INTERNATIONAL JOURNAL OF ADVANCED RESEARCH
ISSN: 3106-7980 : 3106-7972
Volume 3, Issue 3 [July, 2026]

FUSION BASED IDENTIFICATION AND CLASSIFICATION MODEL FOR β - THALASSEMIA DISEASE STAGES

Shahan Yamin Siddiqui[*1], Faryal Tahseen Sadiq[2], Dr. Zoya Rehman[3], Unaiza Rehman[2,4], Muhammad Waleed[5], Abdullah Dar[6], Muhammad Farrukh Khan[7]

[*1]*Department of Computer Science, NASTP Institute of Information Technology, Lahore 54810, Pakistan.*
[2]*Department of Computer Science, Lahore College for Women University, Lahore 54500, Pakistan.*
[3]*Rehman Clinic, Lahore, Pakistan.*
[4]*Department of Software Engineering, NASTP Institute of Information Technology, Lahore 54810, Pakistan.*
[5]*Department of Computer Science, Barani Institute of Sciences, Sahiwal, Pakistan.*
[6]*Global Reporting and Analysis Department, Ibex Global Solution, Lahore, Pakistan.*
[7]*Department of Artificial Intelligence, NASTP Institute of Information Technology, Lahore 54810, Pakistan.*

[*1]drshahan@niit.edu.pk,[2]faryaltahseen20@gmail.com,[3]zoyarehmansheikh@hotmail.com,
[2,4]unaizarehman@niit.edu.pk,[5]muhammadwaleed497@gmail.com,[6]abdullah.dar@ibex.co,
[7]dr.farrukhkhan@niit.edu.pk

DOI: https://doi.org/10.5281/zenodo.22121570

Keywords
β-Thalassemia, Carrier Screening, Feature-Level Fusion, Fuzzy Logic, Support Vector Machine, Convolutional Neural Network, Complete Blood Count.

Article History
Received: 15 May 2026
Accepted: 15 July 2026
Published: 30 July 2026

Corresponding Author: *
Shahan Yamin Siddiqui

Abstract

β-Thalassemia is an autosomal recessive haemoglobinopathy caused by mutations in the HBB gene that reduce or abolish β-globin chain synthesis. Because carriers are clinically silent, undetected carrier couples remain the principal route through which transfusion-dependent thalassemia is transmitted to the next generation. Confirmatory laboratory procedures such as hemoglobin electrophoresis and high-performance liquid chromatography are accurate but costly, slow and unavailable in many low-resource settings, which motivates automated screening from routinely collected data. This study presents a fusion-based decision model that identifies β-thalassemia carriers by combining two complementary modalities. A support vector machine was trained on 2,001 complete blood count reports obtained from the Punjab Thalassemia Prevention Program, and a convolutional neural network was trained on 80 peripheral blood smear images from which 64 color, texture and morphological descriptors were extracted per segmented erythrocyte. The two feature spaces were then combined through a fuzzy-logic feature-level fusion layer. All models were evaluated using a 70:30 training-validation partition. The support vector machine and the convolutional neural network attained validation accuracies of 95.00% and 95.83% respectively, whereas the proposed fuzzy fusion model reached 96.15% accuracy with a 3.85% miss-rate, 95.45% precision and 96.71% recall. The findings indicate that fusing hematological indices with erythrocyte morphology yields more dependable carrier identification than either modality used alone.

INTRODUCTION

Thalassemia is among the most widespread monogenic disorders in the world. It arises when mutations in the globin genes disturb the balanced synthesis of the α- and β-polypeptide chains that together form the hemoglobin tetramer,

producing chronic hemolytic anemia, ineffective erythropoiesis and progressive iron overload. β-Thalassemia, which results from mutations in the HBB gene, is the clinically dominant form and follows an autosomal recessive pattern of inheritance: when both parents carry the trait, each pregnancy carries a one-in-four probability of a transfusion-dependent child [1]. The problem is centered mainly around the Mediterranean basin, Middle East, and South Asia. Pakistan is one of the countries where thousands of patients have been reported who require therapy in private and government health institutions and where the practice of consanguineous marriage leads to a high carrier rate [2], [3].

The problem clinically manifests in that carriers of the β-thalassemia gene are basically asymptomatic. They show up with mild microcytic and hypochromic anemia, which is easily misdiagnosed as iron deficiency anemia, which is probably the most common cause of microcytosis in the world. It is important to distinguish between the two because giving a thalassemia carrier iron supplements would be either useless or even dangerous [4]. Discrimination is possible only through Hb A2 quantitation using HPLC or hemoglobin electrophoresis, together with serum ferritin and transferrin saturation levels [5]. This needs special instruments, trained technicians, and considerable time; and is, therefore, simply unavailable at the primary care level in most areas with a high prevalence. And because prevention through premarital and antenatal counseling depends on identification of carriers before conception, the cost and unavailability of these tests are real bottlenecks in the national program.

