|y-\hat{y}| - \frac{\delta}{2}), & otherwise \end{cases} \tag{11}$$

where y = true value at the last time step
$\hat{y}$ = predicted value at the last time step
δ = transition threshold for quadratic and linear loss

The classification loss, $FL(p, y)$, for labelled data uses a focal loss combined with a margin loss to handle class imbalance and enforce a margin between positive and negative logits (Equation 12).

$$FL(p,y) = \begin{cases} -\alpha(1-p)^{\gamma} log(p), & y = 1 \\ -\alpha(1-p)^{\gamma} log(1-p), & y = 0 \end{cases} \tag{12}$$

where p = predicted probability of the positive class
y = true class label
α = weighting factor to address class imbalance
γ = parameter reducing loss for well-classified cases

These losses are weighted and combined to form the total loss, which is backpropagated to update model parameters via Adam optimization with learning rate scheduling. To ensure robust evaluation and avoid overfitting, training is accompanied by validation on a separate dataset after each epoch. Early stopping is implemented based on validation loss to prevent overtraining, and the best-performing model weights are checkpointed.

2.6 Model Validation

To evaluate the model's ability to detect anomalies, particularly in highly imbalanced datasets where illegal dumping events are rare, several performance metrics are used. These metrics include per-class Precision, Recall, and F1-score, along with overall Accuracy, Macro Average (unweighted mean across all classes), and Weighted Average (mean weighted by class frequency). Given the nature of anomaly detection, particular emphasis is placed on the Precision-Recall Area Under the Curve (PR AUC), which is more informative than standard accuracy in reflecting the model's performance on the minority (anomalous) class. PR AUC evaluates the trade-off between precision (minimizing false positives) and recall (maximizing true positive detections), making it especially suited for imbalanced classification tasks.

2.7 Anomaly Detection

To detect anomalies in illegal dumping activity, the model combined supervised classification and unsupervised reconstruction to compute a hybrid anomaly score for each site and timestamp. The hybrid anomaly score (Equation 13) was defined as the element-wise product of the anomaly probability and the reconstruction error using Huber loss.

$$anomaly\ score = \sigma(logits) \cdot recon.error \tag{13}$$

where $\sigma(logits)$ = anomaly probability (sigmoid output)
k = scaling factor (in dB per degree)

This fusion leveraged both the likelihood of anomalous behavior and deviation from expected feature patterns. Flagging of anomaly was performed by applying a statistical thresholding approach to the model's hybrid anomaly scores. For each site, an expanding mean and standard deviation of the anomaly scores were computed over time to dynamically model typical behavior. Observations falling outside ±2 standard deviations from the expanding mean were flagged as anomalies, capturing significant deviations indicative of unusual dumping activity. This method enabled site-specific anomaly identification while accounting for temporal trends and variability in the data.

2.8 Feature Importance

Feature importance quantifies the influence of each input feature on a model's anomaly detection output, offering interpretability and insight into which variables most significantly contribute to identifying events such as illegal dumping. A gradient-based sensitivity analysis is employed, wherein the gradient of the model's output with respect to each input feature is computed after a forward pass and summed prediction. The absolute gradients are then averaged across batches, spatial nodes, and time steps to yield a normalized importance score for each feature dimension. Features with higher gradient magnitudes are considered more influential. Statistical thresholds based on mean and standard deviation are applied to highlight significantly important features. This method, requiring only a single backward pass, is computationally efficient for large-scale spatiotemporal datasets and allows for clear visualization through bar plots, enabling intuitive interpretation of the most impactful features.

3. Results and Discussion

3.1 Data Diagnostic

The diagnostics in Table 1 and 2 highlight significant data completeness between Sentinel-1 and Sentinel-2 datasets. Sentinel-1 maintained relatively stable coverage, with approximately 35–45% of weeks having at least one node with missing data and short streaks of only 1–2 weeks (Table 1). This indicates that while gaps existed, they were infrequent and dispersed, allowing for relatively reliable temporal analysis.

Diagnostic	Train	Val.	Test
Total No. of Weeks	51	23	25
% of Weeks with Missing Data	35.43	45.28	36.91
Mean Missing Weeks per Node	18.07	10.41	9.23
Std. Missing Weeks per Node	3.02	1.78	1.60
Max. Consecutive Missing Weeks	2	2	2

Table 1. Summary of Weekly Missing Data Diagnostics for Sentinel-1 by Dataset Split

In contrast, Sentinel-2 suffered from much higher data loss, with up to 59% of the weeks affected and prolonged streaks of missing data lasting up to 19 consecutive weeks (Table 2). This points to structural challenges in optical data acquisition, such as cloud cover or acquisition frequency, which are common limitations in tropical monitoring contexts.

