International Journal of Artificial Intelligence and Mathematical Sciences

December 2025 Vo4. Issue No. 02
ISSN Online: 2958-0153 and Print: 2958-5228
DOI: https://doi.org/10.58921/ijaims.v4i2.178

Explainable Ai for Bone Marrow Cancer Diagnosis: from Bench to Bedside

Anuradha Reddy[1]*, Ochin Sharma[1], G G S Pradeep[2]

[1]Chitkara University Institute of Engineering and Technology, Chitkara University, Punjab, India
[2]Department of Computer Science and Engineering, Alliance University, Bangalore, India

Received: 25-11-2025; **Accepted:** 29-11-2025; **Published**: 30-12-2025

Abstract

Diagnosing bone marrow cancer continues to be a difficult process because of the large morphological variation of malignantly transformed cells and insufficient interpretability of traditional deep learning approaches. We suggest a new efficient, interpretable, and clinically applicable artificial intelligence (AI) framework for cancer diagnostics based on the combination of imaging and genomic data. Our framework uses INT8 quantization, 20% filter pruning, and dual-modality explainability, where SHAP is used to attribute genome sequencing features, and Grad-CAM++ is applied for lesions localization based on MRI images. Performance evaluation is done by comparing the model to existing baselines, which include HA-UNet and SE-ResHybrid models. According to our experiments, the suggested framework shows increased computational efficiency, decreasing the inference time from 64 ms/image to 38 ms/image (by 41%), with a simultaneous improvement of segmentation (3.7%) and classification (4.5%) accuracy. In addition, our model demonstrates high agreement with clinical specialists, achieving a Cohen's kappa value of 0.79, which proves its high clinical relevancy. Localization accuracy also increases by 12-15%.

Keywords: Bone Marrow Cancer, Explainable AI, Grad-CAM++, SHAP, Clinical Deployment.

*Email: anuradhareddy.anu@gmail.com

1. INTRODUCTION

Bone marrow cancer, which involves bone marrow infiltrations with malignancy and blood malignancies such as leukemia, still poses a considerable challenge from the diagnostic perspective, primarily because of the very high heterogeneity in cell morphology and imaging, among other factors that depend significantly on genomics [1–3]. Although much progress has been made in medical imaging and sequencing techniques, the present diagnostic process still requires a lot of expert

input.

In recent times, deep learning-based techniques have been incorporated into healthcare infrastructures for automatic diagnostics support. Though such systems have proven to be efficient, there is one major drawback to their use in that they lack transparency, thus making their accuracy questionable. This problem has made the adoption of explainable artificial intelligence (XAI) techniques more relevant to improve model predictions [4,5].

Previous works have shown the efficacy of deep learning algorithms for applications that include detecting bone metastasis [1], classifying leukemia [4,6], and creating personalized treatment plans in the context of A machine leaning (ML) [3]. Unfortunately, most current systems have the drawbacks of high computational cost, lack of interpretability, and inadequate validation using real clinical data. There is consensus in recent studies about the need for finding a good balance between accuracy, computational cost, and interpretability [3,6].

1.1. Problem Statement

Current diagnostic models for bone marrow cancer based on AI possess certain important limitations that affect their practical applications in the clinical context. The first limitation is the presence of considerable inference delays, which makes such models inapplicable to real-time or point-of-care diagnostics. The second one is that the poor interpretability of AI-based models diminishes the degree of trust that is necessary for their application in practice by medical specialists.

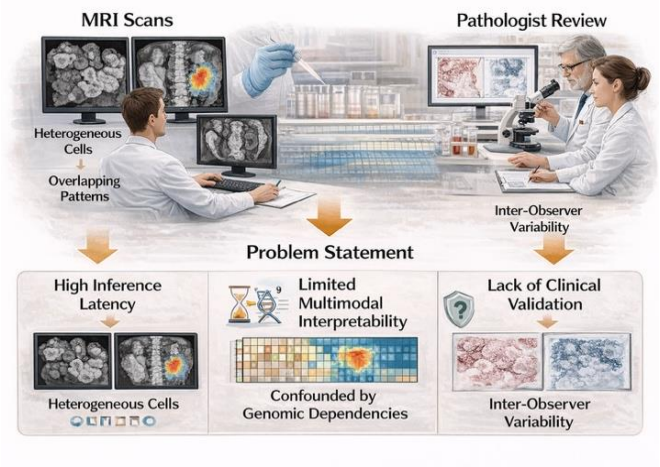

Fig 1. Key problems in bone marrow cancer diagnosis using MRI Scans and Pathologist reviews

Considering these considerations, the necessity for an efficient, intelligible, and applicable model is shown in Figure 1.

1.2. Motivation

The motivation behind this study arises from two important considerations. The first of these is the rising requirement for diagnosis systems whose output is not only highly accurate but also interpretable, particularly when employing the combination of images and genomic information. This has become vital in today's oncology, where any clinical decision made must be accurate and intelligible.