The complete blood count, by contrast, is inexpensive, automated and performed almost universally. This has driven sustained interest in deriving a screening decision directly from red cell indices. Early efforts produced discriminant formulae built on hemoglobin, red cell count, mean corpuscular volume, mean corpuscular hemoglobin and red cell distribution width; more recent work has replaced hand-crafted formulae with supervised learning. Classifiers based on support vector machines, ensemble trees, gradient boosting and deep neural networks have all been reported to separate carriers from non-carriers using hematological indices alone [6], [7], [8], [9]. A parallel line of research treats the problem visually, applying image processing and convolutional networks to peripheral blood smears in order to recognize the abnormal erythrocyte morphologies characteristic of thalassemia [10], [11].

Each modality carries a distinct weakness. Numerical indices summaries the average behavior of the whole red cell population and therefore discard the shape information that a hematologist relies on when examining a slide. Image-based classifiers capture that morphology but are sensitive to staining protocol, illumination and magnification, and are typically trained on small annotated datasets. Multimodal fusion offers a principled way to exploit their complementarity, and reviews of multimodal learning in medical diagnostics consistently report that integrating heterogeneous sources improves robustness relative to any single stream [12]. Fusion may be performed at the feature level, where descriptors from each modality are concatenated into a joint representation before classification, or at the decision level, where the outputs of independently trained models are combined. Fuzzy logic is particularly well suited to the former in a clinical context, because hematological thresholds are inherently graded rather than crisp: a mean corpuscular volume of 79 FL is not categorically different from 81 FL, yet a hard cut-off treats it as such.

Despite this, the great majority of published thalassemia screening systems remain single-modality. Where fusion has been applied to β-thalassemia, it has generally operated on a single feature-based dataset combined at the decision stage [13], leaving the integration of quantitative blood-count indices with erythrocyte morphology largely unexamined. The present study addresses that gap. It develops a fuzzy-logic feature-level fusion model that jointly analyses complete blood count reports and peripheral blood smear images, and benchmarks it against a support vector

machine and a convolutional neural network trained separately on the same underlying data.

The principal contributions of this research are as follows:

- A fusion-based screening model for β-thalassemia carrier identification that unifies hematological indices and blood smear morphology within a single decision pipeline.
- A fuzzy-logic feature-level fusion layer that accommodates the graded nature of clinical reference ranges instead of imposing crisp thresholds.
- Extraction of sixty-four color, texture and morphological descriptors, including distance-angle signature features of the erythrocyte and its central pallor, from segmented single-cell sub-images.
- A controlled comparison in which the support vector machine, the convolutional neural network and the fusion model are evaluated under an identical 70:30 partitioning protocol.
- Benchmarking of the resulting accuracy against state-of-the-art approaches reported in the thalassemia screening literature.

Related Work

Research on automated thalassemia screening has evolved along three lines discriminant formulae from red cell indices, supervised learning on hematological data, and image-based smear analysis with a fourth and more recent strand concerning the fusion of heterogeneous sources. Formula-based discrimination remains the reference point against which learned models are judged. Saha et al. analyzed 21,000 individuals in Eastern India and derived a stepwise discriminant expression combining red cell count, mean corpuscular hemoglobin, hematocrit, mean corpuscular hemoglobin concentration and red cell distribution width, reporting roughly 89% accuracy [14]. Such indices can be computed by hand, but their performance varies with age, sex and the local prevalence of iron deficiency, and none attempts to grade disease severity.

Supervised learning on complete blood count data has produced consistently stronger results. Laengsri et al. developed ThalPred, a publicly accessible web tool that discriminates the thalassemia trait from iron deficiency anemia and reported approximately 95.6% accuracy [15]. Çil et al. tested extreme learning machines, regularized extreme learning machines, support vector machines, logistic regression and k-nearest neighbors for the same discrimination problem and reported that their best configuration reached 94.37% [16]. Sadiq et al. combined random forest, support vector machine and gradient boosting classifiers in an ensemble trained over red cell indices extracted from the Punjab Thalassemia Prevention Program and achieved around 93% accuracy [17]. Fu et al. developed the TVGH-NYCU Thal Classifier on general blood counter indices with support vector machine and validated the model using Monte-Carlo cross-validation and got a considerably lower area under the curve of 0.76, showing how critical is the dependence on sample characteristics [18].

Recent studies scaled these methods up: Schipper et al. predicted hemoglobinopathies on complete blood count data on a large-scale clinical laboratory population [6], Mo et al. trained deep neural network using 8,693 patient records and achieved 96% accuracy [8], and Rustam et al. used principal component analysis followed by a convolutional architecture on 5,066 self-reported cases and got 96% [7]. Das et al. evaluated different machine learning algorithms along with traditional screening formulae for the screening of β-thalassemia trait among antenatal Indian women, reporting 86.6% accuracy and showing the complexity of the issue due to incomplete records [9]. Uçucu and Azik confirmed the same picture for the differential diagnosis of β-thalassemia minor and iron deficiency anemia using automated machine learning on routine counts [19].

Image-based methods focus directly on erythrocyte morphology. Tyas et al. segmented peripheral blood smear images into single-cell sub-images and classified nine types of erythrocytes using multi-layer perceptron trained on geometric parameters, moment invariants and distance-angle signature descriptors of both the erythrocyte and its central pallor; combination of these morphological

properties with color and texture features increased accuracy to 93.77% [10]. This research group had compared convolutional and artificial neural networks for the classification of abnormal red cell images in thalassemia minor before [20]. Khan et al. used obileNetV2 architecture to classify hemoglobin electrophoresis images and reported 95.72% accuracy [11], and Jarujareet et al. used image structure function representation obtained by a compact differential dynamic microscopy instrument to classify β-thalassemia major from HbE/β-thalassemia [21]. These systems demonstrate that morphology carries genuine diagnostic signal, but they depend on image acquisition conditions that are harder to standardize than a blood count.

