The Scientific Temper
Vol. 17, No. 11, November 2026: pp. 7169-7181
ISSN: 0976-8653, E-ISSN: 2231-6396
DOI: 10.58414/SCIENTIFICTEMPER.2026.17.11.2587
A Web Of Science Journal

RESEARCH ARTICLE

From DNA Sequences to Explainable Images: An FCGR-Based Deep Learning Framework for Leukemia Classification

R. Srividhya[1], R. Radha[2]

[1]Department of Computer Science, SDNB Vaishnav College for Women, Chromepet, Chennai-. 600 044, Tamil Nadu, India, vidhyasdnbvc1098@gmail.com*

[2]Department of Computer Science, SDNB Vaishnav College for Women, Chromepet, Chennai-. 600 044, Tamil Nadu, India, radhasundar1993@gmail.com

Corresponding author: vidhyasdnbvc1098@gmail.com

ABSTRACT

Leukemia is a deadly blood cancer and an accurate classification and timely classification is a good way to characterize the disease and support clinical decision-making. This paper proposes a methodology for classification of leukemia based on DNA sequence data that involves frequency-Chaos Game Representation (FCGR), image preprocessing, data augmentation, transfer learning and interpretability through Grad-CAM. Healthy and Leukemia samples, represented by DNA sequences, are transformed to 2D images called frequency-Chaos Game Representation (FCGR) images, which preserve the frequency as well as the spatial representation of nucleotide combinations. Preprocessing and Augmentation followed by classification by fine tuning pretrained ResNet-18 network in which 1,000 class ImageNet output layer was modified to two class output (Healthy, Leukemia) and trained on the target data. Approximately 1,066 samples of the data were used for the experiments, which resulted in a test accuracy of 94.96%. Then, Grad-CAM is used to visualize the spatial regions of FCGR images having the largest influence on each prediction. It is demonstrated that the integration of an interpretable genomic image representation and deep transfer learning can be effective for the classification of leukemia.

Keywords: *Leukemia, DNA Sequences, Genomic Classification, Hematological Malignancy, Sequence Representation, Disease Classification, Clinical Decision Support.*

INTRODUCTION

Leukemia occurs when there is an abnormal overproduction of blood-forming cells and their normal production is impaired, resulting in a cancer of the blood-forming tissues such as bone marrow. Due to its varied types of clinical behaviors, molecular signatures and responsiveness to therapy, early and precise diagnosis is crucial for the treatment planning. Standard diagnosis involves blood counts, microscopy, immunophenotyping, cytogenetic and molecular analysis. These have high information content but need trained personnel and special facilities. With more and more genomic databases, computational techniques to find patterns associated to diseases directly from sequences are becoming possible (Rao et al., 2025).

The sequence of the four nucleotides A (adenine), C (cytosine), G (guanine) and T (thymine) make up the DNA and encode the biological information. The analysis of genomic sequences is usually achieved by hand-coded properties such as nucleotide frequencies, k-mer counts and statistical features. These can be effective but lack the ability to detect intricate relationships and the relative placement of nucleotide

Received: 5th November, 2026 Accepted: 7th November, 2026 Published: 28th November, 2026

combinations, and the feature-engineering process can degrade the quality of the classifier. Representation learning methods that learn the features from the encoded data can also decrease the need for manual feature selection and can provide more expressive inputs to the classifier (A. V. N. et al., 2025).

One of such methods is Frequency Chaos Game Representation (FCGR) which involves creating a two-dimensional numerical image from a DNA sequence. The FCGR is based on the Chaos Game Representation and provides a compressed representation of the nucleotide combination frequencies for a specified length of nucleotide and has intensities in specific locations in space. These images can be used for analysis using image processing and computer vision methods and similar sequence-to-image conversions have been found to reveal sequence information that deep learning classifiers can leverage (Rahman et al., 2026). The FCGR therefore offers a structure which makes the differences between Healthy and Leukemia genomic samples more evident.

Although deep CNNs learn hierarchical representations of images very well, it requires very large amounts of labeled data and a lot of computational resources to train a deep CNN from scratch. Transfer Learning is a solution to this problem that relies on transferring learning knowledge from a larger source dataset for a given task (Nisar et al., 2025). The Residual CNN called ResNet-18 has short-cut connections that help in the training of deep nets and also to propagate features. Patterns learned for diseases can be fine-tuned using ResNet-18 on FCGR images, where the original classification layer is replaced with a classification layer corresponding to the target classes and low and mid-level features trained in previous fine-tuning are kept.

The disadvantage of deep nets is that while they are good at predicting, they are difficult to explain, as they are learned at multiple levels. In biomedical applications, interpretability is particularly critical, to help the user understand why a sample is labeled diseased/healthy. Grad-CAM (gradient-weighted class activation mapping) makes it easier to pinpoint which areas in an input are most important for a prediction. Applied to FCGR images, it serves to visualize the spatial region and the combination of nucleotides used for classification and gives insights about the discriminative genomic patterns (S. & J., 2025).

This study therefore sets out to develop a combined system based on FCGR and transfer learning for leukemia classification from DNA sequences. Sequences are mapped to FCGR images, preprocessed, augmented and then fed into a pretrained ResNet-18 that was originally trained to classify images on ImageNet and whose classification layer has been removed, in order to perform fine-tuning for binary classification of Healthy vs. Leukemia. This framework is implemented and results are presented on an independent test set, and Grad-CAM is included to identify influential regions in the FCGR representations, which are evaluated with about 1,066 samples. The final pipeline is an integration of the genomic sequence transformation, deep visual representation learning, binary disease classification and visual explainability that could be used as a stepping-stone towards finding discriminative genomic signatures for leukemia and toward more interpretable sequence-based disease classification.

The key contributions of this work are as follows. First, the problem of classifying sequences is transformed to an image classification problem by transforming DNA sequences of Healthy and Leukemia samples to FCGR images, without the need for crafted genomic descriptors. Second, adaptation of a pre-trained ResNet-18 is performed to the binary genomic task by swapping the output layers and fine-tuning using preprocessing and label-preserving augmentation. Thirdly, the framework is tested on an independent set for evaluation based on accuracy, confusion matrix and class-based metrics. Fourth, Grad-CAM is used to visualize which regions of the FCGR are used for each prediction, enabling an exploration of the behavior of the classifier rather than considering it as a black box.

