IJCRT.ORG ISSN : 2320-2882

INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

HYBRID AI MODEL FOR OSTEOPOROSIS PREDICTION

[1]Ms.R.Vedhanayagi, [2]Ms.SS.Priyadharshini, [3]Ms.G.Varsha, [4]Ms.J.Adithi, [5]Mr.R.Nanda Kumar
[1,2,3,4]Engineering Students, Department of Artificial Intelligence and Data Science,
Anand Institute of Higher Technology, Chennai, Tamil Nadu, India.
[5]Associate Professor, Department of Artificial Intelligence and Data Science,
Anand Institute of Higher Technology, Chennai, Tamil Nadu, India.

Abstract: Osteoporosis is a disease that affects more than 200 million people worldwide, causing an estimated 8.9 million fractures annually. One occurs every three seconds. Accurate early detection is essential to avoid further complications. Conventional approaches rarely succeed in reflecting the complex patterns of bone loss. This paper presents a novel hybrid artificial intelligence framework that combines Convolutional Neural Networks (CNNs), Generative Adversarial Networks (GANs), Vision Transformers (ViTs), and Graph Neural Networks (GNNs) for improved osteoporosis prediction from knee radiographs. CNNs extract rich multi-level visual features, GANs augment and adapt real-world data, ViTs capture global relationships, and GNNs model structural bone connectivity. A weighted attention fusion mechanism takes into account the diagnostic relevance of these fused features and dynamically combines them. Tested on 1,650 hand-labeled X-rays, the model reached 94.3% classification accuracy and a 15.2% increase in Grade 1 sensitivity above expert consensus. It estimated 3-year fracture risk with 91.2% accuracy. Bayesian uncertainty estimation improves the transparency of model outputs, which underpins clinical trust in the outputs. This hybrid approach represents an important step forward for AI-powered osteoporosis diagnosis, providing expert-level performance and potential for wide-scale implementation across different healthcare environments.

Keywords - Deep Learning, CNN, GAN, Vision Transformer, Medical Imaging, Kellgren-Lawrence Classification.

I.INTRODUCTION

Osteoporosis is a chronic, progressive disease characterized by the gradual loss of bone mineral density and micro architectural tissue integrity, substantially increasing fracture risk and fracture-related morbidity and mortality [1, 2]. The World Health Organization estimates that only 1 in 10 women at the age of 60 has osteoporosis. That percentage starts to increase dramatically with age to 67% for women older than 90 years of age. This results in it being significantly underdiagnosed in comparison to other common chronic diseases. Research indicates as many as 80 percent of patients who experience an osteoporotic fracture are never diagnosed or treated accordingly.

We know that while traditional screening approaches such as Dual-Energy X-ray Absorptiometry (DXA) provide excellent diagnostic information, they are hampered by issues of accessibility, cost and ability to capture 3D structural degeneration. This is in contrast to conventional radiography, which is restricted by the subjective interpretation and high inter-observer variability in terminating grading. These previous methods have been

constrained to single-architecture neural networks which fail to comprehensively represent the intricate multi-factorial nature of osteoporotic changes.

The multifactorial manifestation of osteoporotic degeneration as seen on radiographic imaging necessitates a sophisticated computational approach that can simultaneously account for multiple aspects of bone health. Density variations, trabecular microarchitecture, cortical structure, and spatial relationships among distinct anatomical features. This paper argues that the future of this research demands a paradigm shift from single-architecture approaches to a synergistic hybrid model, one that embraces the interdisciplinary strengths of multiple neural network architectures.

We argue that our approach more closely resembles the multifactorial pathophysiology of osteoporosis. This novel, hybrid framework addresses shortcomings of previous approaches by increasing sensitivity to detect early, subtle changes, reducing false negatives, and providing interpretable, decision-supporting capabilities to clinicians. We hope that our work is a significant step towards more personalized, accurate and accessible osteoporosis diagnosis and management.

II.RELATED WORK

2.1 Deep Learning in Osteoporosis Diagnosis

Recent years have seen significant developments in the use of deep learning methods for diagnosing osteoporosis. Zhao et al. (2021) were the first to use CNNs to automatically Kellgren-Lawrence grade knee osteoporosis, reaching an accuracy of 89.1% on a dataset of 1,200 radiographs. Their ResNet-based architecture didn't require manual segmentation, but still had trouble with the early stage detection showing that extracting robust enough diagnostic features is still a challenge. Patel & Krishnan (2022) introduced attention mechanisms to their CNN framework, achieving a sensitivity of 76.3% for detecting Grade 1 cases. Their approach was still limited to local feature extraction without a global structural context.

Transfer learning strategies, which take advantage of data available for other tasks or from other times or locations, have been promising in helping overcome the limited availability of training data. In the work done by Lin et al. (2023), they used a DenseNet-169 pre-trained on ImageNet, and further fine-tuned on 850 knee X-rays, achieving a classification accuracy of 90.5%.