Strategies that use fusion have been investigated more in areas other than cardiology. Nadeem et al. created a fusion framework at the decision level using fuzzy logic for predicting heart diseases, where they trained SVM models in parallel using two data sets and by merging the results, they achieved 96.23% accuracy [22]. Shankar et al. used fusion of gray-level co-occurrence matrix, local binary pattern and gray-level run-length features for COVID-19 classification from chest radiograph images, which resulted in 95.65% accuracy for multi-class classification [23]. Vijayakumar et al. employed feature-level fusion to achieve 96% accuracy for electrocardiogram interpretation [24], and Khan et al. employed a two-stage fusion classifier using phonocardiogram signatures for detecting and grading coronary artery disease with 88% accuracy [25]. Review articles on multimodal deep learning in medical diagnostic tasks have pointed out that such increases in performance generalize over imaging, omics, and structured clinical data [12].

Within thalassemia specifically, fusion remains sparsely applied. Purwar et al. combined deep image features with clinical features to classify thalassemia patients [26], and Ibrahim et al. proposed a late-fusion model over four machine learning algorithms for β-thalassemia carrier prediction, achieving 96% overall accuracy from a feature-based dataset [13]. Susanto et al. implemented a fuzzy inference system that graded patients as major, intermediate, minor or normal, but validated it on only four complete blood count reports [27]. Akhtar et al. applied linear discriminant analysis to select risk factors for thalassemia prediction, reporting approximately 78% accuracy without disease-level identification [28]. Two observations follow from this body of work. First, no reported system integrates quantitative blood count indices with erythrocyte morphology in a single feature-level fusion layer. Second, the fuzzy formulations that do exist have not been evaluated on datasets of meaningful size. The methodology described next is designed to address both points.

Proposed Methodology

The proposed system is a supervised, fusion-based decision framework for identifying β-thalassemia carriers. It operates in two phases. During the training phase, two learning algorithms are applied in parallel to two distinct modalities: a support vector machine analyses feature-based hematological data, while a convolutional neural network analyses image-based data. The outputs of both branches are then combined through a fuzzy-logic feature-level fusion layer. During the validation phase, the trained model is applied to previously unseen samples and the fused decision is evaluated against the recorded ground truth. All experiments were implemented and simulated in MATLAB.

Each training branch is organized into four layers. The data acquisition layer retrieves raw records from the data source and stores them in the system database. The pre-processing layer resolves missing and noisy values for the feature-based stream, and performs resizing and segmentation for the image-based stream. The application layer executes the classifier, and the performance layer measures accuracy and miss-rate against a predefined criterion; if the criterion is not met, the model is retrained. Results that satisfy the criterion are passed forward to the fusion stage. Figure 1 presents the complete data training architecture.

