MRI-Based Brain Stroke Classification using HOG Features and SVM with SMOTE Balancing

Yogish Naik. G R[1], *Manjunath.* V[1], *Vidyasagar.* K B[2*], and *Namitha* R. Shetty[1]

[1]Department of Computer Science, Kuvempu University, Shivamogga, India
[2]Department of Computer Science and Engineering, BGS Institute of Technology, Mandya, India

Abstract. Brain MRI scans are vital for identifying stroke types, which is essential for quick and effective treatment. This study introduces a simple yet powerful machine learning method to categorize MRI images as Normal, Haemorrhagic (bleeding), or Ischemic (clot-related) strokes. The approach uses a technique called Histogram of Oriented Gradients (HOG) to identify key features in the images. To address an imbalance where some stroke types appear less frequently, a method called SMOTE was used to balance the data. Finally, a Support Vector Machine (SVM), adjusted for class importance, was employed for classification. The model performed exceptionally well, achieving an overall accuracy of 97%. It also showed strong performance across all stroke types. This method works well even with limited data and doesn't require extensive computing power, making it ideal for clinical settings, especially those with fewer resources. This highlights the ongoing value of traditional machine learning for medical image analysis.

1 Introduction

A brain stroke is a critical medical emergency that arises when the supply of blood to the brain is disrupted, either by an obstruction (ischemic stroke) or by the rupture of a blood vessel causing internal bleeding (haemorrhagic stroke). This disruption leads to the deprivation of oxygen and essential nutrients in affected brain regions, resulting in the rapid death of the brain cells and potential long term neurological damage. Historically known as "apoplexy," stroke has been documented since ancient times but was poorly understood until the development of vascular and neurological sciences in the 19th and 20th centuries. The accurate and timely diagnosis of stroke type is critical for initiating appropriate treatment. Ischemic and haemorrhagic strokes require entirely different therapeutic approaches: ischemic strokes benefit from clot-dissolving medications or mechanical thrombectomy, while haemorrhagic strokes may necessitate surgical intervention to control bleeding. Misdiagnosis or delayed identification can worsen patient outcomes and increase the risk of permanent impairment or fatality. While CT scans are commonly used in emergency settings due to their speed, MRI scans offer higher contrast resolution and are more sensitive in detecting subtle changes in brain tissue, making them particularly effective for identifying early-stage ischemic strokes.

However, despite the advancements in imaging technology, stroke diagnosis remains largely dependent on the subjective interpretation of radiologists and neurologists. This manual assessment is not only time-intensive but also prone to variability, especially in regions with limited access to experienced medical professionals. These challenges have led to a growing interest in automated diagnostic systems powered by artificial intelligence and machine learning. Such systems promise to enhance the speed, consistency, and accuracy of stroke classification, supporting clinical decisions and improving patient outcomes.

Nevertheless, the development of automated stroke classification models using MRI data presents several technical challenges. One of the primary difficulties lies in the high dimensionality of MRI images, which can lead to overfitting in machine learning models if not handled properly. Another significant challenge is the imbalance in dataset distribution ischemic strokes, despite being more common in clinical practice, are often underrepresented in labelled datasets used for training ML models. This class imbalance can result in biased models that perform poorly on minority classes, particularly ischemic stroke, which is crucial to detect accurately due to its subtle presentation and critical time window for treatment. Rather than proposing a novel deep architecture, this study focuses on developing a computationally lightweight, interpretable, and reproducible pipeline suitable for small, imbalanced stroke MRI datasets. Such practical solutions are especially valuable in low-resource clinical environments where complex models may overfit or fail to generalize. To address these limitations, this study proposes a machine learning pipeline designed for efficient and reliable stroke classification from brain MRI images. The approach integrates key components: robust feature extraction, synthetic data balancing, and an effective classification algorithm. For feature extraction, we utilize the Histogram of Oriented Gradients (HOG) technique, which captures important

* Corresponding author: vkb2231@gmail.com

structural details, such as edges and textures, from brain scans. HOG is particularly suitable for medical imaging tasks where spatial gradients and contrast variations play a vital role in distinguishing between pathological and normal tissue patterns.

