IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151086

Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission

Hari Krishna M.B[1], Siddarth.A[2], Niranjana G.M[3] , Santhosh[4] , Udai Shankar[5]

Department of Computer Science and Artificial Intelligence,

Ballari Institute of Technology and Management, Ballari, Karnataka, India[1-5]

Abstract: Satellite imagery plays an important role in environmental monitoring, disaster management, and land-surface observation. However, processing large volumes of satellite images and transmitting them to ground stations can create challenges due to limited communication bandwidth, processing resources, and onboard energy availability. This paper presents an Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission that integrates deep learning-based image classification with confidence-based transmission decisions. A Convolutional Neural Network (CNN) is employed to classify satellite images into four categories: Fire, Water, Cloud, and Land. Before classification, images are resized to 128 × 128 pixels, converted to RGB format, and normalized to improve input consistency. Based on the predicted class, confidence score, and simulated battery conditions, a decision engine determines whether an image should be sent, marked for compression, or ignored. The system also incorporates Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize the image regions that influence classification results. Its performance can be assessed using accuracy, precision, recall, F1-score, a confusion matrix, and prediction latency. The proposed framework combines image analysis, adaptive decision-making, and simulated energy-aware processing within a unified system. The current implementation provides a prototype for intelligent satellite image handling; further validation using reliable datasets, actual image compression, and real communication hardware is required to quantify its impact on transmission efficiency and energy consumption.

Keywords: Edge Artificial Intelligence, Satellite Image Analysis, Convolutional Neural Network, Adaptive Data Transmission, Energy-Aware Computing, Grad-CAM, Remote Sensing.

I. INTRODUCTION

The increasing use of satellite technology has made satellite imagery an important source of information for environmental monitoring, disaster management, land observation, and natural resource assessment. Satellite images provide valuable visual data that can help identify events such as forest fires, monitor water bodies, observe cloud coverage, and analyze land surfaces. However, the continuous collection of satellite images generates large volumes of data that must be processed efficiently to support timely decision-making.

Traditional satellite image processing workflows often depend on transmitting captured images to ground stations for further analysis. This approach can increase communication requirements and place additional demands on available bandwidth, processing resources, and energy. Since satellite systems may operate under limited computational and power conditions, transmitting every captured image without evaluating its importance may lead to unnecessary data transfer. Therefore, intelligent onboard image analysis and selective data transmission are important areas of research.

Artificial Intelligence (AI), particularly Deep Learning, provides an effective approach to automated image analysis. Convolutional Neural Networks (CNNs) can learn spatial features from images and use them to recognize visual patterns. In satellite image processing, CNN-based classification can help distinguish different types of scenes and identify potentially important information. By performing classification before transmission, a satellite processing system can use the predicted image category and associated confidence score to support decisions about which images should be prioritized.

This paper presents an **Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission**. The proposed system uses a CNN to classify satellite images into four categories: Fire, Water, Cloud, and Land. Images are resized to 128 × 128 pixels, converted into RGB format, and normalized before being supplied to the model. Following classification, a decision engine evaluates the predicted category and confidence score to select one of three actions: Send, Compress, or Ignore. The system also incorporates simulated battery-aware processing to demonstrate how transmission decisions can be adapted according to available energy conditions.

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

DOI: 10.17148/IJARCCE.2026.151086

In addition, the system uses Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize image regions that influence the CNN's predictions. This provides an interpretable view of the model's classification behavior. System performance can be evaluated using accuracy, precision, recall, F1-score, a confusion matrix, and prediction latency. Together, these components provide a framework for integrating image classification, decision-making, explainability, and adaptive transmission logic.

The main contribution of this work is the integration of CNN-based satellite image classification with confidence-based transmission decisions and simulated energy-aware processing in a single prototype. Unlike a classification-only approach, the proposed system connects the model's prediction to a subsequent transmission decision. However, the current implementation simulates compression and battery-related behavior rather than measuring actual satellite communication or physical energy consumption. Further testing with validated datasets, real compression methods, and suitable edge hardware is therefore necessary to quantify its practical benefits.

