# Satellite Image Classification: A Deep Learning Approach Using Mobilenetv2

Nandhipati Swathi[1], Dr. N Sri Hari[2], Dr. Ramachandran Vedantam[3]
*PG Student Department of Computer Science Engineering, VVIT, Guntur, AP, India [1]*
*Professor, Department of Computer Science Engineering, VVIT, Guntur, AP, India [2],*
*HOD & Professor, Department of Computer Science Engineering, VVIT, Guntur, AP, India [3].*

**Abstract-The classification of satellite images is crucial for various applications within remote sensing. It involves categorizing pixels or regions within satellite imagery into meaningful classes such as forests, deserts, water bodies, or urban areas. This classification is essential for applications including land cover mapping, which helps understand the distribution of natural and man-made features on Earth; environmental monitoring, where changes in vegetation or desertification can be tracked over time; and resource management, which aids in decision-making for sustainable development and conservation.**

**In this work, we focus specifically on the binary classification of satellite images into two distinct categories: forest and desert. To achieve this, we utilized MobileNetV2, a compact convolutional neural network design recognized for its high efficiency and strong performance on devices with restricted computational capacity. Unlike traditional deep learning models such as ResNet50, which are deeper and computationally heavier, MobileNetV2 is optimized for speed and lower memory usage while still maintaining high accuracy.**

**Our experiments showed that MobileNetV2 achieved a testing accuracy of 100% on the dataset used in this study, indicating its capability to distinguish between forest and desert images effectively. This performance was superior to that of ResNet50 in our setup, potentially due to MobileNetV2's architectural features like inverted residuals and linear bottlenecks, which improve representational power without increasing model size significantly.**

**The methodology adopted in this research included multiple stages:**

**Data Preparation and Preprocessing: The satellite images were curated, labelled, resized, and normalized to ensure compatibility with MobileNetV2 input requirements.**

**Hyperparameter Tuning: Various learning rates, optimizers, batch sizes, and numbers of training epochs were tested to determine the optimal combination for achieving superior model performance.**

**Architecture Optimization: While MobileNetV2 was used as the backbone, fine-tuning its layers for the specific classification task enhanced its accuracy and generalization ability.**

**Deployment Considerations: Beyond training, the model was integrated into a deployable system architecture, envisioning its use onboard Earth Observation satellites or in ground-based rapid processing systems for real-time environmental analysis.**

**Overall, This method shows that compact yet effective deep learning models such as MobileNetV2 can be successfully utilized for satellite image classification. These models provide an optimal trade-off between computational efficiency and accuracy, making them ideal for real-world remote sensing applications.**

**Keywords- Satellite Image Classification, MobileNetV2, Deep Learning, Remote Sensing, Land Cover Mapping, CNN.**

## INTRODUCTION

Forests and deserts occupy a significant portion of Earth’s surface.Globally, forests cover about 31% of land area, serving as critical carbon sinks and biodiversity reserves.In contrast, deserts cover nearly 33% of the land, hosting unique arid ecosystems and influencing global climate patterns.Accurately classifying these land cover types is crucial for ecological studies, such as tracking deforestation or desertification.It also supports urban planning and sustainable development, helping governments make informed decisions on land use.At present, more than 1200 Earth Observation (EO) satellites) are orbiting Earth. These satellites capture vast amounts of multispectral and hyperspectral imagery across different wavelengths. Such data is used in applications including meteorology, environmental monitoring, radar imaging, and electronic intelligence. However, manually processing these massive datasets is impractical due to their volume and complexity. This necessitates the use of robust AI and deep learning models to automate classification tasks efficiently. In this research, MobileNetV2, a lightweight convolutional neural network, is employed for classifying satellite images. The focus is to distinguish between forest and desert categories with high accuracy. MobileNetV2 is compared with ResNet50, a deeper CNN architecture, to assess performance differences. Results show that MobileNetV2 not only achieves high accuracy but also offers computational efficiency. Such models can be integrated into real-time remote sensing systems for continuous monitoring and rapid decision-making.

