

Deep Learning–Based Land Use and Land Cover Classification Using the Eurosat Dataset

Sivakumaran Sarvanan

BSc (Hons) in Computer Science (Sri Lanka) MSc Candidate in Data Science and Artificial Intelligence
(France)

DOI: <https://doi.org/10.51583/IJLTEMAS.2026.1502000097>

Received: 13 February 2026; Accepted: 21 February 2026; Published: 19 March 2026

ABSTRACT

Land Use and Land Cover (LULC) classification plays a crucial role in remote sensing applications such as urban planning, environmental monitoring, agricultural analysis, and climate studies. Recent advances in deep learning, particularly convolutional neural networks (CNNs), have significantly improved classification accuracy for satellite imagery. This thesis presents a comparative study of two deep learning approaches for LULC classification using the EuroSAT dataset: a convolutional neural network trained from scratch and a transfer learning model based on a pre-trained VGG-19 architecture. The EuroSAT dataset consists of Sentinel-2 satellite images categorized into ten land cover classes. Experimental results demonstrate that transfer learning achieves superior classification performance compared to training a CNN from scratch, highlighting the effectiveness of pre-trained models for remote sensing image analysis.

Keywords: Remote Sensing, EuroSAT, Land Use and Land Cover, Deep Learning, CNN, Transfer Learning.

INTRODUCTION

Land Use and Land Cover (LULC) classification is a fundamental task in remote sensing that involves identifying and categorizing different land surface types from satellite imagery. Accurate LULC maps are essential for applications such as environmental monitoring, urban growth analysis, deforestation detection, and agricultural planning. Traditionally, LULC classification relied on manual interpretation or classical machine learning techniques using handcrafted features, which are often time-consuming and less scalable.

With the rapid growth of deep learning, convolutional neural networks (CNNs) have emerged as a powerful tool for image classification tasks. CNNs automatically learn hierarchical feature representations from raw image data, eliminating the need for manual feature engineering. In recent years, CNN-based methods have achieved state-of-the-art performance in satellite image classification tasks.

This thesis investigates the application of deep learning techniques for LULC classification using the EuroSAT dataset. Two different approaches are explored: (1) a CNN trained from scratch and (2) a transfer learning model utilizing a pre-trained VGG-19 network. The primary objective is to evaluate and compare their performance and suitability for satellite image classification.

Related Work

Several studies have demonstrated the effectiveness of deep learning for remote sensing image classification. Helber et al. introduced the **EuroSAT dataset** as a benchmark for land use and land cover classification using Sentinel-2 imagery, showing that **CNN-based approaches significantly outperform traditional machine learning methods** [1]. Subsequent research has explored various deep learning architectures, including **AlexNet, VGG, ResNet, DenseNet, and EfficientNet**, for EuroSAT classification [2]. **Transfer learning** has been widely adopted due to its ability to leverage knowledge from large-scale datasets such as ImageNet, with studies consistently reporting higher accuracy and faster convergence compared to training networks from

scratch. This thesis builds upon these works by implementing and comparing a **custom CNN** and a **VGG19-based transfer learning model** on the EuroSAT dataset.

EuroSAT Dataset

The EuroSAT dataset is a publicly available remote sensing dataset derived from Sentinel-2 satellite imagery. It contains approximately 27,000 labeled RGB images with a spatial resolution of 10 meters per pixel. Each image has a size of 64×64 pixels and belongs to one of ten land use and land cover classes:

- A. Annual Crop
- B. Forest
- C. Herbaceous Vegetation
- D. Highway
- E. Industrial
- F. Pasture
- G. Permanent Crop
- H. Residential
- I. River
- J. Sea/Lake

K. The dataset is well balanced and widely used as a benchmark for evaluating deep learning models in remote sensing applications.


```
▶ data.classes_names
```

```
... {'River': 0,  
    'HerbaceousVegetation': 1,  
    'Residential': 2,  
    'Highway': 3,  
    'SeaLake': 4,  
    'AnnualCrop': 5,  
    'PermanentCrop': 6,  
    'Forest': 7,  
    'Industrial': 8,  
    'Pasture': 9}
```


Class name with label

Data Preprocessing

The images are resized to match the input requirements of the deep learning models. Pixel values are normalized, and the dataset is split into training and testing subsets. For transfer learning, ImageNet normalization parameters are applied to ensure compatibility with the pre-trained VGG-19 model.



