

Original Research

Received 28 June 2026

Revised 01 Aug 2026

Accepted 16 Aug 2026

Natural Resources for Human Health 2026, Vol. 6 (16s): 189–217

KEYWORDS:

Hip and Knee X-ray, DMI-Net, Explainable AI (XAI), SHAP, Grad-CAM, Implant loosening detection

An Explainable Deep Hybrid Architecture DMI Net for Knee and Hip Implant Status Classification in X-ray Imaging

Manjit Sandhu^{1*}, Navin Kumar², Ravinder Singh Sawhney³, Lakshay Sanan⁴

¹*Department of Electronics Technology, GNDU Amritsar, India

²Department of Mechanical Engineering, IIT Ropar, India

³Department of Electronics Technology, GNDU Amritsar, India

⁴Department of Orthopaedics, Senior Resident, GMC, Amritsar, India

manjit.ece@gndu.ac.in , nkumar@iitrpr.ac.in , sawhney.ece@gndu.ac.in , lakshaysanan0814@gmail.com

*Corresponding Author (Manjit Sandhu)

ABSTRACT: Objective: Complications in knee and hip joint replacements are implant-related current clinical concerns, and the delay or inaccuracy in diagnosis of these complications may result in serious patient outcomes and revision surgeries. To overcome these issues, this paper proposes a new type of hybrid Deep Multiscale Integration Network (DMI-Net) that can be used to accurately assess the condition of the implants using knee and hip X-ray datasets.

Method: The proposed structure includes a full preprocessing pipeline, image resizing, contrast enhancement using contrast-limited adaptive histogram equalization (CLAHE), normalization, and Gaussian filtering, to enhance visual quality and highlight important structural details. A lightweight parallel depthwise separable convolutional neural network (PD-CNN) is used to effectively extract multiscale features with minimal computational complexity. As an extension of discriminative capability, the architecture combines several high-performance transfer learning models through feature stacking and concatenation to form a strong hybrid system. To gain further insight and confirm feature separability, dimensionality reduction algorithms like PCA, LDA, UMAP, and t-SNE are employed. Also, SHAP and Grad-CAM approaches are used to enhance the interpretability of models, which allows visualizing of decision-relevant regions of the model and enhances clinical trust.

Results: The proposed DMI-Net has shown excellent performance when compared to some of the established models, with an accuracy of 98.92, precision of 98.70, recall of 98.60, F1 score of 98.65 and AUC of 98.80. These results point to the efficacy and validity of the suggested approach, indicating its high potential as a supportive tool to make accurate diagnosis of the implants and well-informed clinical decisions.

Conclusion: The present study is a contribution of a novel hybrid DMI-Net framework for detecting and classifying orthopaedic conditions using combined knee and hip implant X-ray images. Overall, the suggested hybrid DMI-Net is significantly more dependable, powerful, and understandable in computer-aided diagnosis of orthopaedic imaging, particularly the evaluation of knee and hip implants.

1. INTRODUCTION

The associated hip and knee joint disorders of orthopaedic implants have become a critical global healthcare issue, especially with the rising prevalence of degenerative diseases such as osteoarthritis, and the increase in the number of joint replacement surgeries being performed throughout the world. The success of the total hip and knee arthroplasty is well-known and recognized, and it helps to restore mobility and reduce pain; however, the post-surgical complications, including implant loosening, malalignment, infection and periprosthetic fractures have become a major clinical challenge and the main cause of implant failure and revision surgeries (Smith et al., 2022; Brown and Taylor, 2023). The loosening of implants, in particular, is frequent, and may often develop slowly with no noticeable initial symptoms, making it tricky to detect in good time and potentially resulting in severe functional impairment (Khan et al., 2021; Li et al., 2025). Moreover, the patient anatomy, bone density, surgeries, as well as the location of the implants, further complicate the diagnosis, which further illustrates the need to have accurate and reliable methods of assessment (Zhang et al., 2024; Mehta et al., 2026).

Radiographic imaging, especially X-ray imaging is the most common diagnostic tool that is used in the examination of hip and knee implants, due to its accessibility, affordability and effectiveness in the visualization of the bone-Implant interface. X-ray scans are frequently used by clinicians to assess the fixation of the implants, any signs of loosening, and to monitor the post-operative outcomes (Lee et al., 2023; Sharma et al., 2025). However, such images need to be interpreted by a person, which can occasionally be time-consuming and extremely dependent on clinical expertise, which could also be a contributor to inter-observer variability and diagnostic inconsistency (Patel et al., 2022). Especially hard to notice are those indicators of an early failure of the implant, such as minor radiolucent lines or microstructural bone changes (Gupta et al., 2024). These inadequacies demonstrate the desperate need of automated and intelligent diagnostic systems that is capable of producing constant, precise and early detection of implant-related complications.

The last years have witnessed a remarkable improvement in the detection and classification of the implant abnormalities following the integration of the Artificial Intelligence (AI) and deep learning (DL) techniques in the medical imaging process. Conventional neural networks (CNNs) have shown excellent performance in isolating complex and hierarchical features of X-ray images, which can be used to identify more instances of implant loosening, bone loss, and structural irregularities (Wang et al., 2023; Chen et al., 2025). Such AI-based methods do not only enhance the precision of the diagnosis, but also reduce the workload of clinicians, providing more reliable and quicker results. Moreover, recent developments in hybrid architectures, and multiscale feature fusion methods have improved the ability of models to capture local and global patterns, which in turn enhances generalization to diverse datasets (Garcia et al., 2024; Verma et al., 2026).

Learning to adjust representations acquired earlier to new tasks.

Transfer learning (TL) has become one of the most popular deep learning strategies to analyze medical images, especially in those cases when annotated data is scarce. TL allows effective feature extraction by leveraging the knowledge of pretrained models trained on large-scale datasets, at a lower cost of computation and requiring less extensive training data (Anderson et al., 2023; Singh et al., 2024). This method enables pretrained convolutional neural networks to be taught generalized representations and adapt them to domain specific tasks, including implant classification, and abnormality detection in hip and knee X-ray images. The further refinement of these models on special datasets further enhances the learning of discriminative implant-specific features, which leads to high classification accuracy, robustness, and generalization performance on a variety of datasets (Kaur et al., 2025).

More recent developments in deep learning have shown that TL-based models can be both highly accurate in detecting implant-related complications and also efficient and scalable. Architectures based on CNNs, together with the novel method of feature extraction and fusion, give solid representations of complex implant structures. Nevertheless, even in the light of these advancements, the fact that AI models would not be interpretable would continue to hinder their application in clinical practice. To overcome this issue, Explainable Artificial Intelligence (XAI) methods have been proposed to increase transparency in decision-making procedures. The SHAP and Grad-Cam methods make it possible to visualize the key areas that have an impact on model predictions and enhance the interpretability of the models, increase clinician trust, and promote the reliability of clinical decision-making (Kumar et al., 2024; Rao et al., 2026).

Innovative Features and Contributions.

Although there has been significant developments in deep learning (DL) to analyze the conditions of orthopaedic implants based on X-ray images, most of the existing hybrid or ensemble models are of limited design. They generally either use late-stage decision fusion, or simply concatenate features across multiple transfer learning backbones without directly modeling the multiscale nature of hip and knee implants, such as small scale differences at the bone-implant interface, and larger scale patterns in structural alignment. Moreover, much of the existing research lacks the implementation of feature-space visualization and model-agnostic interpretability in the same, cohesive framework, thus limiting both diagnostic transparency and clinical applicability.

To fill these gaps, this paper presents Deep Multiscale Integration Network (DMI-Net) an Explainable Artificial Intelligence (XAI)-driven model customized to the overall evaluation of the hip and knee implant conditions using a combination of X-ray datasets. The proposed architecture is set to concurrently capture (a) fine-grained local features such as micro-textural patterns and boundary irregularities, using a lightweight parallel depth wise separable convolutional neural network (PD-CNN), and (b) high-level contextual representations obtained via three complementary transfer learning models. Parallel to this, dimensionality reduction methods, that is, principal component analysis (PCA), linear discriminant analysis (LDA), uniform manifold approximation and projection (UMAP), and t-distributed stochastic neighbour embedding (t-SNE) are used only to visualise features in the feature space and evaluate qualitatively the separability of classes, thereby increasing interpretability but not affecting the classification process. The fused feature representation is then processed with a stacked classification mechanism, supported by an integrated explainability pipeline that combines Shapley Additive Explanations (SHAP) with Gradient-weighted Class Activation Mapping (Grad-CAM) and variants, to provide easy and clinically relevant explanations of model predictions.

This work has made significant contributions as follows:

New Framework Architecture Design: Design an innovative XAI-enabled hybrid DMI-Net architecture that can seamlessly integrate various transfer learning models with a lightweight PD-CNN to extract and learn multiscale features of both hip and knee X-ray images.

Advanced Feature Fusion Strategy: Introduction of a powerful feature-level fusion mechanism that fuses both local and global representations into a discriminative feature space, which results in enhanced characterization of the implant condition and increased classification accuracy.

Comprehensive Interpretability Integration: Implement dimensionality reduction methods (PCA, LDA, UMAP, t-SNE) to visualize features and SHAP and Grad-CAM-based models to provide understandable, trustworthy, and interpretable explanations of model decisions.

Strong Performance Validation: Extensive comparative analysis and statistical test with several baseline transfer learning models, and superior performance and generalization capability of the proposed DMI-Net framework using the key evaluation metrics.

The assessment of hip and knee implants in clinical practice, where variation between clinicians and a growing number of diagnostic processes can also influence accuracy and efficiency. The need to detect implant-related complications at an early stage and to address loosening and structural degradation of the implants with improved patient outcomes is inherent. The proposed structure is not intended to replace clinical expertise but will be a smart assistive aid, which will enhance uniformity in diagnosis, point out clinically significant areas, and help make confident decisions in a complex or ambiguous case through interpretable and reliable analysis performed by AI.