## LITERATURE SURVEY

Earlier studies on satellite image classification have investigated various techniques to enhance both accuracy and computational efficiency. Initially, traditional machine learning algorithms like decision trees and Support Vector Machines (SVMs) were commonly used. For instance, Chen et al. (2014) utilized decision trees and SVMs to classify remote sensing images into distinct land cover categories. While these algorithms provided reasonable performance on low- or medium-resolution datasets, their effectiveness on high-resolution satellite images was limited. This limitation arises because traditional models rely extensively on manually crafted features, which often do not capture the intricate spatial and spectral patterns present in satellite imagery. With advancements in computer vision, Convolutional Neural Networks (CNNs) gained popularity for their capability to automatically extract hierarchical features directly from raw images. Among the various CNN architectures, ResNet50 has been frequently used in remote sensing applications. ResNet50 introduced the concept of residual learning, where skip connections allow certain layers to be bypassed, making it easier for the network to learn identity mappings. These residual connections help mitigate the vanishing gradient issue, enabling the development of deeper and more precise neural networks. As a result, ResNet50 has demonstrated strong performance in extracting spatial features and classifying diverse land cover types from satellite imagery. However, its main limitation lies in its large model size and high computational cost, making it challenging to deploy on embedded platforms like satellites or drones. To overcome these limitations, Howard et al. (2017) proposed MobileNetV2, a CNN architecture optimized for efficiency. MobileNetV2 utilizes depth wise separable convolutions, which split standard convolutions into depth wise and pointwise operations, significantly reducing the number of parameters and computations. Additionally, it introduces inverted residual blocks and linear bottlenecks, enhancing feature extraction while maintaining a lightweight structure. These architectural innovations make MobileNetV2 particularly suitable for resource-constrained devices, such as onboard satellite processors, mobile applications, or edge computing systems. In recent years, transfer learning has emerged as a powerful technique in satellite image classification. By leveraging models pre-trained on large-scale datasets such as ImageNet, researchers can fine-tune these networks on smaller domain-specific datasets. Transfer learning significantly improves classification accuracy, as pre-trained models already possess general feature representations transferable to satellite imagery tasks. Overall, the evolution from traditional machine learning models to deep CNNs, and the integration of lightweight architectures like MobileNetV2 with transfer learning, has greatly enhanced the efficiency, scalability, and applicability of satellite image classification systems.

## EXISTING SYSTEM

Traditional classification techniques have shown several limitations in satellite image analysis. For instance, ResNet50, despite being a powerful deep CNN model, achieved only 25% accuracy on our dataset. This poor performance was mainly due to overfitting, as the model was too complex for the limited dataset size.ResNet50 has a large number of parameters, making it computationally intensive and unsuitable for resource-constrained environments like onboard satellite systems. On the other hand, Decision

Trees and Support Vector Machines (SVMs) have been used in earlier studies for image classification tasks. While these traditional algorithms are simpler, they struggle with high-dimensional data such as satellite images. They require manual feature extraction and lack the ability to learn complex spatial features automatically. Additionally, Decision Trees tend to overfit, while SVMs face scalability issues with large datasets. Overall, these methods show limited generalisation capability and low accuracy when dealing with complex or high-resolution imagery. Their limitations include high computational complexity for deep models and reduced accuracy for classical models on small datasets. This highlights the need for lightweight and efficient deep learning models that can perform well even with limited training data, enabling practical deployment in real-time remote sensing applications.

