

Illegal Mining Activity Detection Using satellite Images

Shashidhara S and Thouseef Ulla Khan

Department of MCA

Vidya Vikas Institute of Engineering and Technology, Mysuru, India

shashis41297@gmail.com and thouseef.khan@vidyavikas.edu.in

Abstract: *Illegal mining poses a significant threat to ecological balance, biodiversity, and lawful economic systems. Across the globe, unsanctioned mining operations lead to air pollution, water pollution, deforestation, and loss of natural habitats, causing irreversible damage to the environment. Traditional methods for monitoring mining activities—such as physical patrolling or bureaucratic inspections—are often inefficient, slow, and unable to deliver info in real time.*

In order to overcome these constraints, this proposal suggests an AI-powered system capable of detecting illegal mining activities Using satellite photos with deep learning and computer vision techniques. The system integrates a custom Convolutional Neural Network Cnn and transfer learning via DenseNet121 for image classification, distinguishing between forested and illegally mined regions with very high accuracy. To enhance functionality, YOLO (You Only Look Once) is incorporated as a real-time object detection algorithm to identify and localize rivers and mining zones in satellite images. This two-pronged strategy guarantees both classification and spatial localization, enabling authorities to take prompt and informed action. A robust preprocessing pipeline—comprising RGB channel visualization, histogram equalization, Gaussian blurring, Sobel filtering, segmentation, and Canny edge detection—was implemented to enhance image clarity and highlight distinguishing features. The dataset consists of high-resolution satellite images annotated into two categories: forest and illegal mining zones.

The system is deployed using a Flask-based web application, allowing users to upload satellite images and receive real-time predictions with highlighted illegal zones. Designed for environmental agencies, researchers, and government bodies, this solution demonstrates the potential of AI and remote sensing as powerful tools for sustainable environmental monitoring and the mitigation of illegal mining activities..

Keywords: Illegal mining detection, satellite imagery, deep learning, convolutional neural networks cnn, DenseNet121, YOLOv8, object detection, image preprocessing, computer vision, environmental monitoring, Flask web application.

I. INTRODUCTION

Illegal mining has emerged as among the most pressing challenges in environmental conservation and sustainable development. Defined as the extraction of minerals without legal authorization or adherence to environmental standards, it is a worldwide problem that impacts both wealthy and developing nations. Such operations frequently take place in remote and forested areas where monitoring is difficult, enabling destructive practices to continue unchecked. The consequences are severe, including large-scale loss of biodiversity, water pollution, deforestation, and significant socio-economic disruption in local communities.

The ecological damage caused by illegal mining is particularly alarming. Forest ecosystems are destroyed, rivers become contaminated with toxic byproducts such as mercury and cyanide, and air quality deteriorates due to particulate emissions. Beyond environmental degradation, the socio-economic impacts are profound: governments lose revenue from untaxed extraction, local populations face health hazards, and legitimate mining industries suffer from unfair competition. These multifaceted effects highlight the urgency of creating clever systems to keep an eye on and preventing such unlawful activities.



Conventional approaches to monitoring mining—such as field inspections, drone surveillance, and GIS-based mapping—are limited in effectiveness. Manual inspections are resource-intensive, prone to delays, and often cover only small geographical areas. Drone monitoring, while useful, cannot be scaled cost-effectively to vast territories. GIS-based mapping provides insights into land-use changes, but it still relies heavily on manual interpretation and lacks real-time capabilities. As a result, illegal mining operations frequently go hidden until serious harm has already occurred.

Recent advances in satellite imaging technology have created chances for large-scale, real-time environmental monitoring. High-resolution satellite imagery offers wide spatial coverage and temporal snapshots, enabling the detection of subtle changes in land cover over time. These datasets, however, are massive and complex, making manual inspection impractical. This challenge calls for the integration of artificial intelligence (AI) and deep learning to automate analysis and derive actionable insights from satellite data.

Deep learning techniques, particularly Convolutional Neural Networks (CNNs), have revolutionized the field of computer vision. Their ability to extract hierarchical spatial features makes them well-suited for tasks such as image classification and object detection. Transfer learning models like DenseNet121 further improve accuracy by leveraging pre-trained feature representations, while advanced object detection algorithms such as YOLOv8 enable real-time localization of regions of interest. These developments provide a strong foundation for designing systems capable of detecting illegal mining activities with unprecedented efficiency and reliability.

