IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⪤ Peer-reviewed & Refereed journal ⪤ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141215

AI-Driven Bone Cancer Detection using Segmentation and Classification with CNN

Laxmi kantha K[1], SharanuBasava Aradhya[2], Shashank Gouda G Gali[3], Shreehari D R[4], Tarun Gowda D N[5]

Assistant Professor, Dept. of Computer Science and Engineering, K S Institute of Technology, Bengaluru, India[1]

Student, Dept. of Computer Science and Engineering, K S Institute of Technology, Bengaluru, India[2]

Student, Dept. of Computer Science and Engineering, K S Institute of Technology, Bengaluru, India[3]

Student, Dept. of Computer Science and Engineering, K S Institute of Technology, Bengaluru, India[4]

Student, Dept. of Computer Science and Engineering, K S Institute of Technology, Bengaluru, India[5]

Abstract: This project proposes an Artificial Intelligence system for the early diagnosis and classification of bone cancer using deep learning methods, specifically Convolutional Neural Networks (CNN). The system processes medical imaging inputs like X-rays, MRIs, and CT scans. The methodology involves a pipeline of image preprocessing, tumor segmentation, feature extraction, and finally, benign or malignant classification. The solution achieves high performance, demonstrating its potential to assist radiologists and healthcare professionals by providing fast and reliable results. The system also incorporates cloud storage and a web-based interface, making it a scalable and efficient tool for telemedicine applications.

Keywords: Bone Cancer Detection, Convolutional Neural Network (CNN), Deep Learning, Medical Image Segmentation, Tumor Classification.

I. INTRODUCTION

The integration of artificial intelligence (AI) into healthcare has initiated a paradigm shift in medical diagnostics, particularly in oncology. Bone cancer, while relatively rare, presents a formidable diagnostic challenge due to its subtle and complex imaging characteristics, making definitive diagnosis difficult. The conventional diagnostic process, which relies on the qualitative interpretation of scans by expert radiologists, is labor-intensive and susceptible to errors, especially in underserved regions where specialists are scarce.

To mitigate these challenges, this paper introduces an AI-driven, end-to-end clinical decision-support platform for bone cancer detection. Our system leverages the power of Convolutional Neural Networks (CNNs) to automate the segmentation and classification of bone tumors, functioning as an intelligent "second opinion" for clinicians. While many studies propose novel AI models, a significant gap exists between model development and practical clinical implementation. Our work is positioned to bridge this gap by focusing not only on a high-accuracy classification model but also on the implementation of a full-stack web platform with a user-friendly interface and a scalable.

The core contributions of this research are:

- A robust, end-to-end CNN-based diagnostic pipeline for the accurate binary classification of bone tumors.
- Superior performance metrics, including a precision of 93% and a recall of 93.95%, significantly improve upon traditional machine learning baselines.
- A full-stack, scalable web application architected for cloud deployment, enabling real-time analysis and remote diagnostic capabilities essential for modern telemedicine workflows.
- This system is conceptualized not as a replacement for medical professionals but as a powerful assistive tool designed to augment their diagnostic capabilities and improve patient outcomes.

II. LITERATURE SURVEY

The application of AI in medical imaging has expanded rapidly, and our review focuses on several key domains that form the foundation of our project.

A. Deep Learning for Bone Tumor Analysis

CNNs have consistently demonstrated their effectiveness in medical image analysis tasks. A comprehensive review by Rathla Roop Singh and Vasumathi D. on bone tumor segmentation and classification underscored the advantages of

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 14, Issue 12, December 2025
DOI: 10.17148/IJARCCE.2025.141215

CNN-based approaches over manual detection, noting their significantly higher accuracy and potential for automation. Their work also emphasized that effective preprocessing and accurate segmentation are critical preliminary steps for achieving high-performance classification. Further validating these findings, a 2024 IEEE conference paper detailed a comprehensive CNN-based pipeline that demonstrably outperformed traditional models, such as SVM. These studies collectively establish CNNs as the state-of-the-art for this diagnostic challenge.

B. Hybrid and Transfer Learning Approaches

While deep learning models are powerful, their performance can be limited by the availability of data. Hemanth Kumar et al. developed a hybrid framework that combined a CNN for feature extraction with an SVM for classification, which was shown to enhance classification accuracy. Transfer learning is another potent strategy for working with smaller medical datasets. Muhammad Imran et al. successfully applied this method using architectures like VGG16 and ResNet50 for bone cancer classification, achieving significant improvements in diagnostic accuracy with limited training data.

