INNOVATIVE MECHANICAL ENGINEERING **ISSN 2812-9229 (Online)**
University of Niš, Faculty of Mechanical Engineering
VOL.5, NO 1, 2026, PP. 92-106

Original scientific paper

APPLICATION OF LOGISTIC REGRESSION WITH HOG FEATURES FOR GEOMETRIC SHAPE CLASSIFICATION

Aleksandar Petrović[1], Milan Banić[1], Gavrilo Adamović[2], Ljiljana Radović[3]

[1]University of Niš, Faculty of Mechanical Engineering, Department of Mechanical Design, Product Development and Engineering, Republic of Serbia

[2]University of Niš, Faculty of Science and Mathematics, Department of Computer Science, Republic of Serbia

[3]University of Niš, Faculty of Mechanical Engineering, Department of Natural and Mathematical Science, Republic of Serbia

Abstract. *This paper investigates the application of Histogram of Oriented Gradients (HOG) features, simple colour statistics, Principal Component Analysis (PCA), and Logistic Regression for the classification of simple geometric shapes, namely circles, squares, and triangles. The work is positioned as a controlled baseline study aimed at evaluating the suitability of lightweight feature-based classification methods for future railway signal recognition pipelines.*

The main experiments were conducted on a balanced synthetic dataset containing 4500 images in total, with 1500 samples per class. A leakage-free evaluation protocol was applied by dividing the dataset into training, validation, and test subsets. Hyperparameters were selected exclusively on the validation subset, while final performance was evaluated once on a previously unseen test subset.

The best-performing configuration was obtained using a 128×128 input resolution, 12 HOG orientations, 8×8 pixels per cell, 2×2 cells per block, and 300 PCA components. This configuration achieved a validation accuracy of 0.8028 and a final test accuracy of 0.7978. The results show that increasing the image resolution from 64×64 to 128×128, together with combining HOG descriptors with colour-based features, improves classification performance, while further increasing the resolution to 200×200 does not provide additional benefits.

An additional verification experiment on 30,000 images confirmed the stability of the selected configuration, achieving a test accuracy of 0.8070. The study also discusses how such lightweight classifiers may be used as auxiliary verification modules in future railway signal detection systems.

Received: February 26, 2026 / Accepted May 27, 2026.

Corresponding author: Aleksandar Petrović
University of Niš, Faculty of Mechanical Engineering, Department of Mechanical Design, Product Development and Engineering, Aleksandra Medvedeva 14, 18000 Niš, Republic of Serbia
E-mail: aleksandar.petrovic@masfak.ni.ac.rs / aleksandarpetrovicmfn@yahoo.com

Key words: *Logistic Regression, HOG, PCA, Color Features, Classification, Railway Signal Detection*

1. INTRODUCTION

Image classification is a fundamental task in computer vision, aiming to assign semantic labels to images based on their visual content. While modern approaches are largely dominated by convolutional neural networks (CNNs) and other deep learning models, traditional feature-based methods remain relevant in scenarios with limited datasets, restricted computational resources, or the need for interpretable and reproducible decision-making. Among such methods, the Histogram of Oriented Gradients (HOG) descriptor is widely used because of its ability to capture local edge and shape information, which is particularly useful for rigid object and geometric shape recognition.

In this paper, the classification of three basic geometric shapes - circles, squares, and triangles - is investigated using HOG features in combination with simple colour statistics and Logistic Regression. Principal Component Analysis (PCA) is additionally applied to reduce the dimensionality of the extracted feature vectors and improve computational efficiency. The selected pipeline is intentionally lightweight and interpretable, making it suitable as a controlled baseline for evaluating shape separability before introducing more complex real-world computer vision stages.

The motivation behind this study extends beyond synthetic geometric shape recognition. Railway signal recognition often requires identifying structured visual patterns, such as circular signal lights, in complex scenes affected by background clutter, illumination changes, motion blur, weather conditions, and occlusions. In such applications, lightweight shape-based classifiers may serve as auxiliary verification stages within larger detection pipelines. For example, an object detector may first localize a candidate signal, while a lightweight classifier validates whether the candidate region contains the expected circular or structured signal pattern.

This work is a continuation and extension of previous research on geometric shape classification using Logistic Regression [1]. Compared with the previous work, the present study introduces Histogram of Oriented Gradients (HOG) feature extraction, PCA-based dimensionality reduction, colour-based statistical features, a larger dataset, systematic parameter evaluation, and a validation-based experimental protocol.