LITERATURE SURVEY

Leukemia is heterogeneous hematological malignancies marked by abnormal proliferation and accumulation of immature (blast) or dysfunctional cells, which makes early diagnosis, especially of subtypes, difficult. Diagnosis typically combines peripheral blood or bone marrow assessment, immunophenotyping, cytogenetic and molecular studies, but these procedures are affected by observer variability, image quality, cellular similarity and subtype complexity. Advances in machine learning (ML) and deep learning (DL) have enabled automated analysis of blood smear, bone marrow, clinical and molecular data. CNNs learn discriminative morphological features directly from images, and transfer learning makes them practical on limited medical datasets. More recent work has added vision transformers (ViT), ensemble learning, explainable artificial intelligence (XAI), feature selection and data balancing. Nevertheless, differences in datasets, class distributions, acquisition conditions and evaluation protocols still limit generalizability.

Several studies have applied deep and hybrid learning to leukemia diagnosis. A review of DL-based analysis of leukemia types highlighted the value of automated recognition of malignant blood-cell patterns (Akila & Parvathi, 2025). A hybrid strategy combining several learning methods classified multiple leukemia subtypes from bone marrow cell images,

discriminating between visually similar classes (Hossain et al., 2024). MATLAB-based ML detection from blood samples showed that conventional algorithms depend strongly on feature representation, sampling and preprocessing (Sanaullah & Al Noman, 2024). A fine-tuned VGG19 network demonstrated the benefit of transfer learning (Kaur, 2024), and colour-space-aware hybrid CNNs showed that chromatic features complement CNN-learned features for acute lymphoblastic leukemia (ALL) classification (Arya et al., 2025).

Other work has compared conventional ML with deep architectures. SVM and CNN classifiers were compared on peripheral blood smear images to contrast handcrafted and learned features (Kumar et al., 2025). A modified ResNet-101 extracted hierarchical features from leukemic white blood cell images (Choubey et al., 2024), and an adapted Inception V3 used multi-scale feature extraction for blood smear analysis (Pujitha et al., 2024). Nature-inspired feature reduction combined with stacked ensemble learning showed that removing irrelevant or redundant features improves ensemble decisions (Mohapatra et al., 2024). As medical AI grows in importance, explainability has gained attention; an XAI framework using clinical data supported early risk prediction and type classification of leukemia (Shahbaz et al., 2025).

Data quality and representation are also active concerns. Data-balancing techniques were investigated to reduce class imbalance in acute myeloid leukemia (AML) subtype classification (Aby et al., 2025). A ViT-based approach used self-attention to capture global relationships in medical images (Joshi & Mehta, 2025), while a fast CNN model addressed both computational efficiency and diagnostic performance for ALL (Devi et al., 2024). DL has also been used for joint detection and classification of ALL within integrated diagnostic systems (A. A. K. & Raimond, 2025), and LIME-based explanations have shown which image regions drive leukemia predictions (S. K. et al., 2025).

Recent studies have compared deep, transformer, multimodal and ensemble approaches. A comprehensive comparison of DL models emphasized the importance of architecture choice and thorough evaluation (Hasan et al., 2024). ALL-ViT demonstrated attention-based feature learning for ALL detection (Pandey et al., 2025), and gene-expression-based ML showed the value of molecular information as a complementary diagnostic source (Mahmoud et al., 2024). Systematic comparison of DL architectures for B-cell ALL classification (Indraini et al., 2025) and the ensemble-driven DRIXen architecture (Tasnim et al., 2025) further illustrate the benefit of combining classifiers for robust prediction. Despite these advances, gaps remain: class imbalance, high inter-class similarity, dependence on acquisition conditions, limited interpretability and little integration of complementary representations. Moreover, most studies analyze cell images rather than genomic sequences. These limitations motivate frameworks that combine discriminative representations, effective feature learning and interpretable predictions.

METHODOLOGY

The proposed method is an end to end pipeline which is aimed at classification of Healthy and Leukemia DNA sequences with image based representation of the genome and transfer learning (Figure 1). Initially, the DNA sequences are gathered and processed, cleaned and standardized. Then, each sequence is transformed into a FCGR image that spatially represents the frequency of all possible nucleotide combinations. The images are pre-processed and augmented, a ImageNet ResNet-18 is adapted for binary classification and trained on the images. Lastly, the trained network is applied to an independent test set and Grad-CAM is applied to understand which areas of the FCGR images contribute to the classification.

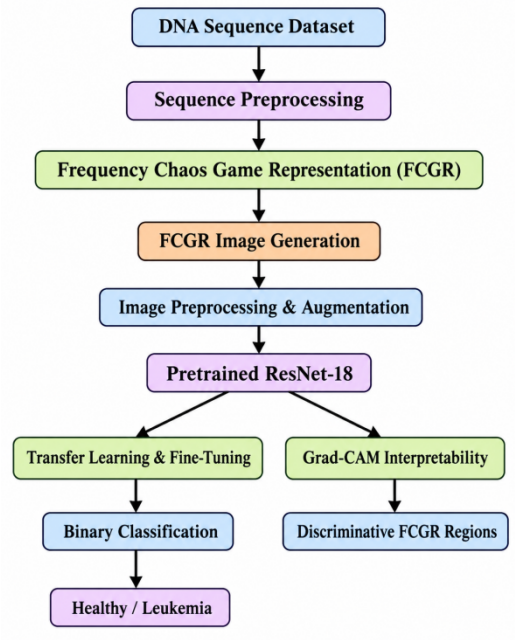

Fig. 1: System Architecture

DNA Sequence Acquisition and Preparation

The sequences of DNA belonging to Healthy and Leukemia classes are acquired, and segregated class-wise in a folder. These are denoted by the 4 base letters A, C, G and T. The sequence order and combination has the biological information so sequences are screened and retained in the original order for use for the FCGR transformation. The sequences are preprocessed to remove the characters which are not invalid and non supported characters to avoid introduction of spurious patterns, treat ambiguous characters consistently and determine if the sequence is healthy or leukemia. This avoids the need for human extraction of man-made genomic features and preserves the genomic sequence, which can be expressed in a structured manner to facilitate learning of features by downstream models.