2. RESEARCH LANDSCAPE AND PREVIOUS ART

Recent developments in computer-aided diagnosis (CAD) systems have significantly improved the automated evaluation of the status of hip and knee implants using X-ray imaging to meet the growing clinical need amid the increasing number of joint replacement surgeries. In orthopaedic diagnostics, precise detection of implant-related complications, such as loosening, malalignment, structural degradation, etc. is also an important challenge. Here, deep learning (DL), specifically convolutional neural networks (CNNs), have proven to be remarkably able in extracting complex radiographic features and providing superior diagnostic performance than conventional methods. As a result, researchers have investigated various datasets, model

architectures, and feature-learning approaches to improve classification accuracy, robustness, and clinical reliability in implant analysis.

Tiwari et al. (2023) utilised a hip arthroplasty radiographic dataset of 3,206 X-ray images and applied transfer learning on DenseNet121 and ResNet50 network architectures. In their study, they found that the exact accuracy of classification of the DenseNet121 network in the task of identifying the implant is 98.23%, which also demonstrates the high effectiveness of deep feature reuse in the task of identifying the implant. Based on this, Chee et al. (2022) examined the classification of knee implants using a dataset of 521 X-ray images, and test a variety of transfer learning models such as VGG16, ResNet50, and DenseNet121. DenseNet121, among them, has proven to be the most successful with a precision of 96.70 exact this time with the implant-specific pattern. All these studies collectively determine the effectiveness of transfer learning in orthopaedic imaging in situations where annotated datasets are not available.

In line with these developments, Yu et al. (2024) suggested a deep learning-based system of total knee arthroplasty using a dataset of 1,412 X-ray images in anteroposterior and lateral views. This study used a ResNet101 architecture to predict implant size, with an exact F1-score of 0.91 with respect to femur components and 0.87 with respect to tibia components, with the prediction accuracy within a range of ± 1 size of 99.00% and 98.00% respectively. The paper emphasizes the relevance of DL to not only the classification but also the planning of surgical procedures. In a similar fashion, Kinger (2024) used the Osteoarthritis Initiative (OAI) dataset that comprised 4,446 knee X-ray images and used InceptionV3 and DenseNet models to assess the condition of the joint. InceptionV3 model was found to have an exact accuracy of 91.00% and this showed its ability to capture structural variations that can be used in making decisions about implants.

Additional developments have been aimed at enhancing the generalization and scalability of models to larger data sets. Zhang et al. (2025) performed a thorough assessment with the use of multiple knee radiographic datasets (more than 10,000 images) and applied complex CNN architectures, such as EfficientNet-B4. The model demonstrated a perfect classification accuracy of 97.82% which demonstrates the advantages of optimized scaling and feature extraction. Parallel to Li et al. (2025) developed an implant loosening detector based on an implant loosening detection framework that is a hybrid CNN-based framework, using a dataset of 2,150 hip and knee X-ray images to achieve an exact accuracy of 97.65%. These results highlight the increasing significance of hybrid implant-like feature-learning methods in treating complex implant-related abnormalities.

More recent research has focused on cross-institutional testing and practical applicability. Yasin et al. (2025) published a dual-center study based on a composite dataset of 5,200 knee X-ray images, using a multi-label classification framework built on a ResNet framework. The proposed model was found to have the exact accuracy of 98.10 which indicates that the model has a high generalization in the heterogeneous clinical environments. Likewise, Verma et al. (2026) proposed a hybrid deep learning architecture that combines multiscale feature extraction and attention mechanisms, which are trained on a dataset of 6,800 combined hip and knee X-ray images. The model has an exact classification accuracy of 98.45, surpassing traditional CNN and transfer learning methods and demonstrating the importance of multiscale representation learning.

All of these studies demonstrate significant advancements in using deep learning and transfer learning methods to the analysis of hip and knee implants. Nevertheless, most current methods mainly use single architectures or traditional ensemble designs, without explicitly considering the multiscale implant properties. Besides, the convergence of feature-space visualization and explainable artificial intelligence (XAI) tools has been limited, thus limiting the model interpretability and clinical trust.

The current study in response to these limitations, proposes a hybrid Deep Multiscale Integration Network (DMI-Net) that can be adapted to analyze hip and knee implants using a combined X-ray dataset. The framework proposed combines a lightweight parallel depthwise separable CNN to effectively extract features across multiple scales using multiple transfer learning models to obtain complementary representations. In addition, dimensionality reduction algorithms like PCA, LDA, UMAP, and t-SNE are used to enable visualization of features in the feature space and evaluate the separability of classes. To improve the level of transparency and clinical applicability, explainability techniques such as SHAP and Grad-CAM are included to point at diagnostically significant areas. The proposed methodology is systematically assessed with the help of standard performance parameters and benchmarked against the state-of-the-art approaches, which is aimed at providing a robust, interpretable, and clinically reliable solution to the problem of implant analysis.

3. PROPOSED METHODOLOGY

3.1 Architectural Blueprint of the Study

This part gives a detailed account of presenting the proposed framework for detecting implants and loosening of the implants in hip and knee radiographs. It not only describes the integrated dataset used in this study, but also systematically describes the end-to-end processing pipeline, including image resizing and preprocessing, multiscale feature extraction, and feature selection. Each of the constituent stages of the framework is well expounded in the subsequent subsections. Fig. 1 shows the proposed DMI-Net network, which is conceptualized as a systematic and end-to-end network to identify and categorize the hip and knee implant conditions using a composite X-ray imaging dataset. The whole methodological pipeline is carefully designed into the following consecutive steps:

Image Refining and Standardisation:

Any radiographic input is to be passed through a detailed preprocessing strategy that is expected to deliver to visual fidelity and offer uniformity across the dataset. The step includes dimensional standardization, region-specific masking where necessary, contrast optimization using contrast-limited adaptive histogram equalization (CLAHE), intensity normalization to provide distributional consistency and Gaussian smoothing to reduce noise, without affecting diagnostically significant bone implant interface characteristics.

Dataset Segregation:

The systematic division of the processed data then allows the effective training of the model and objective judgement. Particularly, 70 percent of the samples will be used as training data, 10 percent will be used as validation data, and the remaining 20 percent will be used as final performance evaluation data.

Multiscale Feature Representation Hierarchical:

The process of feature extraction is implemented through a dual-stream approach that would help to extract local and global contextual features. A lightweight parallel depthwise separable convolutional neural network (PD-CNN) is first used to learn complex structural features, such as boundaries of implants, patterns of trabecularity, and subtle radiographic changes. At the same time, several ready-made transfer learning models are used to generate high-level semantic representations and, therefore, provide a comprehensive understanding of implant conditions using heterogeneous X-ray hip and knee images.

Feature Normalization and Exploration of the Latent Space:

The obtained 220-dimensional representation is statistically normalized to guarantee the consistency of scales and the numerical stability before classification. Such a refined representation is the input to the predictive module. Linear projection methods are used to get reduced-dimensional views to investigate the structural distribution of learned representations. Moreover, nonlinear manifold learning methods, such as Uniform Manifold Approximation and Projection (UMAP) and t-distributed Stochastic Neighbor Embedding (t-SNE) are applied to show intrinsic clustering patterns and inter-class relationships in hip and knee implant radiographic data.

Architecture Formulation:

The discriminative representations that are obtained using transfer learning architectures are smoothly incorporated with those that are obtained by the PD-CNN using a feature-level fusion mechanism. This combined representation is then spread through fully connected layers, and an ultimate decision layer (using the softmax algorithm) to make a multi-class categorization of implant conditions in the combined hip and knee X-ray dataset.

Parameter Optimization and Performance Assessment:

The systematic sweeping of hyperparameters by using the validation subset to enhance the generalization property and minimize overfitting optimize the model structure. The performance of the proposed DMI-Net framework is quantitatively measured in terms of accuracy, precision, recall, F1-measure, macro-averaged indices and the area under the receiver operating characteristic curve (AUC-ROC), which guarantees a sound and comprehensive performance evaluation.

Interpretability and Decision Transparency:

The framework is supplemented with advanced interpretability mechanisms that allow meaningful interpretation and help to support clinical relevance. The methods used to identify the salient areas affecting model predictions include Techniques such as Shapley Additive Explanations (SHAP) and Gradient-weighted Class Activation Mapping (Grad-CAM), and their improved variants. This integration helps to better comprehend the decision-making process and become more confident in the results of automated diagnoses.

3.2 Data Description

The data that will be used in this study will be a unified set of hip and knee radiography images that will be found in Kaggle, an open-source repository of data that is widely recognized. This selected dataset is composed of a total of 6,413 X-ray images in JPG format with a variability in terms of spatial resolution and representation of the anatomy. The images are categorized under four different classes namely, normal, implant, implant-loose, and mask, respectively, which corresponds to the specific clinical condition related to the orthopaedic assessment.




Fig 1: Proposed framework for the detection and classification of Implant

In Fig 2, representative samples of each category are provided, and the picture shows an overall picture of the diversity in the dataset. Moreover, Table 1 and Fig 3 give a comprehensive overview of the dataset composition, the distribution of images per class of the dataset, its respective labels, and frequency per class. Such organized data can be used as a potent base to train and evaluate the future proposed framework in the correct differentiation of normal and implant-related diseases in combined hip and knee radiographic analysis.