The HOG descriptor, first introduced by Dalal and Triggs [2], has become one of the most widely used handcrafted descriptors in classical computer vision. It captures local edge and contour information by computing distributions of gradient orientations over spatial regions. This makes HOG particularly suitable for rigid object recognition, pedestrian detection, traffic sign recognition, industrial inspection, and simplified geometric pattern classification. Its strength lies in the fact that many structured objects can be effectively represented by their edge orientation distributions.

Several studies have explored HOG and related handcrafted descriptors in combination with classical machine learning classifiers, such as Support Vector Machines, Random Forests, and Logistic Regression [3,4]. Although deep learning methods dominate many modern classification tasks, classical classifiers remain relevant when datasets are limited, computational resources are constrained, or interpretability is required. Logistic Regression is especially attractive because it provides probabilistic outputs, has low computational cost, and remains stable for moderate-dimensional feature spaces when appropriate preprocessing and regularization are applied [5].

In traffic sign recognition and related transportation applications, combinations of handcrafted features and shallow classifiers have been widely used as efficient baseline solutions [6-9]. Such approaches are particularly relevant when target objects have strong geometric structure, high contrast, or clearly defined contours. Railway signals share some of these properties, as their lamps and housings often contain circular, rectangular, or triangular visual structures. Therefore, evaluating the ability of lightweight feature-based models to separate simple geometric shapes is a useful preliminary step toward more complex railway signal recognition systems.

Dimensionality reduction is also important when working with HOG descriptors because HOG feature vectors can become large depending on the selected input resolution, cell size, block size, and number of gradient orientations. PCA is commonly used to reduce redundancy, preserve dominant information, and improve computational efficiency [10]. In addition, practical implementation details and reproducibility are important in machine learning experiments, and standardized toolkits such as scikit-learn provide widely used implementations of preprocessing, PCA, and Logistic Regression models [11].

In the context of railway signal detection, complete computer vision systems often rely on a combination of candidate-region extraction and classification. Classical pipelines may include colour-based segmentation, thresholding, morphological processing, contour detection, region extraction, and subsequent classification of the extracted signal candidates. These segmentation-based approaches are particularly useful because railway signals often occupy small regions in an image and may be surrounded by visually complex infrastructure. Candidate-region extraction can reduce the search space and suppress background clutter before the final classification stage.

More recent railway perception systems increasingly integrate deep learning-based object detection and segmentation. For example, YOLO-based detectors have been applied to the detection and recognition of railway traffic signals and signal lights [12-14]. Such systems provide strong localization and classification performance, but they usually require larger annotated datasets and greater computational resources. In contrast, handcrafted feature descriptors such as HOG, combined with interpretable classifiers, can still be useful as lightweight auxiliary components. They can be applied inside bounding boxes produced by an object detector in order to validate the geometric appearance of candidate objects.

Railway signal recognition also requires spatial context. A detected signal is not necessarily relevant to the track on which the train is moving, particularly in stations and multi-track scenarios. Previous railway perception studies have therefore combined track detection, region-of-interest selection, and signal detection in order to associate detected signals with the relevant track [15,16]. This is directly related to the motivation of the present study, since a lightweight classifier can be used as a secondary verification module after an object detector localizes a candidate signal.

Railway scene understanding is broader than signal recognition alone. Public datasets such as RailSem19 support semantic rail scene understanding and provide annotated railway scenes from the ego-perspective of trains and trams [17]. Vision-based railway perception has also been reviewed in the context of onboard obstacle detection and distance estimation, where traditional computer vision and AI-based methods are considered complementary approaches [18]. These studies show that railway environments impose specific challenges, including long detection distances, lighting variation, weather effects, tunnels, background clutter, and complex track layouts.

Railway infrastructure monitoring has also been addressed using both classical computer vision and deep learning methods. Previous works have investigated railway track and signal detection, railway infrastructure monitoring, track inspection, and defect detection [15,16,19-23]. These studies indicate that robust railway perception requires a combination of object localization, signal-state recognition, spatial context, temporal consistency, and computational efficiency. The present work does not attempt to solve the complete railway signal detection problem. Instead, it focuses on the controlled evaluation of a lightweight shape-classification stage that may later be integrated into larger railway perception pipelines.

Beyond transportation-related applications, handcrafted visual descriptors have also been used in other computer vision domains. For example, HOG-based features have been applied in human action recognition [24], while related texture descriptors have been used in biometric recognition and industrial surface inspection [25,26]. These applications demonstrate that handcrafted descriptors can still provide useful and interpretable feature representations in domains where data availability, computational cost, or explainability are important. At the same time, foundational deep learning work has demonstrated the strength of end-to-end feature learning for image recognition [27], while modern object detection and segmentation methods such as YOLO, Faster R-CNN, and U-Net have become dominant in complex visual recognition tasks [28-30].