Diagnostic	Train	Val.	Test
Total No. of Weeks	51	26	26
% of Weeks with Missing Data	49.71	47.42	59.04
Mean Missing Weeks per Node	25.35	12.33	15.35
Std. Missing Weeks per Node	5.12	3.25	3.11
Max. Consecutive Missing Weeks	17	19	19

Table 2. Summary of Weekly Missing Data Diagnostics for Sentinel-2 by Dataset Split

These disparities have implications for model training and interpretation. While nearest-neighbor imputation helped retain continuity, it relied on the assumption of temporal similarity, which becomes problematic during dynamic environmental shifts like seasonal vegetation changes or illegal dumping events. The extended missing periods in Sentinel-2 may have introduced synthetic patterns that obscure true signals, reducing the model's ability to learn subtle trends or detect short-term anomalies. Consequently, the model's performance may be more robust when trained with Sentinel-1, while outputs based on Sentinel-2 should be interpreted with caution, especially in periods reconstructed from lengthy imputations.

Label	Train	Validation	Test
Non-Waste	102	135	47
Waste	34	7	2
Unlabeled	6392	3186	3279

Table 3. Label Distribution for Dataset Splits.

Table 3 reveals a severe label imbalance and extremely sparse positive samples across all dataset splits. In the training set, only 34 of 6528 samples (≈0.52%) are labeled as "waste," while 102 are labeled "non-waste," and the vast majority (over 97%) are unlabeled. Similar trends appear in the validation (7 waste vs. 135 "non-waste" out of 3328 total) and test sets (2 waste vs. 47 "non-waste" out of 3328). This imbalance severely limits the supervised signal available for learning illegal dumping behavior, particularly the "waste" class, which is the primary interest. Such imbalance not only complicates the training of reliable classifiers but also raises concerns about the model's generalization ability and its sensitivity to rare, high-impact events like waste dumping.

While the adoption of a semi-supervised framework, leveraging reconstruction, focal, and margin-based losses, helps mitigate these issues, the extremely low number of positive examples underscores the need for more label-efficient strategies, such as pseudo-labeling, data augmentation, or active learning to improve anomaly detection performance.

3.2 Model Performance Evaluation

The PR AUC curve in Figure 4 shows steady improvement in training performance and a sharp rise in validation AUC after the 10th epoch, stabilizing around 0.80–0.85. This convergence suggests good generalization and effective detection of rare waste anomalies. Figure 5 illustrates a consistent decline in Huber loss for both training and validation, indicating improved reconstruction of normal patterns. The slightly lower validation loss may reflect noise in the training set or the Huber loss's robustness, supporting the model's sensitivity to illegal dumping behavior.

The combined loss curve in Figure 6, reflecting both classification and reconstruction objectives, shows a steady decline for training and validation sets, indicating stable optimization and minimal overfitting. The consistently lower validation loss suggests effective regularization and possibly cleaner data. This pattern highlights the model's strong learning dynamics and reliability for real-world deployment in spatiotemporal illegal dumping detection.

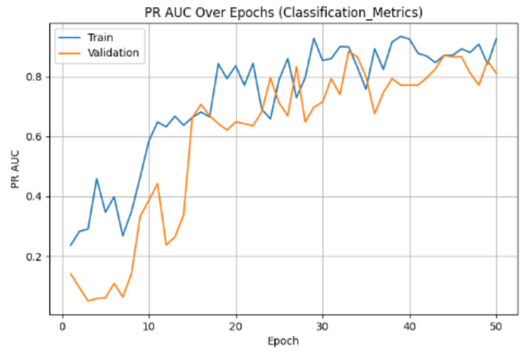

Figure 4. PR AUC for Training and Validation.