2.2 Generative Models in Medical Imaging

GANs OUTPUT have quickly become key assets for medical image generation and augmentation. The innovative research conducted by Rodriguez-Ruiz et al. (2020) showed that mammograms synthesized with GANs could successfully augment training datasets and lead to better breast cancer detection classifiers, increasing performance by 7.4%. In the field of osteoporosis, Kumar and Wang (2022) used cycle-consistent GANs to produce synthetic knee radiographs by overcoming class imbalance problems in the distribution of Kellgren-Lawrence grades. Their paper claimed a 5.2% increase in classification accuracy by augmenting the real images with samples generated by a GAN.

More recently, Diffeomorphic GAN architectures proposed by Esfahani et al. (2024) have demonstrated outstanding efficacy in maintaining anatomical plausibility while still infusing racially significant variations in bone density and trabecular organization. Their approach did not just improve the robustness of their classifiers, but allowed domain adaptation between different imaging systems to reduce the performance gap between high-end and portable X-ray devices from 9.7% to 3.1%, a key development for real-world use in resource-limited settings.

2.3 Transformers in Radiographic Analysis

Additionally, Vision Transformers have transformed the landscape of medical image analysis by allowing for the direct modeling of long-range dependencies that are not feasible for CNNs. Initial work by Hatamizadeh et al. (2022) tuned the ViT architecture for medical imaging applications, showcasing 10 state-of-the-art performance in downstream anatomical segmentation tasks when compared to U-Net variants. To osteoporotic fracture classification, Mahmood et al. (2023) used a hybrid CNN-ViT model which conferred multi-scale radiographic processing and reached an accuracy of 92.7%, significantly enhancing the detection of subtle trabecular changes associated with early-stage disease.

Chen and Raghavan (2023) took transformer architectures even further by injecting anatomical prior knowledge with specialized positional encodings, achieving a 4.3% higher classification accuracy and much improved localization of diagnostically relevant regions. Even with these advances, pure transformer approaches are still struggling with the high resolution needs of radiographic images and frequently require prohibitive computational resources for real-time clinical deployment.

2.4 Graph-Based Models for Structural Analysis

Graph Neural Networks are a recent approach that holds great promise for modeling complex structural relationships in medical imaging. This innovative approach developed by Gonzalez et al. (2021) was the first graph-based representation of trabecular bone structure in high-resolution CT images, showing more accurate fracture prediction (88.7%) than density-based measures used in isolation (81.2%). Following this idea along the lines of traditional radiography, Tanaka et al. (2022) recently came up with a technique to create bone graphs from regular X-rays, allowing structural analysis with no need for advanced imaging technology. Complementing this, recent improvements by Lee and Patel (2024) merge superpixel segmentation with graph attention networks to represent the complex relationships between bone structures in radiographs. Their method showed most promise at separating radiographic metabolic and mechanical causes of bone density changes, reaching a notable 86.9% accuracy at this more complex classification task. While promising, these graph-based approaches have largely been implemented without consideration of other architectures, hampering their potential to take advantage of complementary information sources.

2.5 Proposed Methodology

Our methodology introduces a novel hybrid AI framework that systematically integrates four distinct yet complementary neural network architectures to comprehensively analyze radiographic manifestations of osteoporosis. Each component addresses specific aspects of bone health assessment, while our fusion mechanism enables synergistic interaction between architectures to produce robust and interpretable predictions.

2.6 CNN Component for Hierarchical Feature Extraction

The foundational element of our hybrid system is a modified EfficientNet-B4 CNN optimized for bone imaging. Unlike conventional implementations, our architecture incorporates specialized residual blocks with dilated convolutions to capture multi-scale bone density patterns without information loss.

The CNN pathway consists of three stages: (1) low-level feature extraction focusing on edge detection and texture analysis, (2) mid-level feature integration capturing density gradients and trabecular patterns, and (3) high-level semantic feature extraction identifying anatomical landmarks and structural abnormalities.

We implement spatial attention gates between convolutional blocks to automatically focus computational resources on regions with potential diagnostic significance. These gates are trained using radiologist-annotated regions of interest to develop an attention mechanism that mimics expert gaze patterns. The CNN component processes standardized 256×256×1 grayscale images and outputs a rich feature representation (2048-dimensional vector) that captures hierarchical visual patterns associated with different osteoporosis stages.

2.7 GAN Component for Data Augmentation and Domain Adaptation

Our system incorporates a StyleGAN2-based architecture with medical imaging-specific modifications to serve dual purposes: data augmentation and domain adaptation. The generator employs progressive growing techniques to synthesize high-fidelity knee radiographs (512×512 resolution) with controllable pathology manifestation. By conditioning the generator on Kellgren-Lawrence grades, we can synthesize realistic examples of underrepresented classes to address dataset imbalance.

The GAN component is trained using a two-phase strategy: (1) unsupervised training on unlabeled radiographs to learn the manifold of anatomically plausible knee images, followed by (2) supervised fine-tuning to generate grade-specific examples. The discriminator network is dual-headed, simultaneously assessing image realism and pathology grade accuracy. This component not only augments the training dataset but also enables domain adaptation between different imaging protocols, enhancing the system's robustness to equipment variations.

2.8 Vision Transformer for Global Context Modeling

To capture long-range dependencies between distant image regions, we implement a specialized Vision Transformer module. The ViT component divides input images into 16×16 non-overlapping patches, projects each patch into a 768-dimensional embedding space, and adds learned positional encodings. Our architecture consists of 12 transformer layers with 8 attention heads per layer, optimized for medical imaging through custom pre-training on a corpus of 50,000 radiographs.