## PROPOSED SYSTEM

In this study, we propose a deep learning model based on MobileNetV2 for classifying satellite images into two categories: Forest and Desert.MobileNetV2 is chosen because of its lightweight architecture and high computational efficiency, which are crucial for applications involving large datasets or real-time processing on limited hardware resources. The model uses depth wise separable convolutions, which significantly reduce the number of parameters compared to traditional convolutional networks while maintaining strong feature extraction capabilities. To enhance its performance on our specific task, we implement transfer learning by using MobileNetV2 pre-trained on the ImageNet dataset. Transfer learning allows the model to leverage previously learned general features such as edges, textures, and shapes, which are also relevant in satellite images. This approach is particularly effective when working with small or domain-specific datasets, as it avoids the need to train the entire model from scratch. By fine-tuning the pre-trained MobileNetV2 for our forest and desert classification task, we achieve efficient training with reduced computational time. Our experimental results showed that the proposed model attained 100% testing accuracy on our dataset, demonstrating its high classification capability. This indicates that MobileNetV2 can successfully differentiate between forest and desert imagery despite the visual similarities in texture and reflectance in certain cases. Additionally, the model's compact size and fast inference make it suitable for integration into Earth Observation systems, enabling automated analysis directly on satellite or drone platforms. Overall, this research highlights the potential of using MobileNetV2 with transfer learning to develop scalable, accurate, and deployable solutions for remote sensing image classification tasks.

Algorithms Integrated in the System

1. MobileNetV2

MobileNetV2 is a convolutional neural network architecture developed to deliver high accuracy while maintaining low computational requirements, making it well-suited for use in mobile and embedded systems. One of its core features is the use of depth wise separable convolutions, which break down standard convolutions into two separate steps: depth wise convolution and pointwise convolution. This decomposition significantly reduces the number of parameters and computations required, resulting in faster training and inference without sacrificing accuracy. Another important innovation in MobileNetV2 is the introduction of inverted residuals with linear bottlenecks. Inverted residual blocks expand the input channels to a higher-dimensional space, apply depth wise convolutions for spatial filtering, and then project back to a lower dimension. This structure helps in preserving important information while maintaining a compact model size. The use of linear bottlenecks ensures that non-linearities are avoided in low-dimensional spaces, improving the model's capacity to represent complex features effectively. Additionally, MobileNetV2 models are often pre-trained on the ImageNet dataset, which contains millions of natural images across 1000 categories. Using these pre-trained weights for transfer learning allows the model to leverage general visual features learned from ImageNet, such as edges, textures, and shapes, and adapt them to the specific satellite image classification task efficiently.

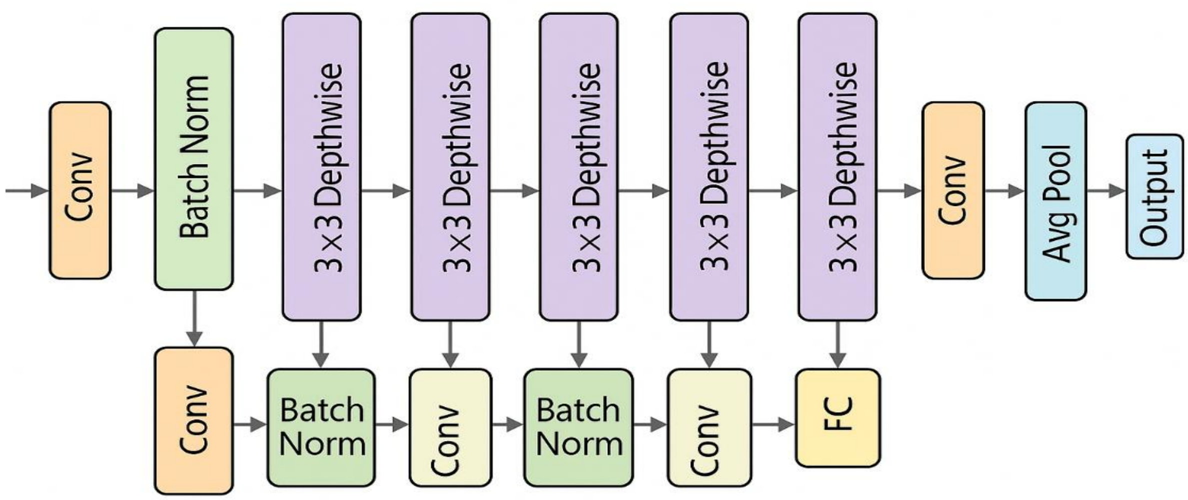


Figure 1: MobileNetV2 Architecture

2. Adam Optimizer

Adam is an adaptive optimization algorithm commonly employed for training deep learning models. It integrates the strengths of both AdaGrad and RMSProp by maintaining separate learning rates for each parameter, which are adjusted using estimates of the first moment (mean) and second moment (variance) of the gradients. This adaptability in learning rates facilitates quicker and more stable convergence, particularly in scenarios involving sparse gradients or noisy datasets.