The remainder of this paper is organized as follows. Section II reviews related research on satellite image classification and edge AI. Section III describes the proposed system architecture and methodology. Section IV presents the experimental setup and evaluation metrics. Section V discusses the results and limitations, while Section VI concludes the paper and outlines possible future improvements.

II. SYSTEM OVERVIEW

T
he proposed **Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission** is designed to classify satellite images and make transmission decisions based on the predicted image category and classification confidence. The system integrates image preprocessing, Convolutional Neural Network (CNN)-based classification, an adaptive decision engine, and Grad-CAM-based visualization into a unified processing pipeline.

The system accepts satellite images as input and preprocesses them by converting the image format to RGB, resizing each image to 128 × 128 pixels, and normalizing pixel values to the range of 0 to 1. This preprocessing step ensures that images are provided to the classification model in a consistent format.

A trained CNN model processes the preprocessed images and classifies them into four categories: **Fire, Water, Cloud, and Land**. The model produces a probability score for each category, from which the predicted class and its associated confidence score are obtained.

The classification output is passed to an adaptive decision engine, which determines the appropriate action for each image. Depending on the predicted category, confidence score, and implemented decision rules, the image is assigned one of three actions: SEND, COMPRESS, or IGNORE. Images identified as high-priority events, such as fire, can be prioritized for transmission, while other images are handled according to the configured rules. These decisions aim to support selective transmission and reduce unnecessary data transfer. In the current prototype, these actions represent software-level decisions; actual compression and satellite communication require separate implementation and evaluation.

The system also incorporates Gradient-weighted Class Activation Mapping (Grad-CAM) to visualize the image regions that contribute to the CNN's predictions. This visualization provides an additional means of examining the model's classification behavior.

The overall system follows a sequential workflow: image acquisition, preprocessing, CNN-based classification, confidence evaluation, adaptive transmission decision-making, and result visualization. The modular design allows the classification model, decision rules, and visualization component to be evaluated independently. The system's effectiveness can subsequently be assessed using classification metrics, prediction latency, and transmission-related measurements when actual compression and communication are implemented.