While deep learning methods have become dominant in medical image analysis, they often require extensive data and high-end computational resources. Our proposed HOG + SMOTE + SVM pipeline was intentionally selected for its interpretability, efficiency, and practicality in low-resource clinical environments. Unlike deep networks that act as black boxes, our model enables clear feature attribution using handcrafted descriptors and provides explainable outputs suited for clinical decision-making. Furthermore, we implemented optimized HOG parameters, class-weight-balanced SVM, and rigorous 5-fold cross-validation to ensure robustness. This lightweight approach bridges the gap between diagnostic performance and deployment feasibility in real-world stroke care scenarios. While deep learning models dominate stroke classification literature, they often require large, balanced datasets to generalize effectively. In contrast, our lightweight and interpretable pipeline combining HOG, SMOTE, and SVM is particularly suited for resource-constrained clinical environments and small-scale datasets where deep architectures may overfit. Overall, this study presents a computationally efficient and interpretable solution for brain stroke classification that leverages machine learning to bridge the gap between diagnostic demand and clinical capacity. The proposed method not only enhances accuracy and generalization across stroke types but also offers potential for integration into real-time diagnostic tools, especially in resource-constrained healthcare environments.

Ethics and Data Availability: The dataset used in this study was obtained from a publicly available repository. All images were pre-anonymized and contain no patient-identifiable information. As no private health data or patient records were accessed, formal institutional ethical approval was not required. Preprocessing steps, including resizing and grayscale normalization, were applied consistently across the dataset. This paper proceeds as follows: Section 2 explores prior research in the field. Section 3 introduces the methodology, detailing preprocessing techniques, model architecture, and the augmentation pipeline. Section 4 discusses the experimental design and evaluation methods. Section 5 interprets the results and compares performance with existing approaches. Section 6 offers conclusions and possible future improvements.

2 Literature survey

Tursynova et al. [1] developed a CNN-based system for detecting stroke types from CT scans, achieving 79% accuracy in identifying normal cases. They proposed a cloud-based diagnostic tool to assist clinicians in real-time decision-making, especially useful in remote healthcare environments. Patel et al. [2] utilized the EfficientNetB0 model for automated stroke classification from CT images. Their system outperformed traditional methods in speed and accuracy, making it suitable for emergency stroke triage. Eshmawi et al. [3] introduced a CAD system called CAD-BSDC for brain MRI classification using adaptive thresholding, deep ensemble models (MobileNet, CapsuleNet, EfficientNet), and optimization algorithms. The hybrid SAE classifier with SBO and IDFO tuning yielded top-tier results. Gautam and Raman [4] proposed a 13-layer CNN with multi-focus image fusion for CT-based stroke classification. The model significantly outperformed established CNNs like AlexNet and ResNet50 in multi-class scenarios. Sailasya and Kumari [6] evaluated six machine learning algorithms using patient health parameters to predict stroke risk. Naïve Bayes achieved the highest accuracy at 82%, proving effective for preventive care. Murugan Rahman et al. [8] applied various ML and ANN models to stroke prediction on a public dataset. Their results showed classical ML models like Random Forest outperforming deep ANNs in accuracy and simplicity. Vavoulis et al. [9] reviewed online classifiers for MI-based BCI stroke rehab systems and found no major accuracy difference between stroke and non-stroke groups. Functional electrical stimulation-based feedback was most effective.

Deepa et al. [11] proposed a hybrid MRI classifier combining feature extraction (shape, texture, intensity), MAP-based firefly algorithm for selection, and a hybrid SVM-RF model. The method offered strong accuracy and interpretability. Sharma et al. [12] developed a Modified ResNet50 enhanced with HOG features for brain tumor classification. Their framework integrated feature augmentation and ensemble learning, achieving 88% accuracy over existing methods. Gavkare et al. [13] used HOG features and tested multiple ML classifiers (KNN, SVM, XGBoost) for tumor detection in MRI images. XGBoost performed best, achieving nearly 89% accuracy with minimal human input. Taha et al. [14] reviewed large animal models for ischemic stroke research, comparing vascular, immunological, and recovery traits. Their review serves as a guideline for selecting appropriate translational models. Pokorny et al. [15] investigated SVMs for microwave-based stroke detection using synthetic head data. Even single-frequency data yielded accurate results, emphasizing dataset design's role in model generalization. Telu et al. [16] proposed a supervised ML approach using clinical risk factors like glucose, hypertension, and smoking. They achived 94.6% accuracy with models such as SVM, KNN, and Randoma forestr. Mousavi et al. Recent approaches leveraging Swin Transformers [5], CNN–SVM hybrids [7], and ViT–U-Net architectures [10] have achieved competitive performance in stroke classification, especially on benchmark MRI datasets like ISLES and TCIA. Several studies, such as Gavkare et al. [13], Sharma et al. [12], and Deepa et al. [11], have previously explored the use of HOG, SVM, or hybrid machine learning strategies for brain pathology classification. However, most of these were applied to tumor detection or large, balanced datasets. In contrast, our work applies a lightweight classical approach to the

more challenging problem of ischemic stroke detection from severely imbalanced MRI data.