This project builds on these advancements by developing an AI-based system for the detection and localization of illegal mining operations from satellite images. A custom CNN and DenseNet121 are employed to classify images into forested and mining categories, while YOLOv8 detects and highlights mining zones and rivers. To ensure robustness, images undergo preprocessing techniques such as histogram equalization, Gaussian blurring, Sobel filtering, and Canny edge detection. Together, these methods enhance accuracy and provide spatial awareness essential for practical intervention.

The developed system is deployed through a Flask-based web application, offering a user-friendly interface where users can upload satellite images and instantly receive predictions with bounding box overlays. This approach ensures that the technology is accessible not only to data scientists but also to environmental authorities, policymakers, and non-technical users. By combining classification with localization, the system empowers stakeholders to take timely and informed actions against illegal mining.

In summary, the proposed solution represents a step forward in the fight against environmental crimes. It demonstrates how AI, remote sensing, and computer vision can converge to address real-world problems of sustainability and governance. By offering a scalable, accurate, and real-time monitoring tool, this work contributes significantly toward safeguarding ecosystems, preserving biodiversity, and supporting regulatory enforcement in regions vulnerable to illegal mining.

II. LITERATURE SURVEY

[1] Hu et al. (2010) proposed one of the earliest frameworks for illegal mining detection through their work Designing an DInSAR-based system for detecting illegal mining. In order to identify ground deformation linked to underground mining operations, the study used Differential Synthetic Aperture Radar Interferometry (DInSAR). This approach worked especially well in areas where optical imaging was limited by vegetation cover or cloud interference. Their findings highlighted that subsidence could be a reliable indicator of hidden mining activity, offering a novel way of monitoring inaccessible terrains. However, the technique required expert interpretation of SAR data and suffered from limited real-time capabilities, making it more suitable as a foundational rather than operational tool.

[2] Balaniuk, Isupova, and Reece (2020) explored modern deep learning techniques in their work Mining and Tailings Dam Detection in Satellite Imagery Using Deep Learning. They designed convolutional neural network (CNN) models trained on high-resolution multispectral imagery to classify mining regions and identify tailings dams. Their study demonstrated that CNNs excel at extracting hierarchical spatial features, enabling accurate differentiation of mining areas from natural landscapes. By stressing the importance of large, well-labeled datasets and multispectral inputs, this work provided a scalable AI-based approach for environmental surveillance and showcased the superiority of deep learning over traditional classification techniques.



[3] Ascencio Cansado et al. (2023) provided a socio- environmental perspective in their study A New Gold Rush: Brazil Fails to Control Illegal Mining and Watches Mechanized Mining Thrive in Amazonia. The authors examined how weak governance and limited enforcement allowed mechanized mining to expand uncontrollably in the Amazon. They emphasized the role of satellite monitoring platforms like MapBiomass and Google Earth Engine in exposing illegal encroachments, highlighting how accessible geospatial intelligence could inform the public and policymakers. Their work underscored the limitations of policy frameworks without technological integration, strengthening the case for automated AI-powered monitoring systems.

[4] Rahman and Hassan (2013) addressed land-use change detection in their study Change Detection and Estimation of Illegal Mining Using Satellite Images. They applied the (NDVI) Normalized Difference Vegetation Index and supervised classification to Landsat imagery to identify vegetation loss caused by mining expansion. Their temporal analysis demonstrated how deforestation patterns serve as indicators of illicit activity. Although constrained by moderate- resolution imagery, their results validated vegetation-based indices as cost-effective tools for detecting changes linked to illegal mining, making NDVI an important baseline for remote sensing applications.

[5] Knyaz and Zheltov (2021) introduced an automated system in their paper Automated Detection of Illegal Nonmetallic Minerals Mining Places According to Sentinel-2 Satellite Images. They utilized spectral signature analysis and morphological filtering to identify disturbed landscapes caused by mining. Their method, while computationally less demanding than deep learning, achieved acceptable accuracy and demonstrated the usability of freely available Sentinel-2 data. This approach provided a low-cost and scalable solution for large-area surveillance, particularly in resource-constrained regions.