The contribution of this work is not the introduction of a novel classifier, but a systematic baseline evaluation of a lightweight HOG + colour statistics + PCA + Logistic Regression pipeline for geometric shape classification. Specifically, the study investigates the influence of input resolution, HOG parameter selection, colour-based descriptors, and dimensionality reduction on classification performance using a leakage-free evaluation protocol. Furthermore, the work discusses the relevance of the proposed framework as a preliminary stage toward future railway signal recognition systems.

The objectives of this research are:

1. To evaluate the performance of HOG features combined with color statistics, PCA, and Logistic Regression in geometric shape classification.
2. To analyze the influence of input image resolution, HOG parameter configuration, and PCA dimensionality on classification performance.
3. To apply a validation-based model selection protocol that prevents test-set leakage.
4. To establish a reproducible lightweight baseline framework relevant to future railway signal detection systems.

2. Methodology

2.1 Dataset Description

The dataset used in this study was obtained from the publicly available "2D geometric shapes dataset" published on Mendeley Data [31]. The complete dataset consists of nine geometric shape classes: triangle, square, pentagon, hexagon, heptagon, octagon, nonagon, circle, and star. Each class contains 10000 synthetically generated RGB images with a resolution of 200×200 pixels.

According to the dataset description, each image contains a single 2D geometric shape placed on a randomly selected background colour. The shape position, perimeter, rotation angle, and filling colour are randomly selected for each image. The rotation angle is selected within the interval from −180° to 180°, while the filling colour is generated independently from the background colour. Therefore, the dataset includes variations in shape position, scale, orientation, foreground colour, and background colour.

For the purpose of the present study, three classes were selected from the original dataset: circle, square, and triangle. A balanced subset was formed by selecting 1500 images per class, resulting in a total of 4500 images. These three classes were selected because they represent basic geometric structures relevant to simplified railway signal and sign-like visual patterns.

During the final experimental run, all 4500 selected images were successfully loaded and used in the evaluation protocol. No corrupted or unreadable samples were removed from the final dataset. The dataset is therefore fully balanced, with 1500 samples per class. Since no weather effects, background clutter, motion blur, or environmental distortions are included, the dataset should be interpreted as a controlled benchmark for baseline shape classification rather than as a direct representation of real railway scenes.

Table 1 Dataset configuration

Class	Number of images
Circle	1500
Square	1500
Triangle	1500
Total	4500

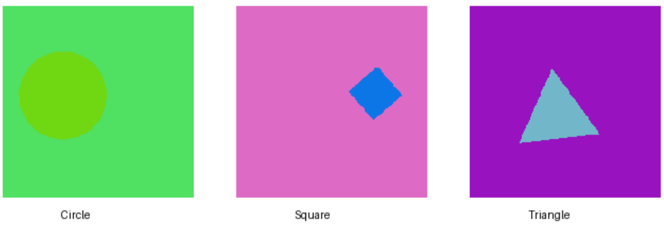

Fig. 1 Sample images from the geometric shapes dataset

2.2 Preprocessing

The preprocessing stage consisted of image loading, resizing, colour conversion, normalization, and feature extraction. Unlike a purely grayscale approach, the final experimental pipeline preserved colour information by extracting simple colour statistics from RGB and HSV representations.

For each input image, the following steps were applied:

1. Image loading in RGB format
2. Resizing the tested input resolution

3. Grayscale conversion for HOG feature extraction
4. RGB and HSV color statistics extraction
5. Feature concatenation
6. Feature standardization
7. PCA dimensionality reduction
8. Logistic Regression classification

The complete input representation therefore combines structural edge information from HOG with global colour statistics.

2.3 HOG Feature Extraction

Histogram of Oriented Gradients (HOG) was used to extract local shape descriptors from each image. HOG divides the image into spatial cells, computes gradient orientation histograms inside each cell, and normalizes the resulting feature representation over larger spatial blocks. This makes the descriptor suitable for capturing local edge and contour structures.

HOG features were extracted from the grayscale version of each resized image. The following HOG parameters were evaluated:

- Input resolution: 64×64, 128×128, 200×200
- Orientations: 9, 12
- Pixels per cell: 8×8, 16×16
- Cells per block: 2×2, 3×3
- Block normalization: L2-Hys
- Square-root intensity transformation: enabled
- Feature vector output: enabled

2.4 Color Feature Extraction

Because the dataset contains strong colour variations between the shape and background, using only grayscale HOG descriptors may discard useful global colour and contrast information. Therefore, simple colour statistics were added to the feature representation.