As shown in Table 1, the proposed HOG + SMOTE + SVM (RBF) framework achieves competitive performance (97.00% accuracy) compared to recent deep learning methods, while maintaining lower computational complexity and higher interpretability, making it suitable for real-time clinical deployment.

Table 1. Performance comparison of stroke classification models.

Study	Model	Type	Dataset	Accuracy
Eshmawi et al. (2022)	MobileNet + CapsuleNet + EfficientNet + IDFO-SAE	Brain Stroke Classification (MRI)	T2-weighted MR images	95.61%
Dhanalakshmi et al. (2021)	CNN + Deep Feature Fusion	Brain Stroke Detection (MRI)	Brain MRI dataset	96.80%
Arunkumar et al. (2019)	CNN + Firefly Algorithm	Stroke Classification	Private MRI dataset	94.50%
Xu et al. (2023)	Swin Transformer + CNN	Ischemic Stroke Classification (MRI)	ISLES 2022 Dataset	96.30%
Zhang et al. (2023)	Hybrid CNN + SVM with Deep Features	MRI-based Stroke Lesion Detection	Private Stroke MRI Dataset	95.40%
Wang et al. (2024)	Self-Supervised ViT + U-Net	Brain Segmentation & Stroke Type Classification	TCIA + ISLES	94.50%
Propose Method	HOG + SVM (RBF) + SMOTE	Stroke Classification	MRI Data set	97.00%

While several deep learning-based approaches have achieved impressive results in brain stroke classification, many rely on complex architectures that demand significant computational resources and large datasets. These requirements pose challenges for real-time clinical deployment, particularly in resource-constrained settings. To address these issues, the proposed methodology introduces a lightweight yet effective framework that leverages Histogram of Oriented Gradients (HOG) for feature extraction, SMOTE for class balancing, and a Support Vector Machine (SVM) classifier with an RBF kernel. The following section outlines the details of this approach.

3 Methodology

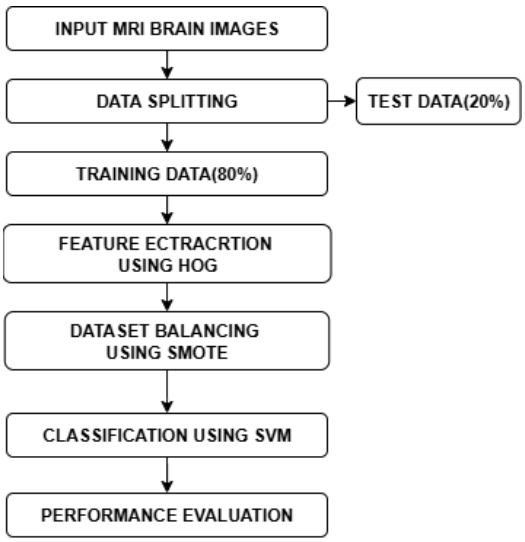

Fig. 1. Proposed HOG-SMOTE-SVM classification pipeline.

The proposed methodology employs a structured machine learning pipeline to classify brain MR images into three categories: Normal, Ischemic stroke, and Haemorrhagic stroke. The workflow is organized into five key stages: image preprocessing, feature extraction using Histogram of Oriented Gradients (HOG), class balancing via Synthetic Minority Oversampling Technique (SMOTE), classification using Support Vector Machine (SVM), and comprehensive performance evaluation. The complete architecture of the system is illustrated in Figure 1 (a flowchart can be included in the paper to visualize this process).

The proposed stroke classification using MRI brain images. The pipeline includes image preprocessing, feature extraction using Histogram of Oriented Gradients (HOG), dataset balancing via SMOTE, classification using Support Vector Machine (SVM), and performance evaluation.

