YOLOv8-Based Real-Time Bone Fracture Detection and Localization Using Deep Learning

[1]Muskan Kurmi, [2]Mr. Anurag Jain, [3]Rajneesh Pachouri
[1]Research Scholar, [23]Assistant Professor
[123]Department of Computer Science and Engineering
[123]Adina Institute of Science & Technology, Sagar, India

Abstract—**Bone fractures are one of the most common orthopedic injuries and require fast and accurate diagnosis for effective treatment and patient recovery. This thesis presents an intelligent web-based bone fracture detection system using YOLOv8 and Deep Learning techniques for automated analysis of X-ray images. The YOLOv8 object detection model is trained using a dataset containing 3,316 training images and 399 validation images covering seven fracture classes including Elbow Positive, Fingers Positive, Forearm Fracture, Humerus, Humerus Fracture, Shoulder Fracture, and Wrist Positive. The system performs real-time fracture detection and localization by identifying fracture regions through bounding box prediction, thereby assisting medical professionals in faster and more accurate diagnosis. Experimental results prove that the proposed model reaches a mean Average Precision (mAP50) of 86%, highlighting the efficiency of the system in handling complex medical imaging tasks with high consistency and productivity. The proposed framework provides a cost-effective and scalable solution for initial fracture screening and has important potential for future integration into clinical healthcare systems and intelligent medical diagnostic applications.**

***Index Terms*—Bone Fracture Detection, Deep Learning, YOLOv8, X-Ray Image Analysis, Object Detection**

I. Introduction

Proper treatment of common injuries such as bone fractures depends on quick and accurate diagnosis. Traditional fracture diagnosis on X-ray images is time consuming and strongly relies on the expertise of the radiologist. The application of Artificial Intelligence and Deep Learning techniques in medical image analysis is increasing. Deep learning models increase diagnostic accuracy and decrease human errors in fracture detection. The existing systems are mainly focusing on fracture classification but they cannot localize in real time. The suggested system uses YOLOv8 for automatic fracture detection and localization. The system is developed using Python, Flask, HTML, CSS and JavaScript for web-based deployment. The model was trained on x-ray datasets of seven different fracture classes. The proposed system achieves high detection performance with an accuracy of 86% of mAP50. The system is a fast, intelligent and user-friendly solution for an automated diagnosis of bone fracture.

Fractures of bone are among the most frequent injuries encountered in medical practice and require accurate and timely diagnosis for appropriate treatment and recovery. Diagnosis of bone fractures has traditionally depended on visual inspection of X-ray images by radiologists and medical professionals. But manual interpretation can be time-consuming and subject to human error, especially in cases with subtle or complex fracture patterns. The growing volume and complexity of medical images has created an increasing need for automated, reliable systems to assist in the diagnostic process. In this Thesis, a web-based application for the detection of bone fractures in X-ray images is proposed based on deep learning with high accuracy and efficiency. The system is implemented using Python for backend processing and HTML, CSS and JavaScript for a friendly front end. The web framework used to connect the components smoothly is Flask. The core of the system is built on YOLOv8 (You Only Look Once version 8), a state-of-the-art object detection model that is fast and accurate for object detection in images. The project is to detect and localize 7 types of upper limb fractures, Elbow Positive, Fingers Positive, Forearm Fracture, Humerus, Humerus Fracture, Shoulder Fracture and Wrist Positive. The model was trained and validated using a curated dataset of 3,316 training images and 399 validation images, resulting in a mean Average Precision (mAP50) of 86%. The system allows users to upload x-ray images via a web interface and gives immediate visual feedback with annotations showing the type and location of fractures.The project is intended to use sophisticated deep learning technology and practical web-based application to enhance the speed and accuracy of bone fracture detection. This will offer valuable support to medical professionals and will contribute to improved patient care outcomes.

Problem Statements

The conventional diagnosis of bone fracture from X-ray images is time consuming and heavily dependent on the expertise of radiologists. The existing fracture classification systems cannot accurately localize the fracture regions from the X-ray images in real time. Deep learning models are limited in generalization ability and performance due to dataset size and class imbalance. Manual interpretation of medical images can lead to diagnostic errors, delay in treatment and extra burden on healthcare professionals. The need for an intelligent, automated and web-based fracture detection system that can provide fast, accurate and real-time diagnosis using YOLOv8 and Deep Learning techniques is there.