Fig 2: Presents representative samples from each category

Table 1: Detailed Sample-wise Distribution of X-ray Dataset

Dataset	Total Images	Image Type	Implant	Implant loose	Mask	Normal
Class			1	2	3	4
X-ray Dataset	6413	X-ray Images	1850	1850	1777	936
Label			0	1	2	3




Fig 3: Dataset Distribution Class-wise

3.3 Image Preprocessing

Image preprocessing is one of the vital elements in improving the visual quality of radiographic images and improving the predictivity of the proposed automated framework. It helps in refining image characteristics, reducing undesired variations, and

providing the availability of standardized inputs to allow effective model learning. A series of preprocessing steps were indexically applied to ensure that the transfer learning models are trained, validated, and evaluated using a consistent and high-quality dataset. These are image resizing to achieve dimensional uniformity, contrast enhancement by use of contrast-limited adaptive histogram equalization (CLAHE), normalization of image intensity to achieve distributional consistency, and Gaussian-based smoothing to suppress noise. Collectively, these procedures improve the visibility of clinically significant bone–implant structures and support robust feature extraction in hip and knee X-ray analysis. The mathematical formulations that represent each of these operations are given in Equations (1)–(4) and the explanation of each of these preprocessing stages is given in the following subsections.

$$I(x, y) = \sum_{i=0}^1 \sum_{j=0}^1 w_{ij} \cdot I(x_i, y_j) \quad (1)$$

The interpolation techniques are used to resize the images in order to have a uniform spatial resolution and structural continuity in the radiographic data. The process of this kind of transformation is the smooth geometric transformation of the image to the necessary size without severe loss of anatomical detail. Where $I(x, y)$ is the original intensity of a given pixel location (x, y) of the input image. The interpolated intensity value is calculated by taking into account the surrounding neighborhood defined by the coordinate to adjacent pixels (x_i, y_j) in which the contribution w_{ij} of each neighboring pixel is controlled by corresponding weighting coefficients, where their relative spatial proximity to the target location determines their weighting coefficient. This weighted interpolation process allows a smooth transition of intensity and preservation of the integrity of bone-implant structures in hip and knee X-ray images.

$$I_{eq}(x, y) = \frac{(L-1)}{MN} \sum_{k=0}^{I(x,y)} h_k \quad (2)$$

Contrast-limited adaptive histogram equalization (CLAHE) is used to enhance contrasts, and in so doing, improve the visibility of fine anatomical tissue in radiographic images. In this expression, $I_{eq}(x, y)$ refers to the transformed intensity at pixel position (x, y) after equalization. The frequency of pixels at a particular gray level k in a local contextual area over which the enhancement is undertaken, is denoted by h_k . MN are used to denote dimensions of the local contextual area over which the enhancement is to be applied. Moreover, L is equal to the sum of the discrete levels of intensity in the image. This technique can be used to effectively enhance contrast, whilst avoiding over-amplification, and in this way critical details of the bone-implant interface can be seen to have been preserved in the hip and knee X-ray images.

$$X_{norm} = \frac{X - X_{min}}{X_{max} - X_{min}} \quad (3)$$

Intensity normalization (also known as intensity normalization) is a technique used to rescale pixel values into a predetermined range, often $(0, 1)$ or $(-1, 1)$, so that there is numerical consistency and that model convergence becomes stable. Here, X will be the original pixel intensity, X_{min} and X_{max} will be the minimum and maximum pixel intensity values in the image, respectively. The normalized output, X_{norm} , is calculated by transforming the original intensities relative to this range, which leads to a standardized distribution which helps reduce the effects of illumination variation and improves the efficiency of feature learning in hip and knee radiography analysis.

$$G(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (4)$$

3.3.1 Image resizing

The pre-processed dataset is uniformly scaled to a fixed size of $224 \times 224 \times 3$ and all radiographic images are assembled to create a uniform dataset at a fixed size of $224 \times 224 \times 3$. This standardization also helps in computational efficiency and preservation of the integrity of key bone-implant interface characteristics necessary to achieve accurate classification. Besides this, to support effective supervised learning, the categorical labels were systematically coded into numerical identifiers 0, 1, 2, and 3 which represent the classes implant, implant loose, mask and normal, respectively, in the proposed system.

3.3.2 Feature Extraction

One of the major goals of this work is the development of a convolutional neural architecture with the ability to produce highly discriminative representations within a modest computational footprint and a limited structural depth. The convolutional network complexity is a paramount issue as too many layers and parameters can greatly increase the computation cost and can even make its practical implementation challenging. To this end, a lightweight parallel depthwise separable convolutional neural

network (PD-CNN) architecture has been designed to provide a good balance between the representational capacity and efficiency, as shown in Fig 4. To combine convolutional operations with fully connected layers, the proposed design will be used to guarantee effective feature abstraction without compromising computational feasibility. The architecture uses several parallel depthwise separable convolutional pathways instead of a series of traditional convolutional layers, which significantly decreases the redundancy of parameters and network complexity. This design paradigm substitutes the conventional stacked design that generally leads to deeper and more resource intensive designs. The architectural structure was rationalized by empirical testing to attain the best performance. In particular, the network combines four depthwise separable convolutional blocks with filter sizes of 128, 64, 32, and 16 to hierarchically extract spatial features at different scales, and to effectively capture multiscale structural patterns. The resulting feature maps are then flattened and concatenated into one unified feature, which is then processed through dense layers with dropout regularization and batch normalization, and followed by a softmax activation mechanism. This pipeline provides a small but expressive feature vector that consists of 220 deep features. The proposed architecture is effective in capturing local and global contextual properties and maintaining low computational complexity to achieve enhanced classification performance and enable a smooth integration with explainable artificial intelligence systems to foster a clinically interpretable deployment.




Fig 4: Architecture of the proposed lightweight PD-CNN network for multiscale feature extraction

3.4 Model Architecture

This section outlines the overall description of the transfer learning (TL) architectures incorporated in the proposed structure, and the development of the hybrid DMI-Net model enhanced with a parallel depthwise separable convolutional neural network (PD-CNN). The architectural design is carefully optimized to allow accurate detection and classification of hip and knee implant conditions based on radiographic images, and with computational efficiency and reduced model complexity. Also, this section describes the computational environment, both in terms of hardware and software settings, as well as the criteria of evaluation used to strictly test the performance of models.

3.4.1 Transfer Learning Models

Transfer learning paradigms in recent years have become one of the most effective methods in the medical image analysis that have shown considerable potential in solving a complex diagnostic problem. Radiographic imaging modalities generate large amounts of high-dimensional visual data the interpretation of which requires complex computational procedures and knowledge of the domain. Transfer learning models alleviate these issues by utilizing knowledge gained on large scale datasets such as ImageNet, thus making it easier to adapt to domain specific tasks and increase the efficiency of learning. In this paper, various transfer learning models are included to enhance feature presentation and augment the classification rates of hip and knee X-ray images.

DenseNet121

DenseNet121 is a deep convolutional neural network architecture consisting of 121 layers which was developed based on the DenseNet framework. The key feature of it is the high density of the connectivity pattern, wherein every layer is connected to all further layers in a feedforward fashion. This architectural approach is able to propagate features significantly, support the efficient flow of gradient which is a common problem in deep networks, and relieves the difficulties of vanishing gradients commonly experienced in deep networks. In addition, feature map reuse between layers can lead to the overall reduction of the

number of parameters as well as enhance representational efficiency. As a result, DenseNet121 is a very effective extractor of rich and hierarchical features of the hip and knee radiographic phenomena.

DenseNet201

DenseNet201 is a state-of-the-art member of the DenseNet family, characterized by its dense architecture (consisting of 201 layers) and its efficient feature propagation mechanism. The network itself is a hierarchical, progressive and ordered process of learning features, where the network is systematically structured into dense block, transition layers, and finally a classification layer. One notable feature of this architecture is its dense connectivity pattern, where each layer is connected to all the other layers of the same dimensionality. Such a design has a great impact on information flow, transmission of gradients across the network is stable, and it effectively eliminates the vanishing gradient problem that is often seen with deep convolutional models. Moreover, the large-scale reuse of feature representations, reduces the redundancy of parameters and enhances computational efficiency, hence allowing the model to extract highly discriminative and context-rich features out of hip and knee radiographic images.

MobileNet

MobileNet is a very efficient deep neural network that is designed to be deployed to computationally limited systems such as mobile and embedded systems. It is designed to be based on depthwise separable convolution, which breaks down conventional convolution into two sequential operations: depthwise convolution to perform channel-wise spatial filtering and pointwise 1×1 convolution to perform feature integration. This decomposition greatly decreases computational cost and number of parameters without a notable loss in representational capability. Due to its compact design, MobileNet can be effectively used to extract meaningful patterns in hip and knee radiographic images with a low computational cost.

MobileNetV2

MobileNetV2 adds architectural improvements to its predecessor by adding inverted residual blocks and linear bottlenecks. The network is implemented as an initial convolutional layer, then a sequence of bottleneck structures that increase feature dimensions by applying pointwise convolution, depthwise convolution to provide efficient spatial filtering and project features back to a lower-dimensional space. The design improves the representation of features without sacrificing efficiency. The residual connections also help to ensure stable gradient propagation and MobileNetV2 is highly effective in solving tasks related to the analysis of medical images in cases involving hip and knee X-ray data.

InceptionV3

InceptionV3 is a state-of-the-art convolutional architecture, which incorporates a variety of design patterns, including factorized convolutions, parallel processing structures and inception modules extracting multi-scale features. The inception modules enable the simultaneous processing of features at the various receptive fields, in effect, capturing the fine and coarse image details. This architecture also depends on the mechanisms of activation optimization as well as the dimensionality reduction algorithm so as to be able to improve the efficiency of computation and convergence behavior. This feature of it in modeling complex spatial and channel-wise relations makes it especially suitable in the detailed analysis of radiographic images.

InceptionResNetV3

InceptionResNetV3 is a hybrid network, which combines the capacity of multi-scale feature extraction Inception modules with the stability of residual learning. The remaining connections are incorporated so that it allows the efficient gradient flow along the deep connections leading to the rapid convergence and elimination of the degradation problem. This combined framework enables the model to learn highly discriminative and hierarchical features and as such, is particularly useful in complex classification tasks involving medical imaging including the evaluation of the conditions of hip and knee implants.

VGG16

VGG16 is a highly evolved architecture of a convolutional neural network that has 16 weighted layers, 13 convolutional layers and 3 fully connected layers. The network has identical 3×3 times 3×3 convolutional filters in all parts of the network, allowing the generation of more hierarchies in the features used, without making the computation infeasible. The max-pooling layers are added periodically to decrease the space characteristics and enhance efficiency. Its simple but strong design allows strong feature extraction, therefore it can be extensively used in image classification tasks, such as in radiographic analysis.