Sample



Sample Data

METHODOLOGY

CNN Trained from Scratch

The first approach involves designing a convolutional neural network trained entirely from scratch. The architecture consists of multiple convolutional layers followed by batch normalization, ReLU activation, and max-pooling layers. These layers progressively extract spatial features from the input images. The convolutional backbone is followed by fully connected layers that perform the final classification into ten classes.

This approach allows the network to learn task-specific features directly from the EuroSAT dataset but requires careful regularization to prevent overfitting.

Transfer Learning Using VGG-19


The second approach employs transfer learning using a pre-trained VGG-19 model. VGG-19 was originally trained on the ImageNet dataset, which contains millions of natural images across 1,000 classes. The convolutional layers of VGG-19 are used as a fixed feature extractor, while the final fully connected layers are replaced with a new classifier tailored to the ten EuroSAT classes.

Transfer learning reduces training time and improves generalization by leveraging previously learned visual features such as edges, textures, and shapes.

```
def build_model():
    vgg19 = models.vgg19(pretrained=True)
    for param in vgg19.parameters():
        param.requires_grad = False
    vgg19.avgpool = nn.AdaptiveAvgPool2d(output_size=(1,1))
    vgg19.classifier = nn.Sequential(
        nn.Flatten(),
        nn.Linear(512, 128),
        nn.ReLU(),
        nn.Dropout(0.2),
        nn.Linear(128, 10))

    loss_function = nn.CrossEntropyLoss()
    optimizer = optim.Adam(vgg19.parameters(), lr=1e-3)

    return vgg19.to(device), loss_function, optimizer
```



VGG 19 model architecture

Training Configuration

Both models are trained using the cross-entropy loss function and optimized with the Adam optimizer. Dropout and batch normalization are applied to reduce overfitting. Training is conducted for a fixed number of epochs, and performance is evaluated on the test dataset.

TABLE I: VGG-19 CNN Architecture for EuroSAT Classification

Layer (Type)	Output Shape	Parameters
Input	[-1, 3, 64, 64]	–
Conv Block 1		
Conv2d (64 filters, 3x3)	[-1, 64, 64, 64]	1,792
ReLU	[-1, 64, 64, 64]	–
Conv2d (64 filters, 3x3)	[-1, 64, 64, 64]	36,928
ReLU	[-1, 64, 64, 64]	–
MaxPool2d (2x2)	[-1, 64, 32, 32]	–
Conv Block 2		
Conv2d (128 filters, 3x3)	[-1, 128, 32, 32]	73,856
ReLU	[-1, 128, 32, 32]	–
Conv2d (128 filters, 3x3)	[-1, 128, 32, 32]	147,584
ReLU	[-1, 128, 32, 32]	–
MaxPool2d (2x2)	[-1, 128, 16, 16]	–
Conv Block 3		
Conv2d (256 filters, 3x3)	[-1, 256, 16, 16]	295,168
ReLU	[-1, 256, 16, 16]	–
Conv2d (256 filters, 3x3)	[-1, 256, 16, 16]	590,080
ReLU	[-1, 256, 16, 16]	–
Conv2d (256 filters, 3x3)	[-1, 256, 16, 16]	590,080
ReLU	[-1, 256, 16, 16]	–
Conv2d (256 filters, 3x3)	[-1, 256, 16, 16]	590,080
ReLU	[-1, 256, 16, 16]	–
MaxPool2d (2x2)	[-1, 256, 8, 8]	–
Conv Block 4		
Conv2d (512 filters, 3x3)	[-1, 512, 8, 8]	1,180,160
ReLU	[-1, 512, 8, 8]	–
Conv2d (512 filters, 3x3)	[-1, 512, 8, 8]	2,359,808
ReLU	[-1, 512, 8, 8]	–
Conv2d (512 filters, 3x3)	[-1, 512, 8, 8]	2,359,808
ReLU	[-1, 512, 8, 8]	–
Conv2d (512 filters, 3x3)	[-1, 512, 8, 8]	2,359,808
ReLU	[-1, 512, 8, 8]	–
MaxPool2d (2x2)	[-1, 512, 4, 4]	–
Adaptive Pooling		
AdaptiveAvgPool2d	[-1, 512, 1, 1]	–
Fully Connected Layers		
Flatten	[-1, 512]	–
Linear (512 → 128)	[-1, 128]	65,664
ReLU	[-1, 128]	–
Dropout (0.3)	[-1, 128]	–
Linear (128 → 10)	[-1, 10]	1,290
Total Parameters	20,091,338	Trainable: 66,954

Note: Total multi-adds: 1.61 G; Estimated total model size: 85.94 MB.