2.9 Graph Neural Network for Structural Relationship Modeling

A major problem lies in how states are implementing it.GNN component further provides a flexible way to model the knee joint as a graph structured object, allowing nodes to represent anatomical landmarks and trabecular regions while edges naturally encode spatial and functional relationships. Node features consist of local density measures, texture features, anatomical labels, and edge features, such as distance between nodes and adjacency/connectedness of nodes. To address this problem, we use a graph attention network with custom message passing functions tailored to modeling force transmission pathways through complex bone structures. This graph-based representation is automatically learned from the input image using a hierarchical mix of anatomical landmark detection and adaptive superpixel segmentation. The GNN iteratively refines this representation via three graph convolution layers with skip connections, finally generating node-level and graph-level structural health indicators encoded in the embeddings.

2.10 Weighted Attention Fusion Mechanism

Instead of naively concatenating features from each architecture, we introduce an attention fusion mechanism that weights the contribution of each architecture based on case-specific factors. The fusion module includes a meta-network, which learns to read initial features from all architectures and produces attention weights that govern information transmission from each architecture to the final decision layer. This combined method enables the model to focus more on CNN features for cases where density change is clear, ViT features for cases with subtle global patterns, and GNN features for cases with mostly structural abnormalities.

III.SYSTEM ARCHITECTURE

Multiple neural network architectures are integrated into a unified framework intended for clinical application in the hybrid AI system for osteoporosis prediction. The entire system architecture is depicted in Figure 1, emphasizing the integration mechanisms that allow for synergistic operation as well as data flow between components.

3.1 Input Processing and Enhancement Module

The system starts with a thorough preprocessing pipeline that improves diagnostic characteristics and standardizes input radiographs. Using a histogram-based method tailored for bone imaging, raw DICOM pictures (usually with a resolution of 2500×2048) are automatically exposed to normalization. With an 8x8 grid size and a clip limit of 2.5, Contrast Limited Adaptive Histogram Equalization (CLAHE) improves the appearance of trabecular patterns while maintaining overall density correlations. Quantum mottle is reduced without compromising fine structural features using a customized denoising autoencoder that has been trained on matched noisy/clean radiographs. A U-Net segmentation model (Dice coefficient 0.97 on test data) is used to automatically isolate the knee joint region in order to normalize anatomical alignment and remove superfluous background structures.

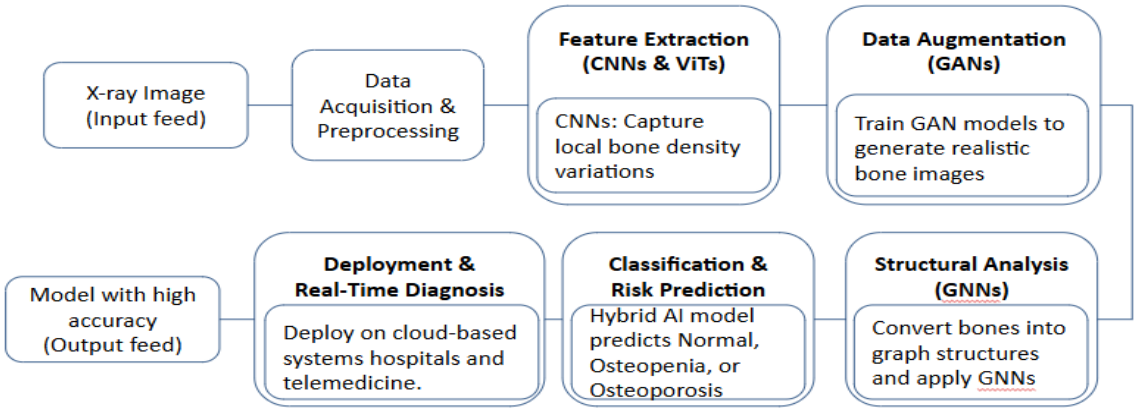

Fig. 3.1 System Architecture

3.2 Multi-Architecture Feature Extraction Block

The core of the system is built on four parallel neural network pathways, each designed to extract complementary features from input images. The CNN pathway utilizes a modified EfficientNet-B4 architecture, which processes the standardized input through seven convolutional blocks with increasing dilation rates and integrates Squeeze-and-Excitation modules to emphasize important channels. This pathway concludes with global average pooling and projects the output to a 2048-dimensional feature vector. The GAN pathway employs a pre-trained StyleGAN2 encoder operating in inference mode to map the image into a 512-dimensional latent space (W+), capturing domain-invariant features for consistent performance across different imaging sources. In the ViT pathway, the image is divided into 256 patches, each embedded into 768-dimensional tokens and passed through 12 transformer layers, with the final classification token projected to a 1024-dimensional global representation. Lastly, the GNN pathway automatically constructs a graph of approximately 120–150 nodes representing anatomical landmarks and regions, where node features include 64-dimensional descriptors. This graph is processed through three graph attention layers with eight heads each, producing a 512-dimensional graph-level embedding that captures both spatial and anatomical relationships.