[6] Rodriguez, Gonzalez, and Lopez (2022) investigated vegetation indices in their study Satellite Imagery Processing Using NDVI for the Detection of Illegal Mining Activities. Their workflow involved NDVI computation, thresholding, and supervised classification to differentiate between healthy forest cover and areas disturbed by mining. The results validated NDVI as a powerful and comprehensible instrument for early detection of mining-induced degradation. Their work highlighted the accessibility of vegetation-based techniques, making them appropriate for integration into governmental GIS platforms for low-cost monitoring.

[7] Torres et al. (2023) presented a radar-centric approach in their study Radar Based Monitoring System to Protect the Colombian Amazon Rainforest. Their model integrated Synthetic Aperture Radar (SAR) with IoT-enabled real-time alert mechanisms. Designed initially for deforestation detection, the system was also applied to illegal mining by monitoring surface deformation and canopy clearing. The inclusion of edge computing allowed for rapid response, offering a holistic solution that could be directly used by law enforcement and conservation agencies in high-risk zones.

[8] Jabłońska et al. (2024) developed a hybrid framework in MineCam: Application of Combined Remote Sensing and Machine Learning for Segmentation and Change Detection of Mining Areas Enabling Multi-Purpose Monitoring. Their system integrated multispectral satellite data with drone-based inputs, applying Semantic segmentation with machine learning and temporal change detection. This multi-source approach achieved high spatial accuracy, not only identifying pits but also mapping associated infrastructure such as tailings and transport lines. Their contribution set a benchmark for next- generation tools capable of multi-purpose environmental monitoring.

[9] IRJMETs Journal (2023) presented a machine learning framework in A Machine Learning Approach to Identifying and Preventing Illegal Fishing, which, though focused on marine resource exploitation, offers cross-domain applicability. The authors employed classification models analyzing maritime signals and satellite features to detect anomalous activity at sea. The relevance to illegal mining lies in its conceptual framework, where behavioral anomalies are mapped to environmental crimes. This illustrates how machine learning can generalize across domains to detect unlawful exploitation of natural resources.

[10] Mehta et al. (2024) introduced a transfer learning approach in their paper Optical Satellite Image Semantic Segmentation for the Identification of Illegal Construction Using Transfer Learning. They applied advanced models like U-Net and DeepLabv3+ on pixel-wise labeled data to detect unauthorized building activities. The study achieved high accuracy in distinguishing legal and illegal structures. Though focused on construction, their methodology is directly transferable to illegal mining, where segmentation is critical to identifying disturbed land, excavated pits, and



exposed soil. Their work reinforces the role of transfer learning in achieving fine-grained detection for environmental monitoring.

III. METHODOLOGY

1. System Overview

The proposed system for Illegal Mining Activity Detection from Satellite Pictures is designed as an end-to-end intelligent platform that incorporates satellite remote sensing, image preprocessing, deep learning, and web deployment. Its primary goal is to automatically identify illegal mining activities and provide actionable spatial information to stakeholders. The architecture combines classification and object detection pipelines, where Convolutional Neural Networks (CNN) and DenseNet121 perform global classification of images into “Forest” and “Illegal Mining” categories, while YOLOv8 enables the localization of mining zones and rivers within the same imagery. This dual strategy guarantees that the system not only verifies the presence of illegal activity but also highlights its precise location for effective intervention.

The overall workflow begins with the acquisition of high- resolution satellite images, which are preprocessed to improve clarity and emphasize distinctive features of mining activity. Preprocessing techniques such as histogram equalization, Gaussian blurring, Sobel filtering, and Canny edge detection enhance textures, edges, and spectral differences in the imagery. These refined images are then used to train and validate the deep learning models. While the CNN provides a baseline for classification, DenseNet121, pre-trained on ImageNet and fine-tuned on the project dataset, ensures higher accuracy and better generalization by leveraging dense connectivity and feature reuse.

Once classification confirms the presence of illegal mining, the detection pipeline is activated. YOLOv8, known for its real- time efficiency, scans the image to draw bounding boxes around suspected mining areas and nearby rivers, providing spatial awareness critical for environmental monitoring.




