UDC 004.4

# FINE-TUNED ALEXNET FOR ROOF SHAPE CLASSIFICATION IN UZBEKISTAN: A TRANSFER LEARNING APPROACH

***Yuldashev S.U.***
ysaidkarim@gmail.com
Digital Technologies and Artificial Intelligence Development Research Institute,
17A, Buz-2, Tashkent, 100125 Uzbekistan.

This study investigates the classification of rooftop shapes using convolutional neural networks (CNNs), with a particular focus on regional adaptation through transfer learning. Initial training of the AlexNet architecture utilizes a publicly available Zenodo dataset comprising satellite imagery of flat, gabled, and hipped roofs. To address generalizability constraints in specific geographic contexts, a custom dataset of Uzbekistan rooftops, sourced from OpenStreetMap and high-resolution Mapbox Static Images API tiles, enables fine-tuning. Experimental outcomes reveal enhanced classification accuracy for flat, gabled, and hipped roofs, underscoring the efficacy of transfer learning in mitigating domain shift due to architectural and environmental variations. Integration of open-source geospatial tools with transfer learning offers a replicable framework for addressing geographic bias in rooftop shape classification, adaptable to other underrepresented regions.

**Keywords:** roof shape classification, convolutional neural networks, transfer learning, AlexNet, satellite imagery.

**Citation:** Yuldashev S.U. 2025. Fine-tuned AlexNet for roof shape classification in Uzbekistan: a transfer learning approach. *Problems of Computational and Applied Mathematics.* 3(67): 133-146.

**DOI:** https://doi.org/10.71310/pcam.3_67.2025.12.

## 1 Introduction

Automated classification of roof shapes has become a vital component in various domains, including urban planning, disaster risk assessment, solar panel installation, and insurance modeling [6, 7]. With the proliferation of satellite imagery and advances in deep learning, particularly convolutional neural networks (CNN), considerable progress has been made in the analysis of built environments from aerial and remote sensing data [1, 8].

However, pretrained CNNs and existing public datasets—such as those from Inria, Wuhan University (WHU), or SpaceNet—are predominantly developed using imagery from Western urban environments. Consequently, these models often exhibit limited generalization when applied to underrepresented regions like Uzbekistan, where architectural styles, building densities, and environmental textures differ significantly [2, 9, 10].

To address this domain gap, we leverage transfer learning techniques to adapt a pretrained AlexNet model [1] to the local context of Uzbekistan. We begin by training the model on the global roof shape classification dataset hosted on Zenodo [2], consisting of flat, gabled, and hipped roofs. To localize the model, we created a new dataset by extracting rooftop geometries from OpenStreetMap [4] and retrieving high-resolution satellite imagery using the Mapbox Static Images API [5]. These localized images were manually reviewed, categorized, and used to fine-tune the CNN model.

The primary objective of this study is to assess how transfer learning improves classification performance on localized satellite imagery. We hypothesize that adapting CNNs to region-specific data will significantly enhance the accuracy of roof shape detection in geographies that are otherwise poorly represented in global datasets.

## 2 Related Work

Recent developments in deep learning and computer vision have significantly advanced the analysis of remote sensing imagery for geospatial tasks [8]. Convolutional neural networks (CNNs) have been particularly effective, with well-established architectures such as AlexNet [1], VGGNet [11], and ResNet [12] achieving strong results in various domains, including building footprint detection, road extraction, and scene classification.

In the specific domain of rooftop analysis, several datasets—such as the Inria Aerial Image Labeling dataset, WHU Building Dataset, and SpaceNet—have been widely used for training and benchmarking models [7, 10]. However, these datasets are largely composed of imagery from high-income and Western urban areas, which introduces significant biases when models are applied to less represented regions such as Central Asia. This geographic imbalance leads to a domain shift, often resulting in poor generalization performance on structurally different urban layouts [9].

To address this issue, the use of transfer learning has gained traction. By reusing pretrained CNN weights and fine-tuning on target-domain data, researchers have successfully adapted models to novel geographic contexts with minimal labeled data [2, 9]. In remote sensing, such strategies have proven effective for land use classification, disaster monitoring, and building detection [8, 13].

Our contribution is novel in two aspects. First, we focus explicitly on rooftop classification in Uzbekistan—a region absent from existing roof shape datasets. Second, we build the fine-tuning dataset using open-source geospatial resources, namely the Mapbox Static Images API [5] and OpenStreetMap vector geometries [4]. This not only ensures reproducibility but also demonstrates how public tools can be combined with transfer learning to enhance model performance in underrepresented regions.

## 3 Dataset and Preprocessing

### 3.1 Zenodo Roof Dataset

The initial model training was conducted using a publicly available dataset hosted on Zenodo [2], which contains labeled satellite imagery of rooftops across various global regions. This dataset includes three primary rooftop geometries commonly found in residential architecture: flat, gabled, and hipped roofs as shown in Figure 1.