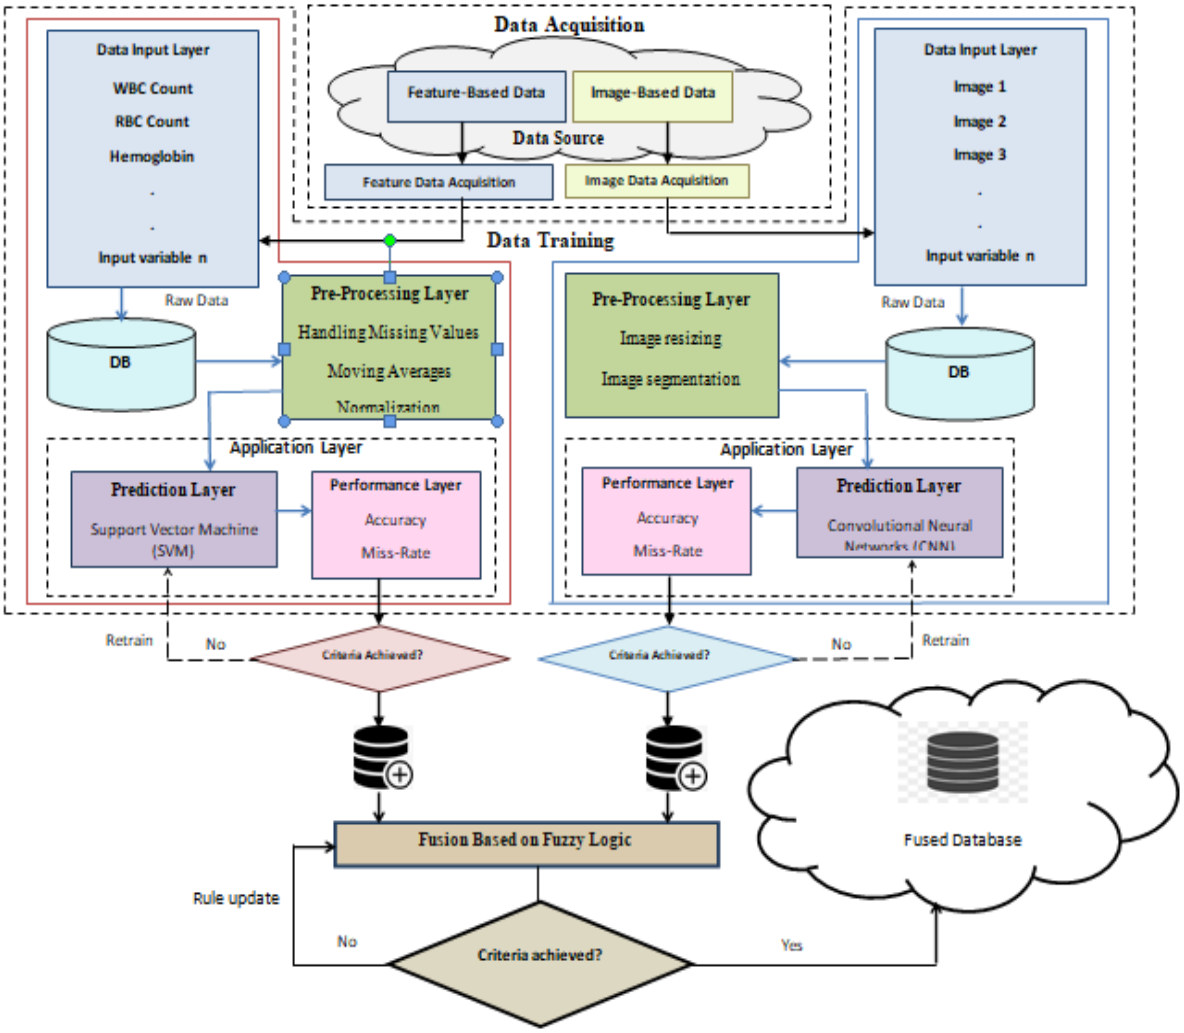

Figure 1: Proposed Model for Identification and Classification of β- Thalassemia Disease

A. Data Sources

Two datasets were used. The feature-based dataset comprises 2,001 complete blood count reports obtained from the Punjab Thalassemia Prevention Program (PTPP), a government initiative established for the prevention and control of thalassemia in Punjab, Pakistan. It contains 1,026 non-carrier samples (class 0) and 975 carrier samples (class 1). The image-based dataset comprises 80 peripheral blood smear images acquired through a microscopic camera and obtained from a public repository, evenly divided into 40 thalassemia-affected and 40 healthy smears. Both were partitioned 70:30 for training and validation.

B. Support Vector Machine Branch

The support vector machine is a supervised classifier that separates classes by constructing the hyperplane which maximizes the margin between the closest points of opposing classes; these points are the support vectors. Once the optimal hyperplane is established, an unseen sample is assigned to a class according to the side of the boundary on which it falls. Eight clinical features were acquired from the automated blood analyzer. Table 1 lists these features together with their units and reference ranges.

Table 1: Clinical Features Obtained from the Blood Analyzer

#	Feature Description	Unit	Normal Range
1	White blood cell count (WBC)	Million/mm³	4.0 – 11.0
2	Red blood cell count (RBC)	Million/mm³	3.76 – 5.70
3	Hemoglobin concentration (HB)	g/dL	M: 14–18; F: 12–15
4	Hematocrit (HCT)	%	33.5 – 52.0
5	Mean corpuscular volume (MCV)	Femtolitre	80 – 100
6	Mean corpuscular hemoglobin (MCH)	Pictogram	28.0 – 32.0
7	Mean corpuscular hemoglobin concentration (MCHC)	g/dL	31.0 – 35.0
8	Platelet count (PLT)	/µL	150,000 – 350,000

Pre-processing of this stream consisted of data cleaning followed by normalization. Cleaning removed duplicate entries by tracking the unique patient identifier, completed incomplete inputs by reference to the source records, and discarded attributes with no bearing on classification, such as patient name and cityNormalization was required due to the fact that the eight features were measured in terms of different units of measure and also some of the reference intervals are sex-based. As a result, sex was included as a binary feature, and all measurements were coded using the ordinal scale with values between 0 and 5, where 0 means that the value is smaller than the reference interval, 5 means that the value is larger than the reference interval, while 1 to 4 refer to the reference interval.

C. Convolutional Neural Network Branch

Microscopic images of peripheral blood smears were resized to 800 × 600 pixels, after which the green channel was extracted from the RGB representation. The green channel was selected because it is comparatively insensitive to variation in brightness and staining intensity. Residual noise was suppressed using a median filter.

Segmentation produced sub-images containing individual erythrocytes through banalization, morphological operations, watershed distance transformation and whole filling. After object detection and labelling, a threshold on object area was applied to distinguish single cells from overlapping clusters, on the assumption that overlapping cells occupy a larger area; the circular Hough transform was then used to separate the touching cells. The result is a collection of single-cell sub-images suitable for feature extraction.