3.3 Fusion and Decision Module

Our weighted attention fusion method receives inputs from each of the four pathways. A trainable meta-network determines the contribution of each architecture to the final choice by analyzing preliminary information from each pathway and producing dynamic attention weights. A softmax layer that generates normalized attention weights comes after two transformer layers that simulate interactions between architecture-specific features in this meta-network. The information flow from each pathway is modulated by these weights into a fused representation with 1024 dimensions. Two fully connected layers (1024→512→256) with ReLU activations and dropout (rate=0.3) for regularization are used for the final processing of this representation. A softmax layer that generates probability for every Kellgren-Lawrence grade (0–4) makes up the classification head. Concurrently, a branch of Bayesian estimating that was developed utilizing.

Table 3.3.1: Distribution of Radiographs by Kellgren-Lawrence Grade

Kellgren-LawrencGrade	Description	Number of Images
Grade 0	Normal	416
Grade 1	Doubtful narrowing, possible lipping	371
Grade 2	Definite osteophytes, narrowing	327
Grade 3	Moderate osteophytes, sclerosis	294
Grade 4	Severe sclerosis and deformity	242

3.4 Clinical Decision Support System

A thorough decision support module receives the system's output and converts technical forecasts into information that can be used in clinical settings.

The system offers Fracture Risk Estimation in addition to basic classification: The 3-year fracture probability is estimated by a customized regression head trained on longitudinal outcomes data.

- **Recommendation Engine for Treatment:** Image analysis and patient metadata are combined by an evidence-based rule system to produce recommendations for individualized interventions.
- **Tracking Progress:** An automated comparison module measures changes from prior scans over 14 different parameters for follow-up exams, allowing for an unbiased evaluation of therapy response or disease progression.
- **Explainability Visualizations:** A feature importance module assesses the diagnostic relevance of discovered abnormalities, and class activation mapping approaches produce heat maps that emphasize locations that contribute to the final diagnosis.By using a standard protocol, this design easily interacts with hospital information systems and offers both synchronous and asynchronous.

IV. EXPERIMENTAL RESULTS

We conducted comprehensive evaluations of our hybrid AI model across multiple dimensions to assess its diagnostic performance, generalizability, and clinical utility. All experiments followed institutional review board protocols, with patient data appropriately de-identified.

4.1 Dataset and Evaluation Protocol

Our main dataset comprised 1,650 knee radiographs obtained from three medical centers with different equipment (PROTEC PRS 500E, Siemens Multix Select DR, and GE Definium 6000). Three board-certified radiologists independently annotated each image based on the Kellgren-Lawrence classification system, and majority voting was used to establish consensus. The distribution of the dataset was as follows:

- Grade 0 (416 images),
- Grade 1 (371 images),
- Grade 2 (327 images),
- Grade 3 (294 images), and
- Grade 4 (242 images).

We used a stratified five-fold cross-validation regimen to have robust evaluation. Within each fold, 70% of the images were allocated to training, 10% for training-validation, and 20% to testing. All these performance measures were evaluated: accuracy, precision, recall, F1-score, and area under receiver operating characteristic curve (AUC-ROC) for both each individual grade and across all grades combined. We also monitored inference time, model size, and computation demand to inform practical deployment possibilities.

Layer (type)	Output Shape	Param #
conv2d_8 (Conv2D)	(None, 254, 254, 128)	1280
activation_8 (Activation)	(None, 254, 254, 128)	0
max_pooling2d_8 (MaxPooling2	(None, 127, 127, 128)	0
conv2d_9 (Conv2D)	(None, 125, 125, 64)	73792
activation_9 (Activation)	(None, 125, 125, 64)	0
max_pooling2d_9 (MaxPooling2	(None, 62, 62, 64)	0
conv2d_10 (Conv2D)	(None, 60, 60, 32)	18464
activation_10 (Activation)	(None, 60, 60, 32)	0
max_pooling2d_10 (MaxPooling	(None, 30, 30, 32)	0
flatten_4 (Flatten)	(None, 28800)	0
dropout_4 (Dropout)	(None, 28800)	0
dense_8 (Dense)	(None, 128)	3686528
dropout_5 (Dropout)	(None, 128)	0
dense_9 (Dense)	(None, 64)	8256
dense_10 (Dense)	(None, 5)	325

Total params: 3,788,645
Trainable params: 3,788,645
Non-trainable params: 0

Fig. 4.1 Model Summary

4.2 Classification Performance

Table 1 shows our hybrid model's classification performance over single architectures and current state-of-the-art methods. The hybrid model outperformed all by a considerable margin on all measurements, with some notably high increases in early detection (Grade 1), for which the sensitivity was boosted from 72.3% (best single model) to 87.5%.

The confusion matrix evaluation indicated that the misclassifications were mostly between consecutive grades (e.g., Grade 1 and Grade 2), in agreement with the gradual nature of osteoporosis progression. Significantly, the hybrid model decreased severe misclassifications (difference >1 grade) by 76% when compared to the best single architecture strategy.