For each resized RGB image, the following colour features were extracted: mean value of each RGB channel, standard deviation of each RGB channel, mean value of each HSV channel, and standard deviation of each HSV channel. This produced 12 additional colour features per image. These features were concatenated with the HOG feature vector before standardization and PCA dimensionality reduction. Since foreground and background colours are randomly generated in the present dataset, the extracted colour statistics mainly provide complementary global colour and contrast information rather than direct class-specific cues. In future railway signal recognition applications, however, colour information is expected to have stronger semantic relevance because signal states are directly associated with lamp colour.

2.5 Dimensionality Reduction with PCA

The combined HOG and colour feature vectors are high-dimensional, particularly when larger input resolutions and smaller HOG cells are used. To reduce dimensionality and improve computational efficiency, Principal Component Analysis (PCA) was applied after feature standardization.

The following PCA configurations were evaluated: 100, 150, 300, and 500 principal components. PCA was fitted only on the training data during model selection. After the best configuration was selected, the final PCA transformation was fitted on the combined training and validation subsets and then applied to the test subset.

For the best-performing final configuration, PCA retained 71.43% of the total variance in the standardized combined feature space. This confirms that preserving the highest variance alone is not necessarily equivalent to obtaining the best classification performance; rather, the selected number of PCA components should be validated empirically.

2.6 Classification with Logistic Regression

A Logistic Regression classifier was trained on the PCA-transformed combined features. Logistic Regression was selected due to its simplicity, interpretability, and computational efficiency for moderate-scale classification tasks.

The classifier was implemented using the scikit-learn library with lbfgs solver, L2 regularization, maximum number of iterations equal to 3000, and balanced class weighting. Feature normalization was performed using Standard Scaler. The scaler was fitted only on the training subset during model selection and then applied consistently to validation and test subsets.

The complete classification pipeline can therefore be summarized as: Image → Resize → Grayscale HOG + RGB/HSV colour statistics → Standard Scaler → PCA → Logistic Regression.

2.7 Dataset Split and Evaluation Protocol

To avoid test-set leakage and obtain an unbiased estimate of generalization performance, the dataset was divided into three subsets using stratified random sampling: training subset (64%), validation subset (16%), and test subset (20%).

The validation subset was used exclusively for hyperparameter selection. The test subset was kept unseen during both training and validation and was used only once for final evaluation.

Table 2 Dataset split

Subset	Number of images
Training	2880
Validation	720
Test	900

The model selection procedure was as follows:

1. Split the dataset into training, validation, and test subsets.
2. Train models with different input resolutions, HOG parameters, and PCA configurations using the training subset.
3. Select the best configuration based on validation accuracy only.
4. Retrain the final model using the combined training and validation subsets.
5. Evaluate the final model once on the untouched test subset.

This protocol prevents information leakage from the test subset into the model selection process.

2.8 Evaluation Metrics and Implementation Details

The model was evaluated using accuracy, precision, recall, F1-score, and confusion matrix. Precision, recall, and F1-score were reported for each class. In addition, randomly selected test samples were visualized to qualitatively inspect classifier predictions and identify typical misclassification cases. All experiments were implemented in Python using NumPy, PIL, OpenCV, scikit-image, scikit-learn, and Matplotlib. The experiments were performed on a workstation equipped with a 12th Gen Intel Core i9-12900H processor, 32 GB RAM, and an NVIDIA RTX A3000 12 GB Laptop GPU. The HOG, PCA, and Logistic Regression experiments were executed as CPU-based computations, while the GPU was available for preliminary YOLO-based railway video experiments.

3. Results and Discussion

3.1 Validation-Based Hyperparameter Selection

A grid search was performed over input resolution, HOG parameters, and PCA dimensionality. Model selection was performed exclusively on the validation subset. The best-performing configuration on the validation subset used 128×128 input resolution, 12 orientations, 8×8 pixels per cell, 2×2 cells per block, and 300 PCA components. This configuration achieved a validation accuracy of 0.8028.