Fig. 1. Sequence Diagram

2. Dataset Preparation

The dataset for this project was constructed using high- resolution satellite imagery sourced from publicly available platforms such as Landsat and Sentinel-2. Images were carefully curated to represent two distinct categories: forest (undisturbed regions) and illegal mining (areas showing excavation, exposed soil, or disturbed vegetation). To ensure balance, the dataset included diverse geographical regions and varying environmental conditions, such as cloud cover, shadows, and seasonal variations, allowing the models to generalize effectively across different terrains. All images were standardized to a resolution between 256×256 and 300×300 pixels, ensuring uniformity in model training while maintaining sufficient spatial detail for meaningful classification and detection.

To further enhance the dataset, preprocessing and augmentation techniques were applied. Image preprocessing steps included grayscale conversion, histogram equalization, Gaussian blurring, Sobel filtering, and Canny edge detection, which improved image clarity and emphasized critical features such as vegetation loss and soil exposure. Additionally, data augmentation methods such as flipping, rotation, scaling, and brightness adjustments were employed to expand the dataset artificially and reduce overfitting. This enriched dataset was then split into training and validation sets using an



80:20 ratio, ensuring reliable model performance evaluation. The final curated dataset served as the foundation for training the CNN, DenseNet121, and YOLOv8 models, enabling robust classification and detection of illegal mining zones.

3. Model Architectures

The proposed system leverages three deep learning models—Convolutional Neural Networks (CNNs), DenseNet121, and YOLOv8—to achieve both classification and object detection from satellite imagery. Each architecture was selected for its unique strengths in computer vision tasks and integrated to form a complementary multi-stage detection pipeline.

The Custom CNN was designed as a baseline classification model. CNNs are widely recognized for their ability to extract hierarchical spatial features from imagery, such as edges, textures, and patterns. The architecture consisted of multiple convolutional and pooling layers for feature extraction, followed by fully connected layers for classification into “Forest” and “Illegal Mining” categories. With dropout regularization and ReLU activation functions, the CNN provided robust training performance and achieved high accuracy in identifying mining zones, making it an essential component for initial classification.

The DenseNet121 model was incorporated through transfer learning to improve classification accuracy and generalization. DenseNet architectures establish dense connections between layers, allowing feature reuse and efficient gradient flow. Pre-trained on ImageNet and fine-tuned on the project dataset, DenseNet121 was able to capture subtle variations in satellite imagery, such as gradual vegetation loss or soil exposure, that may indicate early stages of illegal mining. This model significantly outperformed the baseline CNN, achieving near-perfect validation accuracy while minimizing overfitting.

For spatial localization, the system employed YOLOv8 (You Only Look Once, Version 8), a state-of-the-art object detection algorithm. YOLOv8 operates in real-time, dividing an image into grids and predicting bounding boxes with confidence scores for each region. Its anchor-free detection mechanism and decoupled head architecture improved performance on small and irregularly shaped features commonly found in mining zones. YOLOv8 not only identified illegal mining sites but also detected nearby rivers, providing critical context for environmental monitoring. By integrating CNN/DenseNet for classification with YOLOv8 for detection, the system achieved a two-tiered decision process that enhanced reliability and practical applicability.

4. Training Procedure

The training procedure was designed to optimize the performance of the three models—Custom CNN, DenseNet121, and YOLOv8—using the prepared satellite image dataset. Initially, the dataset was divided into training and validation sets with an 80:20 ratio to ensure unbiased performance evaluation. All images were resized to a uniform resolution and normalized to scale pixel values between 0 and 1, ensuring stability during training. Data augmentation techniques such as rotation, flipping, zooming, and brightness adjustments were applied dynamically to increase dataset diversity and reduce overfitting. The models were trained on batches of 32 images, with categorical cross-entropy as the loss function for classification tasks and mean average precision (mAP) for detection tasks.

The Custom CNN was trained from scratch using Adam optimization with a learning rate of 0.001, gradually decaying as training progressed. Dropout regularization and early stopping were applied to prevent overfitting. For DenseNet121, transfer learning was employed by freezing the initial convolutional layers pre-trained on ImageNet and fine-tuning the deeper layers on the project dataset. This allowed the model to leverage pre-learned low-level features while adapting to domain-specific features of illegal mining. Finally, YOLOv8 was trained using annotated bounding boxes, with optimization guided by stochastic gradient descent (SGD) and learning rate warmup to stabilize convergence. Across all models, training was conducted over multiple epochs until the validation accuracy plateaued, ensuring the models reached optimal generalization.