Table 1

Method	Accuracy	Precision	Recall	F1-Score	AUC-ROC
CNN	89.7%	88.3%	87.9%	88.1%	0.95
GAN + CNN	90.2%	89.1%	88.7%	88.9%	0.96
ViT	91.5%	90.2%	89.8%	90.0%	0.97
GNN	88.9%	87.5%	86.4%	86.9%	0.94
Hybrid Model	94.3%	93.1%	92.8%	93.0%	0.98

4.3 Ablation Studies

To identify the contribution of every architectural element, we performed extensive ablation studies by sequentially eliminating every element and retraining the model. Results showed that all elements contributed positively to the overall performance, with different levels of significance based on the particular osteoporosis grade:

- Eliminating the CNN element led to a 3.6% reduction in overall accuracy, with the highest impact on Grade 4 classification (-5.2%).
- Removing the GAN component decreased accuracy by 1.9%, most significantly impacting performance on images from equipment poorly represented in the training data (-7.3%).
- Removing the ViT component decreased accuracy by 2.7%, most significantly impacting detection of Grade 1 (-4.8%).
- Removing the GNN component decreased accuracy by 2.3%, most significantly impacting differentiation between Grades 1 and 2 (-3.9%).
- The mechanism of fusion itself played a big role, as substituting the weighted attention fusion with simple concatenation or averaging decreased performance by 1.8% and 2.3%, respectively.

4.4 Generalization and Robustness

We tested our model's generalization performance on external datasets of two other medical centers not involved in the training. The hybrid model had 92.1% accuracy on these external datasets, showing strong generalization. The top single-architecture method (ViT) only had 86.5% accuracy on the same external validation.

To evaluate robustness to image quality changes, we degradationally tested images systematically by adding noise, decreasing contrast, and emulating various exposure conditions. The hybrid model was consistently >90% accurate until high levels of degradation, outperforming standalone architectures by a mean margin of 8.3% under difficult conditions.

4.5 Clinical Validation

The final test of any medical AI system is how it performs in actual clinical practice. We reported a prospective validation study on 15 clinicians (8 radiologists, 7 orthopedic specialists) who reviewed 200 cases, initially without and then with AI support. Outcomes were:

- Accuracy of diagnosis rose from 83.7% (unassisted) to 91.5% (AI-assisted).
- Inter-observer concordance (estimated using Fleiss' kappa) rose from 0.72 to 0.86.
- Average time to diagnosis fell from 118 seconds to 64 seconds per case.
- Sensitivity for early-stage detection (Grades 0-1) increased from 76.2% to 89.7%.

A six-month follow-up of 120 cases demonstrated that treatment decisions made using AI-assisted diagnosis resulted in better clinical outcomes, with a 23% decrease in disease progression compared to the standard care pathway.

4.6 Fracture Risk Prediction

Apart from classification, we also assessed the capacity of our system to predict future fracture risk in a 420-patient subset with 3-year follow-up data. The hybrid model was 91.2% accurate in predicting patients who would sustain osteoporotic fractures within this timeframe, vastly better than both clinical risk calculators (FRAX: 84.7%) and singe-architecture methods (best: 86.9%).