VGG19

VGG19 is an extension of VGG architecture with a depth of 19 layers, which includes 16 convolutional layers, and 3 fully connected layers. Like VGG16, it uses a consistent series of $3 \times 3 \times 3$ convolutional filters to allow deeper hierarchical learning and thus better feature representation. Max-pooling is also considered to dimensionality control, and a dropout technique is applied to reduce overfitting and achieve better generalization. The additional layers of VGG19 enable more sophisticated extraction of features, which makes it a trusted model to solve complex image analysis tasks, including the classification of hip and knee implants.

3.4.1 Proposed DMI-Net Model

To perform a systematic evaluation of the performance of the chosen transfer learning architectures, each model was privately implemented and assessed using standard evaluation metrics at the training, validation and testing stages. In the comparative analysis, DenseNet121, MobileNet and InceptionV3 all showed to have better predictive performance as well as better computational efficiency, with lower resource utilization in execution. Therefore, the best performing models were chosen and combined using a stacking and feature-level concatenation scheme to formulate the proposed hybrid DMI-Net architecture. The main goal of this integration is to leverage the complementary advantages of different networks, and thus enrich feature representation and increase classification accuracy. The resulting framework is specifically aimed at the accurate detection and classification of hip and knee implant conditions based on combined X-ray datasets, with robust diagnostic performance and computational efficiency.

3.4.3 Mathematical Description

In order to define the hybrid integration mechanism to be used in accordance with the proposed framework in a formal manner, this subsection presents a mathematical formulation of the feature-level fusion strategy. The main aim is to explain how multiscale representations based on the PD-CNN framework and the choice of transfer learning backbones can be coherently integrated into a single and discriminative representation of feature prior to the classification phase.

Where x refers to an input hip or knee X-ray image. These deep feature representations obtained by DenseNet121, MobileNet, and InceptionV3 are represented as $f_D(x)$, $f_M(x)$, and $f_I(x)$, respectively. Additionally, $f_P(x)$ are the multiscale feature embedding of the radiographic data that is obtained by the PD-CNN module.

$$z(x) = \text{concat} (f_D(x), f_M(x), f_I(x), f_P(x))$$

The integrated feature representation is subsequently propagated through a fully connected classification module for final prediction.

$$\mathcal{Y} = \text{softmax} (g(z(x); \theta))$$

Here, $g(\cdot)$ signifies the parametric fully connected transformations, while θ encapsulates the corresponding learnable weights; the optimization process is conducted via the minimization of the cross-entropy objective function, facilitated through gradient-based backpropagation.

3.4.4 DMI-Net Model Architecture

The proposed hybrid DMI-Net architecture, synergistically integrated with the PD-CNN module, is illustrated in Fig 5 and is meticulously designed to exploit multiscale feature learning, efficient feature integration, and high-precision classification of hip and knee implant diseases based on combined X-ray data. The processing part of the system is the pre-processed radiographic inputs which are processed by the PD-CNN component that consists of a dual-branch structure: a functional branch with parallel depthwise separable convolutional operations and a sequential convolutional pathway. The resulting feature representations are then propagated through fully connected layers and dropout regularization to obtain a small and standardized 220-dimensional feature vector.




Fig 5: Proposed hybrid DMI-Net architecture integrated with PD-CNN for implant prediction

Meanwhile, the three highest-performing transfer learning architectures, namely DenseNet121, MobileNet, and InceptionV3, extract deep hierarchical representations independently, on the same X-ray input. Each network contains a sequence of convolutional operations, pooling mechanisms, flattening layers, and a sequence of fully connected layers (512-1280-2048) to generate feature embeddings of 512, 1280 and 2048 dimensions, respectively. The feature representations that are obtained using the PD-CNN and the transfer learning models are then combined together through feature-level concatenation to form an overall representation that captures both low-level structural attributes and high-level semantic information. This combined representation is then propagated through dense layers and finally a softmax classification layer, which classifies the input into four categories: implant, implant loose, mask and normal.

Additionally, independent applications of the learned feature space to auxiliary interpretability modules, such as Principal Component Analysis (PCA), Linear Discriminant Analysis (LDA), Uniform Manifold Approximation and Projection (UMAP), and t-distributed Stochastic Neighbour Embedding (t-SNE), are each applied to the learned feature space to aid in visualisation and qualitative evaluation of the separability of classes.

The success of the proposed hybrid DMI-Net model is further reinforced in Table 2 where the integration of a custom-crafted PD-CNN with pre-existing transfer learning architectures make it possible to extract both fine-grained texture patterns as well as high-level contextual representations in hip and knee radiographic images. The multiscale fusion approach greatly increases the model ability to effectively distinguish among different implant states, and enhances the overall diagnostic performance.

3.4.5 Hyperparameters for DMI-Net Model

The best set of hyperparameters of the proposed hybrid DMI-Net framework can be summarized in Table 3. The parameters have been obtained very carefully with an extensive experimentation exploration with multiple combinations that lead to a balance between effective convergence and generalizability. The systematization of determining the most suitable parameters setting was through the two-stage optimization plan. The search process was a randomized search over a wide parameter space, including learning rates in the range of 10^{-5} to 10^{-2} , dropout probabilities between 0.2 and 0.6, and different sizes of the batches between 16 and 64 and various configurations of trainable layers within the transfer learning backbones. The approach was useful to determine the fast promising regions within the search space.

Later on, a fine-tuning grid-based search was conducted in the area of the most effective configurations found in the initial phase, allowing a closer fine-tuning of hyperparameters. To further ensure the model generalization and minimize the overfitting, a mechanism of early stopping that was based on the monitoring of the validation loss was regularly used in all

experimental runs. The hybrid optimization method used (that is, random exploration with targeted grid refinement) was an optimal trade off between computational efficiency and robustness to the current application despite alternative optimization strategies also being considered (such as Bayesian optimization). The sTable convergence behavior can be supported by the final parameter configuration, which is indicated in Table 3 and the associated training and validation loss-accuracy curves provided in the results section. The ability to carefully calibrate these parameters is crucial in increasing the ability of the model to extract meaningful representations of the hip and knee radiographic images, thus increasing the reliability of classification and diagnostic accuracy.

Table 2: Significance of the Hybrid DMI-Net Architecture

Aspect	Significance
Multi-Model Integration	The hybrid DMI-Net framework is strategically designed to combine the complementary representational advantages of DenseNet121, MobileNet and InceptionV3, to enhance the ability to generalize, and better the classification strength across a variety of hip and knee radiographic appearances.
PD-CNN Augmentation	The inclusion of the PD-CNN module greatly enhances the extraction of spatial features by making use of depthwise separable convolutions and multi-kernel operations, helping to effectively capture fine-grained structural and contextual variations in X-ray images.
Feature Space Optimization	Linear dimensionality reduction techniques such as PCA and LDA are employed to project the feature space into discriminative subspaces, while nonlinear embedding methods including UMAP and t-SNE are utilized for advanced visualization and qualitative analysis of class separability rather than direct classification.
Regularization and Standardization	The integration of dropout mechanisms alongside feature normalization ensures improved generalization by mitigating overfitting and maintaining consistent feature scaling throughout the training process.
Explainability and Transparency	The architecture is inherently compatible with explainable artificial intelligence methodologies, incorporating SHAP and Grad-CAM to provide interpreTable insights into model predictions and to highlight diagnostically relevant regions within radiographic data.
Clinical Relevance	The proposed hybrid DMI-Net framework offers a computationally efficient, scalable, and interpreTable solution for the accurate detection and classification of hip and knee implant conditions, thereby supporting reliable clinical decision-making in orthopaedic diagnostics.

In the proposed DMI-Net, once the pretrained knowledge was utilized, a fine-tuning approach was adopted in stages to effectively leverage pretrained knowledge and adapt to the target domain. The convolutional layers of each of the pretrained backbones were first frozen and only the classification head was trained by a brief warm-up period of five epochs. This was followed by selective unfreezing to allow the deeper adaptation where the higher-level convolutional blocks would be made trainable, but keep the lower-level feature extractors. In particular, in DenseNet121 the last two dense blocks (which are approximately 50 layers) were fine-tuned; in MobileNet the final eight depthwise separable convolutional layers were unfrozen; and in InceptionV3 the last three inception modules were optimized. This incremental fine-tuning approach enables learning of domain-specific features, with stability and avoiding overfitting. Accordingly, the total number of convolutional layers, summing across the three backbone networks, was almost 50, which were selectively fine-tuned, as outlined in Table 3. This carefully tuned approach allows adapting high-level semantic representations to the domain of hip and knee radiographic imaging, and at the same time, it maintains the strength of low-level features abstractions inherited through large-scale pretraining on ImageNet.

Table 3: Optimal Hyperparameter Configuration of the DMI-Net Model

Hyperparameter	Value
Optimizer	Adam
Learning Rate	0.001
Activation Function	Softmax
Loss Function	Categorical Cross-Entropy

Dropout Probability	0.3
Batch Size	32
Number of Epochs	100
Early Stopping Criterion	Validation Loss
Weight Initialization	Normal
Trainable Transfer Learning Layers	50
Regularization	L2 (0.0005), Batch Normalization

3.4.3 Requirements of Transfer Learning Models

Table 4 below details in detail the architectural specifications and computational needs of the selected transfer learning models, as well as the proposed hybrid DMI-Net framework. The parameters play a crucial role in establishing the practical viability, scalability and computational efficiency of each model in the context of medical image analysis. The comparison describes inherent trade-offs between the complexity of the model, the time to train, and the memory used and predictive accuracy. It is noTable that more complex architectures such as VGG16 and VGG19 have much more parameters and thus require much less time to train and much more memory to fit. On the other hand, the proposed DMI-Net framework is still more effective in terms of classification performance, despite its hybrid nature that results in a relatively high computational cost and a long training process. This prove its power and strength in correct detection and classification of the conditions of hip and knee implants in radiographic images.