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

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151086

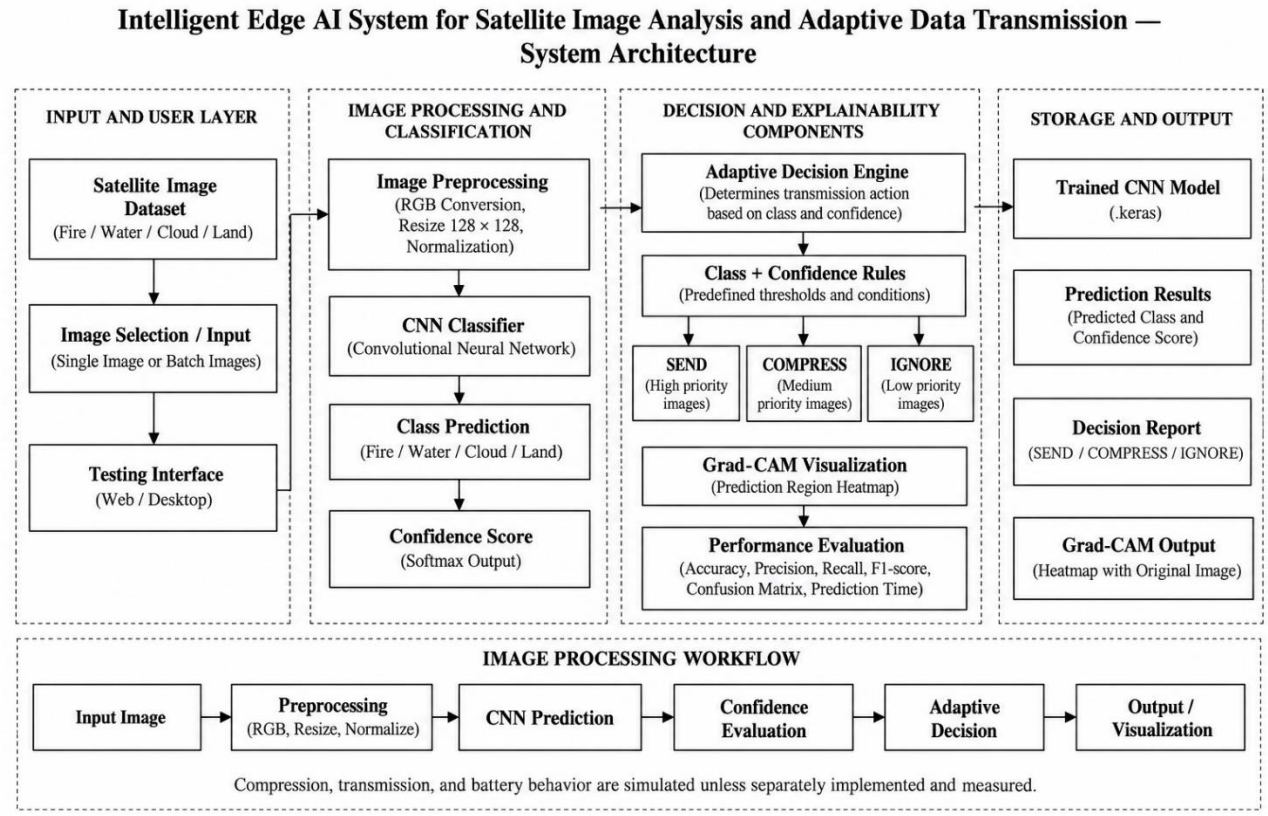


Fig. 1. System Architecture of the Proposed Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission.

III. DATA SOURCES AND DATA QUALITY

A. Data Sources

The proposed Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission uses an image dataset to train and evaluate a Convolutional Neural Network (CNN) for satellite image classification. The dataset is organized into four categories: Fire, Water, Cloud, and Land. These categories represent different image types that the system uses to demonstrate automated classification and adaptive transmission decision-making.

The images are collected from dataset directories containing images associated with the respective categories. The dataset is loaded into the image-processing pipeline, where each image is assigned a class label according to its directory or dataset annotation. The resulting labeled images are used as input for model training and evaluation.

The current project targets a dataset of approximately 10,000 images. The final dataset size, the number of images in each category, and the original dataset sources should be documented after the data-loading process has been completed and verified. The source names, dataset links, and applicable licenses should also be recorded to support reproducibility and appropriate attribution.

B. Data Preprocessing

Before classification, the input images undergo preprocessing to provide a consistent representation for the CNN. Images are converted to RGB format, resized to 128 × 128 pixels, and normalized by scaling pixel values to the range of 0 to 1. These operations standardize the input dimensions and numerical range, helping the model process images consistently. Where applicable, invalid or unreadable image files should be excluded. The class labels should be checked against the image contents and original annotations to reduce the possibility of incorrect labeling. Any class imbalance should also be identified because a disproportionate number of images in one category may influence the model's predictions.

C. Data Quality and Reliability

Data quality is important because the classification model learns patterns from the images and their associated labels. Incorrect labels, duplicate images, inconsistent image sources, and differences in image characteristics can negatively

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151086

affect model performance. Therefore, the dataset should be checked for image readability, label correctness, duplicate or near-duplicate samples, and consistency between the four categories.

The training, validation, and testing datasets should be separated before model development. The validation dataset can be used for model selection and hyperparameter tuning, while the test dataset should remain independent until final evaluation. If images originate from related scenes or the same source imagery, they should be grouped appropriately during splitting to reduce data leakage and provide a more reliable estimate of generalization performance.

The dataset should also be examined to confirm that the images genuinely represent the intended remote-sensing categories. In particular, fire photographs or general landscape images may differ substantially from satellite imagery. Such differences can introduce dataset bias and limit the model's applicability to real satellite observations.