4. Build the Train and Accuracy Functions

```
def train_batch(model, loss_function, optimizer, image, label):
    model.train()
    optimizer.zero_grad()
    prediction = model(image)
    loss = loss_function(prediction, label.long().squeeze())
    loss.backward()
    optimizer.step()
    return loss.item()
```

```
@torch.no_grad()
def accuracy(model, loss_function, image, label):
    model.eval()
    prediction = model(image)
    max_values, argmaxes = prediction.max(-1)
    is_correct = argmaxes == label.long().squeeze()
    return is_correct.cpu().numpy().tolist()
```

```
▶ @torch.no_grad()
def validation_loss(model, loss_function, image, label):
    model.eval()
    prediction = model(image)
    loss = loss_function(prediction, label.long().squeeze())
    return loss.item()
```

Training Accuracy Computation

Training accuracy was computed at each epoch by comparing the predicted class labels with the ground truth labels. The predicted class was obtained by selecting the class with the maximum softmax probability. Accuracy was calculated as the ratio of correctly classified samples to the total number of training samples, as shown in Equation.

$$\text{Training Accuracy} = \frac{\text{Number of Correct Predictions}}{\text{Total Number of Samples}}$$

Training Procedure

The model was trained for 25 epochs using the Adam optimizer with a categorical cross-entropy loss function. During training, both loss and accuracy were computed at each epoch. Training accuracy was calculated by comparing predicted class labels with ground truth labels over the entire training dataset. Validation performance was evaluated after each epoch using a held-out test set, with gradients disabled to ensure unbiased evaluation.

VGG-Inspired CNN (Baseline Custom CNN)

A custom VGG-inspired convolutional neural network (CNN) was developed for the EuroSAT dataset, which contains 64×64 RGB images. The network consists of four convolutional blocks with batch normalization,

ReLU activation, and max pooling, followed by adaptive average pooling and three fully connected layers with dropout for 10-class classification. Designed to be smaller than full-scale pre-trained models like VGG19 or ResNet50, this architecture reduces overfitting, allows efficient training on limited computational resources, and provides flexibility for experimentation while maintaining high classification performance.

TABLE 1
CUSTOM VGG-INSPIRED CNN ARCHITECTURE FOR EUROSAT CLASSIFICATION (MODELS 2)

Layer (Type)	Output Shape	Parameters
Input	[-1, 3, 64, 64]	–
Conv Block 1		
Conv2d (64 filters, 3x3)	[-1, 64, 64, 64]	1,792
BatchNorm2d	[-1, 64, 64, 64]	128
ReLU	[-1, 64, 64, 64]	–
MaxPool2d	[-1, 64, 63, 63]	–
Conv Block 2		
Conv2d (128 filters, 3x3)	[-1, 128, 63, 63]	73,856
BatchNorm2d	[-1, 128, 63, 63]	256
ReLU	[-1, 128, 63, 63]	–
MaxPool2d	[-1, 128, 62, 62]	–
Conv Block 3		
Conv2d (256 filters, 3x3)	[-1, 256, 62, 62]	295,168
BatchNorm2d	[-1, 256, 62, 62]	512
ReLU	[-1, 256, 62, 62]	–
MaxPool2d	[-1, 256, 61, 61]	–
Conv Block 4		
Conv2d (512 filters, 3x3)	[-1, 512, 61, 61]	1,180,160
BatchNorm2d	[-1, 512, 61, 61]	1,024
ReLU	[-1, 512, 61, 61]	–
MaxPool2d	[-1, 512, 60, 60]	–
AdaptiveAvgPool2d	[-1, 512, 1, 1]	–
Fully Connected Layers		
Flatten	[-1, 512]	–
Linear (512 → 128)	[-1, 128]	65,664
ReLU	[-1, 128]	–
Dropout (0.3)	[-1, 128]	–
Linear (128 → 64)	[-1, 64]	8,256
ReLU	[-1, 64]	–
Dropout (0.3)	[-1, 64]	–
Linear (64 → 10)	[-1, 10]	650
Total Parameters	1,627,466	Trainable: 1,627,466

Note: Total multi-adds: 5.82 G; Estimated total model size: 62.09 MB.

VGG-Inspired CNN (Tuned Custom CNN)

The proposed model is a VGG-inspired convolutional neural network designed for classifying EuroSAT satellite images. The network consists of four convolutional blocks, each including a convolutional layer, batch normalization, ReLU activation, and max pooling.