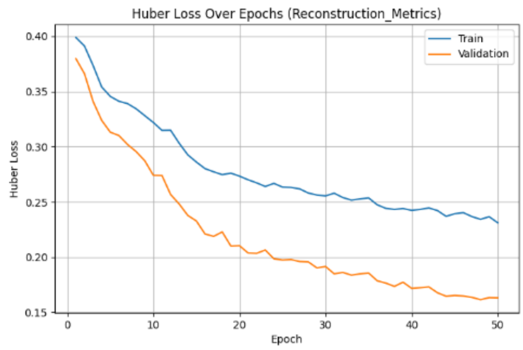

Figure 5. Huber Loss for Training and Validation

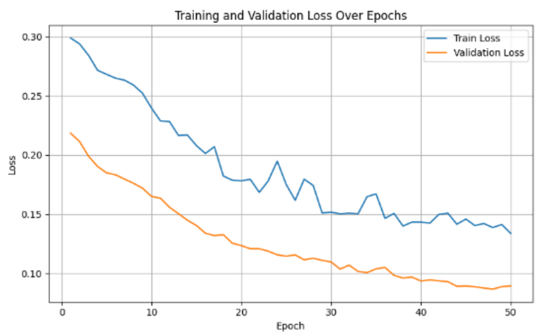

Figure 6. Combined Loss for Training and Validation

3.3 Model Accuracy

As seen in Table 4, the classification results reveal a strong model performance on identifying the non-waste class, with high precision (0.98) and recall (0.89), resulting in a robust F1-score (0.93). This indicates that the model is highly confident and accurate in detecting areas without waste. However, the primary objective of the system is to detect illegal dumping activities, represented by the waste class, which is where the model struggles. With only two positive (waste) samples in the test set, the model managed to correctly identify one, resulting in a recall of 0.50, but its precision was only 0.17, meaning most of the waste predictions were false alarms.

Class	Precision	Recall	F1-Score
Non-Waste	0.98	0.89	0.93
Waste	0.17	0.50	0.25
Accuracy	---	---	0.88
Macro Avg.	0.57	0.70	0.59
Weighted Avg.	0.94	0.88	0.91

Table 4. Classification Performance Metrics

The macro-average metrics (precision = 0.57, recall = 0.70, F1 = 0.59) expose this imbalance clearly. While the overall accuracy is 88%, this number is misleading in an imbalanced context where 47 of 49 samples are “non-waste.” A high accuracy in such settings can be achieved simply by predicting most instances as non-waste. The model appears to lean heavily on the dominant class, likely because it received significantly more examples of “non-waste” during training. As a result, it lacks confidence when generalizing to rare but important waste events. In the context of illegal waste dumping detection, these findings highlight a critical limitation. The ability to detect even a single waste-related anomaly is promising, showing that the model can recognize anomalous patterns. but the low precision suggests high false positives, which can lead to inefficient inspection efforts. A 50% recall means half of the actual waste sites might still go undetected, which undermines the model’s utility in proactive environmental monitoring and enforcement.

3.4 Anomaly Detection

To illustrate the anomaly flagging process, Figure 7 presents the hybrid anomaly scores for the dumpsite in 1st District, Jalajala, Rizal. During the test period, anomalies were flagged on October 11, November 15, and November 22, 2021, each occurring at the peak of sudden spikes in the anomaly score curve. These sharp increases suggest abrupt deviations from typical site conditions, potentially indicating illegal dumping events or related disturbances. Prior to these flagged dates, a gradual upward trend in the scores can be observed, which may reflect increasing site activity or the accumulation of waste over time.

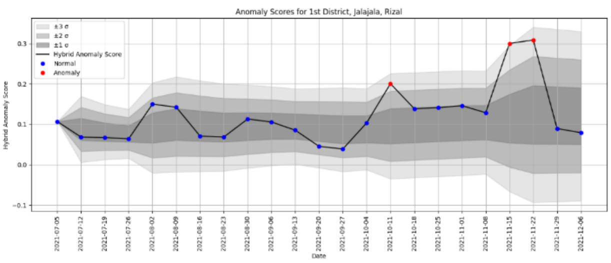

Figure 7. Anomaly Score Plot for 1st District, Jalajala (Rizal)

Conversely, periods with flat or steadily low scores indicate stable conditions, likely representing either inactive dumping behavior or natural site variation. Notable dips following a flagged anomaly may suggest a short-term resolution of the disturbance, possibly due to cleanup efforts, reduced dumping activity, or changes in environmental conditions affecting remote sensing observations. Overall, these patterns highlight how temporal shifts in anomaly scores can reflect the dynamic and episodic nature of waste dumping behavior, offering a valuable basis for early warning and site monitoring.