Pseudo Code:

Input: DNA sequences of healthy and leukemia samples

Output: Preprocessed DNA sequences

BEGIN

Load DNA sequences from the dataset

FOR each DNA sequence DO

Convert sequence characters to uppercase

Remove invalid characters and unwanted symbols

Check for ambiguous nucleotides

Apply the selected strategy to handle ambiguous bases

Associate the sequence with its class label (Healthy or Leukemia)

END FOR

Verify sequence quality and class labels

Store the preprocessed sequences

RETURN preprocessed sequences

END

Frequency Chaos Game Representation Generation

Prepared sequences are converted into FCGR images and show information on a sequence in two dimensions – space. The rationale behind FCGR is that each of the four nucleotides can be assigned to one of the four corners in a 2D space, and each k-mer can be mapped to a specific position. The sequence is not plotted directly but rather FCGR plots the number of times each k-mer appears: high intensities correspond to frequently occurring k-mers and low intensities correspond to infrequent k-mers. If k is the k-mer size, then the output is a 2^k x 2^k matrix containing both compositional and spatial features of the sequence. This is done for all the healthy and leukaemia samples, and is developed as an image-based genomic data set that allows ResNet-18 to automatically learn nucleotide distribution patterns that are relevant to the disease.

More formally, the Chaos Game Representation takes the four nucleotides as the corners of a unit square and places every next nucleotide at the midpoint of the last nucleotide and the corner corresponding to this next nucleotide. All positions in the sequence are then associated with a position that indicates the location of the k-mer preceding it. Normalization guarantees that the images can be compared from sequences with varying length while the grid

arrangement assures that k-mers with similar suffix are located in the vicinity, thus giving the image spatial information that can be utilized by the convolutional filters.

Pseudo code:

Input: Preprocessed DNA sequences and k-mer size

Output: FCGR images

BEGIN

Set the FCGR resolution to $2^k \times 2^k$

FOR each preprocessed DNA sequence DO

Initialize an FCGR frequency matrix with zeros

FOR each valid k-mer in the DNA sequence DO

Determine the matrix position corresponding to the k-mer

Increment the frequency at that position

END FOR

Normalize the frequency matrix, if required

Convert the matrix into an FCGR image

Save the image with its corresponding class label

END FOR

RETURN FCGR images

END

FCGR Image Preprocessing and Augmentation

The FCGR images are then preprocessed for the pretrained ResNet-18. All images are normalized to the input size 224x224x3 of the network and transformed into numerical arrays to keep input values in the range of values that are proper for transfer learning but do not lose spatial information. Random reflection and random scaling are used for augmentation of training images only to vary the training data, but no augmentation is performed on test images to prevent over fitting. No new class labels (Healthy and Leukemia) or new biological samples will be created due to augmentation. This step provides the linkage between the genomic representation and visual classification using transfer learning.

Pseudo code:

Input: FCGR images stored in Train and Test folders

Output: Resized training and testing images

BEGIN

Define the main dataset directory

Load training images from the Train folder

Load testing images from the Test folder

Assign labels using folder names

Verify that Healthy and Leukemia classes are available

Set the target image size to 224 × 224 × 3

Apply resizing to training and testing images

Apply random horizontal reflection and random scaling

to training images only

Preserve the original testing images without augmentation

RETURN augmented training images and resized testing images

END

Transfer Learning Using Pretrained ResNet-18

The classification backbone of the proposed method is pretrained ResNet-18. It was chosen due to its ability to transfer features efficiently through the network via the residual shortcut connections, and its relatively small size. Weights of the model learnt on ImageNet contain low and mid level visual features that can be reused. Fully connected, softmax and classification layers with 1000 original classes are replaced with new layers which have 2 output classes for Healthy and Leukemia for genomic classification. After the new layers, successive residual blocks are trained that extract hierarchical spatial information from FCGR images, with early layers expressing more basic spatial information and deeper layers are expressing spatial information that can differentiate between the target classes. For each epoch, the training samples are randomly selected, fine tuned using SGDM with max 15 epochs, learn rate of 0.0001 and mini batch size of 16. It predicts the probability that the sample is Healthy or Leukemia without handcrafted genomic descriptors

Pseudo code:

Input: Training images and pretrained ResNet-18 model

Output: Trained ResNet-18 model

BEGIN

Load the pretrained ResNet-18 network

Replace the original fully connected layer

with a new layer containing two output classes

Replace the original softmax layer

Replace the original classification layer

Configure training parameters:

Learning rate = 0.0001

Maximum epochs = 15

Mini-batch size = 16

Optimizer = SGDM

Set the testing dataset as validation data

Shuffle training samples at every epoch

Train the modified network using augmented training images

Save the trained network

RETURN trained network

END

Model Training and Binary Classification

During the training stage, the FCGR images and their labels pass through the residual convolutional blocks and the newly developed 2 class layers and the parameters of the model are tuned to learn more distinguishable representations between Healthy and Leukemia images of FCGR. The learned decision function is used to classify every unobserved input into one of the two classes. The model is tested with a test set not used during training and accuracy, sensitivity, specificity, precision and F1-score are from the confusion matrix to see if it generalizes. With around 1,066 samples, this test accuracy of 94.96% indicated that FCGR representations have high discrimination for transfer learning model.