An adaptive average pooling layer reduces spatial dimensions, followed by three fully connected layers with dropout regularization, producing a final 10-class output. The network is trained from scratch using cross-entropy loss and the Adam optimizer with a learning rate of 0.001. This compact architecture balances computational efficiency and classification performance on small 64×64 RGB images.

TABLE 1

VGG-INSPIRED CNN (TUNED CUSTOM CNN MODEL 3) ARCHITECTURE SUMMARY FOR EUROSAT DATASET

Layer (Type)	Output Shape	Parameters
Input	[-1, 3, 64, 64]	–
Conv Block 1		
Conv2d (64 filters, 3x3)	[-1, 64, 64, 64]	1,792
BatchNorm2d	[-1, 64, 64, 64]	128
ReLU	[-1, 64, 64, 64]	–
MaxPool2d	[-1, 64, 63, 63]	–
Conv Block 2		
Conv2d (128 filters, 3x3)	[-1, 128, 63, 63]	73,856
BatchNorm2d	[-1, 128, 63, 63]	256
ReLU	[-1, 128, 63, 63]	–
MaxPool2d	[-1, 128, 62, 62]	–
Conv Block 3		
Conv2d (256 filters, 3x3)	[-1, 256, 62, 62]	295,168
BatchNorm2d	[-1, 256, 62, 62]	512
ReLU	[-1, 256, 62, 62]	–
MaxPool2d	[-1, 256, 61, 61]	–
Conv Block 4		
Conv2d (512 filters, 3x3)	[-1, 512, 61, 61]	1,180,160
BatchNorm2d	[-1, 512, 61, 61]	1,024
ReLU	[-1, 512, 61, 61]	–
MaxPool2d	[-1, 512, 60, 60]	–
AdaptiveAvgPool2d	[-1, 512, 1, 1]	–
Flatten	[-1, 512]	–
Fully Connected Layers		
Linear (512 → 128)	[-1, 128]	65,664
ReLU	[-1, 128]	–
Dropout (0.3)	[-1, 128]	–
Linear (128 → 64)	[-1, 64]	8,256
ReLU	[-1, 64]	–
Dropout (0.3)	[-1, 64]	–
Linear (64 → 10)	[-1, 10]	650
Total Parameters	1,627,466 (Trainable: 1,627,466, Non-trainable: 0)	
Multi-adds (G)	5.82	
Estimated Total Size (MB)	62.09 (Input: 0.05, Forward/Backward: 55.84, Params: 6.21)	

Model Comparison

Model	Name / Type	Description
Model 1	VGG19	Full pre-trained VGG19 (19-layer deep network)
Model 2	VGG-Inspired CNN (2 nd model)	Custom CNN with 4 convolutional blocks, batch norm, ReLU, max pooling, adaptive pooling, 3 FC layers with dropout
Model 3	VGG-Inspired CNN (3 rd model)	Slight variation of Model 2 — e.g., modified number of filters, FC layer sizes, or dropout rates to improve generalization

To evaluate the effectiveness of convolutional neural networks for EuroSAT land use classification, three models were compared: the pre-trained **VGG19**, the **VGG-Inspired CNN (Baseline Custom CNN)**, and the **VGG-Inspired CNN (Tuned Custom CNN)**. VGG19 serves as a standard benchmark with deep pre-trained features, while the two VGG-inspired CNNs are compact, custom architectures designed specifically for the small 64×64 RGB EuroSAT images.

The comparison highlights the trade-offs between model complexity, computational efficiency, and classification performance. The VGG-inspired models require significantly fewer parameters and training resources, reducing the risk of overfitting while maintaining competitive accuracy. Fine-tuning the **Tuned Custom CNN** demonstrates how minor architectural adjustments, such as modified filter sizes or dropout rates, can further improve generalization. Overall, this comparison provides insight into the suitability of pre-

trained versus custom CNN architectures for small satellite image datasets, guiding the selection of an efficient and effective model.

Experimental Results

Evaluation Metrics

The models were evaluated primarily using **classification accuracy**. Additional analyses included **training and validation loss curves**, as well as **class-wise prediction behavior**, to assess model generalization and performance across all categories of the EuroSAT dataset.