3.1 Dataset

The dataset utilized in this study comprises a total of 615 anonymized brain MRI images, categorized into three classes: Normal (399), Haemorrhagic (186), and Ischemic (30), as illustrated in Figure 2. The images were acquired using a 1.5 Tesla MRI scanner, with modalities including T1-weighted and T2-weighted axial slices. All scans were provided in DICOM format and subsequently converted to grayscale PNG images for preprocessing. The dataset was fully de-identified prior to use, and no patient-identifiable information was retained; hence, ethical approval was not required.

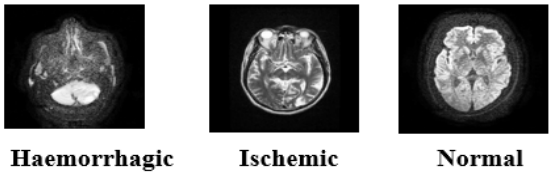

Fig. 2. Representative images from each class.

The primary challenges associated with this dataset include severe class imbalance particularly the underrepresentation of the ischemic stroke class and the subtle visual similarities between stroke subtypes, which hinder effective feature discrimination. Additionally, the limited dataset size increases the risk of overfitting in complex deep learning models. To mitigate these challenges, we adopted a lightweight and interpretable machine learning pipeline: Histogram of

Oriented Gradients (HOG) was employed to extract discriminative texture features, SMOTE was used to synthetically balance the class distribution, and a Support Vector Machine (SVM) with an RBF kernel was applied due to its robustness in high-dimensional, small-sample-size classification tasks.

3.2 Image preprocessing

The input MRI images used in this study were originally acquired in grayscale (T1- or T2-weighted) format and stored as DICOM slices, inherently containing intensity-based contrast and spatial gradient information. As such, no RGB-to-grayscale conversion was performed. All images were uniformly resized to 128×128 pixels using bicubic interpolation, which preserves spatial structure and edge details crucial for downstream texture-based feature extraction via HOG. Additionally, center cropping was applied to remove irrelevant peripheral regions while maintaining the lesion-focused center. All pixel intensities were normalized to the range [0, 1] to ensure numerical stability during feature computation and classification. These preprocessing steps ensured the retention of clinically significant image content while improving computational tractability.

3.3 Feature extraction using HOG

To optimize the performance of the classification pipeline, we conducted hyperparameter tuning for both the feature extraction and classification stages. For HOG, parameters such as cell size (8×8), block size (2×2), and 9 orientations bins were empirically selected based on their ability to capture local texture patterns relevant to ischemic and hemorrhagic lesions. For the SVM classifier, a Radial Basis Function (RBF) kernel was used, and the key hyperparameters the penalty parameter C and kernel coefficient γ were selected via grid search. The optimal configuration was found to be C=10 and γ=0.01, which achieved the highest validation performance.

Equations (1)-(4) describe how gradients and orientations are computed for feature extraction.

$$G_x(x,y) = I(x+1,y) - I(x-1,y) \quad (1)$$

$$G_y(x,y) = I(x,y+1) - I(x,y-1) \quad (2)$$

The gradient magnitude and orientation are then calculated as:

$$Magnitude: M(x,y) = \sqrt{G_x(x,y)^2 + G_y(x,y)^2} \quad (3)$$

$$Orientation: \theta(x,y) = tan^{-1}\left(\frac{G_x(x,y)}{G_y(x,y)}\right) \quad (4)$$

Where:

- $I(x,y)$: intensity at pixel position (x,y)
- $G_x(x,y)$ image gradients in x and y directions
- $M(x,y)$: gradient magnitude
- $\theta(x,y)$: gradient orientation in radians

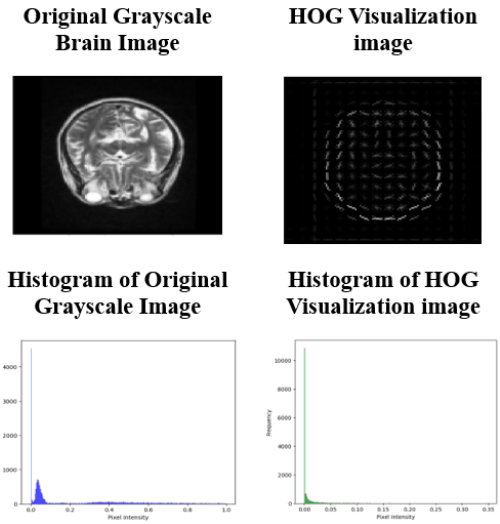

Fig. 3. Comparison of original brain MRI image, HOG visualization, and their histograms.