Pseudo code:

Input: Trained network and testing images

Output: Predicted labels, confusion matrix, and evaluation metrics

BEGIN

Load the trained ResNet-18 model

Predict class labels for all testing images

Obtain prediction scores

Compare predicted labels with actual labels

Calculate:

True Positive (TP)

True Negative (TN)

False Positive (FP)

False Negative (FN)

Calculate classification metrics:

Accuracy = (TP + TN) / (TP + TN + FP + FN)

Sensitivity = TP / (TP + FN)

Specificity = TN / (TN + FP)

Precision = TP / (TP + FP)

F1-score = 2 × Precision × Sensitivity / (Precision + Sensitivity)

Generate the confusion matrix

Display the performance metrics

RETURN predictions, confusion matrix, and metrics

END

Grad-CAM-Based Interpretability and Evaluation

The trained ResNet-18 is finally interpreted using Grad-CAM. A FCGR image is fed into the model which produces a Healthy/Leukemia prediction. Then, the gradient of the predicted class score is passed back to the feature maps of a selected convolutional layer (res5b_relu) and used to weight the feature maps to produce a class-activation heatmap which shows where in the FCGR image the prediction is based. The normalized heatmap is overlaid on the FCGR image for visual explanation. This is not necessarily the nucleotide locations but areas on the FCGR image that represents areas of the transformed genome that have the most significance in the model. Grad-CAM was combined with the 94.96% test accuracy in order to not only show the effectiveness of the model but also the interpretability of the model which is particularly important in the biomedical field.

Pseudo code:

Input: Trained ResNet-18 model and a testing FCGR image

Output: Grad-CAM heatmap and overlay visualization

BEGIN

Select a testing FCGR image

Load the image and its actual class label

Predict the class using the trained network

Resize the image to the network input dimensions

Identify the predicted class index

Select the feature layer "res5b_relu"

Generate the Grad-CAM heatmap

Normalize the heatmap

Display the original FCGR image

Overlay the heatmap on the original image

Display the heatmap separately

Compare the highlighted regions with the predicted class

RETURN heatmap and visualization

END

RESULT AND DISCUSSION

Binary classification was applied to the healthy and Leukemia DNA sequences of ~1,066 samples for

training and testing. ResNet-18 fine tuned, with its original output of 1,000 classes from ImageNet replaced by two classes, and all sequences were mapped to FCGR images and preprocessed and augmented prior to the fine tuning. The accuracy rate of the test set was 94.96% with the fine tuned model, which indicates that there is enough discriminatory information in the FCGR representations to differentiate the two classes. Table 1 shows the experimental configuration.

Table 1. Numerical Configuration of the Experimental Dataset and Classification Task

Parameter	Numerical Value
Approximate dataset size	1066
Number of target classes	2
Original ImageNet classes	1,000
Target classes after modification	2
Test accuracy (%)	94.96

As mentioned in Table 1, the classification layer trained using ResNet-18 has been substituted and the network has been moved to genomic data. The test accuracy of 94.96% reflects that the network has learned patterns that allow discrimination between Healthy and Leukemia FCGR images. This enables the use of convolutional structures, designed for use with image data, to genomic data, where discriminative visual patterns are not prespecified, but instead learned automatically.

Spatially, FCGRs are representations of frequencies of the combination of nucleotides, and thus different DNA sequences will have different distributions of spatial intensity. ResNet-18 uses residual layers to capture hierarchical features from these patterns; transfer learning helps not to build a deep network from scratch and fine tuning involves reusing pretrained parameters to create an initial representation and fine tuning that representation to classify healthy and Leukemia.

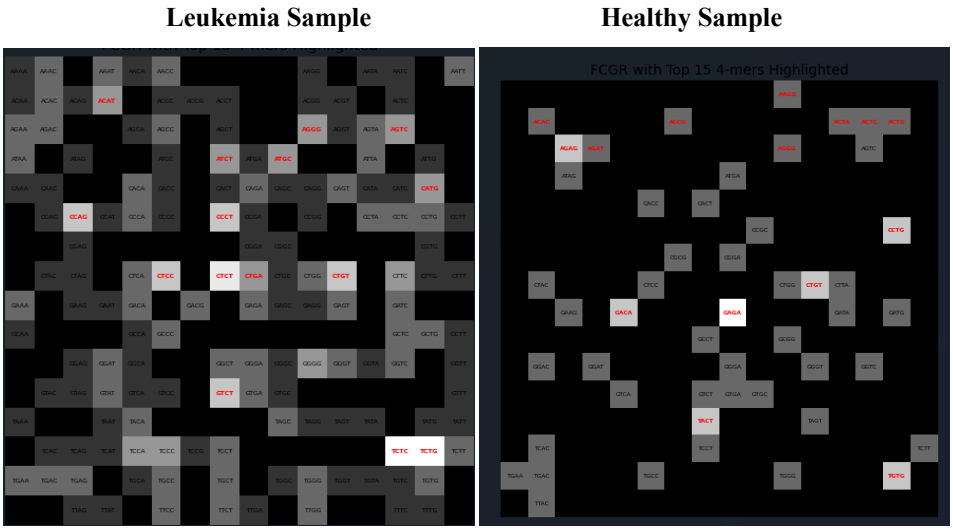

Figure 2. Classification Output of the Proposed FCGR-Based Leukemia Detection Framework

Figure 2 shows the output of the proposed framework for a Leukemia sample and a Healthy sample, illustrating how DNA sequences converted to FCGR images are classified by the fine-tuned ResNet-18 and visually confirming the binary classification setup.

Table 2. Performance of the Proposed Classification Framework

Performance Parameter	Value
Number of classes	2
Test accuracy (%)	94.96
Classification output dimension	2
Original output dimension	1,000
Accuracy error (%)	5.04

Table 2 presents the key performance features. Test accuracy was 94.96% and error rate 5.04%. The modified network has two output classes (instead of 1000 original output classes of ImageNet model), as it is no longer attempting to distinguish the 1000 categories of images but rather the two categories of genomic images.

The findings suggest the transfer learning of FCGRs. ResNet-18 is a compact residual network with shortcut connections that allow information to flow easily and to be fine-tuned. During the fine-tuning process, the pre-trained filters were adjusted to detect spatial patterns in the distributions of the frequencies of the nucleotides in the FCGR images rather than patterns of objects in photographs, and the features were found to be helpful in classifying the two genomic types.