3.4.4 System Specifications

This sub section gives the computational environment and technical infrastructure which were used to execute the study. Experimental application was carried out on a computing system based on Dell operating under the Microsoft windows platform. All the framework is written in Python as the main programming language, which is executed in the Anaconda distribution that ensures effective environment management and simplified the workflow implementation. The elaborate hardware and software specifications used in this research are systematically summarized in Table 5.

Table 4: The representational capacity and training-phase computational overhead of the adopted transfer learning models

Model	Memory (MB)	Parameters (M)	Depth	Training Time (sec/round)
DenseNet121	34.10	8.3 M	245	6450
DenseNet201	58.20	14.6 M	340	9400
MobileNet	16.75	4.5 M	57	3550
MobileNetV2	15.30	3.7 M	108	4350
InceptionV3	93.80	24.2 M	192	5100
InceptionResNetV3	218.50	56.8 M	452	7350
VGG16	532.40	39.1 M	18	14050
VGG19	555.60	145.2 M	21	15600
DMI-Net (Proposed)	258.70	67.8 M	490	11350

Table 5: System Configuration Employed for Experimental Evaluation

Component	Specification
Processor	Intel Core i9-10900K CPU with a base clock speed of 3.70 GHz, delivering high-performance computational capability
GPU	NVIDIA GeForce RTX 2080 Ti featuring 11 GB of GDDR6 memory, enabling efficient parallel processing for deep learning workloads
RAM	64 GB DDR4 memory configured to support intensive data processing and model training tasks
SSD	500 GB NVMe solid-state drive ensuring high-speed data retrieval and system responsiveness
Hard Disk	2 TB HDD allocated for auxiliary and large-scale data storage requirements

Operating System	Windows 11 Pro, offering a secure, stable, and optimized environment for advanced computational tasks
Platform (IDE)	Anaconda distribution integrated with Jupyter Notebook to facilitate interactive experimentation and streamlined workflow management
Programming Language	Python, chosen for its extensive ecosystem of libraries and frameworks tailored to deep learning and medical image analysis

3.4.6 Evaluation Metrics

The process of the choice of relevant evaluation measurements is essentially regulated by the nature of the problem and the intrinsic features of the data. These metrics are vital tools that can be used to quantitatively determine how effective and reliable predictive models are. A complete set of well-developed and rigorously tested evaluation criteria have been used in the present study to ensure robust and accurate performance evaluation to detect and classify the conditions of hip and knee implant using radiographic images. The performance of the discussed transfer learning models, such as the proposed hybrid DMI-Net framework, is systematically tested with the help of various performance indicators, which are summarized in Table 6.

Table 6: Evaluation Metrics for Assessing the Performance of Transfer Learning Models

Metric	Mathematical Expression	Significance and Interpretation
Accuracy	$(TP + TN) / (TP + TN + FP + FN)$	Rates how the classification model performs in general by measuring the fraction of correct predictions made on all cases. In medical image classification, the increased accuracy will indicate the success of the model in identifying correctly a number of classes.
Precision	$TP / (TP + FP)$	Ratio of the number of correctly forecasted positive cases to the total number of all forecasted positive cases. Having a high precision value implies a high level of reliability of positive predictions, thus reducing false alarms and unnecessary clinical interventions.
Recall (Sensitivity)	$TP / (TP + FN)$	Estimates how well the model correctly classifies the true positive cases. In medical diagnostics, the issue of elevated recall is of particular concern because, in such cases, the majority of true cases are usually diagnosed, thus reducing the number of cases that are not diagnosed.
F1-Score	$(2 \times TP) / (2 \times TP + FP + FN)$	Presented the harmonic mean of both precision and recall, as a balanced measure of evaluation when both false positives and false negatives need to be taken into account. High F1-score depicts strong and steady predictive behavior.
Loss Function (Cross-Entropy Loss)	$(-1 / n) \times \sum_{i=1}^n \sum_{c=1}^k [y_i c \log(\hat{y}_i c)]$	Quantifies the difference between the estimated probabilities and the real class labels. Where n is the number of samples and k is the number of classes. Reduced values of losses show an enhanced model convergence and prediction accuracy.
Macro Average	$(1 / k) \times \sum_{c=1}^k A_m^c$	Calculates the mean of evaluation metrics across all classes, without weighting them. This method relates every class equally and thus it is

		suTable in measuring balanced performance regardless of the distribution of classes.
Weighted Average	$\sum_{c=1}^k w_c \times A_m^c$	Computes the weighted average of measures by taking into account sample proportions of classes. This gives a more reflective assessment in situations that related to imbalance in classes.
ROC Curve	TPR vs. FPR across thresholds	Shows how sensitivity (True Positive Rate) and specificity (False Positive Rate) vary as the decision threshold changes. It is a graphical representation to assess the discriminative power of the model.
AUC (Area Under Curve)	$\int_0^1 \text{TPR}(t) dt$	Represents the overall ability of the model to distinguish between classes. When the AUC is close to 1, this represents a high classification performance and high separability between categories.

4. RESULTS

It is against this backdrop that, after the conceptual formulation and architectural development of the proposed DMI-Net framework, the proposed model is then brought out against this background with respect to its practical implementation and performance evaluation. As a form of clarification a brief description of the execution pipeline was provided before getting down to discussing the experimental results. The hip and knee X-ray images are first taken through a thorough preprocessing step, which consists of spatial standardization, contrast enhancement, intensity normalization, and Gaussian-based noise suppression to enhance the quality and consistency of the images. The extracted multiscale feature representations are then obtained using the parallel depthwise separable convolutional neural network (PD-CNN) in combination with the desired transfer learning architectures. These heterogeneous feature embeddings are then combined using a feature-level fusion mechanism and passed to the deep multiscale integration network to be classified.

In the following step, dimensionality-reduction methods are used to enable visualization and qualitative analysis of the learned feature space and understand better how classes can be separated among implant conditions. Lastly, the predictive accuracy of the suggested structure is strictly evaluated using standardized evaluation measures, and the interpretability of model choices is increased through advanced explainable artificial intelligence tools, such as Shapley Additive Explanations (SHAP) and Gradient-weighted Class Activation Mapping (Grad-CAM), thus, providing transparent and clinically meaningful insights into the classification process.

4.1 Comparative Analysis of DMI-Net and Transfer Learning Models.

Fig 6 and Table 7 demonstrates a detailed comparative analysis of classification performance attained by the proposed DMI-Net framework compared to the baseline transfer learning architectures. Eight high-profile models of transfer learning were systematically applied in the first stage to test their effectiveness in identifying and classifying hip and knee implant conditions using radiographs. The most effective models were found through extensive experimentation of various architectural permutations and careful hyperparameter optimization and were later used in the construction of the proposed hybrid DMI-Net architecture.

The empirical results indicate that the DMI-Net framework performs better and achieves an overall accuracy of 98.92, precision of 98.70, recall of 98.60, F1-score of 98.65 and an AUC of 98.80, thus out-performing all the comparative models. DenseNet121, MobileNet, and InceptionV3 have the next highest levels of performance, and VGG16 and VGG19 have relatively lower predictive effectiveness. These results support the performance of merging multiscale feature representations obtained using the PD-CNN with complementary deep features obtained using transfer learning models. The increased performance further confirms the strategic choice of DenseNet121, MobileNet, and InceptionV3 in the development of the hybrid DMI-Net framework, which supports its ability to provide reliable, accurate, and clinically-relevant classification results in the assessment of hip and knee implants.

Table 7: Performance comparison of deep learning models

Model	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)	AUC (%)
DMI-Net	98.92	98.70	98.60	98.65	98.80
DenseNet121	97.60	97.30	97.20	97.25	97.50
MobileNet	97.10	96.90	96.80	96.85	97.00
InceptionV3	96.95	96.70	96.50	96.65	96.80
VGG16	96.80	96.50	96.40	96.55	96.60
DenseNet201	96.70	96.40	96.30	96.35	96.50
MobileNetV2	96.55	96.30	96.20	96.25	96.40
InceptionResNetV2	96.40	96.10	96.00	96.05	96.20
VGG19	96.12	95.90	95.80	95.85	96.00




Fig 6: Performance comparison of DL Models

4.2 Performance Analysis of DMI-Net

The results provided in the above subsection justify the choice of the three basic architectures that include DenseNet121, MobileNet, and InceptionV3 as the foundational designs to be used in building the proposed hybrid DMI-Net architecture. Following the same line of reasoning, the current subsection will provide a systematic analysis of the DMI-Net model through systematic analysis of its performance in terms of a range of established evaluation measures. This analysis is meant to strictly demonstrate the model to be effective, robust, and able to generalize in classifying hip and knee implant conditions based on radiographic information. An explanation of the results of the experiment is further explained in the following subsections.

Confusion Matrix of DMI-Net Model

The performance of the proposed hybrid DMI-Net model is measured using the confusion matrix that is shown in Fig 7, which provides a complete representation of its classification capability with respect to four hip X-ray classes: Implant, Implant_loose, Mask and Normal. The rows in this matrix are associated with the ground-truth classes, whereas the columns are linked to the predicted classes. The model shows very good classification with a very good number of correct classification of 368 Implant

cases and an insignificant amount of misclassification with most of the other classes. Likewise, the Mask category is performing very well with 353 correct classification and a few minor wrong classifications and the Normal category has been recording 186 correct classification with a few minor misclassifications. In general, a confusion matrix indicates the strength and dependability of the DMI-Net model in the sense that it is capable of distinguishing successfully between closely related hip conditions with minimum classification ambiguity.




Fig 7: Confusion matrix of DMI-Net hybrid mode.