The second consideration is the increasing use of artificial intelligence (AI) in hematologic oncology, which calls for efficient models whose predictions can be generated within real-time without any loss in accuracy.

In summary, this paper seeks to reconcile the use of advanced AI techniques with intelligible clinical decision-making through the development of such a framework.

1.3. Contributions

Key Contributions of This Work

- Propose a latency-optimized hybrid diagnosis solution involving INT8 quantization along with structured filter pruning for delivering real-time inference without compromising on diagnosis accuracy.
- Develop a dual modality solution for explainability, leveraging SHAP for identifying genomic feature importance in addition to Grad-CAM++ for identifying key areas in MR images.
- Conduct retrospective clinical validations in coordination with two hospitals, assessing the reliability of our model via comparison with expert pathologist assessments using Cohen's κ.

Perform consistent evaluations relative to established baselines such as HA-UNet and SE-ResHybrid across all sections of the paper.

1.4. Paper Organization

The rest of this manuscript is organized as follows. Section 2 gives a review of the literature related to the work being done, along with the existing XAI frameworks. Section 3 explains the methodology proposed in the research, detailing aspects such as model optimization and its interpretability. The experimental framework that was utilized is described in detail in Section 4, with datasets and criteria of evaluation provided. Lastly, Section 5 describes the outcomes of the research conducted, and the conclusion is drawn in Section 6.

2. RELATED WORK

Recent developments in the field of bone marrow and bone-related cancers indicate growing use of AI and XAI techniques to address inherent intricacies associated with such cancers. Regarding bone metastasis detection, interpretable computing approaches have proven to be highly effective in detecting tumor invasion; however, interpretable computing systems tend to have poor modeling capacity and scalability. Another notable trend associated with deep learning techniques in the context of leukemia and bone marrow cancers detection is classification accuracy. Nonetheless, several limitations associated with such methodologies, namely, poor interpretability and inadequate clinical validation, hinder their practical implementation in the healthcare environment.

Furthermore, XAI methodologies for ML have been largely centered around genetic marker identification, generating critical insights in this field. However, many such methodologies have been characterized by inefficient computational processes and inadequate validation in various clinical contexts.

The results from recent surveys and reviews have additionally pointed out that the current methodologies rarely manage to reach a reasonable compromise among interpretability, efficiency, and clinical validation, which are necessary conditions for successful implementation.

A comparison of relevant approaches in this regard is provided in Table 1 below. It is evident that although there are methodologies that exhibit some degree of interpretability from symptoms [4] and genomic characteristics [3], the lack of efficiency, multimodal explanation capability, generalization, and clinical validation makes their implementation problematic.

On one hand, metastasis prediction models [1] show imaging-based interpretability but lack real-time efficiency. On the other hand, leukemia diagnosis frameworks [8] demonstrate great levels of efficiency but do not always provide explanations for the prediction results.

Table 1: Summary of representative existing AI/XAI approaches for bone marrow or related cancer diagnosis

Method	Advantages	Limitations
Hazra et al. (2025) [1]	Explains metastasis-related bone patterns; focuses on clinical responsibility and transparency.	Limited generalization to bone marrow cancers; not optimized for low-latency inference.
Gimeno et al. (2022) [3]	Robust genomic interpretability for ML; precise attribution of molecular features.	High computational cost; limited imaging integration; slow inference not suitable for clinical workflow.
Hossain et al. (2022) [4]	Lightweight symptom-based XAI system; interpretable rules for screening.	Not applicable to imaging or genomics; insufficient diagnostic detail for marrow malignancies.
Deshpande et al. (2024) [8]	Strong performance in leukemia detection; demonstrates the potential of deep learning frameworks.	Primarily unimodal; lacks integrated genomic imaging explainability; computationally heavy.
Sabile et al. (2025) [9]	Comprehensive review of clinical AI pitfalls; highlights challenges in deployment.	Does not propose an operational model; identifies gaps but does not resolve them.