Preprocessing and augmentation are also important. As compared to the use of photos, the FCGR images of genomic sequences must be standardized to have a uniform input. Augmentation can help improve generalisation because it provides more variation in the training patterns and reduces the dependency on individual training patterns and maintains the labels of Healthy and Leukemia unchanged. But precautions should be taken before transferring augmented genomic representations to a computational to clinical setting.

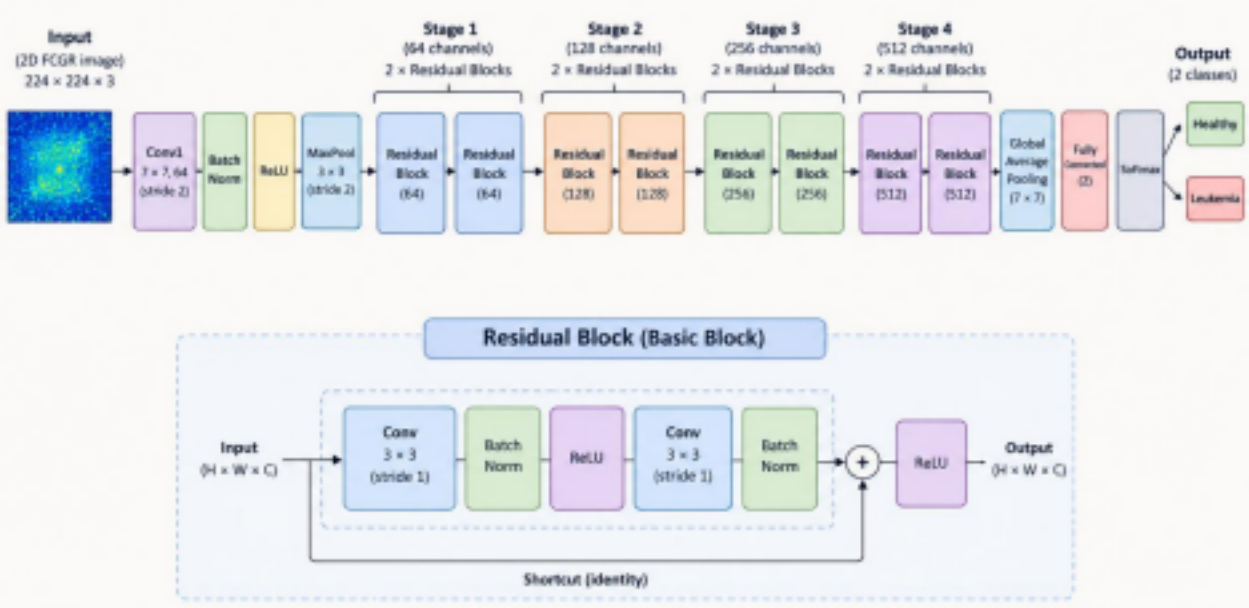

Figure 3. ResNet-18 architecture

ResNet-18 is a 18-layered convolutional neural network (CNN) that contains residual layers as depicted in Figure 3. The initial convolution layer with batch normalisation, ReLU activation and max-pooling, four residual layers, global average pooling and classification layer. Residual layers are built up with basic residual blocks which connect with shortcut to enable gradient flow. The network is pre-trained on ImageNet and its 1,000 class output layer is modified to a two class (Healthy/Leukemia) output layer and then fine-tuned on resized and normalised FCGR images.

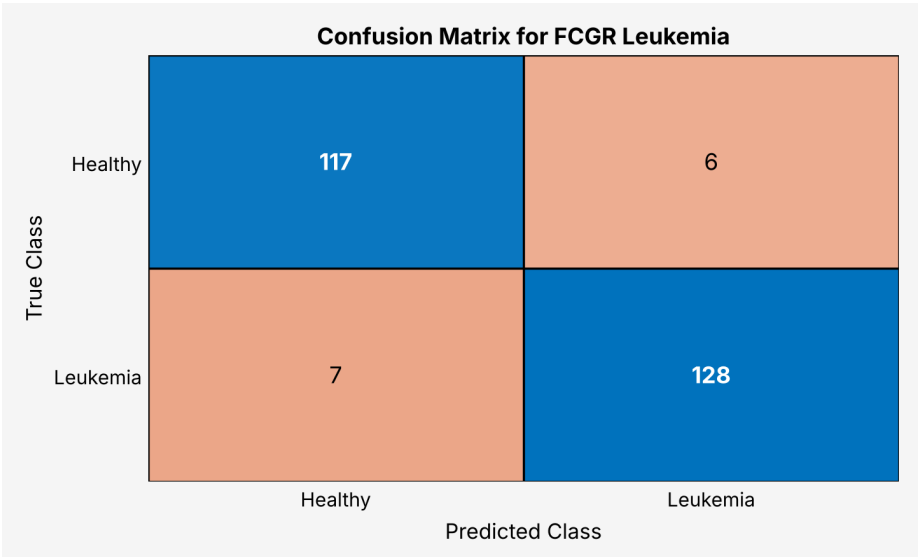

Figure 4. Confusion Matrix

Figure 4 presents the confusion matrix for the FCGR–ResNet-18 framework. Among the 123 healthy test samples, 117 were predicted as healthy, and 6 were incorrectly predicted as leukemia. Of the 135 leukemia test samples, 128 were correctly classified and 7 were misclassified as healthy. Total numbers of correctly classified test samples were 245 out of 258 test samples, for a test accuracy of 94.96%, which is stated here. Large diagonal and few off-diagonal counts (nearly balanced) indicate that the framework is very good at separating Healthy and Leukemia samples, and is not biased toward any class.