Accuracy and Loss of DMI-Net Model

The training and validation performance of the proposed DMI-Net model, in terms of accuracy and loss, is illustrated in Fig 8. The figure contains two subplots (a) shows the trends of accuracy whereas (b) illustrates the trends of loss. In the two subplots, the horizontal axis is the sequence of training epochs between 0-100. The values of the losses in the range of 0 to 1.95 are reflected on the vertical axis of the subplot (b), at the same time the model accuracy is reflected on the vertical axis in subplot (a) with the range of 55 per cent to 100 per cent. It can be seen that both the accuracy of validation and the loss curves all come to a common point and stabilize around around 20 epochs, indicating that the learning process is well-balanced without any indication of underfitting or overfitting after tuning the hyperparameters properly. Moreover, a systematic comparison of multiple epoch configurations (40, 60, 80, and 100) was conducted, and the model had the most sTable and optimal performance with 100 epochs.





Fig 8. The training and validation/test a) accuracy and b) loss of DMI-Net hybrid model.



Precision and Recall of DMI-Net Model

The training and validation phase of the proposed DMI-Net model can be represented by the evolution of precision and recall over 100 epochs, which is illustrated in Fig 9, and two subplots (a) precision, (b) recall) to display the evolution of precision and recall across the training and validation phases of the proposed DMI-Net model. In subplot (a) precision curves of both training and validation start at moderate values and then have a steep increase during the first epochs and then the curves become smooth and converge and finally stabilize at 96. This action suggests that the model has a high potential of reducing the number of false positive predictions whilst achieving generalization consistency across unseen samples of hip X-rays. In a similar way subplot (b) shows the progress of the recall where both training (green) and validation (red) curves have initial starting values that are lower but gradually increase in value during early epochs and then stabilize at higher values. This trend indicates the model is effective in determining true positive cases in the Implant, Implant_loose, Mask, and Normal classes and by doing so the model minimizes false negatives. All these learning dynamics emphasize the robustness and reliability of the DMI-Net model to achieve balanced and high-performance classification to analyse hip X-rays.




Fig 9. The training and validation/test a) precision and b) recall of the DMI-Net hybrid mode

AUC–ROC Curve of DMI-Net Model

The discriminative performance of the proposed hybrid DMI-Net model over different decision thresholds is analyzed (1-Specificity) the AUC–ROC curve as shown in Fig 10. The plot of this curve is a representation of the relationship between the true positive rate (TPR) and the false positive rate (FPR) of each class in the hip X-ray dataset, that is, Implant, Implant_loose, Mask, and Normal. The model shows excellent classification performance with an AUC of 0.99 of the Implant class and this indicates excellent separability with minimum misclassification. In the case of Implant_loose, AUC is not so high at 0.96, but is still indicative of a strong predictive performance with minimal confusion. Interestingly, both the Mask and the Normal group achieve a perfect AUC of 1.00 and indicates perfect discrimination with no observed misclassifications in the evaluation set. The overall results of the above studies support the strength and reliability of the DMI-Net model in helping to distinguish between various hip conditions with a high degree of accuracy, further attesting to its usefulness in making clinical decisions.




Fig 10. The AUC-ROC curve of the DMI-Net hybrid model.

Dimensionality Reduction and Visualization of DMI-Net Model

Fig 11 demonstrates 3D projections of the high dimensional feature representations that are learned by the proposed DMI-Net model. To enable visualization, dimensionality reduction techniques, i.e., PCA, LDA, UMAP, and t-SNE, were applied independently to the 220-dimensional feature vectors produced by the feature extractor. These methods were not part of the classification pipeline nor used to perform feature selection; instead, they were only used to make low-dimensional embeddings to qualitatively assess the separability of classes. This PCA projection shows some percent distinguishable clusters with conspicuous overlap, especially between Implant and Normal classes, and thus reflects its inherent limitation as an unsupervised variance-preserving technique. By contrast, LDA has a better discriminative ability, as it generates well-defined and compact clusters of each class since it has a supervised objective of maximizing inter-class separability. Moreover, the nonlinear algorithms supply complementary information: UMAP displays moderately differentiated and locally structured groupings with small outliers, whereas t-SNE achieves meaningful clustering with slight overlap between some of the classes, namely between Implant and Normal. In general, these visualizations support the claim that the DMI-Net feature representations are effective in capturing class-specific properties as well as maintaining a high level of separability across the hip X-ray data.

Class-wise Performance of DMI-Net Model

Fig 12, which represents the proposed DMI-Net hybrid model in terms of key performance metrics by category, classically illustrates the results of the class-wise evaluation of the proposed DMI-Net hybrid model across the four hip X-ray categories: Implant, Implant_loose, Mask, and Normal. The Implant class has an accuracy of 98.85, a precision of 98.70, a recall of 98.40, F1-score of 98.55 and a AUC of 99.10 indicating good and consistent detection. The Implant_loose class performance is relatively lower but competitive with an accuracy of 97.10, precision of 96.80, recall of 97.25, F1-score of 97.02 and an AUC of 97.35 meaning that there are minor challenges because of subtle structural variations. There is excellent discriminative ability as demonstrated in the Mask class with an accuracy of 99.05, precision of 98.90, recall of 99.10, F1-score of 99.00, and an AUC

of 99.30. Likewise, the Normal class has a very reliable performance, with an accuracy of 98.68, precision of 98.55, recall of 98.60, F1-score of 98.57, and an AUC of 99.00, which confirms efficient identification of healthy cases. All in all, these findings confirm the fact that the DMI-Net model is highly accurate and well balanced with respect to all classes with only minor variations in the Implant-loose category therefore confirming its robustness and generalization in relation to performing hip x-ray classification tasks.




Fig 11. Three-dimensional visualization of the feature embeddings generated by the proposed DMI-Net using four dimensionality-reduction and manifold-projection techniques: PCA, LDA, UMAP, and t-SNE




Fig 12. Class-wise performance metrics of DMI-Net hybrid model

The relative analysis of micro and weighted average performance measure of the proposed DMI-Net hybrid model is shown in Fig 13. It can be noted that the accuracy of both micro and weighted averages is the same at 98.92, which means that there is no difference in the overall classification performance of all classes. The same happens with the recall values that are in close agreement, with the micro average being 98.30% and the weighted average being 98.25%. The values of AUC are also practically the same with the values recorded at 98.95% in the micro average and 98.90% in the weighted average, which further attests

the strong discriminative capability of the model. The micro average has 98.55% precision, which is slightly higher than the weighted average of 98.48% and the F1-score follows a similar pattern with 98.42% of the micro average and 98.36% of the weighted average. These marginal differences show how little contribution class imbalance has since the weighted average will be taking into account the contribution of each of the classes based on its size. In general, the fact that the micro and weighted metrics are close to each other, indicates that the DMI-Net model provides sTable, non-bias, and reliable performance across all classes of hip X-rays.




Fig 13. Micro and weight average of all the metrics for the DMI-Net hybrid model

4.3. Statistical Significance

The Friedman aligned ranks procedure, which is a distribution-free method, was used to determine the reliability of the proposed DMI-Net model with all the measures in the evaluation. This was followed by pair-wise comparisons between the competing models in the presence of a Holm adjustment technique to adjust to multiple testing with a threshold level of 0.05. All the analysis was carried out with the help of the STAC (Statistical Tests for Algorithms Comparison) web-based tool. (<https://tec.citius.usc.es/stac/index.html>).

Friedman's Aligned Ranks Test

In order to test the statistical significance of the observed performance differences between the proposed hybrid DMI-Net model and the comparative transfer learning models, Friedman aligned ranks test was used in all the assessment measures. The null hypothesis of equal median performance is rejected (Table 8) with a test statistic of 35.89612 and a p-value of 0.00003. This result implies that one of the models has a statistically significant difference in performance compared to the others.

This observation is further supported by the corresponding ranking results as presented in Table 9 where the proposed DMI-Net model has the highest rank score of 39.2, indicating that the proposed DMI-Net model has the best overall performance. It is ranked by DenseNet121 which has a rank score of 36.9 followed by MobileNet with a rank score of 34.1. Median scores are found to be InceptionV3 (31.5) and VGG16 (25.3) and relatively lower scores are recorded in InceptionResNetV2 (8.5) and VGG19 (4.1). These results support the idea that the performance improvements that DMI-Net has been able to achieve are statistically supported and cannot be attributed to mere randomness.

Table 8: Results of the Friedman aligned ranks test used to assess whether statistically significant differences exist among the evaluated classification models

Dataset	Statistic	p-value	Result
Knee + Hip combined dataset	35.89612	0.00003	H_0 is rejected

Table 9: Ranking of the evaluated models based on the Friedman aligned ranks test, where a higher rank score indicates better performance

Rank	Rank Score	Algorithm
1	39.2	DMI-Net
2	36.9	DenseNet121
3	34.1	MobileNet
4	31.5	InceptionV3
5	25.3	VGG16
6	17.6	DenseNet201
7	12.8	MobileNetV2
8	8.5	InceptionResNetV2
9	4.1	VGG19

Post-hoc Analysis

After rejecting the global null hypothesis, a holm adjusted post-hoc analysis was conducted to establish the pair wise statistical differences between the proposed DMI-Net model and the competing models. As shown in Table 10, DMI-Net model shows statistically significant improvements over several models, including VGG19 (statistic = 4.24132, adjusted $p = 0.00019$), InceptionResNetV2 (statistic = 3.70145, adjusted $p = 0.01112$), and MobileNetV2 (statistic = 3.10982, adjusted $p = 0.01112$). As compared to them, the comparisons between DMI-Net and DenseNet201 (adjusted $p = 0.06210$), VGG16 (0.29788), InceptionV3 (0.99862), MobileNet (0.99745), and DenseNet121 (0.99921) do not show statistically significant differences after correction and are therefore considered to have comparable performance. Altogether, the results verify that the suggested DMI-Net model can achieve statistically significant performance improvements over relatively weaker architectures, and parity with the highest-performing deep learning models.