To prepare the data for input into the neural network, all images were resized to $256\times\times 256$ pixels, normalized to the range $[0, 1]$, and the class labels were converted to one-hot encoded vectors suitable for categorical classification tasks.

Normalization was performed as follows for each pixel value $p \in [0, 255]$:

$$p_{\text{normalized}} = \frac{p}{255}. \tag{1}$$

This transformation ensures that the pixel values are distributed uniformly and helps the network converge faster during training.

(a) Flat roof (b) Gabled roof (c) Hipped roof

**Figure 1** Examples of different roof types from the Zenodo dataset.

# 4 Uzbekistan Roof Imagery

To assess the generalization capabilities of the model on region-specific data, we created a custom dataset of rooftops located in various urban zones of Uzbekistan as shown in Figure 2.

(a) Flat roof (b) Gabled roof (c) Hipped roof

**Figure 2** Examples of different roof types from the Uzbekistan dataset.

## 4.1 Geolocation Data Extraction

Rooftop geometries were sourced from OpenStreetMap (OSM) using the Quantum Geographic Information System (QGIS) application and the QuickOSM plugin. The extracted data included multipolygon representations of buildings in GeoJSON format, where each rooftop is described as a closed polygon or a set of polygons.

## 4.2 Centroid Calculation

To retrieve satellite imagery centered over each rooftop, the centroid of each polygon was computed. For a polygon defined by $N$ vertices with coordinates $(x_1, y_1), (x_2, y_2), \ldots,$ $(x_N, y_N)$, the centroid $C = (\bar{x}, \bar{y})$ is calculated using the arithmetic mean:

$$\bar{x} = \frac{1}{N} \sum_{i=1}^{N} x_i, \quad \bar{y} = \frac{1}{N} \sum_{i=1}^{N} y_i. \tag{2}$$

This centroid location represents the average spatial position of the rooftop and was used as the anchor point for image retrieval.

## 4.3 Satellite Image Acquisition

Each centroid was passed to the Mapbox Static Images API [5], which returns high-resolution satellite imagery tiles at a specified zoom level. We selected zoom level 19, which provides building-scale resolution. To convert geographic coordinates (latitude $\varphi$, longitude $\lambda$) into tile coordinates $(x, y)$ at zoom level $z$, we used the following standard Web Mercator projection formulas:

$$n = 2^z, \quad x = \left\lfloor \frac{\lambda + 180}{360} \cdot n \right\rfloor, \quad y = \left\lfloor \frac{1 - \log(\tan(\varphi) + \sec(\varphi))/\pi}{2} \cdot n \right\rfloor, \tag{3}$$

where:

- $\tan(\varphi)$ is the tangent of latitude,
- $\sec(\varphi) = \frac{1}{\cos(\varphi)}$ is the secant of latitude,
- $\lfloor \cdot \rfloor$ denotes the floor function.

The resulting tiles were downloaded and saved into class-labeled folders (e.g., flat/, gabled/, hipped/) based on manual inspection. Only those tiles with minimal occlusions, visible rooftops, and correct geometries were retained to ensure label integrity. The final dataset was approximately balanced across the three classes.

### 4.4 Data Augmentation

To enhance the generalization capabilities of the model and mitigate potential overfitting due to limited training samples, we applied data augmentation techniques. These augmentations artificially expand the training set by introducing variations in the input images while preserving semantic content. The following transformations were applied only to the training set:

- Random rotations within ±30°, simulating satellite orientation noise,
- Horizontal flips, mimicking viewing angles from opposite directions,
- Width and height shifts up to 20% of the image size, accounting for positioning errors,
- Zooming and shearing (20%), introducing slight geometric distortions.

These techniques preserve the underlying class semantics while allowing the network to learn invariance to typical image perturbations observed in satellite imagery.

## 5 Methodology

### 5.1 Model Architecture

For the task of rooftop shape classification, we selected AlexNet [1] as the base convolutional neural network (CNN) architecture. Originally introduced by Krizhevsky et al. for the ImageNet Large Scale Visual Recognition Challenge (ILSVRC), AlexNet has proven to be effective in image classification tasks across a variety of domains, due to its ability to extract hierarchical features from moderately sized datasets [1, 14].

The AlexNet architecture comprises eight learnable layers, organized as follows:

- **Five convolutional layers** that apply 2D kernels across the input image to extract local spatial features. Each convolutional layer is followed by ReLU activation to introduce non-linearity, and batch normalization to stabilize and accelerate convergence.
- **Three max-pooling layers**, interleaved between convolutional layers, downsample the feature maps by selecting the maximum value within a local window. This reduces spatial dimensions and provides translation invariance.
- **Two fully connected (dense) layers**, each with a large number of neurons (typically 4096), aggregate features across the entire image and serve as high-level classifiers.
- **A final softmax output layer**, which produces a probability distribution over the three target classes: flat, gabled, and hipped roofs.