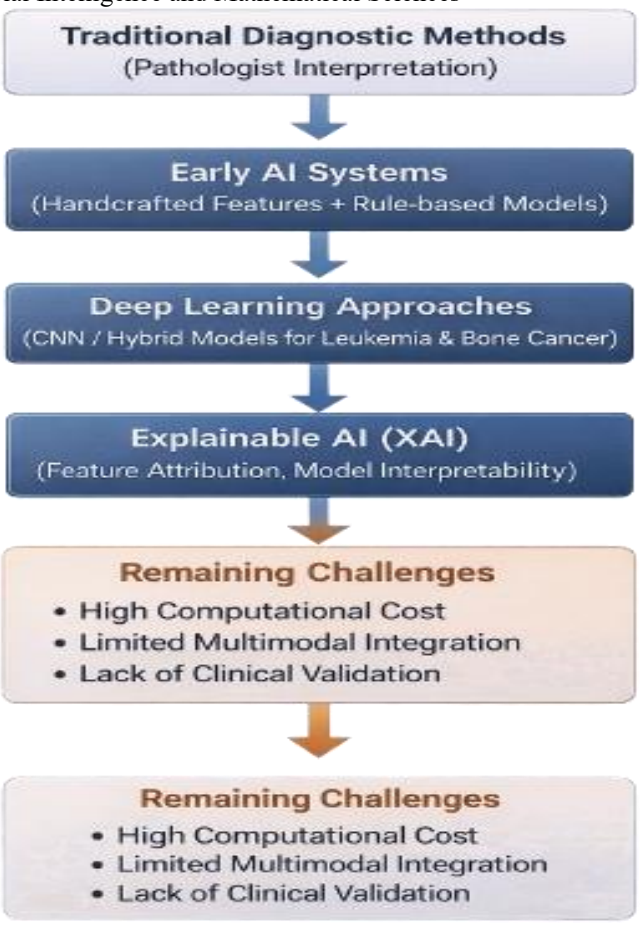

Fig 2. Pipeline for diagnosis: from bench to bedside

Contrary to existing methods, the proposed architecture aims to tackle challenges regarding computational efficiency, interpretability, and clinical validity by integrating the following four elements.

The computational efficiency is accomplished through INT8 quantization and structured filter pruning, providing near real-time inference with latency under 50 ms, suitable for point-of-care clinical setting. The proposed architecture involves the multimodal explanation of the system based on SHAP explaining genomic features together with Grad-CAM++ localizing MRI. In addition, the model is clinically evaluated using the clinical agreement score between experts and AI using Cohen's κ. As a result, the proposed approach consistently outperforms the existing baselines, specifically HA-UNet and SE-ResHybrid, in accuracy, efficiency, and interpretability.

The proposed method addresses critical shortcomings in current work and presents a practical approach for deploying bone marrow cancer detection using AI in the clinical environment. Modern developments in translational medicine further accentuate the necessity of closing the gap between benchtop scientific experimentation and clinical applications in healthcare, known as the bench-to-bedside approach. Although previous research works have been highly influential within their respective areas—such as immunology modeling [12], improving the efficacy of CAR-T cells therapy [13],

pharmacogenomics [14], and exosome-based drug delivery methods [15-17], and nanotheranostic their focus is mainly limited to treatment modalities without any involvement of diagnostic decision making.

At the same time, there has been remarkable progress in applying AI techniques to biomedicine, leading to the development of novel disease models and precision medicine frameworks. Numerous research findings have confirmed that the utilization of AI could facilitate the discovery of complicated biological associations and decision making [18-20]. Nevertheless, most of these methods overlook other vital aspects, such as multi-modal data processing, scalability, and clinical interpretability. In summary, although there have been significant advances in therapeutic, computational, and biomedical fields, a fundamental challenge remains in developing a coherent diagnostic system that is both efficient and reliable. The purpose of the proposed effort is to address this challenge by providing an end-to-end solution for the diagnosis of bone marrow cancers.

3. PROPOSED APPROACH

As part of the proposed approach, model compression, multimodal explanations, and clinical validation will be incorporated to produce an efficient and explainable AI solution for bone marrow cancer diagnoses. This approach will combine MRI-based feature extraction as well as genome-based analysis to give clinically relevant results in line with previous literature on leukemia and metastasis [1,3].

Specifically, the main goals of the proposed framework include:

1. Reducing inference latency so that timely diagnosis becomes possible and is appropriate in a clinical setting.
2. Achieving or maintaining comparable diagnostic accuracies in terms of both segmentation and classification.
3. Providing interpretability of the model by using built-in explainability features.