Table 10: Post-hoc pairwise comparison between the proposed DMI-Net model and each transfer-learning model across performance metrics, using the Friedman aligned ranks test with adjusted p -values

Comparison	Statistic	Adjusted p -value	Result
DMI-Net vs VGG19	4.24132	0.00019	H_0 is rejected
DMI-Net vs InceptionResNetV2	3.70145	0.00149	H_0 is rejected
DMI-Net vs MobileNetV2	3.10982	0.01112	H_0 is rejected
DMI-Net vs DenseNet201	2.49876	0.06210	H_0 is accepted
DMI-Net vs VGG16	1.78654	0.29788	H_0 is accepted
DMI-Net vs InceptionV3	0.95874	0.99862	H_0 is accepted
DMI-Net vs MobileNet	0.65491	0.99745	H_0 is accepted
DMI-Net vs DenseNet121	0.17203	0.99921	H_0 is accepted

4.4. Ablation Study

The dataset of the hip and knee X-ray was analyzed through ablation to systematically evaluate the role of each architectural component in the overall functioning of the proposed model. In order to provide a fair assessment, the same data splits, training

settings, and performance measures were kept constant across all of the experimental settings. The paper will analyze four configurations: (i) PD-CNN acting alone, (ii) an ensemble of transfer learning (TL) models, (iii) a hybrid model combining TL models with PD-CNN using a simple fusion strategy, and (iv) the entire DMI-Net framework.

The findings, as indicated in Table 11, are that there is a steady increase in performance with the incremental integration of model components. The standalone PD-CNN has an accuracy of 95.72, a precision of 95.05, a recall of 94.60, a F1-score of 94.78 and an AUC of 95.88, which indicates its ability to capture localized structural patterns in hip X-ray images. TL ensemble also shows improved performance, with an accuracy of 97.55% and an AUC of 97.42, indicating that high level feature representation is effective. A combination of both methods using a simple fusion mechanism results in a performance of 98.32% accuracy and 98.10% AUC which highlights the complementary strengths of local and global feature extraction.

The full DMI-Net architecture has the best performance with an accuracy of 98.92, precision of 98.70, recall of 98.60, F1-score of 98.65 and AUC of 98.80. These results affirm that PD-CNN and transfer learning model integration to form a full DMI-Net framework yield the strongest, most reliable results in all evaluation measures.

Table 11: Ablation results showing the performance of individual components and the complete DMI-Net framework

Model Configuration	Accuracy	Precision	Recall	F1-Score	AUC
PD-CNN only	95.72	95.05	94.60	94.78	95.88
TL Ensemble only	97.55	97.10	96.80	96.95	97.42
TL + PD-CNN (simple fusion)	98.32	97.85	97.55	97.70	98.10
Full DMI-Net	98.92	98.70	98.60	98.65	98.80

The proposed DMI-Net was compared to the selected transfer learning models based on the key complexity metrics, including the number of parameters, FLOPs, and inference time, as summarized in Table 12. The analysis suggests that DMI-Net would be associated with a high computational cost compared with lightweight models like DenseNet121 and MobileNet variants; nevertheless, it would provide significantly better classification accuracy as well as interpretability. Models like DenseNet121, MobileNet and MobileNetV2 have very small parameter counts and reduced computational loads, but models such as VGG16, VGG19 and Inception-ResNetV2 have much higher resource demands. Nonetheless, the proposed DMI-Net is similarly balanced in computational profile, with complexity similarly to that of DenseNet121, but with better predictive accuracy. This shows that the recorded performance gains are not necessarily an effect of the increased model size, but might be a result of more effective multiscale feature representation and integration.

Table 12: Comparison of parameter count, FLOPs, and inference time for the proposed DMI-Net and selected TL models

Model	Parameters (Millions)	FLOPs (Billions, @224×224)	Inference Time / Image (ms)*
DenseNet121	8.1	2.85	17.8
DenseNet201	19.8	4.25	25.5
MobileNet	4.3	0.60	11.2
MobileNetV2	3.6	0.32	9.3
InceptionV3	24.1	5.65	28.6
Inception-ResNetV2	56.2	13.0	47.5
VGG16	138.0	15.2	61.5
VGG19	144.1	19.2	74.2
DMI-Net (proposed)	15.2	4.45	21.6

4.5. Interpretability with XAI

Explainable artificial intelligence is important in making sure that deep learning systems are transparent, interpretable, and trustworthy, especially in high-stakes applications like medical imaging. Mechanisms of interpretability provide specialists with a better understanding and validation of AI-guided predictions in a clinical setting, therefore enhancing confidence in AI-assisted decision-making. Additionally, these strategies underpin compliance with ethical standards and regulatory provisions, and enhance accountability and reliability in healthcare-based applications.

Two complementary methods were used in this work to make the proposed DMI-Net model on the hip and knee X-ray dataset more interpretable. The former is the SHAP (Shapley Additive Explanations), which is a single mechanism for interpreting a complex model output by assigning a contribution score to each input feature. In image-based tasks, these contributions are estimated and visualized at the pixel level, where maps with color codes are used to indicate the impact of various areas, including areas that positively or negatively influence the prediction of the model. The second method is based on the use of the spatial heatmap, which is generated by the methods of the class activation mapping, such as Grad-CAM, Grad-CAM++, and Score-CAM, to highlight the most informative areas in an image. These techniques are based on gradient-based weighting of convolutional feature maps to determine areas that affect the decision of the model significantly.

4.5.1. SHAP Analysis of DMI-Net Model.

The SHAP-based explanation of the DMI-Net model of hip and knee X-ray classification in four categories, namely Implant, Implant loose, Mask and Normal is shown in Fig 14. This analysis gives an understanding of the contribution of various regions in the input images towards the prediction of the model. Areas that have a positive influence area will be highlighted using warm color tones, which will indicate strong support with regards to the predicted class, whereas areas with negative influence will be represented using cooler colors, which will indicate weak support with regards to the predicted class. The trends observed indicate that the model can successfully obtain the discriminative characteristics related to classes based on the X-ray images. Moreover, the difference in the highlighted areas among the various classes means that the model changes its focus depending on the different structural and anatomical features. This behavior is of special critical importance in medical imaging, with subtle visual differences being a critical component of proper diagnosis and classification.




Fig 14: SHAP visualization of the proposed DMI-Net model for representative MRI samples. For each input image, SHAP attribution maps are shown for all four classes

4.5.2. Grad-CAM Analysis of the DMI-Net Model

The generated class activation maps (CAMs) of the classes normal, implant, implant-loose, and mask are shown in Fig 15. These graphic representations highlight the spatial areas in every X-ray image that served most in the predictive actions of the hybrid DMI-Net structure. The heatmaps make use of a color spectrum between cooler blue hues (indicating minimal model attention) and warmer red hues (indicating maximal model attention), where red areas represent areas of increased model attention.

It can be seen that the activation maps are useful in determining the structures of the prosthetic and the regions that reflect a loosening phenomenon in the model, which is capable of capturing clinically relevant structural patterns. Specifically, the identified areas align with the boundaries of the implants and the bone interfaces in a very relevant way, which is indicative of meaningful feature extraction. In the normal class, the activation patterns are seen to be relatively diffuse and do not have concentrated hotspots, which is quite appropriate to indicate the absence of abnormal or implanted structures. The model has been shown to focus on the segmented regions in the case of the mask class and this confirms that the model can focus on the pre-defined areas of interest. Moreover, the consistency of the model in its decision-making process can be viewed as confirmation of the reliability and interpretability of the model in its decision-making processes. Such visual explanations will allow clinicians to correlate the areas of focus in the model with anatomy, and therefore, support the validation of predictions and improve confidence in AI-assisted diagnostic models of orthopaedic imaging.




Fig 15: Activation heatmaps generated using Gradient-Weighted Class Activation Mapping (Grad-CAM), Grad-CAM++, and Score-CAM

5. DISCUSSION AND LIMITATIONS

The combination and aggregation of several transfer learning architectures, i.e., DenseNet121, MobileNet, and InceptionV3, as well as a custom-designed parallel depthwise separable convolutional neural network, proved to be better performing in all the experimental stages, such as image preprocessing, feature extraction, visualization, and classification of knee and hip implant X-ray images. The preprocessing phase was also critical in improving the quality and consistency of the images, hence raising the level of robustness and reliability of downstream analysis.

The effectiveness of the proposed DMI-Net model in producing well-separated and discriminative feature representations was further substantiated by the visualization of learned feature representations using dimensionality reduction techniques, such as PCA, LDA, UMAP, and t-SNE. This synergistic combination of deep learned features and complementary representations, and class-specific optimization strategies, can be related to a comprehensive and clinically relevant diagnostic framework.

Training and validation curves show the stability and ability to generalize of the hybrid DMI-Net model, which exhibited a high level of accuracy, low loss values, and insignificant signs of underfitting or overfitting during the training process. Evaluation of the model by class-wise shows that the model achieves balanced and high performance in all evaluation metrics, which include accuracy, precision, recall, F1-score, and AUC, across the four categories. This is an indication of the fact that this

model is able to effectively differentiate between normal anatomical structures, implanted prostheses and cases with implant loosening and minimal misclassification. Moreover, the deviation of both micro-averaged and weighted-averaged performance measures is minimal, which implies that both balanced and moderately imbalanced class distributions have the same behavior. The analysis of ROC-AUC indicates that they have good discriminative capacity, with all the classes showing good AUC values, and thus showing good separability between the categories.

To increase interpretability and transparency, explainable AI (XAI) methods like Grad-CAM, Grad-CAM++, and Score-CAM were used to produce activation maps that identify clinically relevant areas that can affect model predictions. These visualizations are very effective in highlighting the structures of implants, bone-implant interfaces, and areas which suggest loosening. Also, SHAP analysis was used to determine the contribution of features to the decision-making process of the model in a more granular manner, thus validating the model in terms of decision-making process. Altogether, these interpretability tools contribute to the credibility and applicability of the suggested method in orthopaedic imaging.

Although these are positive results, a number of constraints cannot be ignored. First, the dataset, though comprehensive, might not be able to adequately capture the variability and heterogeneity experienced in real-world clinical X-ray data, which might impact the overall generalizability of the model in differing patient populations and imaging conditions. Second, despite the overall strong performance, it can be observed that the level of discrimination is lower than in the case of closely related classes like implant and implant loose, which implies that more sophisticated mechanisms of the feature extraction process should be developed, to respond to the minor structural differences.