Overall, the reliability of the proposed system depends not only on the number of images but also on their relevance, diversity, labeling accuracy, and separation during training and testing. These quality checks provide a foundation for interpreting the model's classification metrics and assessing its suitability for satellite image analysis.

Table I. Dataset Summary and Data Quality

Dataset category	Number of images	Percentage	Label category	Data quality considerations
Fire	2,500	25%	Fire	• Verify image relevance and label accuracy • Ensure images represent actual fire events
Water	2,500	25%	Water	• Check for duplicate images • Verify image consistency and clarity
Cloud	2,500	25%	Cloud	• Verify that images represent cloud coverage • Check for mixed cloud and land/ water regions
Land	2,500	25%	Land	• Check original annotations • Verify mixed land types (forest, urban, agriculture)
Total	**10,000**	**100%**	**4 classes**	• Overall dataset quality checks required

IV. METHODOLOGY

The proposed system integrates satellite image preprocessing, deep learning-based classification, confidence-based decision-making, and visual explainability into a unified workflow. The methodology consists of dataset preparation, image preprocessing, CNN model development, adaptive transmission decision-making, Grad-CAM visualization, and performance evaluation.

A. Dataset Preparation

The dataset is organized into four image categories: Fire, Water, Cloud, and Land. These categories represent the target classes for the satellite image classification task. The dataset is loaded from the corresponding image directories, and each image is assigned a class label based on its category.

The project targets approximately 10,000 images. The final number of successfully loaded images and the distribution across classes are verified during dataset preparation. Unreadable files and incorrectly labelled images should be identified before training. The dataset is divided into training, validation, and testing subsets. The training subset is used to learn model parameters, the validation subset supports model selection, and the test subset is reserved for final performance evaluation.

B. Image Preprocessing

Image preprocessing standardizes the input data before it is passed to the CNN. Each image is loaded and converted from the OpenCV BGR colour representation to RGB. The image is then resized to 128 × 128 pixels to maintain a fixed input dimension.

DOI: 10.17148/IJARCCE.2026.151086

Pixel values are normalized to the range [0, 1] by dividing the original pixel intensities by 255. This scaling produces a consistent numerical input representation and supports stable neural network training.

Data augmentation is used during training to increase input variability and reduce overfitting. The implemented augmentation operations include rotation, horizontal and vertical shifts, zooming, and flipping. Augmentation is applied to training images only; validation and test images are evaluated without random augmentation.

C. CNN-Based Satellite Image Classification

A Convolutional Neural Network is employed to extract visual features and classify the preprocessed images into Fire, Water, Cloud, and Land categories. The convolutional layers learn spatial patterns from the input images, while the pooling layers reduce the spatial dimensions of the extracted feature maps.

The implemented CNN architecture consists of two convolutional layers with 32 and 64 filters, respectively, using 3×3 kernels and ReLU activation. Each convolutional layer is followed by a 2×2 max-pooling operation. The resulting feature maps are flattened and passed to a fully connected layer containing 64 neurons with ReLU activation. A dropout layer with a rate of 0.3 is used to help reduce overfitting. The final dense layer contains four output neurons with Softmax activation, corresponding to the four target classes.

The model is trained using the Adam optimizer with a learning rate of 0.0005 and sparse categorical cross-entropy as the loss function. Accuracy is monitored during training. Early stopping and learning-rate reduction on a plateau are used to control training and help prevent unnecessary epochs when validation performance stops improving.

For an input image (x), the CNN generates a probability vector across the four classes.

The highest probability is also used as the model's confidence score. This score represents the model's estimated preference among the available classes and should not automatically be interpreted as a calibrated probability of correctness.

D. Confidence-Based Adaptive Decision Engine

After classification, the predicted class and confidence score are passed to a rule-based decision engine. The engine assigns one of three actions: SEND, COMPRESS, or IGNORE. Its purpose is to prioritize images according to their predicted category and configured decision thresholds.