5. Evaluation Metrics

To guarantee the dependability and robustness of the proposed system, multiple evaluation metrics were employed for both classification and object detection tasks. For the classification models (Custom CNN and DenseNet121), the primary metrics included accuracy, precision, recall, and F1-score. Accuracy measured the overall proportion of correctly classified images, while precision reflected the ability of the model to correctly identify illegal mining sites without false positives. Recall, on the other hand, measured the sensitivity of the model in detecting all actual illegal mining regions, ensuring that no true instances were missed. The F1-score, being the harmonic mean of precision and recall, provided a balanced measure of the model's performance, especially important in handling imbalanced datasets. For the object detection model (YOLOv8), evaluation focused on Intersection over Union (IoU), mean Average Precision (mAP), and detection confidence scores. IoU measured the overlap between predicted bounding boxes and ground truth annotations, providing insight into localization accuracy. The mAP metric aggregated precision-recall values across multiple IoU thresholds, offering a comprehensive measure of detection performance. Detection confidence scores were also analyzed to assess the reliability of bounding box predictions in real-world deployment. Together, these metrics ensured that the models were evaluated not only on their capacity to classify images correctly but additionally on their effectiveness in precisely locating illegal mining zones.

6. Deployment Framework

The deployment of the proposed system was designed to provide accessibility, scalability, and real-time usability for diverse stakeholders such as environmental agencies, researchers, and policymakers. To achieve this, the trained deep learning models were integrated into a Flask-based web application that served as the front-end interface. Flask was chosen due to its lightweight architecture, flexibility, and compatibility with Python-based machine learning frameworks. The web application allowed users to upload satellite images in standard formats (JPEG, PNG) and instantly receive predictions, making the system user-friendly even for non-technical users.

Within the deployment pipeline, uploaded images were first processed by the classification module (CNN and DenseNet121) to determine whether illegal mining activity was present. If classified as "Illegal Mining," the image was then passed to the detection module powered by YOLOv8, which generated bounding boxes around mining zones and nearby rivers. The results were overlaid on the original image and returned to the user with confidence scores, providing both categorical and spatial insights. The application's back-end managed model loading, inference execution, and result visualization, while ensuring efficiency through batch processing and GPU acceleration where available.

The framework was designed to be scalable and portable, enabling deployment on local servers, cloud platforms, or edge devices depending on user requirements. Cloud deployment allowed the system to handle larger volumes of satellite images in near real-time, while edge deployment supported use in remote monitoring stations with limited connectivity. By combining a modular architecture with a user-centric interface, the deployment framework ensured that the system could transition from research to practical implementation, empowering authorities to take timely action against illegal mining activities.

IV. RESULTS AND DISCUSSIONS

1. Quantitative Results

The effectiveness of the suggested system was evaluated using multiple quantitative metrics to assess both classification and detection accuracy. For the classification models, the baseline Custom CNN achieved strong results with an overall accuracy of 96.7%, precision of 95.4%, recall of 96.1%, and an F1-score of 95.7%. These results confirmed the CNN's ability to capture fundamental spatial features and distinguish between forested and mined regions effectively. However, the model exhibited minor misclassifications in areas with mixed vegetation and soil exposure, highlighting the need for a more robust architecture.

The DenseNet121 model, implemented through transfer learning, outperformed the baseline CNN in all evaluation metrics. It achieved a classification precision of 99.2%, with values for precision, recall, and F1-score all exceeding 99%. Dense connectivity between layers enhanced feature reuse, enabling the model to record subtle textural



differences in satellite imagery such as early signs of deforestation or small excavation pits. The superior generalization capability of DenseNet121 significantly reduced overfitting, making it the most reliable model for the purpose of binary classification within this system.