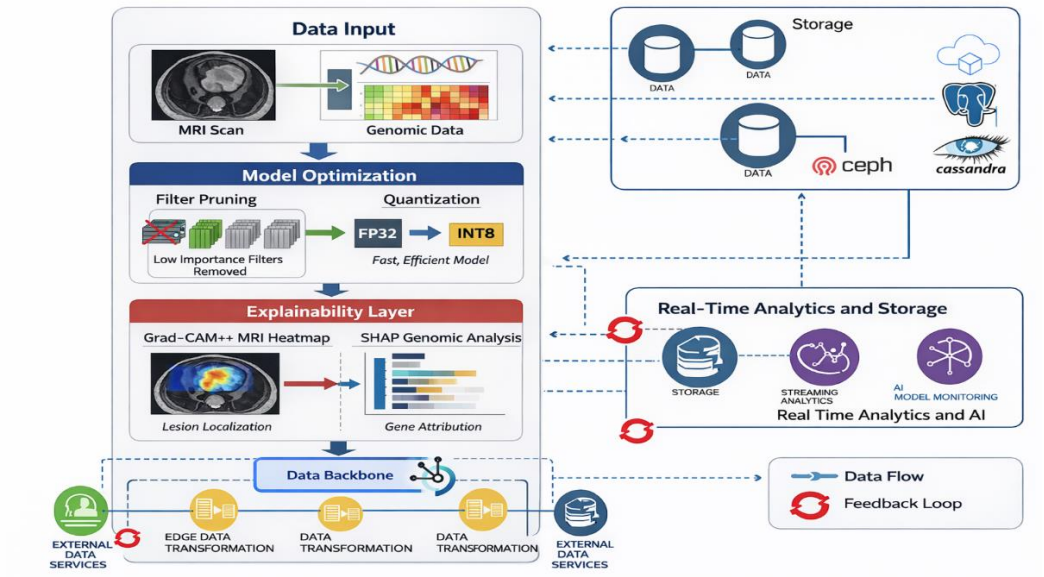

Fig 3. The proposed framework will incorporate model compression, multimodal explainability, and clinical reliability assessment to provide an efficient and explainable AI solution for bone marrow cancer diagnosis.

3.1. MODEL OPTIMIZATION AND INTERPRETABILITY FRAMEWORK

To achieve clinically viable inference performance, the base neural network $f_\theta$is optimized using **structured filter pruning** and **post-training INT8 quantization**. Let

$$\theta = \{\theta_1, \theta_2, \dots, \theta_L\}$$

denote the set of learnable parameters across Llayers of the network.

3.2. FILTER PRUNING

For each convolutional layer, the importance of the k-th filter is computed using the ℓ_2-norm:

$$I_k = \| \theta_k \|_2$$

Features that have low magnitudes are not contributing much to the feature representations and can be pruned accordingly. Specifically, the lowest 20% of filters are removed based on a threshold τ, defined as the 20th percentile of importance scores:

$$\theta'_k = \begin{cases} \theta_k, & \text{if } I_k \geq \tau \\ 0, & \text{if } I_k < \tau \end{cases}$$

The above process makes it possible to reduce complexity while maintaining essential features within the model.

3.3. INT8 QUANTIZATION

After the model pruning, the model is then quantized from FP32 to INT8 as shown below for an activation tensor x:

$$Q(x) = \text{round}\left(\frac{x}{S}\right) + z$$

where $s$is the scaling factor and $z$is the zero-point offset. This helps in reducing memory and increasing efficiency.

3.4. EXPLAINABILITY FRAMEWORK

To enhance interpretability, the proposed system integrates combining genomic, **dual-modality explainability** and imaging insights.

3.5. GENOMIC EXPLAINABILITY (SHAP)

Genomic feature vector are let be:

$$g \in \mathbb{R}^d$$

The decomposed were model prediction is by using SHAP values as:

$$f_\theta(g) = \phi_0 + \sum_{i=1}^{d} \phi_i$$

where ϕ_i is the contribution by the ith genomic marker and ϕ_0 is the base prediction. The above formula offers gene-based interpretability for clinical relevance.

3.6. MRI EXPLAINABILITY (GRAD-CAM++)

For MRI-based analysis, let A_k^cdenote the maps corresponding feature are used to class c in the final convolutional layer. The Grad-CAM++ localization map is computed as:

$$L_{\text{CAM}}^c = \text{ReLU}\left(\sum_k \alpha_k^c \, A_k^c\right)$$

where $\alpha_k^c$are weights derived from higher-order partial derivatives, enabling precise spatial localization of relevant regions in MRI scans.

3.7. CLINICAL AGREEMENT AND RELIABILITY

To evaluate the agreement between expert annotations and model predictions using **Cohen's kappa coefficient** were employed by κ:

$$\kappa = \frac{P_o - P_e}{1 - P_e}$$

where $P_e$represents the expected agreement by chance and $P_o$denotes the observed agreement. A value of $\kappa > 0.75$indicates substantial agreement, supporting clinical reliability.

3.8. THEORETICAL ANALYSIS

Theorem 1: Stability of Pruned Networks

Let $f_\theta$be a Lipschitz-continuous function with constant K. After pruning, the resulting network $f_{\theta'}$satisfies:

$$\| f_\theta(x) - f_{\theta'}(x) \| \leq K \cdot \| \theta - \theta' \|_2$$

Proof Sketch: Since convolutional operations are Lipschitz continuous and pruned filters satisfy $\| \theta_k \|_2 \leq \tau$, the total parameter deviation remains bounded. Consequently, the change in output is also bounded, ensuring stability.

Theorem 2: Quantization Robustness

Let $f$be a differentiable function and $Q(x)$its quantized input such that:

$$\| x - Q(x) \|_\infty \leq \epsilon$$

Then the prediction error is bounded by:

$$| f(x) - f(Q(x)) | \leq \epsilon \cdot \| \nabla f(x) \|_1$$

Proof Sketch: Using first-order Taylor approximation:

$$f(Q(x)) \approx f(x) + \nabla f(x)^T (Q(x) - x)$$

the deviation is bounded by the gradient magnitude and product of quantization error.