Sixty-four descriptors were extracted from each sub-image across three families. Color features were captured by the first four color moments of the green channel: mean, standard deviation, skewness and kurtosis, which describe respectively the average color, the spread of the distribution, the asymmetry of color irregularity and the degree of tapering. Texture features were obtained from gray-level co-occurrence matrices, a second-order statistic that considers relationships between pixel pairs rather than isolated pixel values; contrast, correlation, energy and homogeneity were computed at orientations of 0°, 45°, 90° and 135°. Shape features comprised seven moment invariants together with geometric parameters including area, perimeter, major and minor axis lengths, compactness, eccentricity and solidity. The distance-angle signature was derived by converting the cell to a binary image and plotting the centroid-to-boundary distance as a function of angle, with the starting point fixed at the minimum radius so that the descriptor is rotation-invariant. Because the central pallor adopts characteristically abnormal shapes in affected erythrocytes, its geometry was extracted separately by Otsu thresholding, labelling and erosion. Table 2 summarizes the feature set.

Table 2: Features Obtained Through Feature Extraction

Type	Features Extracted	Count
Colour	Colour moments of green channel (mean, standard deviation, skewness, kurtosis)	4
Shape	Geometric parameters of cell	7
	Moment invariants (Hu's moments) of cell	7
	DAS morphology of central pallor	9
	Geometric parameters of central pallor	8
	DAS morphology of cell	9
Texture	Gray-level co-occurrence matrix feature set	20
Total		64

Within the convolutional network, the pooling layer preserves the captured feature maps and does not participate in backpropagation, so its weights remain fixed. The output of the preceding layers is flattened into a one-dimensional array and passed to the fully connected layer, where the extracted information is consolidated. A softmax layer converts the resulting logits into class probabilities, and the accuracy values are recorded at the final layer.

D. Fuzzy Logic-based Fusion

In the third stage, the feature sets produced by the two branches are combined into a joint representation and a fuzzy inference layer is applied to reach the final decision. The rationale behind using fuzzy logic in this case is that the supporting evidence in medicine is graded and not crisp since red cell indices are not distinct at the borders but continuous, while erythrocyte morphology is a spectrum and not distinct classes. Instead of each modality deciding on which crisp label to use prior to fusion, the fusion stage maps the combined feature vector to membership values of the carrier and non-carrier classes and uses the decision mappings and criterion settings defined during training. Where the criterion is not satisfied, the rule base is updated and the model is retrained. Figure 2 illustrates the structure of the proposed fusion model.

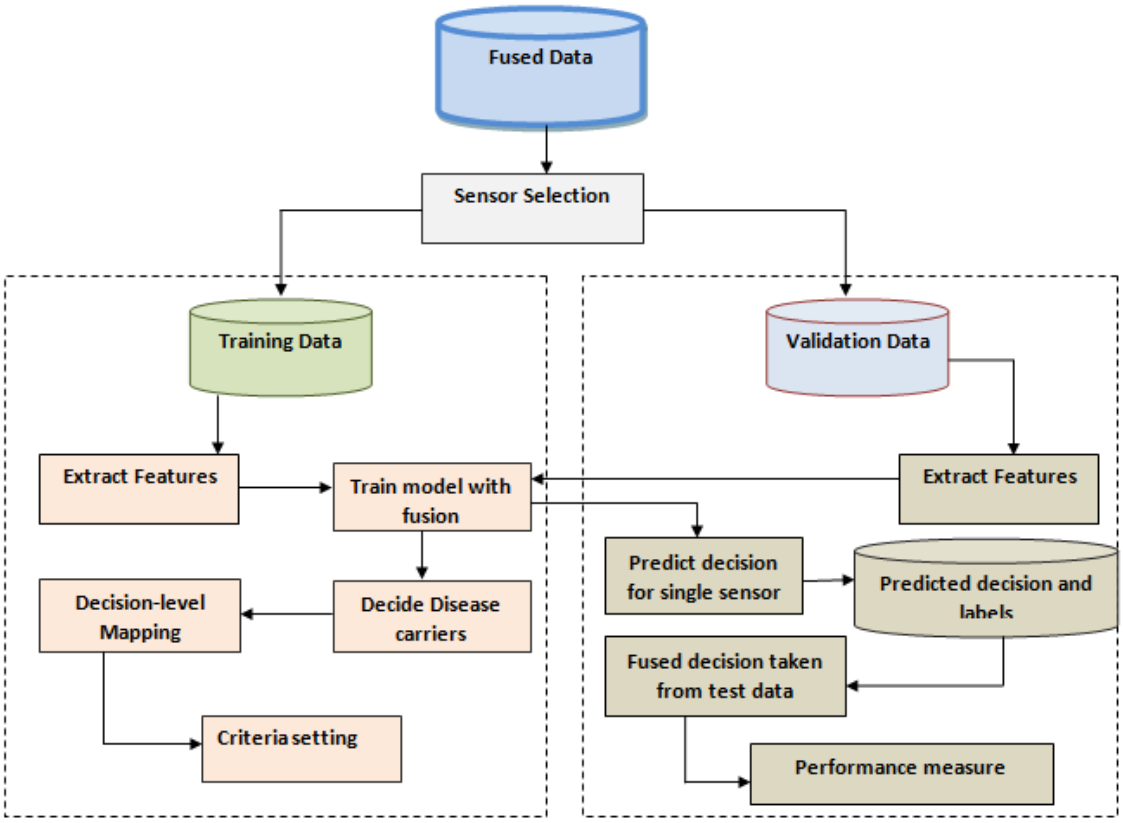

Figure 2: Proposed Model for Fuzzy Logic-Based Fusion

E. Validation Phase

Validation is the decisive stage of the implementation. Thirty per cent of each dataset was withheld from training and used to determine whether the system satisfies the required criterion. For the support vector machine, 600 complete blood count reports were reserved, comprising 308 non-carriers and 292 carriers. For the convolutional neural network, 24 blood smear images were reserved, divided evenly between the two classes. For the fusion model, 624 fused instances were reserved, comprising 320 non-carriers and 304 carriers. The validation pipeline mirrors the training pipeline through data acquisition, input, pre-processing, prediction and performance layers, with the trained model imported at the prediction stage and the fused features assembled before the fuzzy decision is taken.