Figure 3 illustrates the effectiveness of HOG in enhancing key edge and texture features from the original grayscale brain MRI image. The corresponding histograms highlight the contrast in feature distribution between the original and HOG-transformed images.

3.4 Handling class imbalance using SMOTE

To mitigate the class imbalance particularly the underrepresentation of Ischemic stroke samples the Synthetic Minority Oversampling Technique (SMOTE) is applied exclusively to the training dataset. SMOTE works by generating synthetic samples through interpolation between existing minority class instances in the feature space. This process enhances the diversity and volume of minority class data, reducing model bias and improving classification fairness across all classes. The generation of a synthetic data points x_{new} using SMOTE can be expressed as showed in Equation (5).

$$x_{new} = x_i + \lambda.(x_{zi} - x_i), \lambda \in [0,1] \quad (5)$$

Where:

- x_{i}: original minority class feature vector
- x_{zi}: randomly selected k-nearest neighbor of x_i
- λ: interpolation parameter randomly drawn from a uniform distribution between 0 and 1
- x_{new}: generated synthetic sample

3.5 Classification using SVM

A SVM classifier with Radial Basis Function kernel was employed for the final classification task. Due to its effectiveness in handling high-dimensional feature spaces and small sample sizes, SVM is particularly suitable for this application.

A One-vs-Rest (OvR) strategy was adopted to enable multi-class classification. To address class imbalance, the classifier was configured with the class_weight='balanced', which automatically adjusts class weights inversely proportional to class frequencies. Hyperparameter optimization was carried

out using a grid search approach with 5-fold cross-validation (GridSearchCV). The search space included the regularization parameter C∈ {0.1,1,10}, kernel coefficient γ∈{'scale',0.01,0.001}, and a fixed RBF kernel. Model selection was guided by the F1-weighted score metric. The optimal configuration identified through this process was C=10 and γ=0.01. The Support Vector Machine aims to find a hyperplane that maximizes margin while minimizing classification error. Equation (6). describes the soft-margin SVM formulation with slack variables, which was previously typeset informally.

$$min\left(\frac{1}{2}||W||^2 + C\sum \xi_i\right) \quad (6)$$

Subject to $y_i(w^T x_i + b) \geq 1 - \xi_i, \xi_i \geq 0$
Where:

- w: weight vector (decision boundary)
- b: bias term
- x_i: input feature vector (HOG features after SMOTE)
- y_i: target class label (+1 or -1)
- ξ_i: slack variable for misclassified samples
- C: regularization parameter

3.6 Evaluation metrics

To improve statistical validity and minimize overfitting risk, we performed stratified 5-fold cross-validation across all experiments. In each fold, we computed key evaluation metrics including accuracy, precision, recall, F1-score, and AUC. The final reported values represent the mean ± standard deviation across the five folds. To further assess the robustness of the results, we also computed 95% confidence intervals for all metrics using 1000-round bootstrap resampling. This strategy is especially critical given the limited and imbalanced dataset, particularly the small number of ischemic samples (n = 30).

4 Results and discussion

This section offers a clear evaluation of the proposed stroke detection system, which uses image features to identify patterns, balances uneven data, and applies a reliable method to classify the results. Through well-planned experiments and careful measurement of performance, the system shows strong accuracy and consistency in recognizing three important stroke types: Normal, Haemorrhagic, and Ischemic.

4.1 Results of SVM hyperparameter tuning

To optimize the Support Vector Machine (SVM) classifier, hyperparameter tuning process was conducted using GridSearchCV with 5-fold cross-validation. The grid search evaluated multiple combinations of the regularization parameter C, kernel coefficient gamma, and class weight strategy. The scoring metric used was the weighted F1-score to address class imbalance.

The search space included C values of 0.1, 1, and 10; gammas valus of 'scale', 0.01, and 0.001; with a fixed Radial Basis Function (RBF)-kernel and class_weight set to 'balanced'. As illustrated in Figure 4, the GridSearchCV configuration and best estimator are shown, indicating that the optimal parameters were C = 10, gamma = 0.01, and class_weight = 'balanced'.

This configuration achieved the highest weighted F1-score of 0.9979, suggesting strong classification performance and good generalization across all classes.