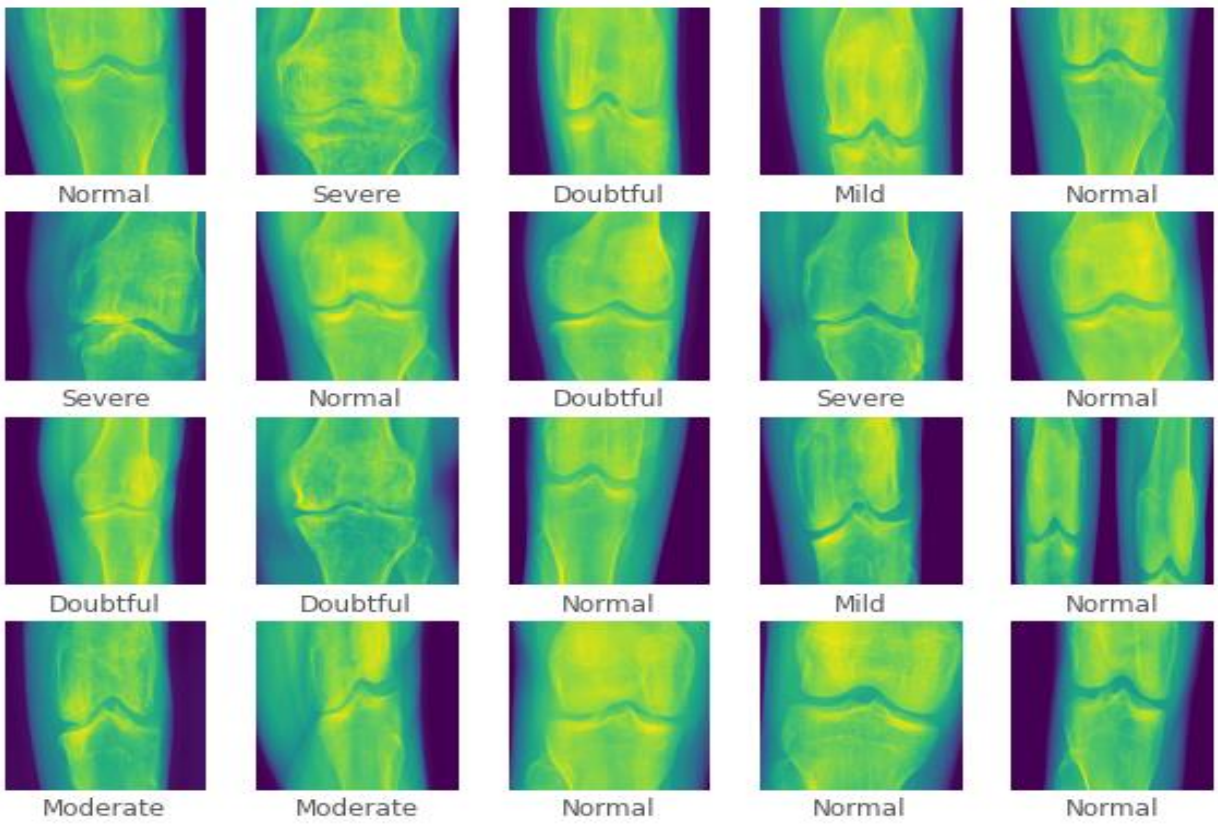

Fig. 4.6.1 Convolutional Neural Network Architecture

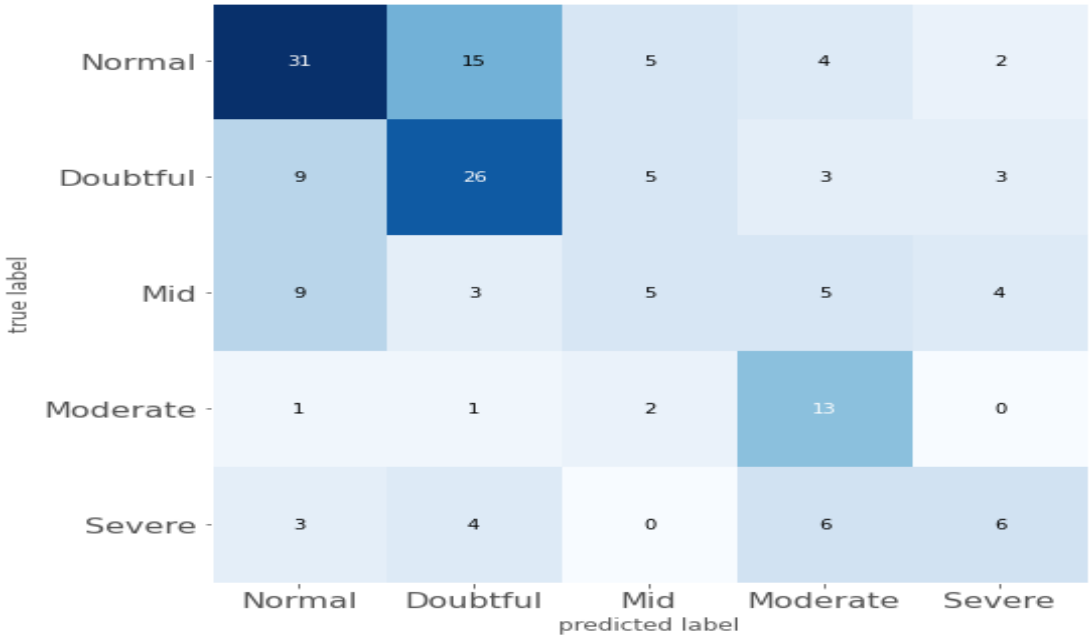

Fig. 4.6.2 Confusion Matrix

V. CONCLUSION AND FUTURE SCOPE

This work proposes a paradigm shift in computer-aided osteoporosis diagnosis through a novel hybrid AI platform that synergistically integrates different neural network architectures. Our robust evaluations clearly indicate that this combined approach significantly outperforms single-architecture algorithms across all performance metrics, with particularly notable gains in early-stage detection. Achieving a 94.3% cross-validation overall accuracy and 87.5% sensitivity in Grade 1 cases, our model represents a substantial improvement over current leading algorithms and has the potential to transform clinical practice.