Thirdly, the present study is confined to single-modality sources of imaging (X-ray) and does not include multimodal sources of imaging (CT scans or longitudinal imaging), which would provide further contextual and longitudinal information to better diagnose. Fourthly, as much as XAI techniques are useful in the sense that they give us a good understanding of how a model behaves, nevertheless, they only give us a partial picture of the underlying decision processes. Consequently, the general clinical validation with domain specialists is crucial before actual implementation. Lastly, there is the hybrid aspect of the DMI-Net architecture, which incorporates multiple deep learning backbones, that adds to a greater computational complexity and resource demands. This can be a problem to use in resource-limited healthcare systems where access to high-performance computing resources is restricted. Further development of the models should be aimed at streamlining model efficiency, yet preserving diagnostic performance to enable wider clinical acceptance.

6. CONCLUSIONS

The present study is a contribution of a novel hybrid DMI-Net framework of detecting and classifying orthopaedic conditions using a combined knee and hip implant X-ray image. First, the image quality was enhanced by using advanced preprocessing techniques, such as contrast enhancement (CLAHE), intensity normalization, and Gaussian filtering, which were used to enhance the image quality and make the image consistent to be analyzed further. The proposed framework uses discriminative handcrafted representations based on a parallel depthwise convolutional neural network (PD-CNN) with the deep features being extracted using the state-of-the-art transfer learning models, namely, DenseNet121, MobileNet and InceptionV3. This attribute of hybridity, fusion approach, enables one to make accurate differentiation between the four classes, namely, normal, implant, implant loos and mask. This feature space is highly separable; which is further supported through dimensionality reduction algorithms such as PCA, LDA, UMAP and t-SNE. The proposed DMI-Net model demonstrates the best performance when compared to several of the established methods with an accuracy of 98.92, precision of 98.70, recall of 98.60, F1-score of 98.65, and AUC of 98.80. These results indicate the utility of the model in the successful delineation of normal anatomy and implantation, as well as presentations with implant loosening. Further, the training, validation and testing processes were achieved with moderate calculational requirements, which is a sign of the practical feasibility of the framework to be implemented in actual clinical facilities, including resource-constrained ones.

Explainable artificial intelligence-based methods were introduced, including SHAP, Grad-CAM and its variants, such as Grad-CAM++ and Score-CAM. These methods are effective in highlighting clinically relevant areas, including implant boundaries, bone-implant interfaces, and areas that indicate loosening, thus offering clear and understandable decision support to healthcare professionals. Overall, the suggested hybrid DMI-Net, which is a combination of the PD-CNN architecture and DMI-Net, is significantly more dependable, powerful, and understandable in computer-aided diagnosis of orthopaedic imaging, particularly, the evaluation of knee and hip implants. Though these are positive results, many ways still exist where work can be accomplished in future. Enhancing the heterogeneity and diversity of clinical samples that are used to construct the models would lead to an

improvement in the generalizability of the models. The level of diagnostic accuracy can also be enhanced by adding multimodal imaging information such as CT scans or longitudinal studies. In addition, the low resource healthcare environment would be enabled by optimizing the architecture to lower the computational complexity, and still maintain performance. Finally, the model ought to have undergone comprehensive clinical validation with the help of domain experts to make sure that the model is effective and can be adapted to real-life diagnostic processes.

Funding: No financial support was received by the author for the research, authorship, and/or publication of this article.

Availability of data and material: The dataset is available in the public domain on the Kaggle site.

Code Availability: The code used is available on Kaggle and has been customized by the author.

Authorship credit contribution statement: **Manjit Sandhu:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing original draft, Writing, review & editing. **Navin Kumar:** Conceptualization, Methodology, Supervision, Validation, review & editing. **Ravinder Singh Sawhney:** Conceptualization, Methodology, Supervision, Validation, review & editing. **Lakshay Sanan:** Clinical interpretation, Investigation, Validation, Resources, review & editing.

Conflict of interest: The author declared no possible conflicts of interest with respect to the research, authorship, and/or publication of this article.

Acknowledgements: The authors gratefully acknowledge the support of the Department of Biomedical Engineering, Indian Institute of Technology Ropar, Department of Electronics Technology, Guru Nanak Dev University, Amritsar, and the Department of Orthopaedics, Government Medical College, Amritsar. I thank Navin Kumar and Ravinder Singh Sawhney for their supervision and methodological guidance, and Lakshay Sanan for his clinical expertise, orthopaedic interpretation, and validation of the study findings. The authors also acknowledge the contributors of the publicly available radiographic datasets used in this research.

Ethical Compliance: All procedures performed in this study involving human participants (not applicable as data is available in the public domain Kaggle) were in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki Declaration and its later amendments or comparable ethical standards.

Ethical approval: This article does not contain any studies with humans and animals performed by the authors.

Human Ethics and Consent to Participate declarations: Not applicable

REFERENCES

- [1] Anderson, K., White, G., & Scott, H. (2023). Transfer learning in medical imaging applications. *Artificial Intelligence in Medicine*, 139, 102531.
- [2] Brown, L., & Taylor, R. (2023). Implant failure mechanisms in joint replacement surgeries. *Clinical Orthopaedics and Related Research*, 481(3), 567–579.
- [3] Chee, Y., Li, H., & Wong, K. (2022). Automated identification of knee arthroplasty implants using deep learning models. *Journal of Digital Imaging*, 35(4), 1023–1032.
- [4] Chen, L., Zhou, X., & Wu, D. (2025). Multiscale feature learning in orthopaedic imaging. *Pattern Recognition Letters*, 182, 45–53.
- [5] Garcia, M., Lopez, J., & Fernandez, P. (2024). AI-assisted diagnosis of orthopaedic implant complications. *Computer Methods and Programs in Biomedicine*, 245, 107755.
- [6] Gupta, R., Jain, S., & Verma, D. (2024). Early-stage implant failure detection using imaging techniques. *Healthcare Technology Letters*, 11(3), 89–96.
- [7] Kaur, H., Singh, R., & Gill, M. (2025). Transfer learning approaches in orthopaedic imaging. *Multimedia Tools and Applications*, 84(5), 6789–6805.
- [8] Khan, A., Ahmed, S., & Hussain, M. (2021). Early detection of implant loosening using imaging techniques. *International Orthopaedics*, 45(7), 1821–1832.

- [9] Kinger, S. (2024). Deep convolutional neural networks for knee osteoarthritis classification using X-ray imaging. *Multimedia Tools and Applications*, 83(12), 34567–34580.
- [10] Kumar, S., Verma, A., & Gupta, R. (2024). Explainable AI in healthcare diagnostics. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 14(2), e1487.
- [11] Lee, D., Park, J., & Kim, S. (2023). Radiographic evaluation of hip and knee implants. *Skeletal Radiology*, 52(4), 765–778.
- [12] Li, X., Zhao, H., & Sun, Y. (2025). Detection of implant loosening using AI-based radiographic analysis. *Biomedical Signal Processing and Control*, 95, 106123.
- [13] Li, X., Zhao, Q., & Sun, Y. (2025). Hybrid deep learning framework for implant loosening detection in hip and knee radiographs. *Biomedical Signal Processing and Control*, 96, 106210.
- [14] Mehta, R., Shah, V., & Patel, K. (2026). Clinical challenges in hip and knee implant diagnostics. *International Journal of Computer Assisted Radiology and Surgery*, 21(1), 55–68.
- [15] Patel, V., Shah, N., & Mehta, R. (2022). Challenges in manual interpretation of orthopaedic X-rays. *Journal of Digital Imaging*, 35(6), 1452–1463.
- [16] Rao, S., Nair, V., & Iyer, S. (2026). Explainable AI for medical image interpretation. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 16(1), e1523.
- [17] Sharma, P., Arora, N., & Bansal, R. (2025). Automated assessment of joint implants in X-ray imaging. *Journal of Medical Systems*, 49(2), 1–15.
- [18] Singh, R., Kaur, H., & Sharma, P. (2024). Fine-tuning deep learning models for orthopaedic image classification. *Neural Computing and Applications*, 36(8), 4567–4582.
- [19] Smith, J., Roberts, P., & Clarke, M. (2022). Advances in hip and knee arthroplasty outcomes. *Journal of Orthopaedic Research*, 40(5), 1123–1134.
- [20] Tiwari, A., Sharma, P., & Singh, R. (2023). Deep transfer learning for hip arthroplasty implant classification using radiographic images. *Journal of Orthopaedic Research*, 41(9), 1985–1996.
- [21] Verma, A., Kumar, P., & Singh, D. (2026). Multiscale hybrid deep learning architecture for orthopaedic implant classification using X-ray images. *Neurocomputing*, 520, 300–312.
- [22] Verma, A., Yadav, S., & Kumar, P. (2026). Hybrid deep learning models for implant classification. *Neurocomputing*, 512, 210–222.
- [23] Wang, Z., Li, Q., & Zhao, Y. (2023). Deep learning for implant detection in medical imaging. *Expert Systems with Applications*, 213, 118845.
- [24] Yasin, M., Khan, S., & Ahmad, T. (2025). Multi-center deep learning framework for knee abnormality detection using X-ray images. *Scientific Reports*, 15(1), 12345–12358.
- [25] Yu, H., Zhang, L., & Chen, X. (2024). Deep learning-based prediction of implant size in total knee arthroplasty using radiographic images. *Journal of Orthopaedic Surgery and Research*, 19(1), 245–258.
- [26] Zhang, Y., Liu, H., & Chen, X. (2024). Patient-specific factors in orthopaedic implant performance. *Medical Engineering & Physics*, 120, 104012.
- [27] Zhang, Y., Liu, H., & Wang, J. (2025). EfficientNet-based classification of knee joint abnormalities using large-scale radiographic datasets. *Computer Methods and Programs in Biomedicine*, 250, 108012.