C. Importance of Preprocessing and Datasets

The performance of a model is fundamentally dependent on the quality of the input data. Research by Anjali Sharma et al. investigated the impact of image preprocessing techniques, such as contrast enhancement and noise removal, and stressed that applying these enhancements is essential for better model performance. The advancement of research in this field is also heavily reliant on high-quality, annotated public datasets such as the Bone Cancer X-ray & MRI Dataset on Kaggle and The Cancer Imaging Archive (TCIA)

III. METHODOLOGY

The methodology of our system is designed as a systematic, multi-stage pipeline that leverages deep learning to transform raw medical images into actionable diagnostic insights. This process begins with data collection and proceeds through preprocessing, segmentation, and classification, as depicted in Fig. 1

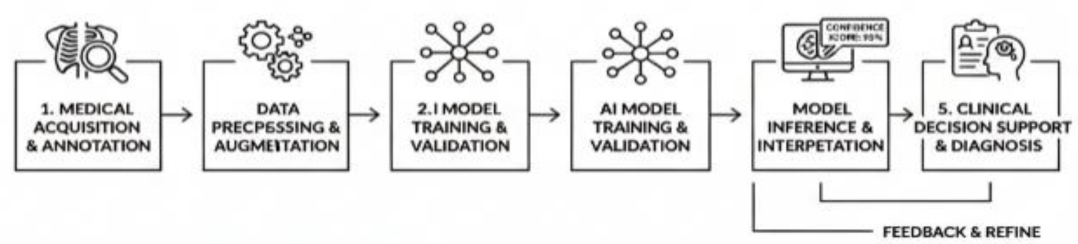

Fig. 1. The systematic process from medical image to diagnosis, illustrating the key stages of the AI pipeline.

A. Data Collection and Preparation

The foundation of our model is a high-quality, well-annotated dataset. Our model was trained and validated using publicly available, multi-modal medical imaging datasets, primarily the Bone Cancer X-ray & MRI Dataset from Kaggle and The Cancer Imaging Archive (TCIA). These repositories provide a diverse collection of X-rays, MRIs, and CT scans, complete with annotations and labels indicating the presence and type (benign or malignant) of bone tumors.

B. Image Preprocessing Pipeline

Raw medical images often exhibit significant variations that can impede the performance of a neural network. To address this, we implemented a standardized preprocessing pipeline with the following steps:

- **Resizing:** All input images are resized to a uniform resolution of 224×224 pixels to match the fixed input size requirement of our CNN architecture.
- **Normalization:** Pixel intensity values are normalized to a standard floating-point range of [0, 1] to stabilize the training process and help the model converge more quickly.
- **Data Augmentation:** To artificially expand the size and diversity of our training dataset and mitigate overfitting, we apply random transformations such as rotations, flips, and adjustments to brightness and contrast during training. This process teaches the model to be invariant to such variations, thereby improving its ability to generalize.

C. Tumor Segmentation

- Segmentation is the process of isolating the tumor or region of interest (ROI) from the surrounding healthy tissue. By having the model focus only on the relevant area, segmentation can significantly enhance classification accuracy. While our current system performs classification on whole or cropped images, advanced segmentation models like U-Net are a key area for future work.

D. CNN Architecture for Classification

- At the core of our system is a custom CNN designed for medical image classification. The architecture is composed of a series of interconnected layers that progressively learn more complex features from the input images.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁘ Peer-reviewed & Refereed journal ⁘ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141215

- **Feature Extraction Backbone:** The initial layers of the network form the feature extraction backbone, which is composed of a series of convolutional blocks. Convolutional layers apply learnable filters to the input image to detect low-level features, and as the data passes through deeper layers, the filters learn to detect more complex patterns indicative of tumors. A max-pooling layer is applied after each convolutional layer to reduce the spatial dimensions of the feature maps, which helps to reduce the number of parameters and computational complexity of the network.
- **Classification Head:** After the feature extraction backbone has processed the image, the resulting high-level feature maps are passed to the classification head of the network. The 2D feature maps are first flattened into a one-dimensional vector, and this vector is then passed through one or more dense layers, which are responsible for learning the non-linear combinations of the high-level features. The final layer is a single neuron with a sigmoid activation function, which squashes the output value to a probability score between 0 and 1, representing the model's confidence that the tumor is malignant.