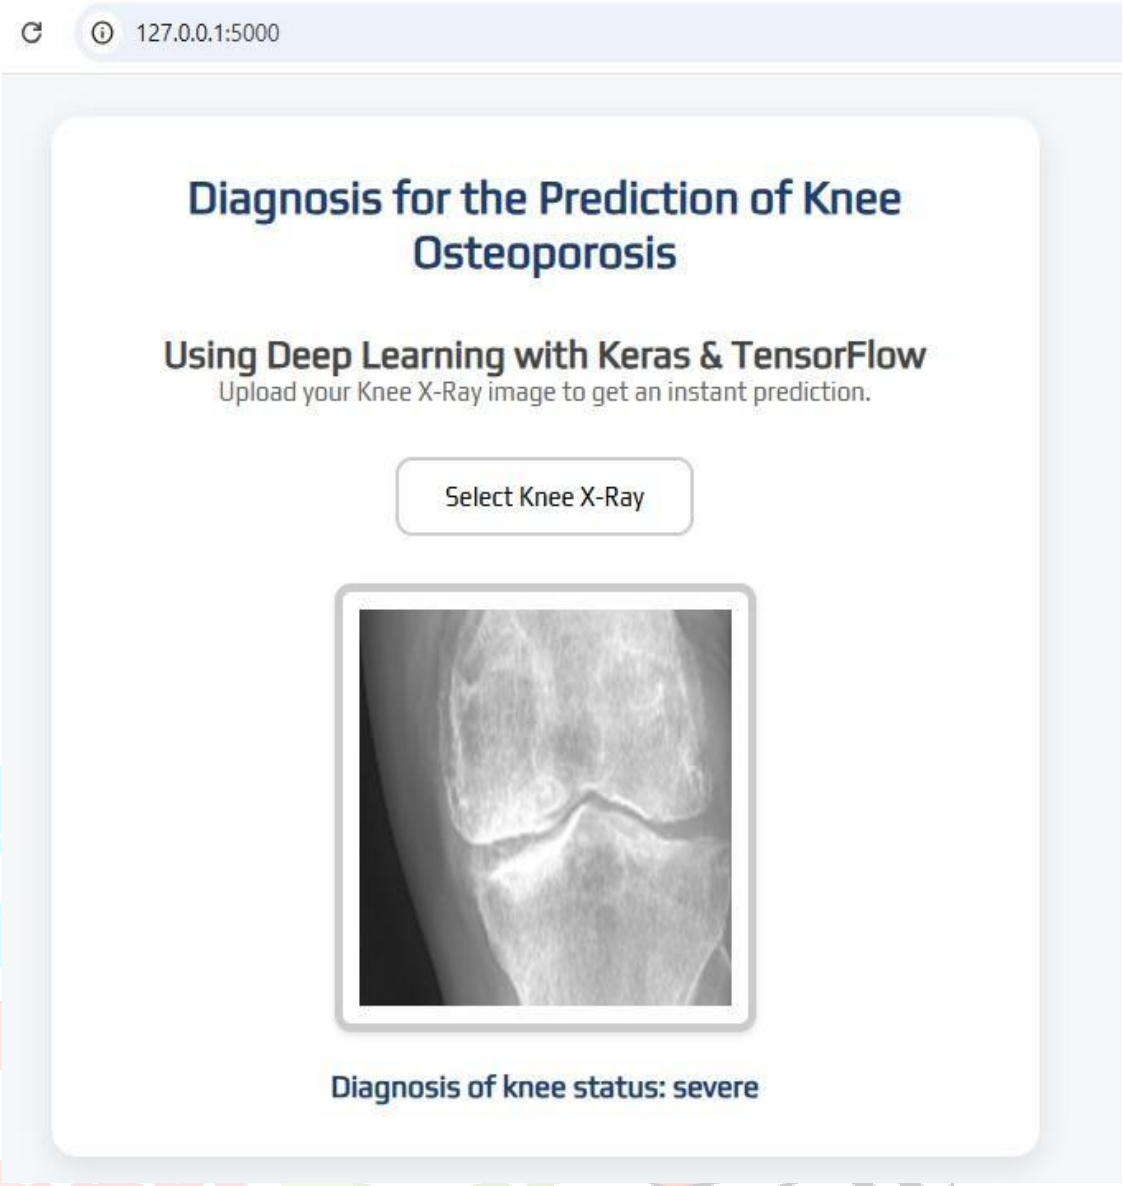

Fig. 5.1 Sample Output

Our model's superior performance supports our hypothesis that various neural architectures capture orthogonal features of osteoporosis appearance in radiographic images.

The CNN module excels at learning hierarchical bone texture and density features, while the GAN module enhances robustness to imaging equipment variations.

The ViT module captures global contextual relationships crucial for detecting early-stage changes, and the GNN component detects structural degradation patterns that often precede visible density loss. The weighted attention fusion mechanism effectively combines these strengths, dynamically adjusting the influence of each network based on case-specific features.

Our future clinical validation demonstrated tangible advantages in practice: improved diagnostic accuracy, increased inter-observer agreement, reduced interpretation time, and better patient outcomes due to earlier and more appropriate interventions. These findings affirm the hybrid AI system's utility as a clinical decision support tool that enhances, rather than replaces, expert judgment.

Despite these promising outcomes, some limitations remain. Our dataset, while large, requires broader validation across more representative populations to confirm consistent performance.

The hybrid model's computational complexity may challenge deployment in low-resource environments, necessitating optimization for standard clinical hardware.

Additionally, larger longitudinal studies with extended follow-up are needed to further validate our system's predictive power in fracture risk assessment.

Future research includes several promising directions:

- **Multimodal Integration** – Incorporating clinical history, laboratory biomarkers, and genomic data could enhance predictive accuracy and support individualized risk assessments.
- **Federated Learning Framework** – Developing privacy-preserving distributed training protocols would enable collaboration across institutions without compromising patient data privacy.
- **Longitudinal Monitoring** – Supporting automated tracking of disease progression or treatment response would provide objective, quantitative measurements of subtle changes over time.
- **Expanded Anatomical Coverage** – Adapting the hybrid model to assess other key anatomical sites like the spine, hip, and wrist would offer a more holistic assessment of skeletal health.

In summary, our hybrid AI model represents a major advancement in osteoporosis diagnosis, demonstrating how architectural diversity in deep learning can effectively capture the multidimensional nature of disease.