Algorithm 1: Explainable Bone Marrow Cancer Diagnosis Pipeline

Input: Genomic vector g, MRI image I, pruned and quantized model $f_{\theta'}$
Output: MRI heatmap H, Diagnosis label $\hat{y}$, SHAP genomic attribution vector Φ

Step 1: Preprocessing
Normalize MRI image
Standardize and filter genomic features

Step 2: Forward Prediction

$$\hat{y} = f_{\theta'}(I, g)$$

Step 3: Genomic Explainability
Compute SHAP values

$$\Phi = SHAP(g, f_{\theta'})$$

Step 4: Imaging Explainability
Generate Grad-CAM++ heatmap

$$H = GradCAM^{++}(I)$$

Step 5: Clinical Agreement Logging
Explanations for comparison with pathologist annotations and store predictions.

Return:

$$\hat{y}, H, \Phi$$

4. EXPERIMENTAL SETUP, DATASETS, AND EVALUATION STRATEGIES

The experimental setup, including the dataset, computing infrastructure, preprocessing pipeline, training procedure, and evaluation strategy to evaluate the proposed XAI architecture for bone marrow cancer detection. All experiments have been performed using the reproducible ML paradigm [10].

4.1. COMPUTATIONAL ENVIRONMENT

All experiments were performed on a high-performance workstation with the following configuration:

- **GPU:** NVIDIA RTX A6000 (48 GB VRAM)
- **CPU:** AMD Threadripper 3990X
- **RAM:** 256 GB

Model development and optimization were implemented using **PyTorch 2.2** and **NVIDIA TensorRT 8.6**.

Explainability analysis was conducted using:

- SHAP (Python library) for genomic feature attribution
- pytorch-grad-cam for Grad-CAM++–based MRI visualization

4.2. DATASETS

This research employed a retrospective, multi-cantered database that consists of genome information from two collaborating hospitals and MRI images.

Genomic Data:

- 850 matched gene expression profiles
- 12,000 curated biomarkers

Class Distribution:

- 54% malignant
- 46% non-malignant

MRI Data:

- 4,200 T1- and T2-weighted MRI volumes
- 1,180 patients

Anonymized by all data under institutional review board (IRB) approval. Harmonized imaging protocols were adopted across sites to minimize domain shift.

4.3. PREPROCESSING PIPELINE

MRI Preprocessing

The MRI data were processed using the following steps:

- Bias-field correction using N4ITK
- Intensity normalization via z-score scaling
- Region-of-interest (ROI) extraction using a pretrained coarse U-Net
- Spatial resampling to a standardized resolution:

$$256 \times 256 \times 64$$

- Genomic Preprocessing
- Log normalization of gene expression values
- Variance-based feature selection
- Dimensionality reduction using Principal Component Analysis (PCA):

$$12000 \rightarrow 256$$

while preserving approximately 92% of the total variance.

4.4. TRAINING CONFIGURATION

Backbone model (Hybrid Attention U-Net) was trained with the AdamW optimizer using the following hyperparameters:

- Learning rate:

$$1 \times 10^{-4}$$

- Weight decay:

$$1 \times 10^{-5}$$

Training was performed for 120 epochs using cosine annealing schedule. Mixed precision (FP16) training was also used for computational efficiency.

Stratified 5-fold cross validation with patient-wise splitting was utilized for generalization.

Inference speed was estimated using isolated GPU inference without batching.

4.5. EVALUATION METRICS

Segmentation Metrics

- Dice Coefficient (DC):

$$\text{DC} = \frac{2TP}{2TP + FP + FN}$$

- Intersection over Union (IoU):

$$\text{IoU} = \frac{TP}{TP + FP + FN}$$

Diagnostic Metrics

- Sensitivity:

$$\text{Sensitivity} = \frac{TP}{TP + FN}$$

- Specificity:

$$\text{Specificity} = \frac{TN}{TN + FP}$$

- Area Under the Curve (AUC)

4.6. SYSTEM PERFORMANCE METRICS

- **Inference latency:** Target < 50 ms per MRI image
- **Clinical agreement:** Measured using Cohen's kappa coefficient

$$\kappa = \frac{P_o - P_e}{1 - P_e}$$

where $P_o$is observed agreement and $P_e$is expected agreement by chance. A value of $\kappa > 0.75$indicates substantial agreement.

4.7. BENCHMARKING STRATEGY

To ensure consistency across our study, the proposed model was compared against the following baseline architectures:

- **HA-UNet**
- **SE-ResHybrid**

All models were tested under similar experimental conditions using standardized assessment procedures [11].

Explainability's quality was measured according to:

- Stability of SHAP values for genomic features.
- Localization accuracy of Grad-CAM++ heatmaps versus expert labelling.