To mitigate overfitting and encourage generalization, we applied *dropout regularization* between the dense layers. In this approach, during training, a proportion $p$ of neuron activations are randomly set to zero. We used a dropout rate of $p = 0.6$, following the empirical recommendations for medium-scale datasets [14].

Additionally, *L2 weight regularization* (also known as weight decay) was employed. This technique adds a penalty term to the loss function that discourages large weight magnitudes, and is expressed as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cross-entropy}} + \lambda \sum_{i=1}^{n} w_i^2, \tag{4}$$

where:

- $\mathcal{L}_{\text{cross-entropy}}$ is the primary classification loss (see Section 12),
- $w_i$ are the learnable weights of the network,
- $\lambda$ is the regularization strength hyperparameter.

This formulation biases the model toward simpler hypotheses, reducing the risk of memorizing noise in the training data. The combination of convolutional filtering, pooling, non-linearity, regularization, and fully connected layers makes AlexNet a powerful yet computationally tractable model for visual recognition tasks involving moderate-resolution satellite images [8, 11].

### 5.2 Training Strategy

The initial training of the AlexNet model [1] was conducted using the global roof shape dataset from Zenodo [2], which served as the source domain. The model was trained from scratch—i.e., all weights were initialized randomly—using the Adam optimizer [14], which combines the advantages of both Adaptive Gradient Algorithm (AdaGrad) and Root Mean Square Propagation (RMSProp).

Adam updates network parameters based on adaptive estimates of lower-order moments and is defined by the following update rule:

$$\theta_t = \theta_{t-1} - \alpha \cdot \frac{\hat{m}_t}{\sqrt{\hat{v}_t} + \varepsilon}, \tag{5}$$

where:

- $\theta_t$ are the model parameters at time step $t$,
- $\alpha$ is the learning rate (set to 0.0001),
- $\hat{m}_t$ and $\hat{v}_t$ are the bias-corrected estimates of the first and second moments of the gradient,
- $\varepsilon$ is a small constant to prevent division by zero.

The loss function used was categorical cross-entropy, suitable for multi-class classification where labels are one-hot encoded. It is computed as:

$$\mathcal{L}_{\text{CE}} = -\sum_{i=1}^{C} y_i \log(\hat{y}_i), \tag{6}$$

where:

- $C$ is the number of classes,
- $y_i$ is the ground truth label (1 if the class is correct, 0 otherwise),
- $\hat{y}_i$ is the predicted probability for class $i$.

Training was performed for up to 50 epochs. To prevent overfitting and optimize computational efficiency, we applied early stopping, which halts training if the validation loss does not improve for five consecutive epochs.

In addition, we implemented learning rate scheduling via ReduceLROnPlateau. When the validation loss plateaued, the learning rate was reduced by a factor of 0.5. A minimum learning rate threshold of $\alpha_{\min} = 1 \times 10^{-6}$ was enforced to prevent excessively small updates.

This strategy ensures both stable convergence and adaptive training dynamics, which are critical when training CNNs on moderately sized remote sensing datasets [8].

### 5.3 Class imbalance handling.

To address the significant class imbalance in the Uzbekistan roof shapes dataset—where *hipped* roofs constituted the majority class—we computed class weights using the `compute_class_weight` function from `scikit-learn`. These weights were incorporated into model training via the `class_weight` parameter in the `fit()` function to penalize underrepresented classes more heavily and promote balanced learning.

The class distribution in the training set was as follows:

- **Flat**: 104 samples
- **Gable**: 263 samples
- **Hipped**: 742 samples

The resulting class weights were:

`{0: 3.55, 1: 1.41, 2: 0.50}`.

This reweighting strategy helped mitigate bias toward the dominant class and improved the model's ability to correctly identify underrepresented roof types, particularly *flat* and *gable* categories.

### 5.4 Evaluation on Uzbekistan Data

To assess the model's generalization capability, the pretrained network—trained solely on the Zenodo dataset—was evaluated on the custom dataset of rooftop images from Uzbekistan, introduced in Section **??**. This step was crucial for simulating a domain shift scenario, where the model is exposed to imagery from a different geographic and architectural context.

We employed four standard metrics for multi-class classification evaluation:

- **Accuracy**: the proportion of correctly predicted samples.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}. \tag{7}$$

- **Precision (per class)**: the proportion of true positive predictions over all predicted positives.

$$\text{Precision} = \frac{TP}{TP + FP}. \tag{8}$$

- **Recall (per class)**: the proportion of true positive predictions over all actual positives.

$$\text{Recall} = \frac{TP}{TP + FN}. \tag{9}$$

- **F1-score**: the harmonic mean of precision and recall.

$$F1 = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}. \tag{10}$$

The results highlighted a significant drop in performance, confirming the presence of a strong domain gap due to differences in roof geometries, illumination, and texture between the Zenodo and Uzbekistan datasets [9]. This underperformance motivated the use of fine-tuning, detailed in Section 12.