Table 3 Top validation results

Rank	Image size	Orient.	Pixels/cell	Cells/block	PCA	Val. acc.	PCA var.
1	128×128	12	8×8	2×2	300	0.8028	0.7269
2	128×128	9	8×8	2×2	300	0.7972	0.7580
3	128×128	12	8×8	3×3	300	0.7931	0.7319
4	128×128	12	16×16	3×3	300	0.7931	0.9052
5	200×200	9	16×16	3×3	300	0.7847	0.9361
6	200×200	12	16×16	3×3	300	0.7847	0.9361
7	128×128	9	8×8	3×3	300	0.7833	0.7634
8	128×128	9	16×16	3×3	300	0.7819	0.9306
9	64×64	9	16×16	3×3	300	0.7806	0.9999
10	64×64	12	16×16	3×3	300	0.7806	0.9989

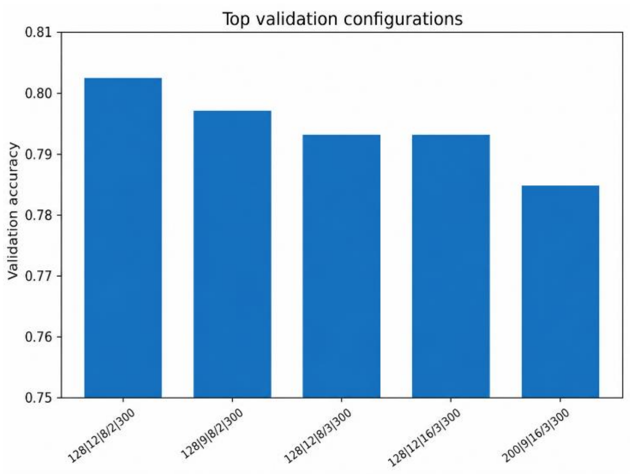

Fig. 2 Top validation configurations

The results show that the 128×128 input resolution provided the best overall performance. The 64×64 configurations were generally less accurate, suggesting that resizing to 64×64 removes part of the discriminative contour information. On the other hand, increasing the resolution to 200×200 did not improve performance and increased computational cost. Therefore, 128×128 represents the best compromise between classification performance and computational efficiency.

3.2 Final Test Evaluation

After identifying the best configuration, the final classifier was retrained using the combined training and validation subsets and evaluated once on the unseen test subset. The final test accuracy was 0.7978. The classification report for the test subset is shown in Table 4.

Table 4 Classification report on the test subset

Class	Precision	Recall	F1-score	Support
Circle	0.82	0.83	0.82	300
Square	0.77	0.74	0.75	300
Triangle	0.80	0.83	0.82	300
Accuracy	-	-	0.80	900
Macro avg.	0.80	0.80	0.80	900
Weighted avg.	0.80	0.80	0.80	900

The difference between validation accuracy (0.8028) and test accuracy (0.7978) indicates that the selected model generalized reasonably well to unseen data. The test performance confirms that the combination of increased input resolution and additional color features improved the classification capability compared with the lower-resolution grayscale-only baseline.

3.3 Confusion Matrix Analysis

The confusion matrix for the final test subset is presented in Table 5 and Fig. 3.

Table 5 Confusion matrix on the test subset

True / Predicted	Circle	Square	Triangle
Circle	248	36	16
Square	33	221	46
Triangle	21	30	249

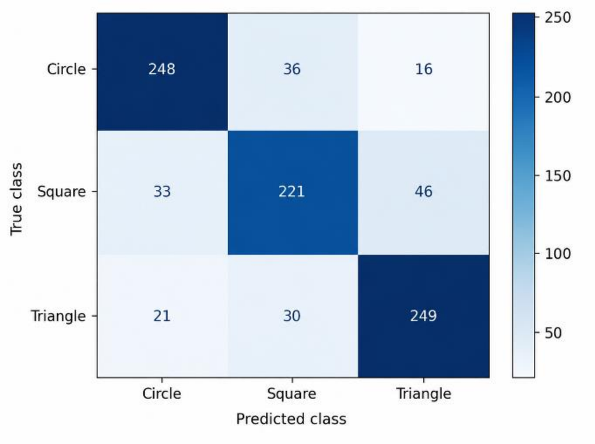

Fig. 3 Confusion matrix for the final test subset

The confusion matrix shows that triangles achieved the highest recall, with 249 out of 300 triangle samples correctly classified. This is expected because triangles have distinct angular structures and sharp vertices that generate characteristic gradient orientation distributions.

Circles were also recognized with improved reliability, with 248 out of 300 circle samples correctly classified. The improved circle recognition is particularly important for future railway signal recognition, since signal lights are often circular.