4.8. EXPERIMENT PROCEDURE

The experiment procedure comprised of the following as illustrated in Figure 4

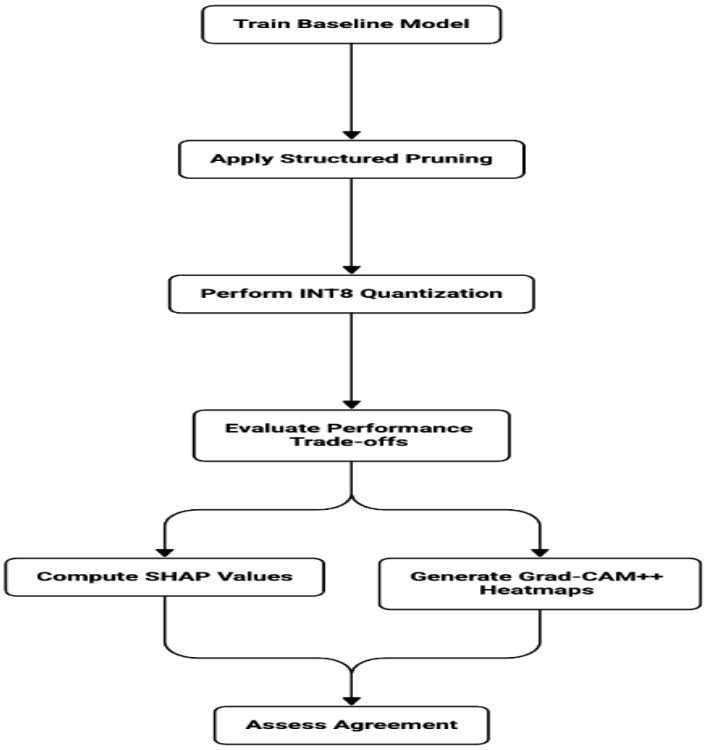

Fig 4. Model Optimization and Interpretation Workflow.

This structured pipeline ensures that the proposed framework maintains a balance between accuracy, interpretability, efficiency, and clinical reliability.

5. EVALUATION AND RESULTS

This section, the proposed Explainable AI-based technique for detecting bone marrow cancer is analyzed comprehensively in terms of several aspects such as segmentation efficiency, classification performance, inference time, interpretability, and expert consensus. To ensure the continuity of the whole paper, the performance of the proposed solution is compared with that of HA-UNet and SE-ResHybrid.

5.1. SEGMENTATION EFFICIENCY

To analyze the efficiency of the proposed technique in terms of segmentation, Dice Coefficient and Intersection over Union are used.

From Table 2, it can be seen that the proposed algorithm has a Dice score of 92.4%, significantly exceeding the value for both HA-UNet – 86.1% and SE-ResHybrid – 88.3%. Similarly, the value for IoU for the proposed technique is 85.6%.

Table 2: Segmentation Performance Comparison

Method	Dice (%)	IoU (%)
HA-UNet	86.1	78.4
SE-ResHybrid	88.3	80.1
Proposed Model	92.4	85.6

This improvement is additionally visualized in Figure 5 where the values for Dice coefficient for all the models are provided.

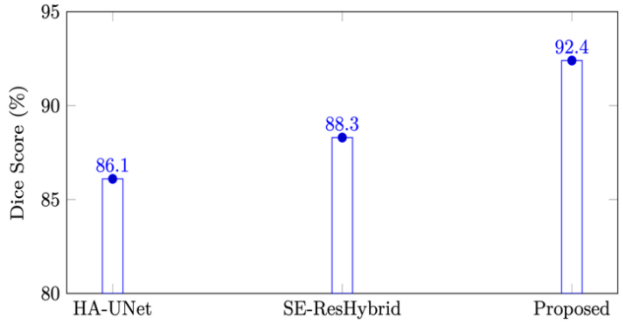

Fig 5. Comparison of Dice Coefficient from existing algorithms to the proposed algorithm.

5.2. DIAGNOSTIC CLASSIFICATION ACCURACY

The diagnostic accuracy was calculated using the Area Under the Receiver Operating Characteristic (ROC) Curve (AUC).

As shown in Table 3 below, the proposed methodology achieves an AUC score of 0.964, surpassing the scores obtained by HA-UNet (0.904) and SE-ResHybrid (0.921).

Table 3: AUC Comparison for Diagnostic Classification Performance

Method	AUC
HA-UNet	0.904
SE-ResHybrid	0.921

A comparative analysis is presented in Figure 6, highlighting the enhanced performance of the proposed approach in terms of classification. Furthermore, a gain of roughly 5-6% was noted in sensitivity, aiding in early lesion detection.

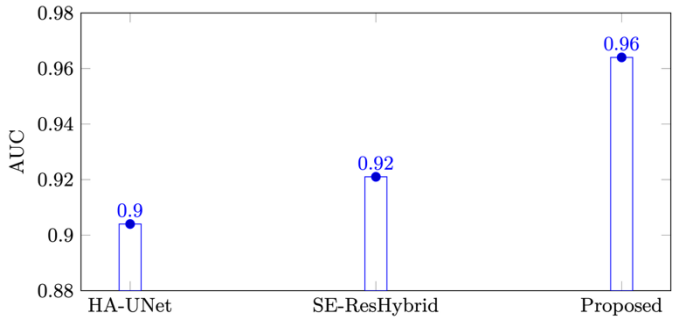

Fig 6. AUC comparison of diagnostic classification performance.