### 5.5 Fine-tuning on Uzbekistan Roofs

To overcome the limitations of domain shift, we employed a transfer learning strategy by fine-tuning the pretrained AlexNet model [1] on the Uzbekistan specific rooftop dataset. Transfer learning allows the reuse of knowledge from a source domain (global roof imagery) and adaptation to a target domain (Uzbekistan architectural context), even when the latter has fewer labeled samples [9, 13].

The fine-tuning procedure included the following key steps:

- **Loading pretrained AlexNet weights:** The model trained on the Zenodo dataset served as the initialization point. The early convolutional layers, which extract general low-level features (edges, textures), were retained.
- **Freezing convolutional layers:** To preserve learned general features, all convolutional layers were frozen, meaning their weights were not updated during fine-tuning.
- **Replacing fully connected layers:** The top classifier layers were removed and replaced with new trainable dense layers, customized for the three-class roof classification task. This part of the network learns features specific to the Uzbekistan dataset.
- **Using a reduced learning rate:** A small learning rate of $\alpha = 1 \times 10^{-5}$ was used to prevent catastrophic forgetting and ensure gradual adaptation. The model was fine-tuned for 25 epochs, with early stopping and learning rate reduction enabled (see Section 12).

This approach leverages the feature extraction capabilities of a globally trained CNN while allowing selective learning of domain-specific representations, in line with established practices in geospatial deep learning [6, 7].

### 5.6 Training Curve Analysis

To monitor training progress and assess model behavior, we recorded and visualized both training and validation metrics—accuracy and loss—across all epochs for the baseline and fine-tuned models. The results are shown in Figures 3 and 4.

The training curves show marked differences in learning dynamics:

**Baseline Model (Figure 3)**

- **Overfitting is evident:** while training accuracy steadily improves and surpasses 85%, validation accuracy remains significantly lower and fluctuates throughout the epochs.
- **Validation instability:** the validation accuracy curve shows abrupt changes and lacks a consistent upward trend, indicating poor generalization and possible domain mismatch with the test data.
- **Loss convergence:** although both training and validation loss decrease smoothly and nearly identically, the discrepancy between accuracy metrics highlights the model's limited applicability to local conditions.

Training line graph of the beaseline model

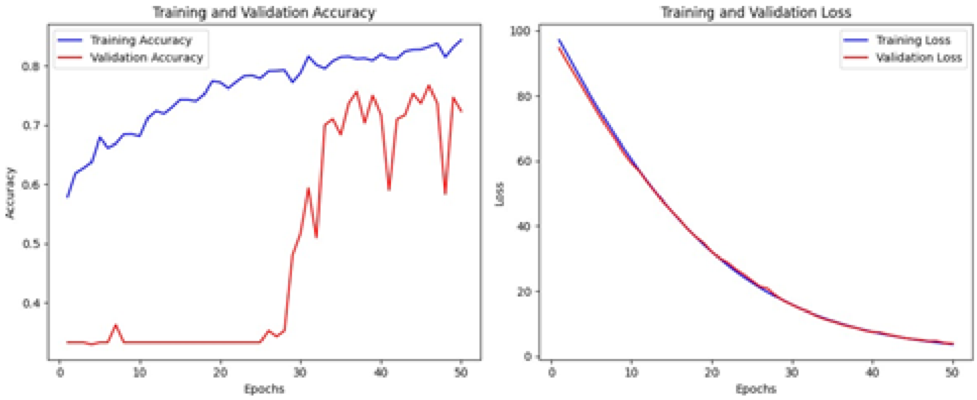


**Figure 3** Training and validation accuracy and loss curves of the baseline AlexNet trained exclusively on Zenodo data.

### Fine-tuned Model (Figure 4)

- Rapid adaptation: the model demonstrates fast convergence, with training accuracy increasing sharply within the first few epochs—leveraging the general knowledge gained from global data.
- Improved generalization: both training and validation accuracies stabilize around 60–65%, considerably outperforming the baseline model on local Uzbekistan rooftop imagery.
- Stable learning behavior: the parallel decline of training and validation loss curves indicates consistent optimization and reduced overfitting on the region-specific dataset.