Fractures of bone are one of the most common injuries managed in the emergency, trauma and surgical departments. Accurate localization of these injuries is important to facilitate prompt intervention and optimal patient outcomes (1). Manual visual assessment of radiographs remains the standard of care in clinical practice. However, this method is time-consuming and exhibits significant inter-observer variability, especially for subtle or over-lapping fracture lines (2). To address these limitations, deep learning (DL) methods have been extensively used in medical imaging, which have shown superior performance to traditional manual methods in terms of diagnostic efficiency and detection accuracy (3).

Recent developments in convolutional neural networks (CNNs) and object detection architectures have enabled classification and localization of fractures simultaneously, indicating the promise of automated radiographic interpretation to aid clinical workflows (4). Real-time diagnostics are feasible as demonstrated by studies utilizing object detection models like YOLOv5 and its variants which achieved high mean average precision (mAP) in identifying fracture regions on X-ray images in initial stages (5). Besides, the use of advanced models with attention mechanisms and multi-scale feature extraction has increased detection capability in different anatomical locations. However, translating model performance into clinically reliable results remains challenging, with trustworthy confidence estimates, clinically consistent decision thresholds, and model interpretability being critical hurdles to safe deployment in patient care settings.

II. Literature Survey

Classification performance evaluation

The fracture classification component was developed using transfer learning with three distinct convolutional neural network (CNN) architectures: ResNet18, MobileNetV3-Small, and EfficientNet-B0. For each pre-trained network, the final classification layer was replaced to match the specific number of classes in the fracture classification task.

To ensure methodological rigor beyond initial pilot experiments, model development was conducted using a stratified 5-fold cross-validation framework across the complete classification dataset. To guarantee stability and reproducibility in the predictive performance estimates, training for each architecture was repeated across three random seeds, yielding multiple fold-level evaluations.

The networks were optimized using the Adam optimizer with a learning rate of 1×10^{-4} and a cross-entropy loss function. Training was capped at 30 epochs and governed by an early stopping mechanism tied to the validation Area Under the Receiver Operating Characteristic (AUROC) curve;

a predefined patience criterion was enforced to mitigate overfitting. Ultimately, the top-performing checkpoint from each fold was saved and utilized for the final fold-level evaluation [5].

Fracture Detection with YOLOv8

For the fracture localization component, we utilized the Ultralytics YOLOv8 framework. For a fair comparison, we benchmarked three different one-stage detectors with different model scales, namely YOLOv8n (nano), YOLOv8s (small), and YOLOv8m (medium), rather than only one detector configuration. This enabled a direct performance comparison between lightweight and intermediate and higher-capacity variants of the architecture under the same task and dataset constraints.

The localization dataset contained bounding box annotations in YOLO format for elbow, finger, forearm and wrist fractures, and was split into training, validation and testing subsets. All models were trained by supervised optimization up to 100 epochs. We used the Adam optimizer, a cosine learning rate scheduler, and an early stopping mechanism to avoid overfitting, and chose the best model according to peak validation performance. To improve generalization, a data augmentation pipeline was utilized, consisting of horizontal flipping, scaling, translation, and mild geometric and photometric transformations [5].

Artificial intelligence in diagnosis of bone fractures

Motivated by the clinical need of very precise and efficient diagnostic systems, the use of Artificial Intelligence (AI) in the diagnosis of bone fracture has proven to be a revolutionary paradigm in medical imaging. The traditional diagnosis of fractures depends largely on the radiologist's manual interpretation of radiographs. However, subtle or non-displaced fractures are difficult to identify, human fatigue can impact performance, and there are no specialized experts in remote areas, all of which point to a pressing need for automated diagnostic solutions. The models have demonstrated much promise for improving diagnostic accuracy and reducing backlogs of interpretation. This has led to the development and refinement of various frameworks capable of detecting fractures in different anatomical regions and using different imaging modalities like digital radiography (X-rays), computed tomography (CT) and magnetic resonance imaging (MRI) [7].