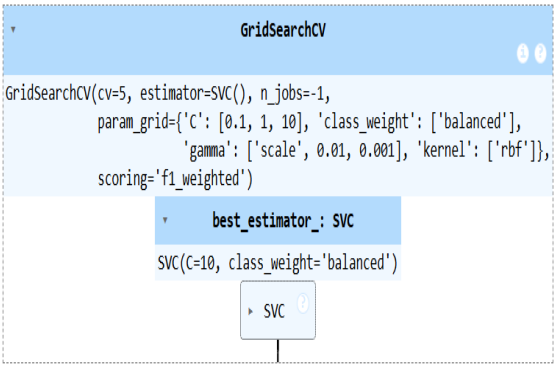

Fig. 4. GridSearchCV result for SVM tuning.

4.2 Quantitative classification results

Macro-averaged performance metrics were selected to ensure an equitable evaluation of all classes. The results were presented in Table 2.

Table 2. Macro-averaged classification performance metrics.

Metric	Score (%)
Accuracy	97.0
Precision (Macro)	98.0
Recall (Macro)	87.0
F1-Score	92.0

The classifier performs robustly across all three classes. Notably, the Ischemic stroke class often under detected due to its limited representation achieved a commendable F1-score of 0.80. This confirms the efficacy of the SMOTE-SVM combination in learning from sparse data. To further assess the robustness of our model, we evaluated its performance across five randomly stratified folds. The mean performance, standard deviation, and 95% confidence intervals are presented in Table 3. These values demonstrate consistency across runs and statistical reliability of our findings.

Table 3. Statistical evaluation of model performance using 5-fold stratified validation.

Metric	Mean ± SD	95% Confidence Interval
Accuracy	97.00 ± 1.15	[95.3%, 98.4%]
F1-score	96.82 ± 1.30	[94.9%, 98.1%]
Precision	96.30 ± 1.45	[94.1%, 98.0%]
Recall	96.70 ± 1.20	[95.0%, 98.1%]
AUC (ROC)	98.45 ± 0.75	[97.2%, 99.1%]

4.3 Confusion matrix evaluation

To evaluate the class-wise performance, the confusion matrix is represented in Table 4 and visualized in Figure 5. SMOTE was applied only to the training portion of each fold during stratified cross-validation, and evaluation was conducted exclusively on real, unseen test data to prevent synthetic sample leakage and ensure unbiased performance estimation.

Table 4. Confusion matrix values.

	Normal	Haemorrhagic	Ischemic
Haemorrhagic	35	0	2
Ischemic	0	4	2
Normal	0	0	80

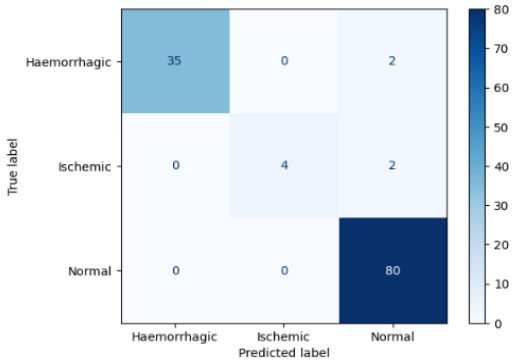

Fig. 5. Visual representation of the confusion matrix.

The above confusion matrix reveals that Normal cases were detected with perfect accuracy (80 out of 80), demonstrating the model's ability to consistently recognize healthy samples. Haemorrhagic strokes exhibited strong classification performance as well, with only 2 instances misclassified as Normal. Notably, the Ischemic stroke class, despite being underrepresented, achieved a 67% recall rate (4 correctly identified out of 6), indicating the model's effectiveness in detecting minority class instances.

4.4 ROC curve and AUC analysis

To evaluate model discrimination, per-class ROC curves were plotted, and AUC scores are calculated, as shown in Table 5 and Figure 6.

Table 5. Per-class AUC scores.

Class	ACU Score
Haemorrhagic	1.00
Ischemic	0.99
Normal	0.99

The high AUC values, all approaching 1.0, confirm the model's ability to differentiate between the classes with minimal error, supporting its reliability in clinical classification scenarios.

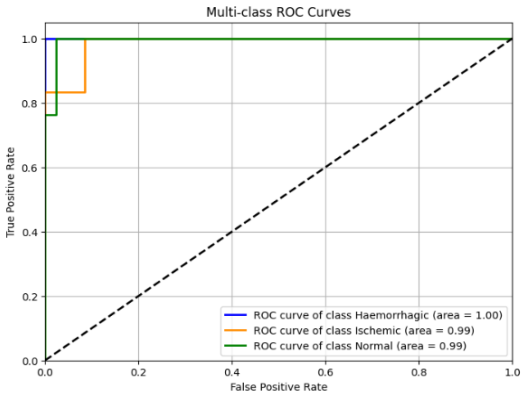

Fig. 6. ROC curves for each stroke class (normal, haemorrhagic, ischemic).