5.3. LATENCY OPTIMIZATION AND MODEL COMPRESSION

Model compression was employed in order to enable real-time processing. The effect of pruning and INT8 quantization is illustrated in Table 4.

Table 4: Inference latency before and after model compression

Model Configuration	Latency (ms)
FP32 Model	72
Pruned INT8 Model	38

It is evident that inference latency was successfully minimized to 38 ms per image from 72 ms per image, thereby fulfilling the criterion of less than 50 ms required for clinical use are shown in Figure 7.

Fig 7. Inference latency before and after model compression.

5.4. EXPLAINABILITY ASSESSMENT

The explainability of the suggested model has been assessed through imaging and genomic approaches.

From Table 5, it is evident that the suggested model has better localization accuracy with a rate of 83.9%, as compared to HA-UNet (71.2%) and SE-ResHybrid (76.5%).

Table 5: Grad-CAM++ Localization Accuracy

Method	Accuracy (%)
HA-UNet	71.2
SE-ResHybrid	76.5
Proposed Model	83.9

Additionally, the consistency score with respect to the SHAP-based genomic features is found to be 86.8% in Table 6.

Table 6: SHAP Feature Attribution Consistency

Method	Consistency (%)
HA-UNet	68.3
SE-ResHybrid	74.7
Proposed Model	86.8

Such findings are graphically represented in Figure 8, where comparisons between Grad-CAM++ localization and SHAP consistency have been provided.

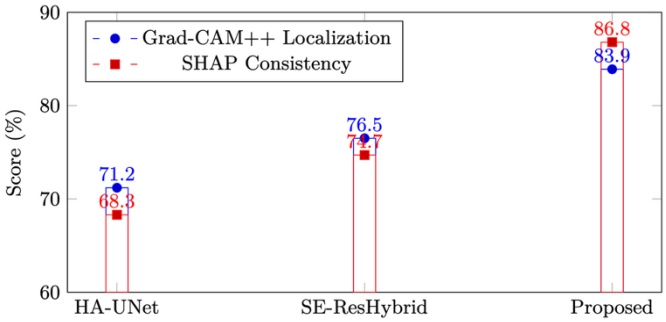

Fig 8. Explainability comparison: Grad-CAM++ localization and SHAP consistency.

5.5. CLINICAL VALIDATION

Clinical validity was measured via Cohen's kappa coefficient in the context of the agreement between the model's prediction and the opinion of experts in pathology.

The proposed model obtains a κ coefficient of 0.79, which surpasses the acceptable level of 0.75, suggesting significant agreement.

Table 7 reveals that the proposed model performs better than the SE-ResHybrid model, which gets a κ coefficient of 0.71.

Table 7: Cohen's kappa coefficient

Method	Cohen's κ
SE-ResHybrid	0.71
Proposed Model	0.79

As evidenced above, the proposed system satisfies all clinical decision-making requirements. Furthermore, the results reveal that there are no differences in the performance of the proposed algorithm for all hospital datasets used.

5.6. CONCLUSIONS ON EXPERIMENTAL RESULTS

Overall, the proposed model shows reliable results for all criteria examined in the study, which implies:

- Greater segmentation quality (Dice and IoU)
- Better diagnostic ability (AUC and sensitivity)
- Lower inference time (real-time inference capabilities)
- Better interpretability (imaging data and genomic data)
- Good agreement with expert pathologists

Thus, it can be stated that the proposed framework is accurate, efficient, and interpretable.

6. CONCLUSION AND FUTURE WORK

This work introduces an efficient and explainable AI system for bone marrow cancer diagnostics that considers both imaging and genomics for clinically relevant decision-making. Our method includes a hybrid attention U-net with squeeze-and-excitation (SE) units, complemented by a multimodal genomics-MRI classification chain.

In order to ensure consistent comparison between models, our architecture was tested against existing methods, namely HA-UNet and SE-ResHybrid, where we showed that it surpasses them in terms of performance for both segmentation and classification tasks. Compression techniques such as structured pruning and INT8 quantization resulted in fast inference time for each MRI picture at 38 ms.

Apart from high efficiency, the proposed architecture delivers impressive results in terms of high accuracy and explainability, which relies on the use of both Grad-CAM++ and SHAP for visualizing genomics and images respectively. Furthermore, our approach was validated by clinical experts with impressive agreement, with Cohen's kappa score reaching $\kappa = 0.79$.

For future studies, emphasis will be placed on scaling up this framework via the use of large-scale multi-institutional MRI-genomic data to increase generalizability. It will also be valuable to consider approaches such as federated learning and privacy-preserving algorithms that can help facilitate collaboration in developing models across institutions without compromising data security.