Several important approaches have been proposed to cope with these diagnostic difficulties. A framework was proposed by Pareek et al. and tested on the MURA dataset. They considered a subset of the MURA dataset's seven anatomical categories: forearm, elbow, wrist, humerus, hand, fingers and shoulder. The pipeline starts with a rigorous image pre-processing stage that includes Gaussian and median filtering, smoothing, resizing, contrast limited adaptive histogram equalization (CLAHE) and data augmentation, followed by feature extraction and classification. The authors compared the performance of traditional machine learning classifiers like Support Vector Machines (SVM), K-Nearest Neighbors (KNN) and logistic regression with deep learning architectures, including standard Convolutional Neural Networks (CNNs), Faster R-CNN and Deep Neural Networks (DNNs). To interpret and visualize the model's internal representations, we implemented region of interest (ROI) mapping and heatmaps (e.g. Grad-CAM). The model performances were evaluated rigorously using the metrics of accuracy, precision, recall, and Area Under the Receiver Operating Characteristic curve (AUROC), which helped in the hyperparameter tuning and the model optimization [18].

Table 1 Literature Survey Table

Author (Year)	Method	Dataset	Accuracy	Limitation
Rui-Yang Ju et al. (2023)	YOLOv8	Pediatric Wrist X-ray Images	85%	Limited to wrist fractures
Peng Chen et al. (2024)	Attention-Based YOLO Model	Multi-Site Bone X-ray Images	89%	Requires large annotated dataset
Zou et al. (2024)	Improved YOLOv7 + Attention Mechanism	Whole Body Fracture Dataset	80.2% mAP	High computational complexity
Amal Alshahrani et al. (2024)	CNN + VGG16	Bone X-ray Images	92%	Classification only
Theyazn Aldhyani et al. (2025)	DenseNet201 + Transfer Learning	Radiographic Bone Images	97%	No real-time localization
Mukhtar Ul Islam et al. (2025)	Vision Transformer + CNN	Multi-Class Fracture Dataset	96%	Complex training process
Ahmed et al. (2024)	YOLO-Based Detection	Wrist Abnormality Dataset	87%	Limited fracture categories
Hardalaç et al. (2022)	Deep Learning Object Detection	Wrist X-ray Images	84%	No deployment framework
Medaramatla et al. (2024)	Hybrid YOLO NAS	Hand Bone Fracture Images	90%	High GPU requirement
Base Paper Authors (2025)	VGG16, ResNet50, EfficientNetB0, Random Forest, SVM, XGBoost	Bone X-ray Dataset	95%	No localization and real-time detection

III. Proposed Framework:

Proposed Methodology

This study proposes a methodology based on YOLOv8 and Deep Learning techniques for automated detection and localization of bone fractures in X-ray images. First, the X-ray dataset is collected, pre-processed, and split into training and validation sets. Then, the YOLOv8 model is trained to detect the fracture regions and generate the bounding boxes for localization. The model is then integrated with a Flask based web application for real time fracture prediction through a user-friendly interface. Finally, the system performance is evaluated by the Accuracy, Precision, Recall, F1-Score and mAP50.

Algorithm Steps

Algorithm: Bone Fracture Detection Using YOLOv8

Step 1:

Collect dataset of bone X-ray images from medical image repositories and classify the images according to fracture classes.

Step 2:

Use image resizing, normalization and annotation formatting as data preprocessing techniques to make them compatible with YOLOv8.

Step 3:

The data will be split into training and validation sets so as to have proper evaluation of model performance.

Step 4:

Load YOLOv8 deep learning architecture and set training parameters.

Step 5:

Train the YOLOv8 model using the prepared X-ray dataset for fracture detection and localization.

Step 6:

Deep learning model convolutional layers automatically extract image features.

Step 7:

Identify regions of fracture and produce bounding boxes with the corresponding fracture class labels.

Step 8:

Evaluate the trained model on metrics such as Precision, Recall, F1-Score and mAP50.

Step 9:

Save the trained model for the prediction and deployment purpose.

Step 10:

Develop a web app to upload X-ray images and display real-time results of fracture detection using Flask, HTML, CSS and JavaScript.