The observed anomaly patterns provide actionable insights for policymakers and local authorities by identifying the timing and locations of illegal dumping events. This enables prompt and targeted interventions, such as scheduling inspections and organizing cleanup efforts immediately after spikes in dumping activity are detected. Tracking the return to stable conditions also allows for evaluation of enforcement effectiveness and community initiatives. Integrating this dynamic monitoring into waste management strategies can improve resource allocation, support the implementation of stricter zoning and penalties in high-risk areas, and facilitate adaptive policy management to enhance environmental compliance.

However, without actual field validation, the model’s outputs remain hypothetical, demonstrating potential patterns and behaviors but lacking confirmation of their real-world accuracy. While the anomaly scores may align with plausible dumping events or environmental changes, there is no definitive evidence that these detections correspond to true instances of illegal dumping. This limitation restricts the ability to assess the model’s practical reliability and precision.

Upon inspecting the few false positives in the test dataset (Figure 8), historical satellite imagery showed no visible signs of illegal dumping, suggesting that these anomalies were likely triggered by sudden, short-term changes, such as land clearing, vegetation loss, or surface disturbance, that temporarily deviated from normal patterns but did not reflect actual waste activity. In some cases, changes in surface reflectance or texture may have mimicked dumping behavior, particularly in exposed or recently altered areas. The relatively low anomaly logits (0.009–0.37) support this interpretation, reflecting low model confidence and highlighting its sensitivity to minor, non-waste fluctuations. These findings point to the need for refining thresholding strategies and incorporating field validation or contextual layers to improve the model’s precision and interpretability.


Figure 8. Sample areas flagged as False Positive © Google Earth.

3.5 Feature Importance

The gradient-based feature importance analysis (Figure 9) reveals that specific spectral bands and remote sensing indices are key to detecting illegal waste dumping. Notably, shortwave infrared bands (B11, B12) and vegetation-sensitive bands (B6, B9) ranked high in importance due to their ability to capture changes in moisture, biomass, and soil, often altered by dumping. Water-related indices like MNDWI and NDWI also emerged as significant, indicating shifts in surface moisture potentially caused by vegetation loss or the presence of impervious waste materials. Vegetation and structural indices such as NDI45, NDTI, and RVI were among the top features, reflecting disturbances in plant cover and land surface texture. These changes are consistent with the environmental impact of dumping, where waste often replaces vegetation and leads to patchy or degraded land. The high ranking of the standard deviation of RVI suggests repeated disturbances or fluctuating vegetation, typical in active or expanding dumpsites.

Radar-derived features, especially the variability in corrected VH polarization backscatter and terrain angle statistics, further contributed to anomaly detection. These radar metrics capture surface roughness and structural heterogeneity—key traits of dumping areas. Together, these findings emphasize the value of combining optical and radar data to detect complex, irregular patterns associated with illegal dumping more effectively.

https://doi.org/10.5194/isprs-annals-X-5-W4-2025-105-2026 | © Author(s) 2026. CC BY 4.0 License.

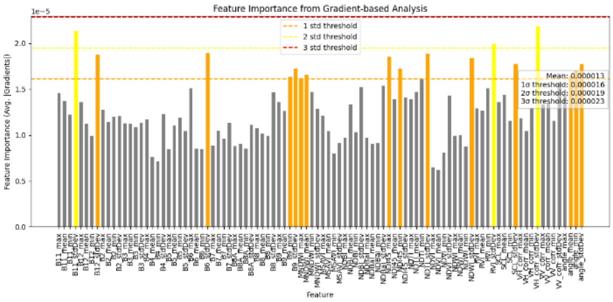

Figure 9. Feature Importance from Gradient-based Sensitivity Analysis

4. Conclusion and Recommendations

This study developed and validated a semi-supervised graph transformer neural network for detecting illegal dumping in the CALABARZON region using multivariate time-series remote sensing data. The model demonstrated strong learning and generalization capabilities, shown by the convergence of training and validation PR AUC curves around 0.80–0.85 and the consistent decline in Huber and total loss. These patterns highlight the model's capacity to model typical site behavior and detect deviations associated with waste dumping, even under severe class imbalance.

Nevertheless, performance on the minority "waste" class was limited by the scarcity of positive samples (34 out of 6,528), leading to high recall but low precision and a high rate of false positives. This trade-off could reduce the efficiency of inspection efforts and impact the system's operational utility. Sentinel-2 data gaps also hindered long-term temporal reliability, emphasizing the need for more robust imputation strategies. Still, sensitivity analyses revealed the relevance of shortwave infrared bands (B11, B12), vegetation-sensitive bands (B6, B9), and indices such as MNDWI, NDWI, and RVI, along with radar-derived texture and backscatter features. These results affirm the value of multi-temporal, multi-sensor fusion in detecting complex, irregular surface changes linked to illegal dumping. Overall, the model contributes toward environmental hazard monitoring aligned with SDG 11 and SDG 13.