Results of the Proposed Model

This section reports the outcome of all three implementations. The architecture classifies each subject into one of two categories: class 0, denoting a non-carrier, and class 1, denoting a β-thalassemia carrier. Under the 70:30 protocol the support vector machine was trained on 1,401 reports and validated on 600; the convolutional neural network was trained on 56 images and validated on 24; and the fusion model was trained on 1,765 fused instances and validated on 624. The class composition of each partition is given in the confusion matrices that follow.

Table 3 presents the training confusion matrices. The support vector machine returned 1,370 correct predictions from 1,401 samples, a training accuracy of 97.79% with a 2.21% miss-rate. The convolutional neural network returned 55 correct predictions from 56 images, giving 98.21% accuracy and a 1.79% miss-rate. The fusion model

returned 1,730 correct predictions from 1,765 fused samples, a training accuracy of 98.02% and a 1.98% miss-rate.

Table 3: Training Confusion Matrices

Model (Training)	Actual	Predicted NC	Predicted C
SVM (N = 1,401)	Non-carrier (718)	701	17
	Carrier (683)	14	669
CNN (N = 56)	Non-carrier (28)	27	1
	Carrier (28)	0	28
Fusion (N = 1,765)	Non-carrier (1,054)	1,031	23
	Carrier (711)	12	699

Table 4 presents the corresponding validation confusion matrices. On unseen data the support vector machine misclassified 16 non-carriers and 14 carriers, yielding 95.00% accuracy and a 5.00% miss-rate. The convolutional neural network misclassified a single non-carrier image, giving 95.83% accuracy and a 4.17% miss-rate. The fusion model misclassified 14 non-carriers and 10 carriers, for a validation accuracy of 96.15% and a 3.85% miss-rate.

Table 4: Validation Confusion Matrices

Model (Validation)	Actual	Predicted NC	Predicted C
SVM (N = 600)	Non-carrier (308)	292	16
	Carrier (292)	14	278
CNN (N = 24)	Non-carrier (12)	11	1
	Carrier (12)	0	12
Fusion (N = 624)	Non-carrier (320)	306	14
	Carrier (304)	10	294

Two patterns are worth noting. First, although the convolutional neural network records the highest training accuracy, it does so on only 56 images, and a single validation error costs a full 4.17 percentage points; that estimate is therefore far less stable than those from the larger datasets. Second, the fusion model is the only configuration whose validation accuracy exceeds both individual classifiers, and it does so while reducing the false-negative count relative to the support vector machine — the more consequential error in screening, since a missed carrier may proceed to an at-risk marriage without counselling.

Performance Evaluation of Classification Model

Model performance was quantified using accuracy, miss-rate, precision, recall (sensitivity), specificity and the F1-score, computed from the confusion matrix entries in the conventional manner, where TP and TN denote correctly identified carriers and non-carriers and FP and FN denote the corresponding errors:

$$Accuracy = (TP + TN) / (TP + TN + FP + FN)$$
$$Miss\text{-}Rate = 1 - Accuracy$$
$$Precision = TP / (TP + FP)$$
$$Recall = TP / (TP + FN)$$
$$Specificity = TN / (TN + FP)$$
$$F1\text{-}Score = 2 \times (Precision \times Recall) / (Precision + Recall)$$

Table 5 reports these metrics for the validation partition of each model.

Table 5: Validation Performance of the Three Models

Metric	SVM	CNN	Proposed Fusion
Accuracy	95.00%	95.83%	96.15%
Miss-rate	5.00%	4.17%	3.85%
Precision	94.56%	92.31%	95.45%
Recall (sensitivity)	95.21%	100.00%	96.71%
Specificity	94.81%	91.67%	95.63%
F1-score	94.88%	96.00%	96.08%

The fusion model returns the highest accuracy, precision, specificity and F1-score of the three configurations. The convolutional neural network records perfect recall, but this reflects the absence of any false negative among only twelve carrier images rather than a genuinely superior sensitivity, and its precision is correspondingly the lowest of the three. The support vector machine is the most stable of the two single-modality models by virtue of its far larger sample, yet it is outperformed by the fusion model on every metric. The improvement from 95.00% to 96.15% is modest in absolute terms, but it is obtained together with a reduction in both error types, which suggests that the two modalities are contributing genuinely complementary information rather than duplicating the same signal.