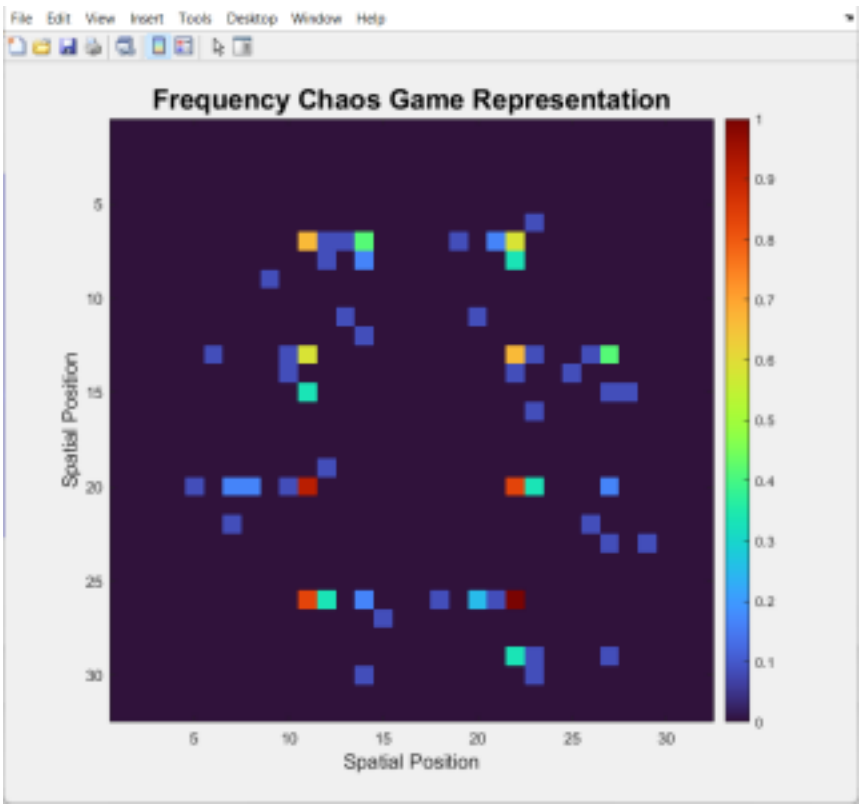

Figure 5. Frequency Chaos Game Representation of DNA Sequence

An example FCGR image, which is passed as an input to ResNet-18 classifier, is shown in Figure 5 to illustrate how sequential nucleotide data can be converted into a structured spatial representation amenable for convolutional layers.

The quantitative evaluation is followed by an explainability analysis. The trained ResNet-18 was analysed using Grad-CAM to find the spatial locations in the FCGR images on which the decision relies. This is important because while a classification is great, without understanding of where the decision was made in the input, a good classification is not enough. Grad-CAM leverages class specific gradients to compute the importance of regions in the convolutional feature map and overlay it on the FCGR image to highlight important regions which impact on the predicted class. These regions are a part of the transformed genomic representation and may not be correlated with specific nucleotides.

It should therefore be viewed as an explanatory process of interpretation. Highlighted region in FCGRs indicates that there is a difference in spatial distribution of certain nucleotide combination frequencies and is not sufficient evidence for a leukemia associated biomarker; biological relevance of such a difference should be confirmed. But visualization helps to determine if the model is sensitive to important regions of the genomic representation rather than random locations, making the classifier more transparent, and facilitating further exploration of unique sequence patterns.

Unlike the studies reviewed in Section II which typically involve classifying images of blood smears or bone marrow, the proposed framework processes images of DNA sequences. This is a framework that complements the morphology-based approaches in that it uses genomic information and, just as morphology-based approaches use transfer learning, can be reused successfully for a specialized biomedical task. As compared to the gene-expression or clinical-data models, the FCGR representation doesn't have any manually selected features, and is of fixed image size, irrespective of the length of the sequence.

Table 3. Parameters of the Proposed FCGR–ResNet-18 Framework

Parameter	Numerical Value
Dataset scale	1066
Input representation dimension	2D
Original ResNet-18 output classes	1,000
Modified output classes	2
Target categories	2
Reported test accuracy (%)	94.96
Accuracy error (%)	5.04

Table 3 summarizes the main parameters of the framework: DNA sequences are converted into 2D FCGR representations, the 1,000-class output of the pretrained network is replaced by a two-class layer, and the final system reaches a test accuracy of 94.96%.

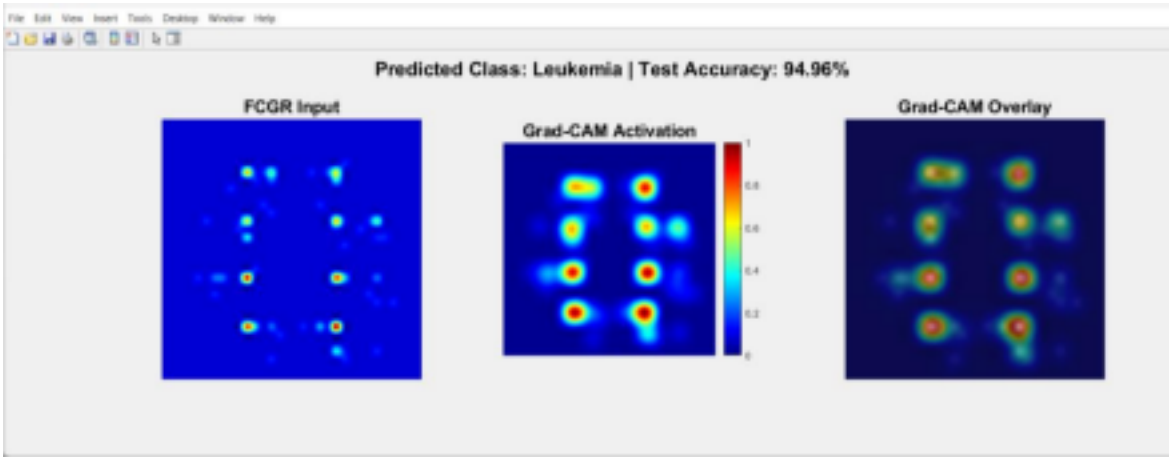

Figure 6. Grad-CAM-Based Visualization of Influential FCGR Regions

The Grad-CAM visualization of a prediction made by the trained ResNet-18 is shown in Figure 6, alongside the original FCGR image and the activation map of influential spatial regions. This map gives a visual understanding of which parts of the image are most relevant to the prediction class, thereby enriching the numerical accuracy with an understandable representation of the model's decision process.