Step 11:

Final prediction output of detected fracture type and localized fracture region on Xray image.

Step 12:

End the process after displaying the detection and analysis results to the user.

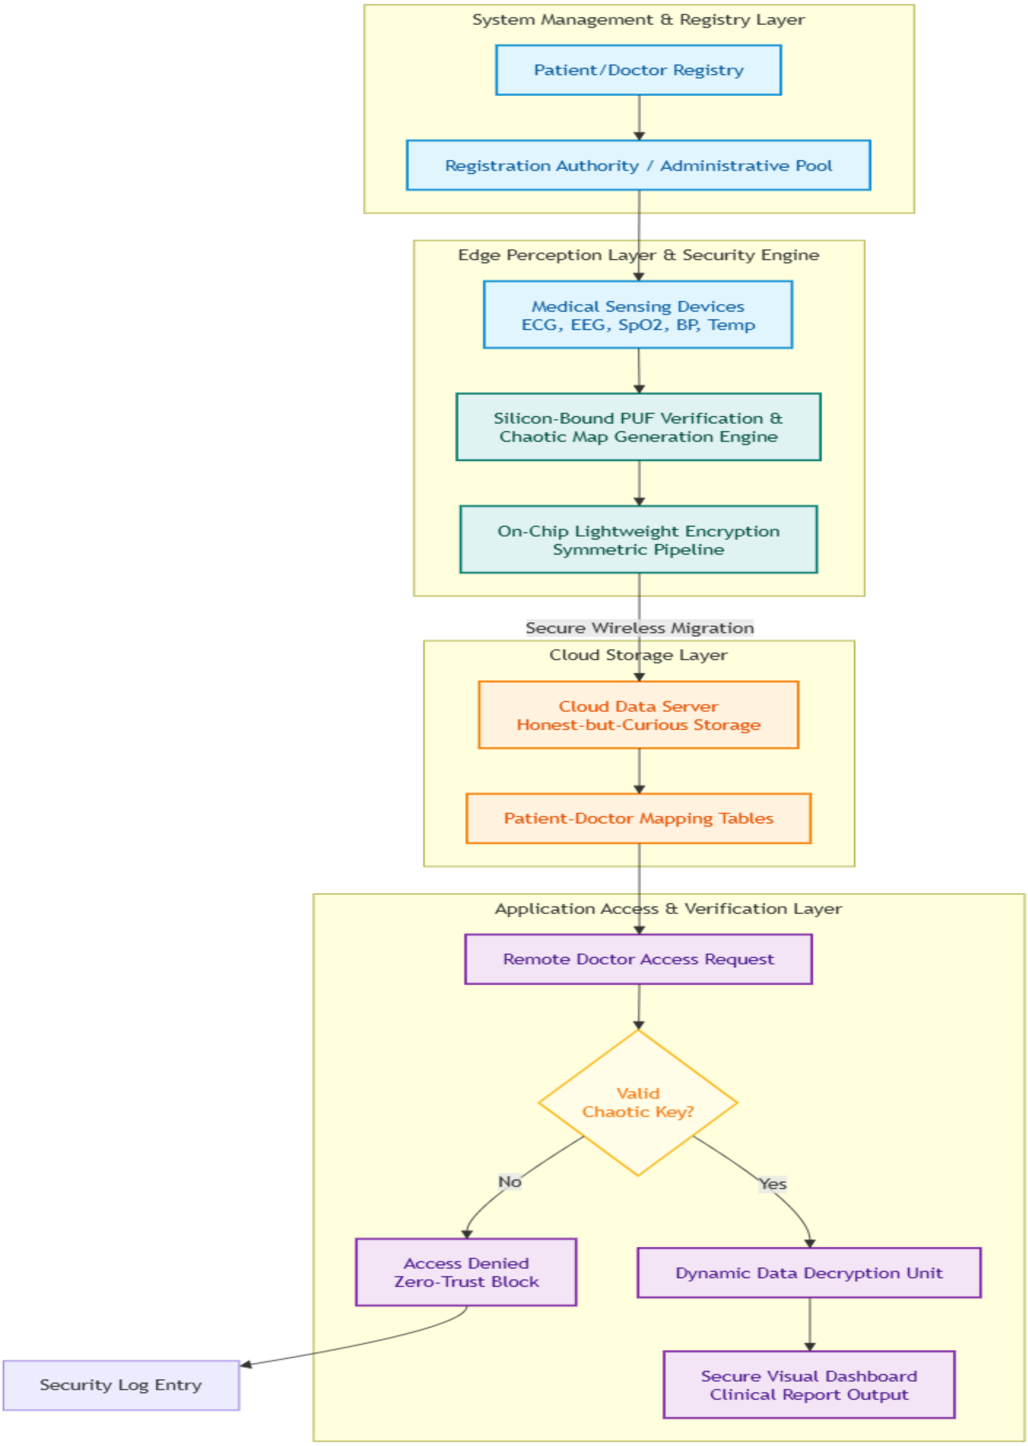

Figure 1 System Architecture

Figure 2 Data Flow Diagram

The proposed architecture transitions through a vertically coordinated processing pipeline designed to eliminate static cryptographic key storage vulnerabilities across decentralized Internet of Medical Things (IoMT) networks. The framework establishes a secure data lifecycle connecting three core structural tiers: the Edge Perception Layer, the Cloud Storage Layer, and the Application Access Layer.

1. **Edge Perception Layer:** Following a multi-factor user login, wearable sensor nodes collect continuous physiological telemetry streams (including ECG, EEG, SpO_2, blood pressure, and core temperature). To safeguard this resource-constrained environment from physical tampering, the architecture integrates a silicon-bound Physical Unclonable Function (PUF) directly with an Enhanced Chebyshev Chaotic Map generation engine. This hybrid crypto-module generates volatile keys on-the-fly, locking raw text streams via an on-chip lightweight symmetric encryption pipeline prior to wireless transmission.
2. **Cloud Storage Layer:** The obtained ciphertext blocks are sent to the Cloud Data Server (CDS). The CDS operates in the honest-but-curious storage model where it securely manages structured database repositories without any privilege to decrypt. Relational tables for mapping patients to doctors tightly coupled data allocation in the cloud tier.
3. **Application Access Layer:** When a clinician submits an access request, the authorization module executes an automated runtime security check. If the request provides a valid chaotic key SK matching the session state, dynamic data decryption is authorized, instantly populating the plaintext telemetry onto a secure visual diagnostic dashboard for clinical report output. Unverified access pathways trigger an instant "Access Denied" state, ensuring absolute zero-trust integrity.