Performance Comparison

Experimental results show that the **transfer learning model based on VGG19** outperforms the CNN trained from scratch in both accuracy and convergence speed. The pre-trained model achieves faster convergence and demonstrates superior generalization on unseen test data. These findings are consistent with existing literature, which reports that transfer learning models generally outperform custom CNNs for satellite image classification tasks [1], [3].


$$\text{Accuracy} = \frac{\sum_{i=1}^N 1(\hat{y}_i = y_i)}{N}$$



Visualize the Train Loss/Accuracy and the Test Loss/Accuracy (VGG19)




Visualize the Train Loss/Accuracy and the Test Loss/Accuracy (VGG-Inspired CNN (2nd model))



Visualize the Train Loss/Accuracy and the Test Loss/Accuracy (VGG-Inspired CNN (3rd model))

- **VGG-19:** Accuracy ~61–62%, F1 ~0.60, overfitting evident.
- **VGG-Inspired 2nd model:** Accuracy ~91%, F1 ~0.90, good generalization.
- **VGG-Inspired 3rd model:** Accuracy ~93%, F1 ~0.92, best overall performance.



DISCUSSION

The experimental results underscore the benefits of **transfer learning** for satellite image classification tasks, particularly for small datasets like EuroSAT. While training a CNN from scratch provides architectural flexibility and simplicity, it often requires larger datasets and longer training times to reach competitive performance. In contrast, transfer learning leverages pre-trained feature representations from large-scale datasets, leading to improved accuracy, faster convergence, and better generalization on unseen test data. This aligns with prior findings in computer vision applications for remote sensing, where pre-trained CNNs effectively capture features relevant to land use and land cover classification [1], [3].

Despite these advantages, pre-trained models such as **VGG19** have higher computational and memory requirements, which may limit their use in resource-constrained environments. Additionally, standard CNN architectures, including the custom VGG-inspired models proposed in this study, are designed primarily for RGB image classification and may need adaptations to fully exploit multi-source satellite data, such as multispectral imagery or temporal sequences.

From an applied perspective, the insights learned from EuroSAT classification models have **direct implications for economic activity estimation**. Accurate land use and land cover classification can serve as proxy features for modeling infrastructure density, urbanization, and agricultural productivity, which are key

indicators of local economic activity. For example, correctly identifying industrial, residential, or agricultural areas enables the construction of derived features that could feed into higher-level economic models without requiring extensive ground-truth data.

Future work could explore more **efficient architectures** such as EfficientNet or Vision Transformers to maintain high accuracy while reducing computational costs. Integration of segmentation-based features (e.g., buildings, roads, vegetation) could further enhance interpretability and provide actionable insights for economic analyses. Additionally, leveraging **active learning or self-supervised pretraining** could mitigate the limitations of small labeled datasets, a challenge common to both remote sensing classification and economic modeling tasks.

In summary, the results indicate that **transfer learning and model fine-tuning are critical for accurate land use classification on small satellite datasets**, and that these models can provide **informative feature representations** for downstream applications, including economic activity estimation. By bridging feature extraction from satellite imagery with potential economic indicators, the proposed approach demonstrates a pathway to efficient and interpretable remote sensing analysis.

Applications

The proposed deep learning approaches can be applied to various real-world remote sensing tasks, including:

- Urban expansion monitoring
- Agricultural land analysis
- Environmental change detection
- Water resource management
- Disaster assessment

CONCLUSION

This study presented a comparative analysis of two deep learning approaches for land use and land cover classification using the EuroSAT dataset: a CNN trained from scratch and a **transfer learning model based on VGG19**. Experimental results demonstrate that transfer learning significantly improves classification accuracy and training efficiency. The findings confirm the effectiveness of pre-trained deep learning models for remote sensing image classification and provide a strong foundation for future research in this domain.

REFERENCES

1. P. Helber, B. Bischke, A. Dengel, and D. Borth, "EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 12, no. 7, pp. 2217–2226, 2019.
2. J. Terven, A. Smith, and L. Johnson, "Deep learning approaches for satellite image classification," *Remote Sensing*, vol. 13, no. 10, pp. 1987, 2021.

Source for review

1. <https://colab.research.google.com/drive/1ieouHAQ7KDjYpHH6ZwPq16gOBe2bUtGr> -VGG 19
2. <https://colab.research.google.com/drive/1XBIEUjy9RpMRITce-sUKRFRWzVcJ5bM1>-VGG-Inspired CNN (2nd model)
3. <https://colab.research.google.com/drive/1hAKkpWDWsfBvuNI2h36L3QoHdpBInexi> - -VGG-Inspired CNN (3rd model)