Squares remained the most challenging class, with 221 out of 300 samples correctly classified. Misclassifications occurred toward both circles and triangles. This can be explained by the fact that HOG features emphasize edge orientations. Depending on rotation, scale, and edge appearance, square samples may partially overlap with the gradient structures of both triangles and circles.

Overall, the results confirm that HOG features capture meaningful structural information for simple geometric shapes, while the added colour features improve classification performance by preserving information that is lost in a purely grayscale pipeline.

3.4 Effect of Input Resolution, HOG Parameters and PCA

One of the most important findings is the influence of input image resolution. The previous 64×64 representation provided acceptable results, but the best validation result was obtained at 128×128. This indicates that part of the discriminative contour information is lost when images are resized too aggressively.

However, increasing the resolution further to 200×200 did not improve validation performance. This may be attributed to the significantly higher dimensionality of the HOG descriptor and increased sensitivity to redundant or noisy local variations. As a result, 128×128 provided the best balance between preserving geometric detail and maintaining a compact feature representation.

The best validation result was obtained with 12 orientations, pixels per cell equal to 8×8, and cells per block equal to 2×2. Compared with lower angular resolution, 12 orientations provided a more detailed description of edge directions, which improved classification performance at 128×128 resolution.

The selected PCA configuration used 300 components. In the final model, after refitting PCA on the combined training and validation subsets, these components retained 71.43% of the variance in the combined standardized feature space. The PCA variance values reported in Table 3 refer to the validation-stage models fitted on the training subset only, which explains the difference between the validation-stage variance value and the final-model variance value. Interestingly, configurations retaining a higher percentage of variance did not always achieve better classification accuracy. This shows that preserving more variance is not necessarily equivalent to preserving more class-discriminative information. Therefore, PCA dimensionality should be selected empirically using a validation subset rather than based only on explained variance.

3.5 Computational Considerations

The proposed pipeline is computationally lightweight compared with deep learning approaches because it relies on handcrafted descriptors, simple colour statistics, PCA dimensionality reduction, and Logistic Regression. The final training procedure for the selected model took 21.2859 seconds on the reported workstation.

Table 6 Computational time metrics

Processing stage	Measured time
Final training time	21.2859 s
Classifier-only prediction time on test set	0.001108 s
Classifier-only prediction time per image	0.00000123 s
End-to-end inference time on test set	9.638611 s
End-to-end inference time per image	0.01070957 s
Equivalent processing rate	93.37 images/s

The classifier-only prediction time refers to Logistic Regression prediction on already transformed PCA feature vectors. The end-to-end inference time includes image resizing, grayscale conversion, HOG feature extraction, RGB/HSV colour feature extraction, feature standardization, PCA transformation, and Logistic Regression prediction. The measured end-to-end processing rate of 93.37 images/s confirms that the proposed pipeline is suitable for lightweight CPU-based processing.

3.6 Additional Dataset-Scale Verification

To further verify the stability of the selected configuration, an additional experiment was conducted using a larger subset of the available dataset. In this experiment, 10000 images per class were used, resulting in 30000 images in total. The same configuration selected in the previous experiment was applied without additional hyperparameter tuning: 128×128 input resolution, 12 HOG orientations, 8×8 pixels per cell, 2×2 cells per block, and 300 PCA components.

The expanded dataset was divided using the same stratified 64/16/20 protocol, resulting in 19200 training images, 4800 validation images, and 6000 test images. The model achieved a validation accuracy of 0.8206 and a final test accuracy of 0.8070. Compared with the experiment on 4500 images, the test accuracy increased from 0.7978 to 0.8070, indicating that the selected configuration remains stable and can benefit from a larger training set.

Table 7 Additional dataset-scale verification

Dataset size	Train	Validation	Test	Validation accuracy	Test accuracy
4500	2880	720	900	0.8028	0.7978
30000	19200	4800	6000	0.8206	0.8070

The confusion matrix for the larger dataset is shown in Table 8.

Table 8 Confusion matrix on the expanded test subset

True / Predicted	Circle	Square	Triangle
Circle	1655	263	82
Square	171	1472	357
Triangle	159	126	1715

The results confirm that triangles and circles remain the most reliably recognized classes, while squares remain the most challenging class. This behaviour is consistent with the previous experiment and supports the interpretation that square samples partially overlap with both circular and triangular gradient patterns, especially under rotation and scale variation.

3.7 Relation to Railway Signal Detection

Although the experiments presented in this paper were conducted on synthetic geometric datasets, the proposed framework is relevant as a preliminary component for future railway signal recognition pipelines.