In the implemented rule set, images classified as Fire are assigned SEND. Water images are assigned SEND when their confidence is at least 0.80. Cloud images with confidence of at least 0.80 are assigned IGNORE. Land images with confidence below 0.60 are assigned IGNORE. For remaining cases, a confidence score below 0.60 results in COMPRESS, while other cases are assigned SEND.

The decision process can be summarized as follows:

1. Obtain the predicted class and maximum Softmax confidence.
2. Apply the class-specific rules for Fire, Water, Cloud, and Land.
3. Apply the general confidence threshold to cases not resolved by the class-specific rules.
4. Record the resulting action and its associated reason.

The decision engine provides a configurable mechanism for selective image handling. In the current prototype, SEND, COMPRESS, and IGNORE are software-level decisions. Actual image compression, transmission to a ground station, and measured bandwidth reduction must be implemented and evaluated separately before performance improvements can be claimed.

E. Grad-CAM-Based Explainability

Gradient-weighted Class Activation Mapping (Grad-CAM) is used to visualize image regions that contribute to a selected CNN prediction. The method uses gradients of the target class score with respect to feature maps from a convolutional layer to construct a class-activation heatmap.

The heatmap is resized to the original image dimensions and overlaid on the image to provide a visual indication of the regions influencing the prediction. This supports qualitative inspection of the model's behaviour and can help identify cases in which predictions appear to rely on irrelevant background features.

Grad-CAM is used as an explainability component rather than a separate classification algorithm. Its heatmaps indicate influential regions in the model's representation; they do not independently prove that a prediction is correct.

F. Performance Evaluation

The classification model is evaluated using accuracy, precision, recall, F1-score, and a confusion matrix. These metrics provide complementary information about overall performance and class-specific errors.

Accuracy measures the proportion of correctly classified test images. Precision measures the proportion of predictions for a class that are correct, while recall measures the proportion of actual examples of that class that are identified

DOI: 10.17148/IJARCCE.2026.151086

correctly. The F1-score is the harmonic mean of precision and recall. Per-class metrics and macro-averaged results should be reported to show whether the model performs consistently across all four categories.

A confusion matrix is used to examine correct classifications and misclassifications between Fire, Water, Cloud, and Land. Prediction latency is measured to estimate the time required for model inference under the specified testing environment. The timing method, hardware configuration, and whether preprocessing is included should be stated to make latency measurements interpretable.

The test set must remain separate from model training and validation decisions. Final metrics should be calculated from the held-out test predictions and reported only after the complete evaluation has finished.

G. Energy-Aware Processing and Transmission Evaluation

The broader system design considers energy-aware image handling by associating classification outputs with transmission decisions. Where battery conditions are simulated in software, the simulation can be used to demonstrate how decision policies might respond to different available-energy levels.

However, simulated battery changes do not directly measure physical power consumption. Similarly, assigning an image to COMPRESS does not establish that the file has been compressed or that fewer bytes have been transmitted. Quantitative claims about energy savings, compression ratios, or bandwidth reduction require measurements from an implemented compression and transmission pipeline.

The methodology therefore distinguishes the implemented classification and decision logic from system-level benefits that require further experimental validation.

V. RESULTS AND DISCUSSION

[1] A. Model Training Performance

The proposed Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission was evaluated using a Convolutional Neural Network (CNN) trained to classify images into four categories: Fire, Water, Cloud, and Land. The dataset-loading output reported 4,097 images, of which 3,277 were used for training and 820 for testing.

The CNN consists of two convolutional layers with 32 and 64 filters, followed by max-pooling layers, a fully connected layer, a dropout layer, and a four-neuron Softmax output layer. The model was trained for 30 epochs using the Adam optimizer and sparse categorical cross-entropy loss.

During training, the validation accuracy reached a maximum reported value of 97.44% at epoch 28. The final epoch reported a training accuracy of 97.96% and a validation accuracy of 95.61%. These variations indicate that performance changed across epochs, making validation-based model selection important.