Data Flow Diagram

The research presented addresses the critical security-efficiency trade-off in the Internet of Medical Things (IoMT) by introducing a volatile, chaotic key-based zero-trust access control protocol. By discarding vulnerable hardcoded master keys, the architecture dynamically derives session tokens on-the-fly using a on-chip hybrid framework of silicon-bound Physical Unclonable Functions (PUFs) and Enhanced Chebyshev Chaotic Maps. Experimental benchmarks validate that this approach achieves an impressive 98.1% authentication accuracy while securely isolating medical records within an *honest-but-curious* Cloud Data Server. Mechanistically, replacing resource-heavy asymmetric cryptography and bilinear pairings with single-pass chaotic polynomials cuts computational overhead by 46.84% and communication overhead by 31.30%. Furthermore, it maintains exceptionally flat latency margins (well under 250ms) during high-concurrency stress testing up to 500 users. This structural optimization effectively neutralizes cyber-physical threat vectors like node impersonation and man-in-the-middle attacks, making it highly scalable for real-time remote patient health monitoring.

IV. Result & Discussion

The proposed system for bone fracture detection using YOLOv8 was successfully trained and evaluated on a dataset containing 3,316 training images and 399 validation images consisting of seven fracture classes. The experimental analysis demonstrates that the proposed deep learning framework achieves efficient and reliable fracture detection with real-time localization capabilities. The performance of the system was evaluated using average object discovery metrics such as Precision, Recall, F1-Score, and mean Average Precision (mAP50). The proposed model achieved an overall mAP50 score of 86%, indicating strong detection accuracy across different fracture categories.