Table 6: Comparative Analysis with Existing Approaches

Study	Year	Technique	Accuracy
Akhtar et al. [28]	2020	Linear discriminant analysis	78.00%
Das et al. [9]	2022	Machine learning + screening formulae	86.60%
Saha et al. [14]	2020	Stepwise discriminant analysis	89.00%
Sadiq et al. [17]	2021	Ensemble (RF + SVM + GBM)	93.00%
Tyas et al. [10]	2020	Multi-layer perceptron (images)	93.77%
Çil et al. [16]	2020	ELM / RELM / SVM / k-NN	94.37%
Laengsri et al. [15]	2019	ThalPred (SVM)	95.59%
Khan et al. [11]	2022	MobileNetV2 (electrophoresis images)	95.72%
Rustam et al. [7]	2022	PCA + convolutional network	96.00%
Ibrahim et al. [13]	2024	Late fusion of four ML models	96.00%
Mo et al. [8]	2023	Deep neural network	96.00%
Proposed model	2026	Fuzzy-logic feature-level fusion (CBC + smear)	96.15%

Table 6 situates these results within the published literature. The comparison should be read with the usual caution, since the studies listed differ in cohort, class balance, geography and validation

protocol, and accuracy figures are therefore not strictly commensurable. Nevertheless, the proposed model compares favorably with reported approaches to β-thalassemia carrier identification, and in particular improves substantially on the linear discriminant analysis baseline that motivated this work.

Conclusion

β-Thalassemia remains a substantial public health burden in South Asia, and the cost and scarcity of confirmatory laboratory testing continue to limit the reach of prevention programs that depend on identifying asymptomatic carriers. This study developed a fusion-based decision model that combines two complementary sources of evidence for carrier screening. A support vector machine was trained on 2,001 complete blood count reports from the Punjab Thalassemia Prevention Program, and a convolutional neural network was trained on 80 peripheral blood smear images from which 64 color, texture and morphological descriptors were extracted per segmented erythrocyte. The two feature spaces were then integrated through a fuzzy-logic feature-level fusion layer.

Under an identical 70:30 partition, the support vector machine achieved 97.79% training and 95.00% validation accuracy, and the convolutional neural network achieved 98.21% and 95.83% respectively. The proposed fusion model achieved 98.02% training and 96.15% validation accuracy, with a 3.85% miss-rate, 95.45% precision, 96.71% recall and a 96.08% F1-score. Fusing hematological indices with erythrocyte morphology therefore produced a more dependable decision than either modality used alone, and reduced both false positives and false negatives relative to the feature-only classifier. Several limitations qualify these findings. The image dataset is small, at 80 smears, so the convolutional branch is the weakest-supported component of the system and its validation estimate carries wide uncertainty. Both datasets are drawn from a single regional context, which restricts generalization to populations with different mutation spectra and different background rates of iron deficiency. The model also performs binary carrier identification; grading severity across the major, intermediate and minor categories requires labelled data at those levels, which were not available here. Future work is, thus, recommended to increase the size of the image dataset, perform prospective validation of the developed model using data from several centers, take into consideration hemoglobin A2 levels, which will help differentiate thalassemia major from iron-deficiency anemia, and conduct multi-class severity classification. It will also be beneficial to calculate confidence intervals and perform statistical significance testing between models.

REFERENCES

[1] R. Galanello and R. Origa, "Beta-thalassemia," Orphanet Journal of Rare Diseases, vol. 5, art. 11, 2010.

[2] T. Tanveer, H. Masud, and Z. A. Butt, "Are people getting quality thalassemia care in twin cities of Pakistan? A comparison with international standards," International Journal for Quality in Health Care, vol. 30, no. 3, pp. 200–207, 2018.

[3] M. U. Nasir, M. Zubair, M. T. Naseem, T. Shahzad, A. Saeed, K. M. Adnan, and A. H. Gandomi, "Multiclass classification of thalassemia types using complete blood count and HPLC data with machine learning," Scientific Reports, vol. 15, art. 26379, 2025.

[4] J. J. Hoffmann, E. Urrechaga, and U. Aguirre, "Discriminant indices for distinguishing thalassemia and iron deficiency in patients with microcytic anemia: A meta-analysis," Clinical Chemistry and Laboratory Medicine, vol. 53, no. 12, pp. 1883–1894, 2015.

[5] V. Brancaleoni, E. Di Pierro, I. Motta, and M. D. Cappellini, "Laboratory diagnosis of thalassemia," International Journal of Laboratory Hematology, vol. 38, no. S1, pp. 32–40, 2016.

[6] A. Schipper, M. Rutten, A. van Gammeren, et al., "Machine learning-based prediction of hemoglobinopathies using complete blood count data," Clinical Chemistry, vol. 70, no. 8, pp. 1064–1075, 2024.

[7] F. Rustam et al., "Prediction of β-thalassemia carriers using complete blood count features," Scientific Reports, vol. 12, art. 19999, 2022.

[8] D. Mo, Q. Zheng, B. Xiao, and L. Li, "Predicting thalassemia using deep neural network based on red blood cell indices," Clinica Chimica Acta, vol. 543, art. 117329, 2023.

[9] R. Das, S. Saleh, I. Nielsen, et al., "Performance analysis of machine learning algorithms and screening formulae for β-thalassemia trait screening of Indian antenatal women," International Journal of Medical Informatics, vol. 167, art. 104866, 2022.