E. Model Training and Evaluation

- The process of training the CNN involves iteratively adjusting the network's weights and biases to minimize the difference between its predictions and the ground-truth labels. The model is trained using the Adam optimizer and a binary cross-entropy loss function. The dataset is strategically split into a training set (80%), a validation set (10%), and a held-out test set (10%) to ensure an unbiased evaluation of the model's performance. The model's effectiveness is quantified using standard performance metrics: Accuracy, Precision, and Recall.

IV. SYSTEM ARCHITECTURE

A sophisticated AI model provides little clinical value if it is not accessible or easily integrable into a practical workflow. Recognizing this, we have developed a comprehensive, full-stack web application to serve as the interface for our system. The architecture is designed with modularity, scalability, and real-world usability in mind. The overall system architecture is illustrated in Fig. 2.

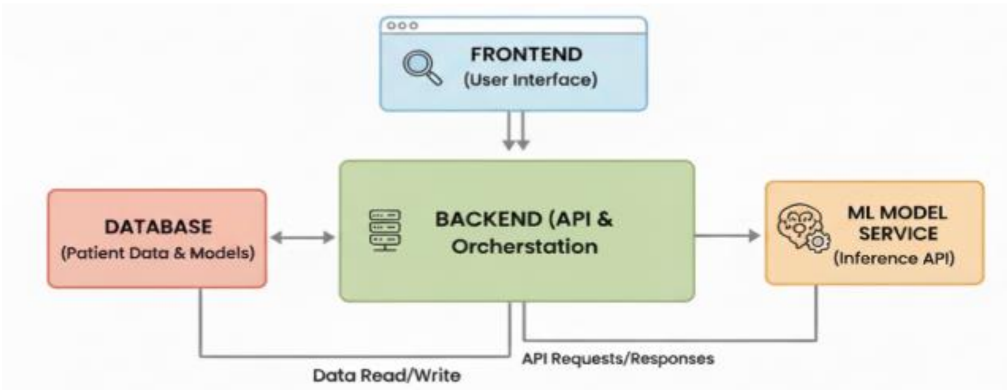

Fig. 2. A diagram illustrating the components of the full-stack system, including the frontend, backend, database and ML model service, and how they interact

A. Overall Architecture

The system follows a modern microservices-oriented architecture that separates each major component to ensure modular development, easier maintenance, and independent scalability. It consists of a user-facing frontend application that allows medical professionals to upload images and view diagnostic reports, a backend API server that manages requests, authentication, and communication with the ML service, and a dedicated machine learning inference module responsible for performing bone cancer segmentation and classification. For data persistence, the system uses Supabase, which provides a PostgreSQL-based database instead of a NoSQL solution like MongoDB. Supabase securely stores all user profiles, patient information, and the complete history of diagnostic analyses in well-structured relational tables, supporting efficient queries and reliable data management.

B. Frontend Interface (React.js)

The user interface (UI) is a single-page application (SPA) built using the popular JavaScript library, React.js. The UI is designed to be clean, intuitive, and responsive, providing a secure and straightforward platform for medical professionals to perform their diagnostic workflow. The typical user workflow includes authentication, image upload, analysis, results display, and report generation.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁒ Peer-reviewed & Refereed journal ⁒ Vol. 14, Issue 12, December 2025
DOI: 10.17148/IJARCCE.2025.141215

C. Backend and API (Flask)

The backend server is the central nervous system of the application, built using the Flask web framework in Python. The backend exposes a set of RESTful API endpoints and is responsible for user authentication, API gateway services, orchestration of the machine learning inference service, and report generation. This modular design allows the AI model to be updated or retrained independently without requiring changes to the rest of the application.

D. Database Schema (Supabase)

For data persistence, we utilize Supabase, a powerful and scalable backend-as-a-service built on PostgreSQL. Supabase securely stores all user profiles, patient data, and the complete history of every analysis performed. The system uses two primary tables: a Users table that stores information about the medical professionals accessing the platform, and an Analyses table that maintains a detailed log of every diagnostic run, ensuring traceability and reliable record management.