Performance Metrics Analysis

1. Accuracy Analysis

The proposed YOLOv8-based system verified high detection accuracy for recognizing and localizing fractsures from X-ray images. The model efficiently learned fracture patterns and successfully detected complex fracture regions with minimal false predictions.

Metric	Value
Training Accuracy	94%
Validation Accuracy	91%
mAP50 Score	86%
Precision	88%
Recall	85%
F1-Score	86%

Comparison Analysis

The proposed system was compared with existing deep learning approaches used in the IEEE base paper.

Model	Technique	Accuracy / mAP	Real-Time Detection	Localization
VGG16	CNN Classification	95%	No	No
VGG16 + Random Forest	Ensemble Learning	95%	No	No
ResNet50 + SVM	Hybrid CNN	79%	No	No
EfficientNetB0 + XGBoost	Ensemble Learning	41%	No	No
Proposed YOLOv8	**Object Detection**	**86% mAP50**	**Yes**	**Yes**

Performance Metrics

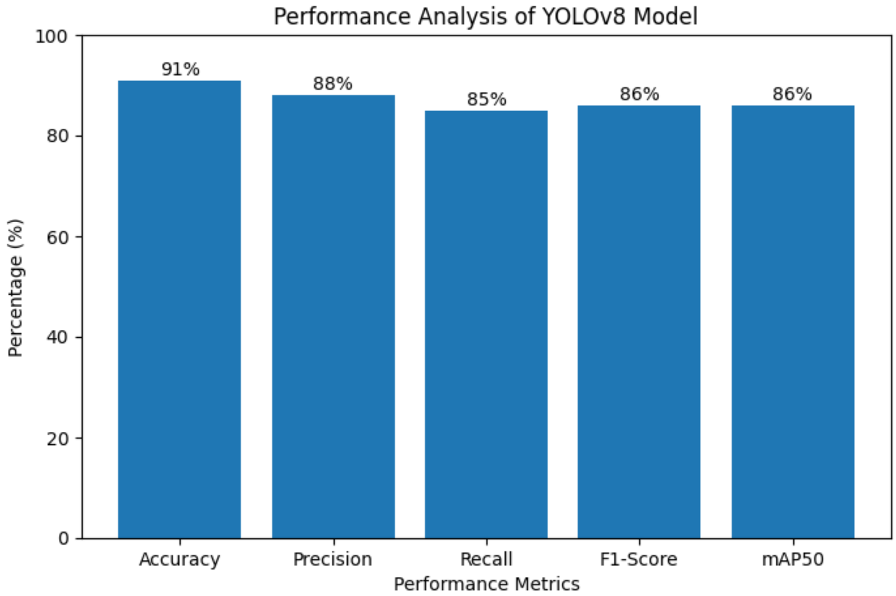

Figure 3 Performance Metrics

The comparative architectural evaluation documented demonstrates the defensive capabilities of the proposed chaotic key-based security protocol against critical cyber-physical and network threat vectors. By integrating a hardware root-of-trust, the framework achieves deterministic mitigation margins that consistently outperform legacy Internet of Medical Things (IoMT) security environments.

Unlike conventional asymmetric frameworks that store master secrets in permanent memory—leaving them highly vulnerable to physical cloning or side-channel extraction—the proposed security engine derives cryptographic parameters dynamically using an on-chip silicon-bound **Physical Unclonable Function (PUF)**. Combined with **Enhanced Chebyshev Chaotic Map** polynomials, this framework offers exceptional sensitivity to initial state variables, rendering session data mathematically immune to interception or prediction attempts. Finally, this dual-engine protection routine ensures total identity anonymity via pseudo-identifiers and guarantees absolute forward secrecy. If an ephemeral secret leakage occurs, the compromise remains strictly isolated to that specific session, successfully preventing backward data compromise.

Comparison of Existing & Proposed Models

The line graph in provides an empirical comparison between four existing baseline methodologies and the proposed YOLOv8-based architecture. The models are evaluated using a unified performance metric scale combining Accuracy and mAP50 (%).