  Training line graph of the fine-tuned model

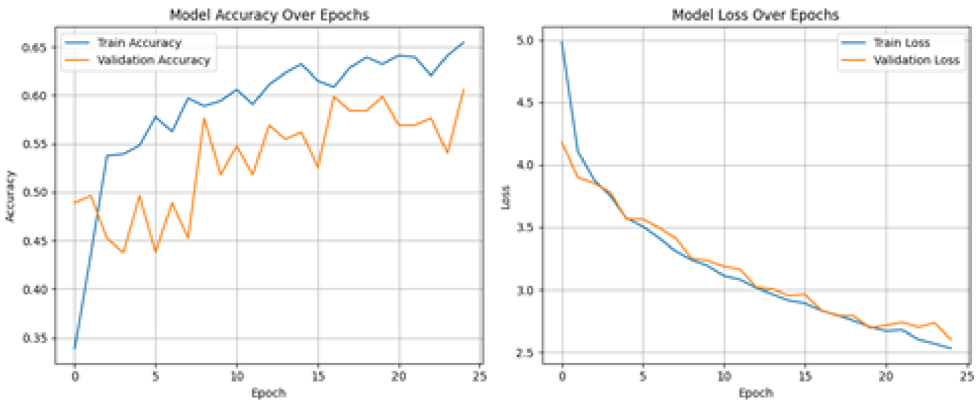


**Figure 4** Training and validation curves of the fine-tuned AlexNet model on the Uzbekistan dataset.

These patterns confirm that fine-tuning significantly improves the model's adaptability and performance on regional data, validating the use of transfer learning for underrepresented geographies [8].

## 6 Results

### 6.1 Initial Evaluation on Uzbekistan Dataset

To assess the baseline generalization capability of the model, we evaluated the original AlexNet architecture trained exclusively on the global Zenodo dataset (see Section 12) against the Uzbekistan specific test set. This test was designed to simulate a cross-domain generalization scenario, where the model encounters data from an unseen geographic and architectural context.

As expected, the model's performance degraded significantly. The evaluation yielded low accuracy and weak class-level performance metrics, as shown in Table 1. These results confirm the presence of a domain shift [9]—differences in rooftop morphology, image resolution, and environmental features between the global training data and local Uzbekistan images—leading to model misclassification and overgeneralization.

**Table 1** Baseline evaluation performance on Uzbekistan test set

| Metric | Value |
|---|---|
| Accuracy | 46.80% |
| Precision | 40.05% |
| Recall | 46.80% |
| F1-score | 43.11% |

### 6.2 Fine-tuning Performance

To address the above limitations, the model was fine-tuned using the Uzbekistan training set, following the strategy outlined in Section 12. During fine-tuning, the convolutional base was frozen, while the classifier layers were retrained with a reduced learning rate to capture region-specific features.

After 25 epochs, the model achieved the following performance on the held-out Uzbekistan test set:

**Table 2** Fine-tuned model performance on Uzbekistan test set

| Metric | Value |
|---|---|
| Accuracy | 70.21% |
| Precision | 73.87% |
| Recall | 70.21% |
| F1-score | 71.68% |

The results demonstrate a dramatic improvement in all evaluation metrics, affirming that transfer learning with regional adaptation is critical to improving prediction accuracy for satellite imagery in underrepresented regions [6, 13].

### 6.3 Comparison of the Baseline and Fine-tuned Models

To quantify the impact of transfer learning, we conducted a side-by-side comparison of the baseline model and the fine-tuned version. Both models were evaluated on the same Uzbekistan test set. Performance metrics for each are summarized in Table 3.

**Table 3** Comparison of baseline and fine-tuned models on Uzbekistan test set

| Model | Accuracy | Precision | Recall | F1-score |
|---|---|---|---|---|
| Baseline | 46.80 | 40.05 | 46.80 | 43.11 |
| Fine-tuned | 70.21 | 73.87 | 70.21 | 71.68 |

Confusion matrices for visualizing class-specific prediction accuracy are shown in Figure 5.