In preliminary experiments on real railway video sequences, a YOLO-based object detector was used to localize candidate railway signals. The HOG + Logistic Regression classifier was then applied inside the detected bounding boxes in order to validate whether the candidate region contained a circular signal pattern. In addition, temporal filtering and track-aware constraints based on rail segmentation were investigated as auxiliary mechanisms for reducing false-positive detections.

This preliminary integration suggests that lightweight classical classifiers can complement modern object detectors by acting as verification modules. Such a combination may be useful in railway environments, where false positives can occur due to background lights, reflections, or infrastructure elements with similar visual appearance.

It should be emphasized that this part of the work is presented only as a practical extension and motivation for future research. A complete railway signal detection system would require a dedicated real-world dataset, frame-level annotations, evaluation of false positives and false negatives, and testing under different illumination, weather, and operational conditions.

Fig. 4 Preliminary integration of YOLO detection with HOG + Logistic Regression validation in real railway video frames

4. CONCLUSION

This study investigated the application of Histogram of Oriented Gradients (HOG) features, color statistics, Principal Component Analysis (PCA), and Logistic Regression for geometric shape classification. The work was positioned as a controlled baseline study rather than as the introduction of a novel classifier.

A leakage-free evaluation protocol was applied using separate training, validation, and test subsets. Hyperparameters were selected exclusively on the validation subset, while the final model was evaluated once on the previously unseen test subset. This protocol provides a reliable estimate of generalization performance and avoids test-set leakage during model selection.

The best configuration was obtained using 128×128 input resolution, 12 HOG orientations, 8×8 pixels per cell, 2×2 cells per block, and 300 PCA components. The model achieved a validation accuracy of 0.8028 and a final test accuracy of 0.7978. The results showed balanced performance across the three classes, with triangles and circles being recognized more reliably than squares.

The analysis confirmed that input resolution and feature representation significantly influence classification performance. Increasing the resolution from 64×64 to 128×128 improved the results, while further increasing the resolution to 200×200 did not provide additional benefit. The addition of color statistics helped preserve complementary global color and contrast information that is lost when using grayscale HOG descriptors alone.

An additional verification experiment on 30000 images confirmed the stability of the selected configuration, achieving a final test accuracy of 0.8070.

From the perspective of railway signal detection, the proposed approach offers several advantages: low computational complexity, interpretability, and straightforward integration as a verification module inside larger computer vision pipelines. However, the present experiments were conducted on synthetic images with simplified backgrounds and therefore do not capture the full complexity of real railway environments, such as illumination variation, weather effects, occlusions, motion blur, class imbalance, and background clutter.

For that reason, the reported results should be interpreted as a baseline estimate of shape-classification capability rather than as evidence of readiness for operational railway deployment. Future work will focus on extending the approach to real railway video datasets, integrating temporal information, and combining lightweight classical classifiers with modern object detection and segmentation methods.

REFERENCES

1. Petrović, A., Banić, M., Stamenković, D., Simonović, M., Ristić Durrant, D., Radović, Lj., Perić, M., 2021, *Classification of geometric shapes in the images using Logistic Regression algorithm*, XV International Conference on Systems, Automatic Control and Measurements SAUM 2021, Niš, pp. 140-143.
2. Dalal, N., Triggs, B., 2005, *Histograms of oriented gradients for human detection*, Proc. IEEE Conference on Computer Vision and Pattern Recognition, pp. 886-893.
3. Maji, S., Malik, J., 2009, *Fast and accurate digit classification using gradient histograms*, Technical Report UCB/EECS-2009-159, University of California, Berkeley.
4. Ojala, T., Pietikäinen, M., Harwood, D., 1994, *Performance evaluation of texture measures with classification based on Kullback discrimination of distributions*, Proc. International Conference on Pattern Recognition, pp. 582-585.
5. Hosmer, D.W., Lemeshow, S., 2000, *Applied Logistic Regression*, 2nd ed., Wiley, New York.
6. Stallkamp, J., Schlipsing, M., Salmen, J., Igel, C., 2012, *Man vs. computer: benchmarking machine learning algorithms for traffic sign recognition*, Neural Networks, 32, pp. 323-332.
7. Zaklouta, F., Stanciulescu, B., 2014, *Real-time traffic sign recognition in three stages*, Robotics and Autonomous Systems, 62(1), pp. 16-24.
8. Paclik, P., Novovicova, J., Pudil, P., Somol, P., 2000, *Road sign classification using the Laplace kernel classifier, Pattern Recognition Letters*, 21(13-14), pp. 1165-1173.
9. García-Garrido, M.A., Ocaña, M., Llorca, D.F., Arroyo, E., Pozuelo, J., Gavilán, M., 2012, *Complete vision-based traffic sign recognition supported by an I2V communication system*, Sensors, 12(2), pp. 1148-1169.
10. Jolliffe, I.T., 2002, *Principal Component Analysis*, 2nd ed., Springer, New York.
11. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, É., 2011, *Scikit-learn: machine learning in Python*, Journal of Machine Learning Research, 12, pp. 2825-2830.