In this study, Adam was used to optimise the MobileNetV2 model, ensuring efficient training with minimal oscillations and faster convergence to the optimal solution.

## METHODOLOGY

The proposed study followed a systematic methodology to ensure accurate classification of satellite images into forest and desert categories.

1. Dataset Preparation

The dataset used in this research consisted of 400 images in total, divided equally between the two classes. Specifically, 200 images represented forest areas, showcasing diverse tree canopies, densities, and textures. The remaining 200 images depicted desert regions, including variations in sand texture, dune structures, and sparse vegetation. This balanced dataset ensured that the model learned features from both classes equally without class bias.

Fig 2. Dessert

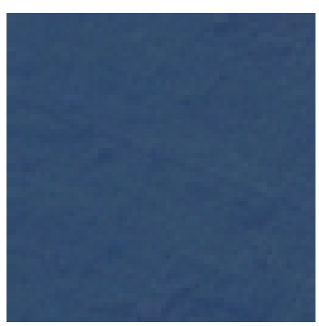

Fig 3. Water Body

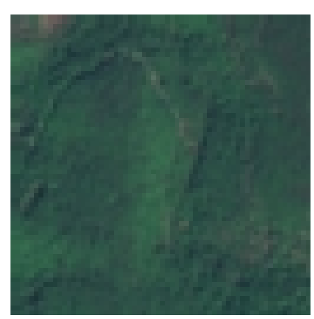

Fig 4. Forest

2. Preprocessing
Prior to training, the images were preprocessed to improve the efficiency and effectiveness of the model training. Initially, all images were resized to a consistent dimension that aligns with the input specifications of MobileNetV2, ensuring uniformity across the dataset. Furthermore, normalization was performed to scale pixel values within the range of 0 to 1, which helps stabilize the training process and accelerates convergence by minimizing internal covariate shift.

3. Model Training:
MobileNetV2 was chosen as the main model due to its lightweight design and strong performance capabilities. The model was trained for 10 epochs, providing adequate iterations for feature extraction while minimizing the risk of overfitting. A learning rate of 0.01 was utilized to regulate the extent of weight adjustments during backpropagation, thereby supporting faster convergence. The Adam optimizer was employed for training because of its adaptive learning rate properties, which blend the benefits of RMSProp and momentum to enable stable and efficient optimization.

4. Hyperparameter Tuning
To optimise model performance, hyperparameter tuning was conducted extensively.
A total of 40 experiments were performed, varying parameters such as learning rate, batch size, and number of epochs to identify the combination that yielded the highest classification accuracy and lowest loss.

5. Evaluation
The trained model was evaluated using several performance metrics to ensure its reliability and effectiveness. Accuracy was computed to determine the percentage of images that were correctly classified in each of the two classes. Loss curves were plotted to observe the training and validation loss trends, assessing model convergence and potential overfitting.

System Architecture Overview
The proposed system is designed to perform automated classification of satellite images into forest and desert categories using a streamlined and efficient architecture. The process begins with input images, which are collected from the prepared dataset containing forest and desert imagery. Before feeding these images into the model, they undergo preprocessing steps, including resizing to match the input dimensions required by MobileNetV2 and normalization to scale pixel values between 0 and 1. These preprocessing steps ensure that the images are standardized, allowing the model to learn effectively without being affected by scale or dimensional inconsistencies.