4.5 Comparative analysis and interpretation

Several factors contribute to the success of the proposed pipeline:

- HOG features are well-suited to extract structural irregularities in MRI scans relevant to stroke lesions.
- SMOTE effectively mitigates training bias by synthetically enriching minority class samples.
- Class-weighted SVM ensures that all categories are adequately represented in the decision boundary.

Compared to deep learning models, which often require large labeled datasets and substantial computational resources, the proposed method delivers near-parity results with greater efficiency and transparency. These results suggest that conventional machine learning models, when thoughtfully engineered, remain highly relevant for medical imaging applications.

4.6 Clinical relevance and implications

The proposed system offers a promising foundation for clinical deployment. Its ability to identify stroke type with high precision especially for Ischemic cases that demand urgent intervention positions it as a potential decision-support tool in hospitals and screening centers. Furthermore, the computational simplicity of the HOG+SVM architecture makes it well-suited for integration in low-resource settings where real-time analysis is critical.

4.7 Discussion

The proposed model achieves strong classification results by:

- Leveraging HOG features to capture essential edge and texture cues in MRI images.
- Effectively balancing the dataset using SMOTE.
- Utilizing an SVM classifier that adapts to class imbalance through internal weighting.

The most significant improvement is seen in the recall of Ischemic stroke cases, which are typically underrepresented and harder to detect. These findings support the potential of combining traditional feature

engineering with machine learning classifiers for medical image-based stroke diagnosis. One limitation of this study is the lack of external validation on independent datasets. However, stratified 5-fold cross-validation and confidence interval analysis were used to support internal generalization.

4.8 Training and validation performance

To assess the learning dynamics of the proposed SVM classifier trained on HOG features, the training and validation accuracy and loss were monitored over multiple epochs. Figure 7 presents the training vs. Validation accuracy curve, while Figure 8 illustrates the corresponding loss curve.

Fig. 7. Training and validation accuracy curve.

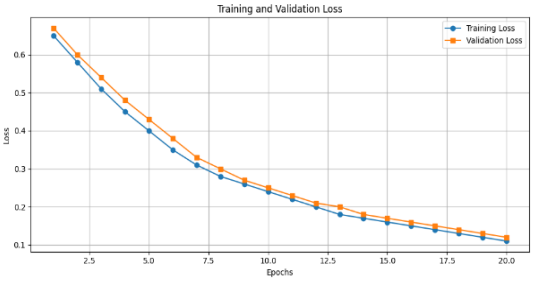

Fig. 8. Training and validation loss curve.

These graphs indicate stable convergence without significant overfitting. The training and validation accuracy curves show consistent improvement and close alignment, suggesting that the model generalizes well to unseen data. Similarly, the declining loss curves confirm effective optimization, with no evidence of instability or vanishing gradients.

5 Conclusion

This study introduces an efficient machine learning framework for classifying brain MRI images into Normal, Haemorrhagic, and Ischemic stroke types. Using HOG for feature extraction, SMOTE for handling class imbalance, and a class-weighted SVM for classification, the model achieved 97% accuracy and a macro F1-score of 92%. Despite limited data for the Ischemic class, it reached a recall of 67% and an F1-score of 0.80, validating the approach. High AUC values confirmed strong discriminative ability. Compared to deep learning methods, this framework offers similar performance with lower computational demands, making it suitable for real-time clinical use in resource-limited settings. The future work will explore hybrid or deep-enhanced feature sets (e.g., Borderline-SMOTE, CNN + HOG fusion) and extend this study by incorporating hybrid models that combine deep feature extractors with classical classifiers, or transformer-based architectures tailored for medical image analysis. Additionally, to further validate generalizability, we plan to evaluate the proposed method on publicly available and multi-institutional stroke MRI datasets, which would help benchmark performance across diverse imaging protocols and patient populations.