[10] D. A. Tyas, S. Hartati, A. Harjoko, and T. Ratnaningsih, "Morphological, texture, and color feature analysis for erythrocyte classification in thalassemia cases," IEEE Access, vol. 8, pp. 69849–69860, 2020.

[11] A. S. Khan et al., "Deep learning assisted automated assessment of thalassemia from haemoglobin electrophoresis images," Diagnostics, vol. 12, no. 10, art. 2405, 2022.

[12] A. M. Al-Zoghby, A. Ismail Ebada, A. S. Saleh, et al., "A comprehensive review of multimodal deep learning for enhanced medical diagnostics," Computers, Materials & Continua, vol. 84, no. 3, pp. 4155–4193, 2025.

[13] M. Ibrahim, S. Abbas, A. Fatima, T. M. Ghazal, M. Saleem, M. Alharbi, F. M. Alotaibi, M. A. Khan, M. Waqas, and N. Elmitwally, "Fuzzy-based fusion model for β-thalassemia carriers prediction using machine learning technique," Advances in Fuzzy Systems, vol. 2024, art. 4468842, 2024.

[14] S. Saha, S. Ghosh, K. Basu, and M. Bhattacharyya, "Prevalence of β-haemoglobinopathies in Eastern India and development of a novel formula for carrier detection," Journal of Hematopathology, vol. 13, pp. 159–164, 2020.

[15] V. Laengsri, W. Shoombuatong, W. Adirojananon, C. Nantasenamat, V. Prachayasittikul, and P. Nuchnoi, "ThalPred: A web-based prediction tool for discriminating thalassemia trait and iron deficiency anemia," BMC Medical Informatics and Decision Making, vol. 19, art. 212, 2019.

[16] B. Çil, H. Ayyıldız, and T. Tuncer, "Discrimination of β-thalassemia and iron deficiency anemia through extreme learning machine and regularized extreme learning machine based decision support system," Medical Hypotheses, vol. 138, art. 109611, 2020.

[17] S. Sadiq, M. U. Khalid, S. Ullah, W. Aslam, A. Mehmood, G. S. Choi, and B. W. On, "Classification of β-thalassemia carriers from red blood cell indices using ensemble classifier," IEEE Access, vol. 9, pp. 45528–45538, 2021.

[18] Y.-K. Fu et al., "The TVGH-NYCU Thal-Classifier: Development of a machine-learning classifier for differentiating thalassemia and non-thalassemia patients," Diagnostics, vol. 11, no. 9, art. 1725, 2021.

[19] S. Uçucu and F. Azik, "Artificial intelligence-driven diagnosis of β-thalassemia minor and iron deficiency anemia using machine learning models," Journal of Medical Biochemistry, vol. 43, no. 1, pp. 11–18, 2024.

[20] D. A. Tyas, T. Ratnaningsih, A. Harjoko, and S. Hartati, "The classification of abnormal red blood cell on the minor thalassemia case using artificial neural network and convolutional neural network," in Proc. International Conference on Video and Image Processing, 2017, pp. 228–233.

[21] U. Jarujareet, K. Wiratchawa, T. Petiwathayakorn, P. Koonyosying, S. Hantrakool, S. Srichairatanakool, and T. Intharah, "Classification of beta-thalassemia major and HbE/beta-thalassemia via deep learning of image structure function image," Biomedical Signal Processing and Control, vol. 102, art. 107265, 2025.

[22] M. W. Nadeem, H. G. Goh, M. A. Khan, M. Hussain, M. F. Mushtaq, and V. A. Ponnusamy, "Fusion-based machine learning architecture for heart disease prediction," Computers, Materials & Continua, vol. 67, no. 2, pp. 2481–2496, 2021.

[23] K. Shankar, E. Perumal, P. Tiwari, M. Shorfuzzaman, and D. Gupta, "Deep learning and evolutionary intelligence with fusion-based feature extraction for detection of COVID-19 from chest X-ray images," Multimedia Systems, vol. 28, no. 4, pp. 1175–1187, 2022.

[24] T. Vijayakumar, R. Vinothkanna, and M. Duraipandian, "Fusion based feature extraction analysis of ECG signal interpretation – A systematic approach," Journal of Artificial Intelligence, vol. 3, no. 1, pp. 1–16, 2021.

[25] M. U. Khan, S. Aziz, K. Iqtidar, G. F. Zaher, S. Alghamdi, and M. Gull, "A two-stage classification model integrating feature fusion for coronary artery disease detection and classification," Multimedia Tools and Applications, vol. 81, pp. 13661–13690, 2022.

[26] S. Purwar, R. Tripathi, R. Ranjan, and R. Saxena, "Classification of thalassemia patients using a fusion of deep image and clinical features," in Proc. 11th International Conference on Cloud Computing, Data Science & Engineering, 2021, pp. 410–415.

[27] E. R. Susanto, A. Syarif, K. Muludi, R. R. W. Perdani, and A. Wantoro, "Implementation of fuzzy-based model for prediction of thalassemia diseases," Journal of Physics: Conference Series, vol. 1751, art. 012034, 2021.

[28] F. Akhtar, A. Shakeel, J. Li, Y. Pei, and Y. Dang, "Risk factors selection for predicting thalassemia patients using linear discriminant analysis," in Proc. Prognostics and Health Management Conference, 2020, pp. 1–7.