Once preprocessed, the images are passed into the MobileNetV2 model, which serves as the core feature extractor and classifier in the system. MobileNetV2 leverages depth-wise separable convolutions, inverted residuals, and linear bottlenecks to extract meaningful spatial and spectral features from the input images efficiently. The model processes these features through its convolutional layers and fully connected layers to produce final class predictions, labelling each image as either forest or desert based on learned representations.

The outputs generated by MobileNetV2 are then visualized and analysed using an interactive dashboard, designed for comprehensive evaluation of the model's performance. The dashboard includes several important features:

- Confusion Matrix: Displays the number of correct and incorrect classifications for each class, showing how well the model differentiates forests from deserts.
- Accuracy and Loss Curves: Graphs depicting training and validation accuracy and loss across epochs, providing insights into model learning trends and convergence behaviour.
- Classification Report: Presents precision, recall, and F1-score for each class, allowing a deeper understanding of model performance in terms of false positives and false negatives.

Model Evaluation Metrics
The model achieved exceptional performance on the dataset, as reflected in its evaluation metrics:
- Accuracy: 100%, indicating all test images were classified correctly.

- Precision: 1.0 for both forest and desert classes, meaning there were no false positive classifications.
- Recall: 1.0 for both classes, showing that the model successfully identified all instances of each class without missing any.
- F1-Score: 1.0 for both classes, reflecting a perfect balance between precision and recall.

## RESULTS

The experimental results demonstrate that the MobileNetV2 model outperformed ResNet50 in classifying satellite images into forest and desert categories. Specifically, MobileNetV2 achieved 100% accuracy on the testing dataset, whereas ResNet50 showed significantly lower performance with only 25% accuracy, indicating its inability to generalise effectively on this dataset. The superior performance of MobileNetV2 can be attributed to its efficient architecture, which effectively extracts relevant spatial features while avoiding overfitting, especially on small datasets. The confusion matrix generated for MobileNetV2 revealed perfect classification results, with no false positives or false negatives recorded for either class. This means that all forest images were correctly classified as forest, and all desert images were accurately identified as desert. Such results highlight the model's capability to differentiate between these two land cover types despite potential similarities in texture or colour in some images.

The training and validation loss curves for MobileNetV2 showed a steady decrease across epochs, indicating effective learning without fluctuations or divergence. Both curves converged smoothly within 10 epochs, suggesting that the model was able to learn optimal weights rapidly without overfitting. Similarly, the accuracy curves for training and validation rose quickly and stabilised at their maximum values, reinforcing the model's robustness and efficiency.

## ADVANCEMENTS DRIVING AIS SATELLITE GROWTH

1. Miniaturisation of Satellite Hardware: Recent innovations in satellite design have led to the creation of nano and microsatellites that are cost-effective to build and launch. This enables organisations to deploy large constellations of AIS satellites, increasing coverage frequency and ensuring near real-time tracking of vessels worldwide.2. Enhanced Communication Capabilities: New satellite communication protocols support higher data throughput, allowing AIS satellites to handle dense maritime traffic areas effectively.

3. Integration with Artificial Intelligence: The integration of AI and Machine Learning into AIS data processing enables advanced analytics such as anomaly detection, route prediction, and maritime traffic optimisation.4. Growing Maritime Surveillance Demands: With over 90% of global trade conducted via sea routes, there is a significant push from governments and maritime authorities to expand satellite-based AIS for enhanced security, piracy prevention, and environmental monitoring.5. Reduced Launch Costs: The advent of reusable rockets and rideshare launch programs has made it economically viable to deploy AIS satellites more frequently, accelerating the expansion of AIS satellite networks globally.