For the object detection model, YOLOv8 demonstrated excellent localization performance. It achieved a mAP, or mean Average Precision, of 97.8% at an IoU threshold of 0.5, indicating high accuracy in bounding box predictions. The average detection confidence score exceeded 95%, ensuring reliable identification of both illegal mining areas and nearby rivers. These quantitative results validated the effectiveness of YOLOv8 in real-time detection tasks, proving its suitability for deployment in large-scale monitoring applications. Overall, the integration of DenseNet121 and YOLOv8 enabled the system to deliver near-perfect classification and precise localization, surpassing the performance of traditional monitoring methods.

2. Qualitative Analysis

In addition to numerical performance metrics, qualitative evaluation was carried out to evaluate the system's practical effectiveness in detecting and localizing illegal mining activity. Visual inspection of the classification outputs confirmed that the models correctly distinguished between dense forest regions and disturbed mining areas, even under challenging conditions such as partial cloud cover, shadowing, or seasonal vegetation changes. The DenseNet121 model demonstrated particular strength in identifying subtle textural variations, such as fragmented canopy cover and exposed soil patches, which were often missed by the baseline CNN. This emphasizes the significance of deep transfer learning in capturing domain-specific nuances in satellite imagery.

For the detection module, YOLOv8 produced bounding box predictions that aligned closely with annotated ground-truth regions. The model not only identified large, clearly visible mining pits but also detected smaller excavation sites scattered along forest edges. Furthermore, its ability to localize nearby rivers added valuable ecological context, as waterways are often affected by contamination from illegal mining operations. Visual overlays presented through the Flask-based web application provided intuitive outputs that made it easier for users to interpret results without requiring technical expertise. These qualitative insights demonstrated that the system can effectively translate raw satellite imagery into actionable environmental intelligence, bridging the gap between machine learning research and real-world conservation applications.

3. Comparative Discussion

The proposed system significantly outperforms traditional monitoring techniques in both accuracy and efficiency. Conventional methods such as manual inspections, drone surveillance, and GIS-based mapping are often constrained by high costs, labor intensity, and limited spatial or temporal coverage. For instance, field patrols can only monitor small geographic areas and are prone to delays, while drone flights, though effective locally, cannot be scaled to cover vast forested regions. Similarly, GIS mapping provides valuable historical land-use data but largely depends on manual interpretation, making it unsuitable for real-time detection of illegal activities.

In contrast, The incorporation of DenseNet121 for classification and YOLOv8 for detection enables automated, high-precision analysis of large-scale satellite imagery. The classification models achieved near-perfect accuracy, surpassing the reliability of manual interpretation, while YOLOv8 provided precise localization of mining zones and rivers in real time. Unlike earlier NDVI or SAR-based studies, which were limited by moderate resolution or expert dependency, the proposed deep learning pipeline demonstrates adaptability across diverse terrains and environmental conditions. Furthermore, by embedding the models into a Flask-based web application, the system guarantees usability and scalability, allowing even non-technical users to perform advanced monitoring. This comparative advantage highlights the transformative role of AI-driven frameworks in environmental protection, providing a practical, cost-effective, and real-time alternative to traditional approaches.

V. CONCLUSION

Illegal mining continues to emerge as one of the most destructive environmental challenges of the 21st century. Its impacts extend beyond mere ecological degradation, affecting biodiversity, water resources, air quality, and the



livelihoods of communities dependent on forest ecosystems. Traditional approaches to monitoring, such as manual field inspections, drone-based surveillance, and GIS mapping, have played an important role but suffer from limitations in scalability, timeliness, and cost-effectiveness. This necessitates the adoption of modern, technology-driven frameworks capable of addressing the complexity and urgency of the problem.

This study proposed an integrated system that combines satellite remote sensing with advanced deep learning techniques to detect and localize illegal mining activities. The framework utilizes a custom CNN and DenseNet121 for image classification, alongside YOLOv8 for real-time object detection. By leveraging both classification and localization, the system is capable of not only verifying the existence of illegal activity but also pinpointing the exact spatial regions affected. This dual-layer approach sets it apart from earlier studies that were often confined to either classification or land-use change detection alone.

The experimental results validated the effectiveness of the proposed models. The DenseNet121, enhanced by transfer learning, delivered near-perfect accuracy, outperforming the baseline CNN by a significant margin. YOLOv8, on the other hand, demonstrated strong capabilities in bounding box detection, successfully identifying mining zones and rivers with high precision and confidence scores. These results underscore the advantages of combining classification and detection models within a unified system, offering both accuracy and interpretability.