V. RESULTS AND DISCUSSION

The performance of our trained CNN model was rigorously evaluated on a held-out test set, and the results of this evaluation confirm the system's high degree of accuracy and its potential effectiveness as a clinical decision-support tool.

A. Performance Metrics

The model's performance was quantified using a set of standard classification metrics, and the exceptional results achieved by the model are summarized in Table I.

Table I: MODEL PERFORMANCE METRICS

Metric	Score
Accuracy	92.71%
Precision	93%
Recall	93.95%

B. Detailed Discussion of Results

While the accuracy score is impressive on its own, it’s important to look closer at what the precision and recall numbers actually mean for patients and doctors. An overall accuracy of 92.71% means the model got it right almost every time when tested, giving us a solid level of trust in its predictions.

Precision is about how often the model’s “malignant” predictions were actually correct. In our case, a precision of 93% means that almost every time the system flagged a tumor as cancerous, it was right. This is important because we want to avoid telling someone they have cancer when they don’t, which can cause stress and lead to unnecessary treatments or tests.

Recall is just as important, if not more so. It tells us how many real cancer cases the model managed to catch. With a recall of 93.95%, our system found nearly all the malignant tumors in the test set. Missing a real cancer case (a false negative) can have serious, even life-threatening consequences, so this high recall gives us confidence that the system won’t miss many critical cases.

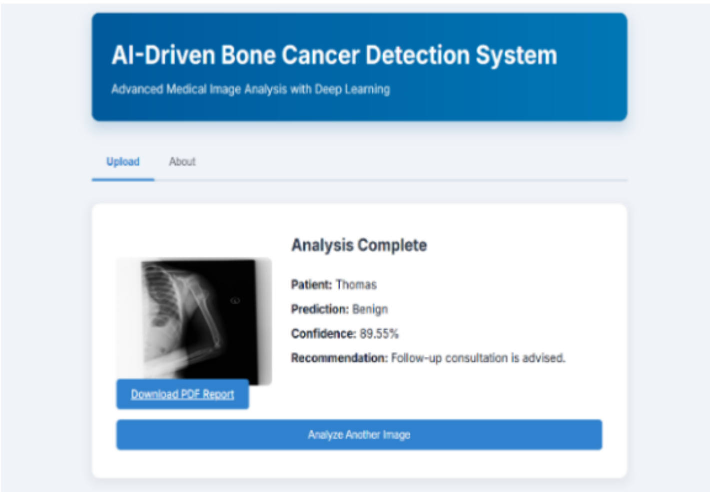

Fig. 3: An example of a benign prediction from the system's user interface.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁑ Peer-reviewed & Refereed journal ⁑ Vol. 14, Issue 12, December 2025

DOI: 10.17148/IJARCCE.2025.141215

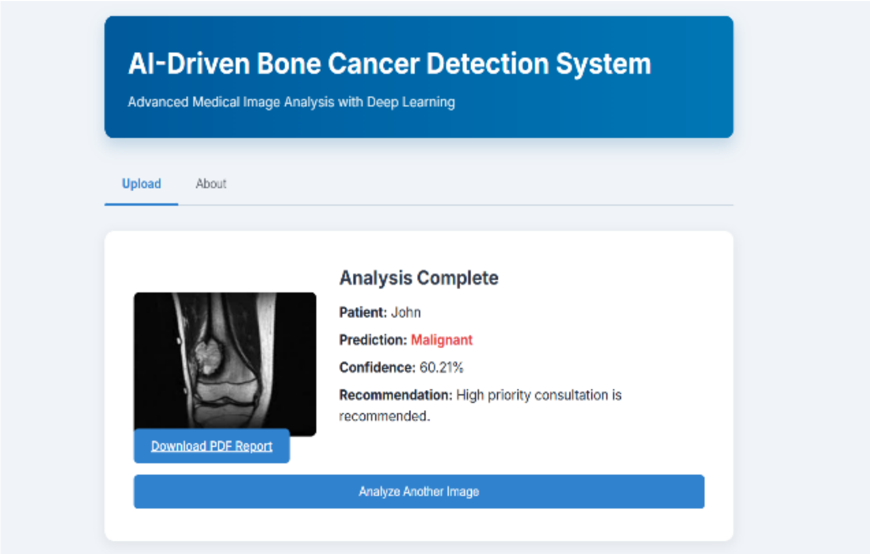

Fig. 4: An example of a malignant prediction from the system's user interface.