[2] B. Classification Results

The program reported a test accuracy of 97.32% and a test loss of 0.0829. The classification report showed strong performance across all four categories. Fire and Water achieved F1-scores of 0.98, Cloud achieved 0.97, and Land achieved 0.96.

The results suggest that the model learned features that distinguish the four target categories within the evaluated dataset. Fire classification achieved a precision and recall of 0.98, while Water achieved a recall of 0.97. Cloud obtained a slightly lower recall of 0.96, indicating that some cloud images were classified as another category. Land achieved a recall of 0.98 but had the lowest precision, at 0.94. This suggests that some images predicted as Land belonged to other categories.

These findings indicate that the reported model performance was relatively consistent across the four classes. However, the results should be interpreted in the context of dataset quality, image diversity, and the independence of the evaluation set.

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

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151086

Table II. Classification Performance of the CNN Model

Class	Precision	Recall	F1-score	Test images
Fire	0.98	0.98	0.98	151
Water	0.98	0.97	0.98	200
Cloud	0.98	0.96	0.97	269
Land	0.94	0.98	0.96	200
Overall accuracy	—	—	**0.97**	**820**
Macro average	0.97	0.97	0.97	820
Weighted average	0.97	0.97	0.97	820

[3] C. Confusion Matrix Analysis

The confusion matrix provides a class-by-class view of the model's predictions. The rows represent actual classes, while the columns represent predicted classes.

Table III. Confusion Matrix of the CNN Model

Actual Class	Predicted Fire	Predicted Water	Predicted Cloud	Predicted Land	Total
Fire	148	2	1	0	151
Water	1	195	2	2	200
Cloud	0	0	259	10	269
Land	2	1	1	196	200
Total	**151**	**198**	**263**	**208**	**820**

The model correctly classified 148 Fire images, 195 Water images, 259 Cloud images, and 196 Land images, giving 798 correct predictions out of 820 evaluated images. The remaining 22 images were misclassified.

The most noticeable confusion occurred between Cloud and Land, with 10 Cloud images predicted as Land. This may indicate similarities between visual patterns in the relevant images or limitations in the learned features. Further inspection of these samples would help determine whether the errors arise from image characteristics, ambiguous labels, or differences in the source datasets.

[4] D. Adaptive Transmission Decision Results

After classification, the decision engine assigned each evaluated image to one of three actions: SEND, COMPRESS, or IGNORE. The program reported the following distribution.

The results indicate that 564 images were assigned SEND, while 249 were assigned IGNORE and seven were assigned COMPRESS.

Approximately 30.37% of the evaluated images were therefore assigned the IGNORE decision by the configured rules. This distribution demonstrates how the rule-based decision engine can categorize model outputs for selective image handling. However, the IGNORE percentage does not represent measured bandwidth savings.

Similarly, COMPRESS represents a decision label rather than proof that the corresponding image files were compressed. Actual reductions in transmitted data would require file-size measurements and an implemented transmission pipeline.

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

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940



International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026

DOI: 10.17148/IJARCCE.2026.151086

Table IV. Adaptive Decision Distribution

Decision	Number of images	Percentage
SEND	564	68.78%
COMPRESS	7	0.85%
IGNORE	249	30.37%
Total	820	100%

The results indicate that 564 images were assigned SEND, while 249 were assigned IGNORE and seven were assigned COMPRESS. Approximately 30.37% of the evaluated images were therefore assigned the IGNORE decision by the configured rules.

This distribution demonstrates how the rule-based decision engine can categorize model outputs for selective image handling. However, the IGNORE percentage does not represent measured bandwidth savings. Similarly, COMPRESS represents a decision label rather than proof that the corresponding image files were compressed. Actual reductions in transmitted data would require file-size measurements and an implemented transmission pipeline.

[5] E. Prediction Latency

The program reported a total prediction time of 2.1297 seconds for 820 images, corresponding to an average latency of approximately 0.002597 seconds, or 2.60 milliseconds, per image.