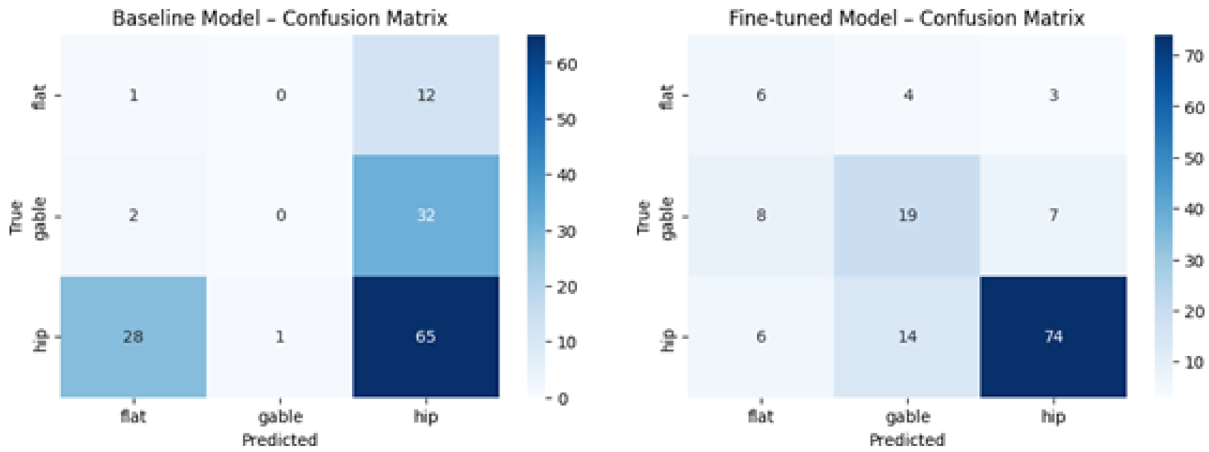


(a) Baseline model

(b) Fine-tuned model

**Figure 5** Confusion matrices of the baseline and fine-tuned models on the Uzbekistan test set.

The baseline model, trained solely on the global Zenodo dataset, achieved moderate performance with an F1-score of 43.11%. While it was relatively effective in recognizing hipped roofs (65 correct predictions), it failed to generalize to flat and gable roofs, misclassifying nearly all of them as hipped. This reflects a substantial class imbalance in predictions and a lack of adaptation to the local architectural patterns of Uzbekistan.

In contrast, the fine-tuned model, which was further trained on a region-specific dataset of Uzbekistan rooftops, showed clear improvements across all evaluation metrics. It achieved an F1-score of 71.68%, with balanced performance across the three classes. Particularly notable is the improvement in the gable category, where the model went from zero correct predictions (baseline) to 19 correctly identified instances (fine-tuned).

The confusion matrices in Figure 5 highlight these differences:

- The baseline confusion matrix is dominated by misclassifications, especially for the flat and gable classes, which are overwhelmingly predicted as hipped.
- The fine-tuned matrix exhibits strong diagonal dominance, indicating significantly improved class-wise accuracy and a more reliable model for deployment in real-world scenarios.

These findings reinforce the critical role of domain adaptation in computer vision for remote sensing. Fine-tuning on geographically relevant data proves essential for achieving robust and balanced performance, particularly in underrepresented regions such as Central Asia [12].

## 7 Applications and Use Cases

The fine-tuned AlexNet model developed in this study has direct utility across a range of real-world domains, particularly in geospatial intelligence, urban planning, and disaster management. Its ability to accurately classify rooftop shapes using satellite imagery enables scalable, automated extraction of structural information in contexts where manual surveys are infeasible.

Key application areas include:

- **Urban planning and infrastructure monitoring.** Accurate mapping of rooftop types aids city authorities in assessing residential density, building typology, and zoning compliance. For instance, gabled and hipped roofs may correlate with older housing stock, while flat roofs often appear in newer or industrial buildings. Automated classification can support infrastructure development and housing policy formulation [15].

- **Renewable energy deployment.** Identifying flat rooftops is particularly valuable for solar panel placement and optimization of photovoltaic potential. Using the model's outputs, energy agencies or private firms can prioritize solar installations in densely built-up zones without requiring costly site visits [16].
- **Disaster response and resilience planning.** Roof shapes are indicators of vulnerability to natural hazards. For example, hipped roofs generally perform better in high-wind events than flat roofs. The model can also be used to compare pre- and post-disaster satellite imagery, enabling rapid damage assessment and resource allocation following earthquakes, floods, or cyclones [17, 18].
- **Geospatial mapping and open-source cartography.** Platforms such as OpenStreetMap often lack rooftop attribute data in developing regions. This model can assist contributors by generating automated roof annotations, enhancing the richness of public geospatial datasets and supporting crowdsourced mapping efforts [18, 19].
- **Insurance and structural risk analysis.** In the insurance industry, roof geometry is a known risk factor in property underwriting. Flat roofs, for example, are more prone to water pooling and structural stress. The proposed classification system can be integrated into risk assessment pipelines to automate premium estimation and claim validation [21, 22].

By leveraging openly accessible data sources (Mapbox, OpenStreetMap) and lightweight CNN architectures, this system demonstrates how AI and remote sensing can be combined to create scalable and reproducible pipelines for urban analytics, particularly in underrepresented and low-resource regions.

# 8 Conclusion

This study presents a transfer learning-based approach for rooftop shape classification using satellite imagery, with a specific emphasis on regional adaptation to the architectural landscape of Uzbekistan. Starting from a pre-trained AlexNet model trained on a globally distributed dataset, we fine-tuned the network using a custom-curated dataset derived from OpenStreetMap and high-resolution Mapbox imagery. The results demonstrate that fine-tuning significantly improved the model's performance on previously unseen, geographically specific data, with accuracy rising from 46.80% to 70.21%, and substantial gains in precision, recall, and F1-score.

These findings underscore the importance of domain adaptation in remote sensing tasks. Models trained on global datasets often underperform in underrepresented regions due to differences in architectural typology, data quality, and environmental context. By leveraging open-source data and lightweight CNN architectures, our approach provides a scalable solution for accurate rooftop classification in data-scarce environments.

Despite the strong performance, several limitations warrant further attention. The manually curated Uzbekistan dataset remains limited in geographic and temporal scope, and the resolution of Mapbox imagery restricts the capture of fine-grained architectural details. The reliance on a relatively dated model architecture (AlexNet) also suggests room for improvement through more advanced networks such as EfficientNet or Vision Transformers. Additionally, the model operates purely on 2D imagery, omitting potentially valuable 3D data such as elevation or slope.