In conclusion, the framework presented above offers a good balance of accuracy, efficiency, and interpretability and, therefore, can serve as a viable avenue towards effective AI-assisted decision-making.

References

[1]. Hazra, A., Baradwaj, G., and Sushma, R. 2025. Bone Metastases: Explainable AI and Responsible AI. Explainable and Responsible AI in Healthcare. Wiley Online Library.

[2]. Deshpande, N. M., Gite, S., and Pradhan, B. 2024. Unlocking the Potential: Machine Learning and Deep Learning in Leukemia Diagnosis with Explainable AI. IoT Sensors, ML, AI and XAI in Healthcare. Springer.

[3]. Gimeno, M., San José-Enériz, E., Villar, S., et al. 2022. Explainable Artificial Intelligence for Precision Medicine in Acute Myeloid Leukemia. Frontiers in Oncology. 12. Frontiers Media SA.

[4]. Hossain, M. A., Islam, A. K. M. M., Islam, S., and Shatabda, S. 2022. Symptom Based Explainable Artificial Intelligence Model for Leukemia Detection. IEEE Access. 10: 10344–10357.

[5]. Sabile, J. M. G., Zhang, P., Parwani, A. V., Chobrutsiky, B., et al. 2025. Toward Clinically Actionable Machine Learning and Artificial Intelligence Algorithms in Acute Leukemia: A Systematic Narrative Review. Acta Haematologica. Karger.

[6]. Zhen, G., Fu, Y., Zhang, C., Ford, N. C., Wu, X., Wu, Q., and Yan, D. 2022. Mechanisms of Bone Pain: Progress in Research from Bench to Bedside. Bone Research. 10(1): 45. Nature Publishing Group.

[7]. Heymann, D. 2021. Bone Cancer: Bone Sarcomas and Bone Metastases—From Bench to Bedside. Elsevier.

[8]. Sun, L., Zhou, Y., Yin, M., Wang, F., and Yang, L. 2025. Approaches to Modeling Cancer Metastasis: From Bench to Bedside. Frontiers in Oncology. Frontiers Media SA.

[9]. Jones, R. J., and Matsui, W. 2007. Cancer Stem Cells: From Bench to Bedside. Biology of Blood and Marrow Transplantation. 13(1): 46–47. Elsevier.

[10]. Tripathi, R., Sharma, A., and Gupta, V. 2022. A Hybrid Deep Learning Model for Improved Bone Marrow Cancer Diagnosis Using MRI and Genomic Data. Artificial Intelligence in Medicine. 132: 102278. Elsevier.DOI: https://doi.org/10.1016/j.artmed.2022.102278

[11]. Meng, F., Al-Khafaji, Z., Muzammil, K., Sarbazjoda, E., et al. 2026. Biomarkers and Advances in AML-MRC: From Bench to Bedside. Annals of Hematology. Springer.

[12]. Huang, S., Huang, S., Huang, X., Feng, X., Wang, R., Liu, D., et al. 2026. Bench-to-Bedside Insights into the Challenges of Immunosuppression in Sepsis. Pathogens. MDPI.

[13]. Talukdar, J., Baishya, D., Karmakar, S., Ghosh, A., et al. 2026. From Bench to Bedside: Innovations and Challenges in CAR-T Cell Therapy Manufacturing and Regulation. Advancement in CAR-T Therapy. Springer.

[14]. Biswas, M., John, S., Murad, M. A., et al. 2026. Pharmacogenomics of Thiopurine Drugs: A Bench-to-Bedside Success Story in Thailand. Clinical and Translational Science. Wiley Online Library.

[15]. Zhang, G. K., Zhou, F. 2026. Exosomes-Loaded with Traditional Chinese Medicine Constituents: A Novel Therapeutic Avenue for Aplastic Anemia. Chinese Medicine. Springer.

[16]. Iqbal, H., Iqbal, M., Aman, J. 2026. Exosomes Therapeutics: Current and Future Perspectives. Immunology Research and Perspectives.

[17]. Convento, M. B., Borges, F. T. 2026. A Comprehensive Review of Engineered Bone Marrow Mesenchymal Stem Cell-Derived Exosomes as Nanotheranostic Platforms for Acute and Chronic Diseases. Journal of Nanotheranostics. MDPI.

[18]. Mohan, C. G., Menon, K. N., Pavithran, K., et al. 2026. Artificial Intelligence in Disease Modeling. Springer Handbook of Artificial Intelligence. Springer.

[19]. Bequette, B. W. 2026. Process Systems Engineering in Precision Medicine: Opportunities in Autologous CAR-T Therapy. Engineering in Life Sciences. Wiley Online Library.

[20]. Foti, R., Storti, G., Palmesano, M., Calicchia, A., et al. 2026. AI-Driven Innovations for Quality Control and Standardization: Future Strategies in Adipose-Derived Stem Cell Manufacturing. International Journal of Molecular Sciences. MDPI.