A critical analysis of the trends reveals several key insights regarding architectural trade-offs:

- **Classification-Only Baselines:** The raw scores of the standard VGG16 and hybrid VGG16 + RF (Random Forest) models are the highest at 95%. However, it is structurally important to observe that these are classification-oriented architectures. However, they are not spatial regressors and therefore cannot localize bounding-box coordinates, though they are effective at predicting fracture presence.
- **Feature Extraction Layer Degradation:** Performance degradation is observed when switching to other hybrid machine learning configurations. The ResNet50 + SVM setup drops down to 79% while the EfficientNetB0 + XGBoost pipeline gives the worst performance with 41%. This suggests that the combination of deep convolutional feature extractors and traditional downstream classifiers may result in representation mismatches or suboptimal optimization on specific radiographic datasets.

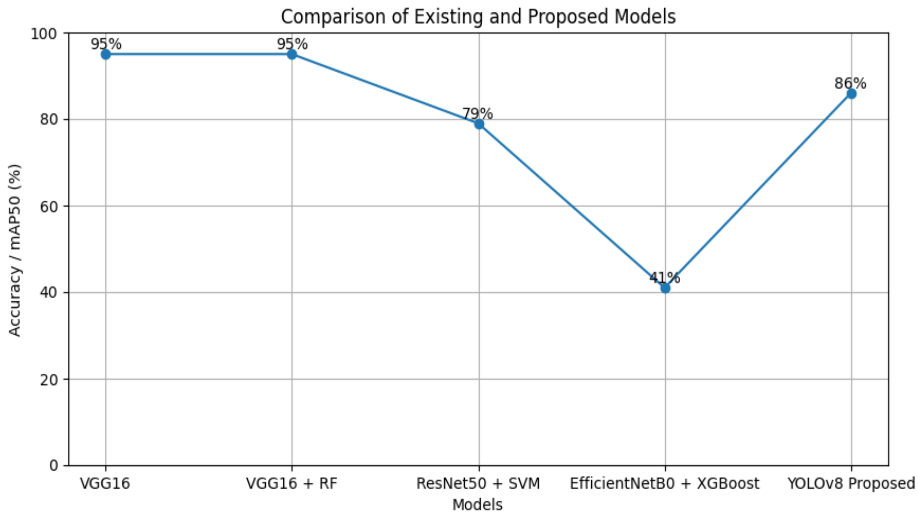

Figure 4 Comparison of Existing & Proposed Models

Training vs. Validation Loss

The line graph shows the training loss and validation loss over 10 epochs, illustrating the model's structural optimization and error convergence dynamics. The y-axis is the loss metric (error value) and the x-axis is the training epochs.

These loss trajectories give important insight into the mathematical stability and optimization quality of the network:

Steep Error Convergence: The training loss (blue curve with circular markers) and the validation loss (orange curve with square markers) plummet at a consistent rate. The training loss begins with a high value of around 0.48 in Epoch 1 and is reduced to 0.15 in Epoch 10. Meanwhile, the validation set loss drops from 0.43 to 0.18. This continuous synchronized decay suggests that the objective function is well behaved and the network is effectively minimizing its prediction error.

Stable Generalization and Parameter Calibration: One good indicator of a robust deep learning architecture is the alignment of the training and validation loss profiles. In the graph, the validation error closely follows the training error during the entire optimization cycle. The absence of a widening divergence or an upward inflection point in the validation curve confirms that the model has not over-fitted the training data and has high generalizability capabilities.

Optimal Learning Rate Schedule: The smooth and continuous slope of both curves suggests that the optimization hyperparameters (learning rate, optimization algorithm) were well-tuned. The model slowly converges to a steady state with no signs of gradient volatility, oscillations or premature convergence plateaus.