Future work should consider expanding the dataset across multiple regions and urban contexts, integrating elevation models or LiDAR data, and employing more recent deep learning architectures. Addressing label noise and incorporating semi-supervised or crowdsourced labeling strategies could further enhance model robustness. Operational deployment will also require investigation into system latency, scalability, and integration with GIS infrastructures.

In summary, this research demonstrates the feasibility and effectiveness of transfer learning for rooftop classification in underrepresented regions, while highlighting technical and practical pathways for enhancing generalizability and real-world applicability in future work.

## References

[1] Krizhevsky A., Sutskever I., Hinton G.E. 2012. *Imagenet classification with deep convolutional neural networks.* Advances in Neural Information Processing Systems, 25,

[2] Roof shape classification dataset `https://zenodo.org/record/3986721`,

[3] Roof shapes of Uzbekistan dataset, `https://drive.google.com/file/d/1jccNLjz3Tb3oMrrgb_HGHl_TIJprAq72/view?usp=sharing`, OSM

[4] OpenStreetMap contributors. Planet dump retrieved from https://planet.openstreetmap.org, n.d.

[5] OpenStreetMap contributors 2007. Planet dump retrieved from https://planet.openstreetmap.org. https://www.openstreetmap.org

[6] Chollet F. 2017. Deep Learning with Python. *Manning Publications.* Greenwich.

[7] Simonyan K. and Zisserman A. 2015. *Very deep convolutional networks for large-scale image recognition.* volume arXiv:1409.1556, https://arxiv.org/abs/1409.1556

[8] Yuldashev S. 2025. *Roof Shape Classification in Uzbekistan Using Fine-tuned AlexNet.* Google Colab Notebook, https://colab.research.google.com/drive/1tlQBzS7ALSqkhe12q8yYW_1AGHmUT-Sla?usp=sharing

[9] He K., Zhang X., Ren S., Sun J. 2016. *Deep residual learning for image recognition.* Proc. IEEE Conf. Comput. Vis. Pattern Recognit. – P. 70–778.

[10] Pritt M., Chern G. 2020. *Satellite image classification with deep learning.* arXiv preprint,

[11] Simonyan K. and Zisserman A. 2015. *Very deep convolutional networks for large-scale image recognition. arXiv preprint.* arXiv:1409.1556,

[12] He K., Zhang X., Ren S. and Sun J. 2016. *Deep residual learning for image recognition.* In Proc. IEEE Conf. Comput. Vis. Pattern Recognit., – P. 770–778.

[13] Chroni A., Vasilakos C., Christaki M. and Soulakellis N. 2024. Fusing multispectral and lidar data for cnn-based semantic segmentation in semi-arid mediterranean environments. Remote Sensing, 16(15): 2729,

[14] Chollet 2017. *Deep Learning with Python.* Manning Publications, Greenwich, CT,

[15] Ni J., Chen J., Wu Y., Chen Z., Liang M. 2022. Method to determine the centroid of non-homogeneous polygons based on suspension theory. ISPRS Int. J. Geo-Inf., 11(4): 233,

[16] Zhao Y., Li X., Wang Z., Chen H. and Liu J. 2025. Deep-learning-based evaluation of rooftop photovoltaic deployment in tianjin, china. ISPRS International Journal of Geo-Information, 14(3):101,

[17] Jie Gu, Zhen Xie, Jiawei Zhang, and Xiaohong He. 2024. *Advances in rapid damage identification methods for post-disaster regional buildings based on remote sensing images: A survey.* Buildings, 14(4): 898,

[18] Haoyu Hao, Sreeharsha Baireddy, Ethan R. Bartusiak, Robert T. Collins, and Michael K. Hinders. 2020. *An attention-based system for damage assessment using satellite imagery.* arXiv preprint arXiv:2004.06643,

[19] Jeremy D. Castagno and Ella M. Atkins. 2018. *Roof shape classification from lidar and satellite image data.* Sensors, 18(11):3960,

[20] John E. Vargas-Muñoz et al. 2019. *Correcting rural building annotations in openstreetmap using convolutional neural networks.* arXiv preprint arXiv:1901.08190,

[21] Shuochuan Meng, Mohammad Hesam Soleimani-Babakamali, and Ertugrul Taciroglu. 2023. *Au- tomatic roof type classification through machine learning for regional wind risk assessment.* arXiv preprint arXiv:2305.17315,

[22] He K., Zhang X., Ren S., and J. Sun. 2016. *Deep residual learning for image recognition.* In Proc. IEEE Conf. Comput. Vis. Pattern Recognit., – P. 770–778.

[23] Robins Kaplan LLP. 2023. Ai's impact on property insurance cov- erage. https://www.robinskaplan.com/newsroom/insights/ ais-impact-on-property-insurance-coverage, Accessed: 2025-05-24.