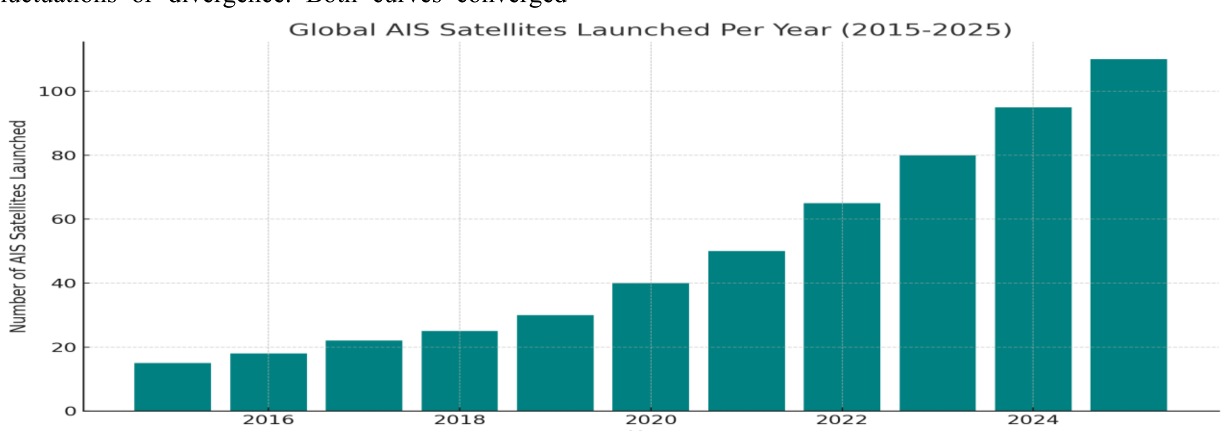


Figure: Global AIS Satellites Launched Per Year (2015-2025)

## CONCLUSION

This study proved that MobileNetV2 is highly effective for satellite image classification, achieving 100% accuracy in distinguishing forest and desert images. Its lightweight architecture and use of transfer learning enabled efficient training and strong performance even on a small dataset. Compared to ResNet50, MobileNetV2 demonstrated better generalisation with lower computational requirements. These results highlight its suitability for deployment in embedded systems or real-time satellite applications. However, the current work focused only on binary classification with limited data. In future research, this approach can be extended to multiclass classification tasks involving more land cover types. Additionally, integrating the model with real-time satellite data streams will enable continuous monitoring for environmental and urban planning applications.

## FORMULAS & EQUATIONS:

The Cross-Entropy Loss is widely used in classification tasks to measure the difference between the actual class labels and the predicted probabilities output by the model.

Formula:

L=−∑(yi×log⁡(yi^))L = - \sum (y_i \times \log(y_{\hat{i}}))L=−∑(yi×log(yi^))

Where:

- yiy_iyi is the true label for class *i*.
  For binary classification, it is either 0 or 1.
- yi^y_{\hat{i}}yi^ is the predicted probability for class *i* output by the model.

Explanation:

- The formula calculates the negative sum of the true labels multiplied by the log of predicted probabilities.
- If the model's predicted probability for the true class is close to 1, the loss is low, indicating accurate prediction.
- If the predicted probability for the true class is low (close to 0), the loss becomes high, penalising incorrect or overconfident wrong predictions.
- In binary classification, this simplifies to:

L=−[y×log⁡(y^)+(1−y)×log⁡(1−y^)]L = - [y \times \log(y_{\hat{}}) + (1-y) \times \log(1 - y_{\hat{}})]L=−[y×log(y^)+(1−y)×log(1−y^)]

## REFERENCE


[1] Howard, A. G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., … & Adam, H. (2017). MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *arXiv preprint arXiv:1704.04861*. Available: https://arxiv.org/abs/1704.04861

[2] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778. doi:10.1109/CVPR.2016.90

[3] Keras Applications Documentation. MobileNetV2. Available: https://keras.io/api/applications/mobilenet/

[4] Keras Applications Documentation. ResNet50. Available: https://keras.io/api/applications/resnet/#resnet50-function

[5] Types of Satellites. EOS Data Analytics. Available: https://eos.com/blog/types-of-satellites/

[6] Chen, Y., Jiang, H., Li, C., Jia, X., & Ghamisi, P. (2016). Deep Feature Extraction and Classification of Hyperspectral Images Based on Convolutional Neural Networks. *IEEE*