Additional performance metrics for the FCGR–ResNet-18 framework are presented in Table 4. These are directly calculated from the values of the confusion matrix in Figure 4 (TP=128, TN=117, FP=6, FN=7), where Leukemia is considered the positive class. The model achieved 94.96% accuracy, 95.52% precision (meaning that most samples predicted to be Leukemia were indeed Leukemia), 94.81% recall (meaning that most leukemia samples were detected), 95.12% specificity (meaning that most healthy samples were correctly predicted), and 95.17% F1-score (meaning that the model strikes a good balance between precision and recall). All these measures give more detailed evaluation of the framework than accuracy alone and further support the effectiveness of the proposed framework.

Table 4. Corresponding Performance Measures

Performance Measure	Proposed FCGR–ResNet-18 (%)
Accuracy	94.96
Precision	95.52
Recall / Sensitivity	94.81
Specificity	95.12
F1-Score	95.17

Overall, the combination of FCGR transformation, preprocessing, augmentation, pre-trained ResNet-18 (fine-tuned) and Grad-CAM constitutes a powerful computational framework for leukaemia detection based on DNA sequences with test accuracy of 94.96%. The innovation lies in the fact that genomic sequential data can be visualized as a 2D frequency matrix that can be used as input for visual deep learning architecture that has been trained on a different task. Grad-CAM adds an interpretability layer for the exploration of influential FCGR regions. The results are computational and do not imply clinical validation. The spatial patterns seen need to be tested for generalisability on independent datasets and external cohorts, and also statistically confirmed and biologically investigated to establish that they are real leukemia patterns and that they do not change under various sequencing conditions and leukemia genomes.

There are a number of restrictions, however. The experiments were conducted on a single dataset comprising approximately 1,066 samples and there is no information at this time about the performance of this framework on other data sources and sequencing platforms from other populations. False positives and false negatives were also found, but while there are 6 false positives in the confusion matrix, 7 false negatives are most serious errors in a screening setting and should be also confirmed. Augmentation operations used for natural images may not be biologically relevant when applied to genomic data. Lastly, Grad-CAM explanations are qualitative and indicate where representation has an impact, but not where genes/mutations.

CONCLUSION

The study proposed a computational pipeline for leukemia classification using DNA sequences based on frequency CGR, image preprocessing, transfer learning and Grad-CAM interpretability. It allows to use a pretrained deep learning network for discriminative feature learning, for genomic sequence classification, where there is no need to solely depend on hand-crafted sequence features. This modified ResNet-18 is an effective classifier with a test accuracy of 94.96% and Grad-CAM helps in the interpretability of the model by showing important areas in the genomic representation. In future studies, the regions identified by explainability tools will be validated against biological data, the model will be trained on larger and more diverse datasets of genomic sequences, validation cohorts external to our training set will be explored, and different types of genomic data (e.g., multimodal data) and deep learning models will be explored to improve model robustness, generalizability, interpretability and potential clinical usefulness.

Acknowledgment

This work was supported by the scheme of "Minor Research Project" grants from Shrimathi Devkunvar Nanalal Bhatt Vaishnav College for Women, Chromepet, Chennai, to Srividhya R (Grant No. 26MRPE012).

Conflict of interest: None

Source of support: None

REFERENCES

A. A. K., & Raimond, K. (2025). Enhancing acute lymphoblastic leukemia classification and detection using deep learning. In *2025 Second International Conference on Cognitive Robotics and Intelligent Systems (ICC-ROBINS)* (pp. 363–368). IEEE. https://doi.org/10.1109/ICC-ROBINS64345.2025.11086223

A. V. N., J. N., T. M., & Subramaniam, K. (2025). Certain investigation on deep learning algorithms for leukemia classification. In *2025 7th International Conference on Innovative Data Communication Technologies and Application (ICIDCA)* (pp. 1392–1397). IEEE. https://doi.org/10.1109/ICIDCA66325.2025.11280469

Aby, A. E., S., S., K. K., A., & Rajan, T. (2025). Optimizing acute myeloid leukemia subtype classification through data balancing techniques. In *2025 2nd International Conference on Trends in Engineering Systems and Technologies (ICTEST)* (pp. 1–4). IEEE. https://doi.org/10.1109/ICTEST64710.2025.11042354

Akila, S., & Parvathi, R. (2025). A comprehensive review of leukemia and its types using deep learning techniques. In *2025 7th International Conference on Artificial Intelligence and Speech Technology (AIST)* (pp. 1–6). IEEE. https://doi.org/10.1109/AIST68591.2025.11441613

Arya, M., Arya, S., & Arya, S. (2025). Color space-aware hybrid CNN architectures for acute lymphoblastic leukemia classification. In *2025 IEEE 16th Annual Ubiquitous Computing, Electronics & Mobile Communication Conference (UEMCON)* (pp. 240–246). IEEE. https://doi.org/10.1109/UEMCON67449.2025.11267768

Choubey, A., Bharath, D., Reddy, G. M., Shekar, K. A., Choubey, S. B., & Gumudavelli, S. (2024). Modified ResNet-101 model for enhancing classification of leukemia white blood cell images. In *2024 IEEE International Conference on Intelligent Signal Processing and Effective Communication Technologies (INSPECT)* (pp. 1–5). IEEE. https://doi.org/10.1109/INSPECT63485.2024.10896237

Devi, J. R., Kadiyala, P. S., Lavu, S., Kasturi, N., & Kosuri, L. (2024). Enhancing acute lymphoblastic leukemia classification with a rapid and effective CNN model. In *2024 Third International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICDCECE60827.2024.10549348

Hasan, J., et al. (2024). Transforming leukemia classification: A comprehensive study on deep learning models for enhanced diagnostic accuracy. In *2024 IEEE International Conference on Power, Electrical, Electronics and Industrial Applications (PEEIACON)* (pp. 1–6). IEEE. https://doi.org/10.1109/PEEIACON63629.2024.10800693