Future work should address data imbalance with label-efficient approaches like pseudo-labeling, active learning, or augmentation. Improving anomaly precision may require refining loss functions, adding contextual features, or post-processing. Low precision for the waste class could be mitigated with advanced time-series data augmentation (e.g., dynamic time warping or sequence permutation) and adaptive thresholding such as the Peaks Over Threshold (POT) approach, which could reduce false positives and better capture temporal variability. Ablation studies are recommended to evaluate each architectural component's contribution, enhancing interpretability and justifying model complexity.

Although this study relies on annotated datasets for model training and validation, integration with field verification remains a key opportunity. Future collaboration with DENR or LGUs could involve joint field assessments to confirm model predictions, contribute additional ground-truth data, and support iterative model refinement for operational deployment. Finally, exploring the model's transferability to other regions will test its adaptability and potential as a scalable solution for nationwide illegal dumping surveillance.

References

Agrawal, S., & Khairnar, G. B., 2019. A Comparative Assessment of Remote Sensing Imaging Techniques: Optical, SAR and LIDAR. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-5/W3, 1–6. doi.org/10.5194/isprs-archives-XLII-5-W3-1-2019.

Al-Harbi, M., Tawabini, B., Al-Farhan, M., Soupios, P., Farooqui, M., & Basaleh, A., 2024. Mapping of uncontrolled dumpsites in arid regions through remote sensing and image processing. *Waste Management & Research: The Journal for a Sustainable Circular Economy*, *43*(4), 616-624. doi.org/10.1177/0734242X241257098.

Bhumika, Das, D., 2022. UApredictor: Urban Anomaly Prediction from Spatial-Temporal Data using Graph Transformer Neural Network. *2022 International Joint Conference on Neural Networks (IJCNN)*, Padua, Italy, 1–8. doi.org/10.1109/IJCNN55064.2022.9892885.

Cheng, H., Tan, P.N., Potter, C., Klooster, S., 2009. Detection and Characterization of Anomalies in Multivariate Time Series. *Proceedings of the 2009 SIAM International Conference on Data Mining*, 413–424. doi.org/10.1137/1.9781611972795.36.

Cicala, L., Gargiulo, F., Parrilli, S., Amitrano, D., Pigliasco, G., 2024. Progressive Monitoring of Micro-Dumps Using Remote Sensing: An Applicative Framework for Illegal Waste Management. *Sustainability*, *16*(13), 5695. doi.org/10.3390/su16135695.

Deng, A., Hooi, B., 2021. Graph Neural Network-Based Anomaly Detection in Multivariate Time Series. *Proceedings of the AAAI Conference on Artificial Intelligence*, *35*(5), 4027–4035. doi.org/10.1609/aaai.v35i5.16523.

Fraternali, P., Morandini, L., Herrera González, S.L., 2024. Solid waste detection, monitoring and mapping in remote sensing images: A survey. *Waste Management*, *189*, 88–102. doi.org/10.1016/j.wasman.2024.08.003.

Kong, L., Yu, J., Tang, D., Song, Y., & Han, D., 2023. Multivariate Time Series Anomaly Detection With Generative Adversarial Networks Based on Active Distortion Transformer. *IEEE Sensors Journal*, *23*(9), 9658–9668. doi.org/10.1109/JSEN.2023.3260563.

Mohammed, O.I., Saeidi, V., Pradhan, B., Yusuf, YA., 2014. Advanced differential interferometry synthetic aperture radar techniques for deformation monitoring: a review on sensors and recent research development. *Geocarto International*, *29*(5), 536–553. doi.org/10.1080/10106049.2013.807305

Shahab, S., Anjum, M., 2022. Solid Waste Management Scenario in India and Illegal Dump Detection Using Deep Learning: An AI Approach towards the Sustainable Waste Management. *Sustainability*, *14*(23), 15896. doi.org/10.3390/su142315896.

Sharma, K., Sood, M., 2024. Monitoring, classification and analysis of waste disposal sites using Machine Learning. *Procedia Computer Science*, *235*, 1558–1567. doi.org/10.1016/j.procs.2024.04.147.