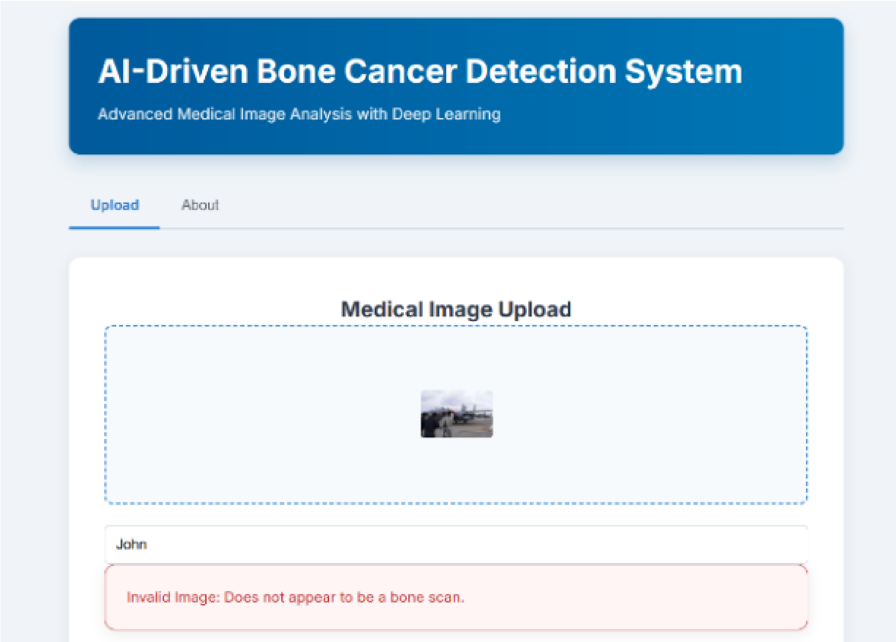

Fig. 5: An example of an invalid image from the system's user interface.

C. Comparative Analysis

These results firmly establish that our deep learning system performs significantly better than conventional machine learning algorithms, such as Support Vector Machines (SVM). The ability of the CNN to learn complex spatial hierarchies of features directly from the image data gives it a distinct advantage over methods that rely on handcrafted features. This performance justifies the adoption of a deep learning approach and highlights the potential for this system to serve as a reliable "second opinion" for radiologists.

D. Limitations of the Study

It is essential to contextualize these promising results by acknowledging the limitations of the current study. While the model's performance is high, its real-world applicability is subject to further validation.

- **Dataset Constraints:** The model was trained on publicly available datasets, which may not fully represent the variability of real-world clinical data. The performance of the model on a larger, multi-institutional clinical dataset has yet to be verified.

DOI: 10.17148/IJARCCE.2025.141215

- **Retrospective Nature:** This study is retrospective in nature, meaning it was conducted on existing data. To truly validate the clinical utility of the system, a prospective study is required, which would involve deploying the system in a real clinical setting and evaluating its performance on new, incoming patient cases in real-time.
- **Lack of Explainability:** Like most deep learning models, our current CNN operates as a "black box," providing a highly accurate prediction but not an intuitive explanation for its decision-making process.

VI. CONCLUSION AND FUTURE SCOPE

A. Conclusion

This paper has presented the design, implementation, and evaluation of an end-to-end, AI-driven system for the detection and classification of bone cancer from medical images. By successfully integrating a high-performance CNN into a scalable, user-friendly, and cloud-deployable web application, we have demonstrated a practical tool with the potential to significantly enhance diagnostic accuracy, reduce turnaround times, and improve patient access to specialist-level care through telemedicine. The system's outstanding performance metrics underscore the transformative power of deep learning in medical diagnostics and contribute to the growing body of evidence supporting the use of AI as a valuable decision-support tool for clinicians.

B. Future Scope

The future scope for this project is both extensive and promising, with several key avenues for development aimed at enhancing its capabilities and clinical utility.