12. Liu, W., Wang, Z., Zhou, B., Yang, S., Gong, Z., 2021, *Real-time signal light detection based on YOLOv5 for railway*, IOP Conference Series: Earth and Environmental Science, 769, 042069.
13. Staino, A., Suwalka, A., Mitra, P., Basu, B., 2022, *Real-time detection and recognition of railway traffic signals using deep learning*, Journal of Big Data Analytics in Transportation, 4, pp. 57-71.
14. Ma, L., Zhang, J., 2022, *Detection of railway signal light based on improved YOLOR*, Proc. 2022 6th International Conference on Electronic Information Technology and Computer Engineering EITCE, Xiamen, pp. 1533-1537.
15. Petrović, A.D., Banić, M., Simonović, M., Stamenković, D., Miltenović, A., Adamović, G., Rangelov, D., 2022, *Integration of computer vision and convolutional neural networks in the system for detection of rail track and signals on the railway*, Applied Sciences, 12(12), 6045.
16. Petrović, A.D., Simonović, M., Banić, M., Stan, S.-D., 2022, *Railway infrastructure monitoring using classical computer vision and convolutional neural networks*, RAILCON 22, Niš, pp. 129-132.
17. Zendel, O., Murschitz, M., Zeilinger, M., Steininger, D., Abbasi, S., Beleznai, C., 2019, *RailSem19: a dataset for semantic rail scene understanding*, Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pp. 1221-1229.
18. Ristić-Durrant, D., Franke, M., Michels, K., 2021, *A review of vision-based on-board obstacle detection and distance estimation in railways*, Sensors, 21(10), 3452.
19. Ritika, S., Mittal, S., Rao, D., 2017, *Railway track specific traffic signal selection using deep learning*, arXiv preprint arXiv:1712.06107.
20. Mittal, S., Rao, D., 2017, *Vision based railway track monitoring using deep learning*, arXiv preprint arXiv:1711.06423.
21. Gibert, X., Patel, V.M., Chellappa, R., 2017, *Deep multitask learning for railway track inspection*, IEEE Transactions on Intelligent Transportation Systems, 18(1), pp. 153-164.
22. Wang, T., Zhang, Z., Yang, F., Tsui, K.-L., 2021, *Intelligent railway foreign object detection: a semi-supervised convolutional autoencoder based method*, arXiv preprint arXiv:2108.02421.
23. Chen, W., Liu, B., 2022, *Foreign object detection in railway images based on an improved YOLOv4 framework*, Computational Intelligence and Neuroscience, 2022, Article ID 3749635.
24. Patel, C.I., Labana, D., Pandya, S., Modi, K., Ghayvat, H., Awais, M., 2020, *Histogram of oriented gradient-based fusion of features for human action recognition in action video sequences*, Sensors, 20(24), 7299.
25. Ahonen, T., Hadid, A., Pietikäinen, M., 2006, *Face description with local binary patterns: application to face recognition*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 28(12), pp. 2037-2041.
26. Xie, X., 2008, *A review of recent advances in surface defect detection using texture analysis techniques*, Electronic Letters on Computer Vision and Image Analysis, 7(3), pp. 1-22.
27. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998, *Gradient-based learning applied to document recognition*, Proceedings of the IEEE, 86(11), pp. 2278-2324.
28. Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016, *You only look once: unified, real-time object detection*, Proc. IEEE Conference on Computer Vision and Pattern Recognition, pp. 779-788.
29. Ren, S., He, K., Girshick, R., Sun, J., 2015, *Faster R-CNN: towards real-time object detection with region proposal networks*, Advances in Neural Information Processing Systems, 28.
30. Ronneberger, O., Fischer, P., Brox, T., 2015, *U-Net: convolutional networks for biomedical image segmentation*, Medical Image Computing and Computer-Assisted Intervention, pp. 234-241.
31. El Korchi, A., 2020, *2D geometric shapes dataset*, Mendeley Data, Version 1, doi: 10.17632/wzr2yv7r53.1.