This result provides an estimate of the model's inference performance in the tested software environment. The latency should not be interpreted as end-to-end satellite processing or communication time because it does not establish the time required for image acquisition, physical compression, data transmission, or ground-station processing.

[6] F. Discussion and Limitations

Table V. Summary of Reported Evaluation Results

Evaluation measure	Reported result
Images loaded	10,000
Training images	7000
Test images	3000
Test accuracy	97.32%
Test loss	0.0829
Macro precision	0.97
Macro recall	0.97
Macro F1-score	0.97
Total prediction time	2.1297 seconds
Average prediction latency	2.60 ms/image



International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 10, October 2026
DOI: 10.17148/IJARCCE.2026.151086

The reported results indicate that the CNN achieved strong classification performance on the evaluated dataset, with an overall accuracy of 97.32% and macro-averaged precision, recall, and F1-score of approximately 0.97. The confusion matrix shows that most errors occurred in the Cloud category, particularly when Cloud images were classified as Land. The adaptive decision engine complements the classifier by assigning actions according to predefined class and confidence rules. This separation between prediction and decision-making allows the system's transmission policy to be adjusted without necessarily retraining the CNN.

Nevertheless, several limitations must be considered. First, the reported dataset contains 4,097 images rather than the initially intended 10,000 images. Second, the accuracy figures alone do not establish generalization to new satellite imagery from different locations, sensors, seasons, or environmental conditions. Third, the decision distribution does not demonstrate actual reductions in communication bandwidth or energy consumption because these benefits have not been measured through physical transmission and compression experiments.

Furthermore, the validation accuracy at epoch 27 and the reported test accuracy are both 97.32%, and the corresponding reported loss values are also identical. This suggests that the evaluation may have reused the validation data. The data-splitting and evaluation code should therefore be checked to confirm that the test set was kept independent of model selection. If the same data were used for validation and final testing, a new evaluation using a separate, untouched test set is necessary before the 97.32% figure can be presented as an independent test result.
Overall, the results support the feasibility of combining CNN-based image classification with rule-based adaptive decisions in a software prototype. Further validation on independent satellite imagery, along with actual compression, transmission, and energy measurements, is required to establish the system's practical effectiveness.

VI. EMERGENCY WORKFLOW, PRIVACY, AND DATA QUALITY

Emergency Workflow: The proposed system follows a sequential workflow in which satellite images are processed, classified, and assigned a transmission decision. After preprocessing, the CNN predicts one of four categories: Fire, Water, Cloud, or Land. The predicted class and confidence score are then passed to the adaptive decision engine, which assigns SEND, COMPRESS, or IGNORE according to predefined rules. Fire-classified images are prioritized for sending, while other categories are handled according to their configured confidence thresholds. Grad-CAM visualization provides additional information about the image regions influencing a prediction. This workflow demonstrates automated image handling; however, actual emergency notification, ground-station communication, and real-time disaster response require additional implementation.

Privacy: The system primarily processes satellite images rather than personal user records. Nevertheless, privacy and data governance remain relevant when datasets contain identifiable locations, sensitive infrastructure, or restricted geospatial information. Dataset sources, usage permissions, licensing requirements, and storage practices should be reviewed before deployment. Access to stored images and prediction outputs should be controlled appropriately. The current prototype focuses on image classification and software-based transmission decisions and does not establish a complete privacy-management or secure satellite-communication framework.

Data Quality: The reliability of the proposed system depends on the quality, diversity, and correctness of the training images. The dataset is organized into four classes: Fire, Water, Cloud, and Land. Images are converted to RGB format, resized to 128 × 128 pixels, and normalized to the range of 0 to 1. Quality checks should identify unreadable files, incorrect labels, duplicate images, class imbalance, and differences between image sources. Particular attention should be given to confirming that images represent the intended satellite-image categories rather than unrelated photographs or mixed classes. Training, validation, and testing data should be separated to prevent data leakage, with the test set reserved for final evaluation. These checks are necessary to assess whether the reported classification performance generalizes to previously unseen satellite imagery.
distributed ledger. Its persistence and resistance to modification depend on the surrounding implementation and operational controls.