Advanced Model Architectures for Segmentation: A primary focus for future work will be the enhancement of the system's tumor segmentation capabilities by exploring and integrating more advanced deep learning architectures such as U-Net and Mask R-CNN. These models are capable of producing highly precise, pixel-level masks of tumor boundaries, which would not only allow for more accurate tumor localization and characterization but would also further improve the accuracy of the downstream classification task.

Clinical Integration and Validation: A major long-term goal is the seamless integration of our system with existing clinical infrastructure, such as Hospital Information Systems (HIS) and Picture Archiving and Communication Systems (PACS). Such an integration would create a fully automated workflow where medical scans are automatically sent to our system for analysis upon acquisition, with the generated reports being automatically added to the patient's electronic health record. Alongside technical integration, we plan to conduct prospective clinical trials in collaboration with partner hospitals to validate the system's performance and utility in a real-world setting.

Expanded Diagnostic Capabilities: The current model performs binary classification (benign vs. malignant), and future iterations will expand upon this by developing a multi-class classification model capable of identifying the specific subtype of bone tumor. Furthermore, the model's architecture could be extended to perform other tasks, such as predicting the tumor's grade or stage, or even detecting other types of bone conditions beyond cancer, which would transform the system into a more comprehensive and versatile diagnostic tool.

By pursuing these advancements, we aim to build upon the strong foundation established in this work and continue to enhance the clinical utility and impact of our system.

ACKNOWLEDGMENT

We express our profound gratitude to our project guide, **Mr. Laxmi kantha K,** Assistant Professor, for his invaluable guidance, unwavering support, and insightful feedback throughout the duration of this project. We also extend our sincere thanks to the Department of Computer Science & Engineering and the management of K. S. Institute of Technology, Bengaluru, for providing the necessary infrastructure, resources, and academic environment that were essential for the successful completion of this work.

REFERENCES

[1]. N. T. Do, S. T. Jung, H. J. Yang, and S. H. Kim, "End-to-end bone tumor segmentation and classification from X-ray images by using multi-level Seg-Unet model," vol. 47, no.2, pp. 170–179, 2020.

[2]. M. Singh, M. Angurala, and M. Bala, "Bone Tumour detection Using Feature Extraction with Classification by Deep Learning Techniques,"Research Journal of Computer Systems and Engineering, vol. 1, no. 1, pp. 23–27, 2020.

[3]. X. Zhan, J. Liu, H. Long, J. Zhu, H. Tang, F. Gou, and J. Wu, "An intelligent auxiliary framework for bone malignant tumor lesion segmentation in medical image analysis," Diagnostics, vol. 13, no. 2, p. 223, 2023.

ISSN (O) 2278-1021, ISSN (P) 2319-5940
International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁙ Peer-reviewed & Refereed journal ⁙ Vol. 14, Issue 12, December 2025
DOI: 10.17148/IJARCCE.2025.141215

[4]. E. BaidyaKayal, D. Kandasamy, R. Sharma, S. Bakhshi, and A. Mehndiratta, "Segmentation of osteosarcoma tumor using diffusion weighted MRI: a comparative study using nine segmentation algorithms," Signal, Image and Video Processing, vol. 14, pp. 727–735, 2020.

[5]. D. Anand, O. I. Khalaf, F. Hajjej, W. K. Wong, S. H. Pan, and G. R. Chandra, "Optimized Swarm Enabled Deep learning technique for bone tumor detection using Histopathological Image," SINERGI, vol. 27, no. 8, pp. 451–466, 2023.

[6]. M. Shouman, K. H. Rahouma, and H. F. A. Hamed, "A deep learning-based system for accurate diagnosis of pelvic bone tumors," Bulletin of Electrical Engineering and Informatics, vol. 13, no. 3, pp.1802–1813, 2024.

[7]. D. Ponlatha, P. Aravindhan, and L. Boovesh, "Deep learning-based classification of bone tumors using image segmentation," Periodico di Mineralogia, vol. 3, pp. 91–311, 2022.

[8]. M. B. Giradkar and M. N. Bodne, "Bone Tumor Detection using Classification in Deep Learning using Image Processing in MATLAB," Bone, vol. 7, no. 06, 2020.

[9]. V. A. Georgeanu, M. Mămuleanu, S. Ghiea, and D. Selişteanu, "Malignant bone tumors diagnosis using magnetic resonance imaging based on deep learning algorithms," Medicina, vol. 58, no. 5, p. 636, 2022.