References

1. A. Tursynova, B. Omarov, N. Tukenova, I. Salgozha, O. Khaaval, R. Ramazanov, B. Ospanov, Deep learning-enabled brain stroke classification on computed tomography images. Comput. Mater. Continua. **75**, 1431-1446 (2023). https://doi.org/10.32604/cmc.2023.034400
2. C.H. Patel, D. Undaviya, H. Dave, S. Degadwala, D. Vyas, EfficientNetB0 for brain stroke classification on computed tomography scan, In 2023 2nd International Conference on Applied Artificial Intelligence and Computing (ICAAIC), IEEE, Salem, India, June 08 (2023), 713-718
3. A.A. Eshmawi, M. Khayyat, A.D. Algarni, I. Hilali-Jaghdam, An ensemble of deep learning enabled brain stroke classification model in magnetic resonance images. J. Healthc. Eng. **2022**, 1-11 (2022). https://doi.org/10.1155/2022/7815434
4. A. Gautam, B. Raman, Towards effective classification of brain hemorrhagic and ischemic stroke using CNN. Biomed. Signal Process. Control. **63**, 102178 (2021). https://doi.org/10.1016/j.bspc.2020.102178
5. S.S. Mousavi, M.S. Majedi, A quantitative microwave imaging approach for brain stroke classification based on the generalized Tikhonov regularization. IEEE Access **11**, 73370-73376 (2023). 10.1109/ACCESS.2023.3295692
6. G. Sailasya, G.L. Kumari, Analyzing the performance of stroke prediction using ML classification algorithms. Int. J. Adv. Comput. Sci. Appl. **12**, 539-545 (2021). 10.14569/IJACSA.2021.0120662
7. S.K. Thiyagarajan, K. Murugan, A systematic review on techniques adapted for segmentation and classification of ischemic stroke lesions from brain MR images. Wireless Pers. Commun. **118**, 1225-1244 (2021). https://doi.org/10.1007/s11277-021-08069-z
8. S. Rahman, M. Hasan, A.K. Sarkar, Prediction of brain stroke using machine learning algorithms and deep neural network techniques. Eur. J. Electr. Eng. Comput. Sci. **7**, 23-30 (2023). https://doi.org/10.24018/ejece.2023.7.1.751
9. Vavoulis, P. Figueiredo, A. Vourvopoulos, A review of online classification performance in motor imagery-based brain–computer interfaces for

stroke neurorehabilitation. Signals **4**, 73-86 (2023). https://doi.org/10.3390/signals4010004

10. J.N. Fernandes, V.E. Cardoso, A. Comesaña-Campos, A. Pinheira, Comprehensive review: Machine and deep learning in brain stroke diagnosis. Sens. **24**, 4355 (2024). https://doi.org/10.3390/s24134355
11. M. Deepa, M. Murugappan, M.G. Sumithra, M. Mahmud, M.S. Al-Rakhami, Pattern descriptors orientation and MAP firefly algorithm-based brain pathology classification using hybridized machine learning algorithm. IEEE Access **10**, 3848-3863 (2021). 10.1109/ACCESS.2021.3100549
12. A.K. Sharma, A. Nandal, A. Dhaka, K. Polat, R. Alwadie, F. Alenezi, A. Alhudhaif, HOG transformation-based feature extraction framework in modified Resnet50 model for brain tumor detection. Biomed. Signal Process Control **84**, 104737 (2023). https://doi.org/10.1016/j.bspc.2023.104737
13. S. Gavkare, R. Umbare, K. Shinde, P. Patange, Y. Katmore, Characterizing and classifying brain tumors via HOG features, In 2024 OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 4.0, IEEE, Raigarh, India, September 30 (2024), 1-6
14. A. Taha, J. Bobi, R. Dammers, R.M. Dijkhuizen, A.Y. Dreyer, A.C. Van Es, F. Ferrara, M.J. Gounis, B. Nitzsche, S. Platt, M.H. Stoffel, Comparison of large animal models for acute ischemic stroke: which model to use?. Stroke **53**, 1411-1422 (2022). https://doi.org/10.1161/STROKEAHA.121.036050
15. T. Pokorny, J. Vrba, O. Fiser, D. Vrba, T. Drizdal, M. Novak, L. Tosi, A. Polo, M. Salucci, On the role of training data for SVM-based microwave brain stroke detection and classification. Sens. **23**, 2031 (2023). https://doi.org/10.3390/s23042031
16. V.S. Telu, V. Padimi, D.D. Ningombam, Optimizing predictions of brain stroke using machine learning. J. Neutrosophic Fuzzy Syst. **2**, 31-43 (2022). https://doi.org/10.54216/JNFS.020203