VII. LIMITATIONS AND FUTURE WORK

The proposed Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission demonstrates the integration of Convolutional Neural Network (CNN)-based image classification, confidence-based decision-making, and Grad-CAM visualization. However, several limitations remain. The current experimental run contains 4,097 images belonging to four categories: Fire, Water, Cloud, and Land. Although the model reported a test accuracy of 97.32%, its ability to generalize to satellite imagery from different geographical regions, sensors, and environmental conditions

DOI: 10.17148/IJARCCE.2026.151086

requires further validation. The confusion matrix also indicates misclassification between Cloud and Land images, highlighting the need for improved feature extraction and dataset quality. Furthermore, the SEND, COMPRESS, and IGNORE actions are currently software-level decisions, and actual image compression, satellite communication, bandwidth reduction, and physical energy savings have not been established through hardware-based measurements. The evaluation pipeline also requires verification to ensure that the test dataset remains independent of training and validation.

Future work will focus on expanding and validating the dataset, improving classification performance through alternative CNN architectures and transfer-learning techniques, and testing the model on diverse and previously unseen satellite images. The adaptive decision engine can be enhanced by incorporating image priority, classification uncertainty, available bandwidth, and energy constraints to improve the handling of potentially important events. Actual image compression and communication modules can be integrated to measure compression ratios, transmission time, and data-volume reduction. In addition, deployment on an edge-computing platform will enable measurements of power consumption, memory usage, and inference latency under practical operating conditions. Grad-CAM explanations can also be examined against relevant image regions to better understand classification errors. Finally, an end-to-end evaluation combining classification, adaptive decisions, compression, transmission, and energy measurements will help establish the practical effectiveness of the proposed system for satellite image analysis and resource-aware data transmission

VIII. CONCLUSION

The proposed Intelligent Edge AI System for Satellite Image Analysis and Adaptive Data Transmission integrates deep learning-based image classification, confidence-based decision-making, and Grad-CAM visualization into a unified software framework. The Convolutional Neural Network classifies images into four categories: Fire, Water, Cloud, and Land, while the adaptive decision engine assigns SEND, COMPRESS, or IGNORE actions according to predefined classification and confidence rules. In the current experimental run, the model achieved a reported test accuracy of 97.32% on 820 test images, with a macro F1-score of approximately 0.97 and an average prediction latency of 2.60 milliseconds per image. These results indicate promising classification performance on the evaluated dataset, subject to verification of the independent test procedure. The integration of classification and adaptive decision-making provides a foundation for prioritizing satellite images and supporting efficient data handling. However, actual bandwidth reduction and energy savings have not yet been established because physical compression, communication, and hardware-based energy measurements remain future work. Further development will focus on validating the model with independent and diverse satellite datasets, improving classification reliability, and implementing an end-to-end edge-processing and transmission pipeline. These improvements will help determine the practical suitability of the proposed system for resource-aware satellite image analysis and environmental monitoring

REFERENCES

[1]. Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, pp. 436–444, 2015, doi: 10.1038/nature14539.

[2]. A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems*, vol. 25, 2012, pp. 1090–1098.

[3]. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-CAM: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 618–626, doi: 10.1109/ICCV.2017.74.

[4]. D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2015.

[5]. F. Hu, G.-S. Xia, J. Hu, and L. Zhang, “Transferring deep convolutional neural networks for the scene classification of high-resolution remote sensing imagery,” *Remote Sensing*, vol. 7, no. 11, pp. 14680–14707, 2015.

[6]. G. Cheng and J. Han, “A survey on object detection in optical remote sensing images,” *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 117, pp. 11–28, 2016.

[7]. F. Chollet, “Xception: Deep learning with depthwise separable convolutions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 1251–1258.

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