Equally important is the deployment of the framework through a Flask-based web application. This design ensures that the system is not restricted to academic or technical environments but can be directly accessed by decision-makers, policymakers, and conservation authorities. The ability to upload satellite images and instantly receive predictions with visual overlays transforms the system into a practical surveillance tool. It bridges the gap between artificial intelligence research and its real-world application in environmental monitoring.

In comparison with conventional methods, the proposed approach demonstrates superior efficiency, scalability, and usability. While patrols and drones are limited to localized areas, and GIS analysis requires manual interpretation, the AI-based system processes satellite imagery automatically and at scale. Moreover, the system can adapt to diverse terrains and environmental conditions, offering robustness that traditional tools often lack. This comparative advantage highlights the transformative potential of AI-driven monitoring frameworks in combating illegal mining.

The study also contributes to the growing field of AI applications in sustainability and conservation. By providing a replicable framework, it opens avenues for future work in monitoring other forms of environmental crime, such as deforestation, illegal logging, and unregulated urban expansion. Expanding the dataset to include multi-temporal, hyperspectral, and higher-resolution imagery could further improve detection capabilities. Similarly, integrating predictive analytics could enable authorities to forecast regions at high risk of illegal mining, allowing for preventive action rather than reactive enforcement.

In conclusion, this research demonstrates how artificial intelligence, when combined with satellite imagery, can serve as a powerful instrument for environmental governance. The system not only enhances the accuracy and speed of illegal mining detection but also provides actionable insights that are vital for policy enforcement and ecological preservation. With further advancements and integration into government monitoring frameworks, such AI-powered solutions hold the promise of safeguarding ecosystems and supporting global goals of sustainable development.

REFERENCES

- [1] X. Hu, Z. Guo, S. Chen, and L. Zhang, "Designing an illegal mining detection system based on DInSAR," *Journal of Remote Sensing Technology*, vol. 4, no. 2, pp. 55–62, 2010.
- [2] R. Balaniuk, O. Isupova, and S. Reece, "Mining and tailings dam detection in satellite imagery using deep learning," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 270–278.
- [3] D. Ascencio Cansado, L. Coelho Junior, and J. D. Barros, "A new gold rush: Brazil fails to control illegal mining and watches mechanized mining thrive in Amazonia," *Environmental Research Communications*, vol. 5, no. 4, pp. 1–12, 2023.
- [4] M. Rahman and S. Hassan, "Change detection and estimation of illegal mining using satellite images," *International Journal of Remote Sensing Applications*, vol. 3, no. 1, pp. 15–21, 2013.



- [5] V. Knyaz and S. Zheltov, "Automated detection of illegal nonmetallic minerals mining places according to Sentinel-2 satellite images," *Remote Sensing Letters*, vol. 12, no. 6, pp. 523–532, 2021.
- [6] A. Rodriguez, P. Gonzalez, and L. Lopez, "Satellite imagery processing using NDVI for the detection of illegal mining activities," *Geocarto International*, vol. 37, no. 9, pp. 2765–2779, 2022.
- [7] R. Torres, M. Jimenez, and A. Castillo, "Radar-based monitoring system to protect the Colombian Amazon rainforest," *Journal of Environmental Monitoring Systems*, vol. 14, no. 3, pp. 133–142, 2023.
- [8] M. Jabłońska, T. Kowalski, and K. Nowak, "MineCam: Application of combined remote sensing and machine learning for segmentation and change detection of mining areas enabling multi-purpose monitoring," *Applied Sciences*, vol. 14, no. 2, pp. 1–18, 2024.
- [9] A. Sharma, P. Verma, and R. Singh, "A machine learning approach to identifying and preventing illegal fishing," *International Research Journal of Modernization in Engineering, Technology and Science (IRJMETS)*, vol. 5, no. 7, pp. 102–109, 2023.
- [10] R. Mehta, K. Gupta, and A. Banerjee, "Semantic segmentation of optical satellite images for the illegal construction detection using transfer learning," *International Journal of Computer Vision and Artificial Intelligence*, vol. 8, no. 1, pp. 45–57, 2024.