Hossain, M. S., et al. (2024). Hybrid AI classification techniques for multi-class leukemia sub-types classification with bone marrow cells. In *2024 6th International Conference on Sustainable Technologies for Industry 5.0 (STI)* (pp. 1–6). IEEE. https://doi.org/10.1109/STI64222.2024.10951147

Indraini, P. N., Yasmin, A., & Suandi, D. (2025). Comparative analysis of deep learning models for acute lymphoblastic leukemia cancer B-cell classification based on image data. In *2025 4th International Conference on Digital Transformation and Applications (ICDXA)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICDXA69105.2025.11329929

Joshi, V. C., & Mehta, M. A. (2025). Vision Transformer (ViT)-based approach for accurate detection of acute lymphoblastic leukemia. In *2025 IEEE 6th India Council International Subsections Conference (INDISCON)* (pp. 1–6). IEEE. https://doi.org/10.1109/INDISCON66021.2025.11254778

Kaur, A. (2024). Revolutionizing leukemia diagnosis: Fine-tuned VGG19 CNN for enhanced classification accuracy. In *2024 Global Conference on Communications and Information Technologies (GCCIT)* (pp. 1–5). IEEE. https://doi.org/10.1109/GCCIT63234.2024.10862090

Kumar, P. A., Kumar, P. N., Kumaran, S., & Rufus, N. H. A. (2025). Machine learning-based classification of leukemia from peripheral blood smear images using SVM and CNN. In *2025 International Conference on Recent Innovation in Science Engineering and Technology (ICRISET)* (pp. 1–9). IEEE. https://doi.org/10.1109/ICRISET64803.2025.11252235

Mahmoud, A., Talpur, K. R., Saini, S., Talpur, B. A., Shah, A., & Zaki, J. (2024). Machine learning-based leukemia classification using gene expression for accurate diagnosis. In *2024 IEEE 9th International Conference on Engineering Technologies and Applied Sciences (ICETAS)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICETAS62372.2024.11120204

Mohapatra, S. K., S. K. P., & Mohanty, M. N. (2024). Optimizing leukemia classification through nature-inspired feature reduction and stacked ensemble learning algorithms. In *2024 International Conference on Electronics, Computing, Communication and Control Technology (ICECCC)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICECCC61767.2024.10593969

Nisar, H., Lee, K.-F., & Ali, S. A. (2025). Optimizing leukemia classification with fine-tuned SENet variants: A comparative study on binary and multi-class tasks. In *2025 4th International Conference on Smart Cities, Automation & Intelligent Computing Systems (ICON-SONICS)* (pp. 51–56). IEEE. https://doi.org/10.1109/ICON-SONICS67145.2025.11309239

Pandey, P., RG, R. T., Pati, R. P., & Singh, S. (2025). ALL-ViT: A novel approach for detection of acute lymphoblastic leukemia. In *2025 12th International Conference on Emerging Trends in Engineering & Technology—Signal and Information Processing (ICETET-SIP)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICETETSIP64213.2025.11156790

Pujitha, M. V., Srikantam, S., Sahithi, C., & Timmasarti, A. (2024). Automated leukemia diagnosis from microscopic blood smear images using adapted Inception V3. In *2024 International Conference on Big Data Analytics in Bioinformatics (DABCon)* (pp. 1–6). IEEE. https://doi.org/10.1109/DABCon63472.2024.10919430

Rahman, M. M., Mollick, M. A. A., & Shova, A. (2026). Advanced transfer learning framework for the automated detection and classification of acute lymphoblastic leukemia. In *2026 IEEE 2nd International Conference on Quantum Photonics, Artificial Intelligence & Networking (QPAIN)* (pp. 1–6). IEEE. https://doi.org/10.1109/QPAIN69676.2026.11546231

Rao, V. S., Sateesh, N., Sneha, Y., Moganti, V. S. L. P. N., Prasad, K. S., & Shruthi, K. (2025). DeepL-DDCL: Deep learning-driven detection and classification of leukemia using microscopic blood smear images. In *2025 IEEE 2nd International Conference on Information Technology, Electronics and Intelligent Communication Systems (ICITEICS)* (pp. 1–5). IEEE. https://doi.org/10.1109/ICITEICS64870.2025.11341392

S., R., & J., V. (2025). Multi-modal classification of pediatric leukemia using PCA-enhanced gene expression and clinical data with XGBoost. In *2025 Innovations in Power and Advanced Computing Technologies (i-PACT)* (pp. 1–6). IEEE. https://doi.org/10.1109/i-PACT65952.2025.11307808

S. K., B. G. S., N. A. K., & D. T. (2025). Explainable AI for hematological diagnostics: Integrating image processing and LIME for leukemia

detection. In *2025 7th International Conference on Signal Processing, Computing and Control (ISPCC)* (pp. 846–851). IEEE. https://doi.org/10.1109/ISPCC66872.2025.11039577

Sanaullah, M., & Al Noman, A. (2024). Automated leukemia detection from blood samples using machine learning in MATLAB. In *2024 IEEE International Conference on Computing, Applications and Systems (COMPAS)* (pp. 1–6). IEEE. https://doi.org/10.1109/COMPAS60761.2024.10795893

Shahbaz, M., Basharat, A., Yasmin, R., Ahmad, N., Abbasi, R., & Hussain, S. (2025). Leveraging explainable AI for early risk prediction and type classification for leukemia: Insights using clinical data from Pakistan. *IEEE Journal of Biomedical and Health Informatics, 29*(7), 4672–4679. https://doi.org/10.1109/JBHI.2025.3531294

Tasnim, H., Ahmed, M., Shafin, M. S., & Rahman, T. (2025). DRIXen architecture: An ensemble-driven classification of acute lymphoblastic leukemia. In *2025 International Conference on Quantum Photonics, Artificial Intelligence, and Networking (QPAIN)* (pp. 1–6). IEEE. https://doi.org/10.1109/QPAIN66474.2025.11171747