[24] Yuldashev 2025. *Roof shape classification in Uzbekistan using fine-tuned alexnet* Google Colab Notebook.

*Received May 30, 2025*

УДК 004.4

# ТОНКАЯ НАСТРОЙКА ALEXNET ДЛЯ КЛАССИФИКАЦИИ ФОРМ КРЫШ В УЗБЕКИСТАНЕ: ПОДХОД С ИСПОЛЬЗОВАНИЕМ ТРАНСФЕРНОГО ОБУЧЕНИЯ

***Юлдашев С.У.***

ysaidkarim@gmail.com

Научно-исследовательский институт развития цифровых технологий и искусственного интеллекта,

100125, Узбекистан, г. Ташкент, Мирзо-Улугбекский р-он, м-в Буз-2, д. 17А.

В данной работе исследуется задача классификации форм крыш с использованием сверточных нейронных сетей (CNN), с особым вниманием к региональной адаптации с помощью переноса обучения. Первоначальное обучение архитектуры AlexNet проводится на общедоступном наборе данных Zenodo, содержащем спутниковые изображения плоских, двускатных и вальмовых крыш. Для устранения ограничений обобщающей способности в специфических географических контекстах создан пользовательский набор данных крыш Узбекистана, полученный из OpenStreetMap и тайлов высокого разрешения через API Mapbox Static Images. Точная настройка на этом наборе данных повышает точность классификации плоских, двускатных и вальмовых крыш, подтверждая эффективность трансферного обучения в смягчении доменного сдвига, вызванного архитектурными и экологическими различиями. Интеграция инструментов геопространственного анализа с открытым исходным кодом и трансферного обучения предоставляет воспроизводимую основу для устранения географической предвзятости в задачах классификации форм крыш, адаптируемую для других недостаточно представленных регионов.

**Ключевые слова:** классификация форм крыш, сверточные нейронные сети, трансферное обучение, AlexNet, спутниковые снимки.

**Цитирование:** *Юлдашев С.У.* Тонкая настройка AlexNet для классификации форм крыш в Узбекистане: подход с использованием трансферного обучения // Проблемы вычислительной и прикладной математики. – 2025. – № 3(67). – С. 133-146.

**DOI:** https://doi.org/10.71310/pcam.3_67.2025.12.

№ 3(67) 2025

ISSN 2181-8460

# HISOBLASH VA AMALIY MATEMATIKA MUAMMOLARI

# ПРОБЛЕМЫ ВЫЧИСЛИТЕЛЬНОЙ И ПРИКЛАДНОЙ МАТЕМАТИКИ

# PROBLEMS OF COMPUTATIONAL AND APPLIED MATHEMATICS

RAQAMLI TEXNOLOGIYALAR VA
SUN'IY INTELLEKTNI RIVOJLANTIRISH
ILMIY-TADQIQOT INSTITUTI

# ПРОБЛЕМЫ ВЫЧИСЛИТЕЛЬНОЙ И ПРИКЛАДНОЙ МАТЕМАТИКИ

## № 3(67) 2025

Журнал основан в 2015 году.
Издается 6 раз в год.

**Учредитель:**
Научно-исследовательский институт развития цифровых технологий и искусственного интеллекта.





Журнал зарегистрирован в Агентстве информации и массовых коммуникаций при Администрации Президента Республики Узбекистан.
Регистрационное свидетельство №0856 от 5 августа 2015 года.

**ISSN 2181-8460, eISSN 2181-046X**





**Адрес редакции:**
100125, г. Ташкент, м-в. Буз-2, 17А.
Тел.: +(998) 712-319-253, 712-319-249.
Э-почта: `journals@airi.uz`.
Веб-сайт: `https://journals.airi.uz`.

**Дизайн и вёрстка:**
Шарипов Х.Д.

Отпечатано в типографии НИИ РЦТИИ.
Подписано в печать 30.06.2025 г.
Формат 60x84 1/8. Заказ №5. Тираж 100 экз.

# PROBLEMS OF COMPUTATIONAL AND APPLIED MATHEMATICS

## No. 3(67) 2025

The journal was established in 2015.
6 issues are published per year.

**Founder:**
Digital Technologies and Artificial Intelligence Development Research Institute.



The journal is registered by Agency of Information and Mass Communications under the Administration of the President of the Republic of Uzbekistan.
The registration certificate No. 0856 of 5 August 2015.

**ISSN 2181-8460, eISSN 2181-046X**



**Address:**
100125, Tashkent, Buz-2, 17A.
Tel.: +(998) 712-319-253, 712-319-249.
E-mail: journals@airi.uz.
Web-site: https://journals.airi.uz.

**Layout design:**
Sharipov Kh.D.

DTAIDRI printing office.
Signed for print 30.06.2025
Format 60x84 1/8. Order No. 5. Print run of 100 copies.

# Содержание

# Contents