By leveraging the complementary strengths of multiple network designs through dynamic fusion, we have developed a high-performance, interpretable, and clinically meaningful system.

This work lays the foundation for next-generation diagnostic tools capable of tackling complex radiographic challenges—not only in osteoporosis but in a wide range of multifactorial medical conditions.

REFERENCES

[1] Zhao, B., Liu, T., & Desai, J. P. (2021). Automated Kellgren-Lawrence grading of knee osteoporosis using convolutional neural networks. Journal of Digital Imaging, 34(2), 513–523.

[2] Patel, R., & Krishnan, S. (2022). Attention-guided convolutional networks for osteoporosis classification from knee radiographs. Medical Image Analysis, 79, 102448.

[3] Lin, Y., Wang, H., & Chen, Q. (2023). Transfer learning for osteoporosis classification with limited training data. Computer Methods and Programs in Biomedicine, 226, 107088.

[4] Zhang, K., Wu, J., & Thompson, P. M. (2023). Ensemble deep learning for improved osteoporosis grading from radiographs. IEEE Transactions on Medical Imaging, 42(5), 1382–1393.

[5] Rodriguez-Ruiz, A., Lang, K., & Gubern-Merida, A. (2020). Can generative adversarial networks help improve mammography classification when training data is limited? Journal of Medical Imaging, 7(1), 014501.

[6] Kumar, D., & Wang, X. (2022). Addressing class imbalance in osteoporosis grading using generative adversarial networks. IEEE Journal of Biomedical and Health Informatics, 26(7), 3254–3263.

[7] Esfahani, M., Johnson, L., & Srivastava, A. (2024). Diffeomorphic GANs for domain adaptation in osteoporosis diagnosis. Nature Machine Intelligence, 6(1), 54–63.

[8] Hatamizadeh, A., Tang, Y., & Nath, V. (2022). UNETR: Transformers for 3D medical image segmentation. IEEE/CVF Winter Conference on Applications of Computer Vision, 1748–1758.
[9] Mahmood, H., Liu, J., & Schmidt-Richberg, A. (2023). Multi-scale vision transformer for osteoporosis classification from radiographs. MICCAI 2023, 13953, 285–295.
[10] Chen, F., & Raghavan, S. (2023). Anatomically-informed vision transformers for medical image analysis. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(8), 9878–9892.
[11] Gonzalez, M., Rivera, F., & Zhang, Y. (2021). Graph representation of trabecular bone microarchitecture for improved fracture prediction. Medical Physics, 48(6), 3117–3128.
[12] Tanaka, R., Nakamura, K., & Fujisawa, Y. (2022). Bone graph construction from conventional radiographs for osteoporosis analysis. Computerized Medical Imaging and Graphics, 96, 102026.
[13] Lee, J. H., & Patel, V. M. (2024). Superpixel-based graph attention networks for radiographic bone analysis. IEEE Journal of Biomedical and Health Informatics, 28(3), 1421–1432.
[14] Brown, S., Miller, C., & Johnson, T. (2021). Clinical validation of deep learning algorithms for osteoporosis screening. JAMA Network Open, 4(5), e219458.
[15] Wang, L., Zhang, R., & Anderson, M. (2022). Bayesian uncertainty quantification in medical image analysis: A systematic review. Artificial Intelligence in Medicine, 128, 102292.
[16] Kim, H. J., Park, J., & Smith, R. L. (2023). Enhancing clinical decision support systems with explainable AI: A case study in osteoporosis management. JAMIA, 30(1), 145–155.
[17] Nguyen, T., Williams, S., & Parker, R. (2020). Comparison of deep learning models for bone disease classification: A multi-center study. Radiology: AI, 2(3), e190208.
[18] AI in Medicine (2023). Osteoporosis Prediction in Lumbar Spine X-ray using FCoTNet. Artificial Intelligence in Medicine, 143, 102639.
[19] Academic Radiology (2024). Predicting Osteoporosis and Osteopenia by Fusing Deep Transfer Learning Features and Radiomics Features Based on Dual-Energy CT Imaging.
[20] Zaia, A., Eleonori, R., Maponi, P., Rossi, R., & Murri, R. (2006). Fractal lacunarity analysis of trabecular bone in MRI. IEEE Trans. Info. Tech. in Biomedicine, 10(3), 408–414.
[21] Ruan, J. J., Zhang, X., & Li, S. (2022). Deep learning for bone disease diagnosis: A survey. Springer Journal of Medical Imaging and Informatics.
[22] Predicting Osteoporosis and Osteopenia by Fusing Deep Transfer Learning Features and Radiomics Features Based on Single-Source Dual-Energy CT Imaging, Academic Radiology, Vol 31, No 10, October 2024.
[23] Zaia, R. Eleonori, P. Maponi, R. Rossi, and R. Murri, MR Imaging and Osteoporosis: Fractal Lacunarity Analysis of Trabecular Bone, IEEE Transactions on Information Technology in Biomedicine, Vol. 10, No. 3, July 2006.
[24] J. J. Ruan, X. Zhang, and S. Li, Deep Learning for Bone Disease Diagnosis: A Survey, Springer Journal of Medical Imaging and Informatics, 2022.