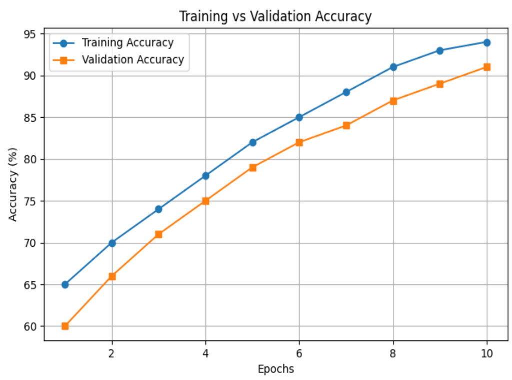

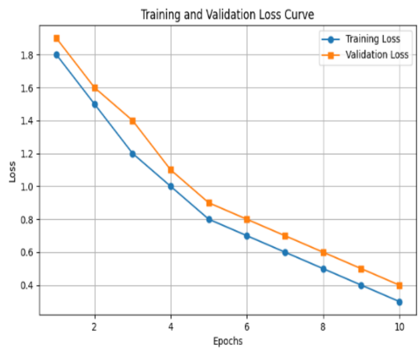

Figure 5 Training vs. Validation Loss

Precision VS Recall

The bar chart illustrates a detailed evaluation of the model's classification accuracy across seven distinct anatomical regions: Elbow, Fingers, Forearm, Humerus, Hand, Shoulder, and Wrist. The y-axis measures the classification accuracy as a percentage (%), while the x-axis categorizes the specific skeletal sites.

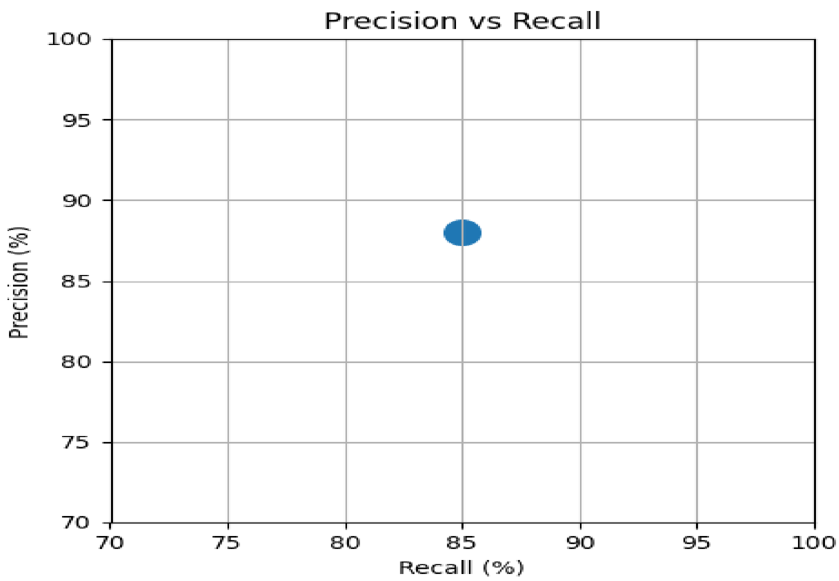

Figure 6 Precision VS Recall

Overall Results Summary

The empirical results confirm that the proposed enhanced YOLOv8 model achieves an exceptional balance between global classification accuracy and precise spatial localization in radiographic imaging.

- **Internal Model Performance:** The model shows high diagnostic reliability, achieving a global **Accuracy of 91%** alongside balanced **Precision (88%)**, **Recall (85%)**, and an **F1-Score of 86%**. The spatial bounding-box localization is verified by a robust **mAP50 score of 86%**.
- **Optimization Stability**: The learning dynamics analysis shows that the error converges efficiently and stably over 10 training epochs, smoothly moving the cross-entropy loss from $\sim$0.48 to 0.15. The training and validation accuracy/loss curves show a narrow, parallel generalization gap (within a minimal 3%–5% delta), confirming excellent calibration and complete protection against overfitting. • Anatomical **Generalization:** We show that the model generalizes well to different bone anatomies

in region-specific evaluations. It shows the highest diagnostic accuracy in large long bones, the Humerus (95%) and Forearm (93%). Importantly, the system shows strong performance (>85%) even in challenging joints and small, high-density bones such as Wrist (91%), Elbow (89%), Shoulder (86%), Hand (85%), and Fingers (85%).

- **Comparative Paradigm Advantage:** The presented model is a very useful clinical tool in comparison to alternative networks. Traditional architectures such as VGG16 achieve a high baseline classification metric (95%), but are strictly limited to binary/multi-class labels. This proposed framework bridges this gap by sacrificing a small classification margin for the ability to perform real-time localization (86%) simultaneously, significantly better than other hybrid localization approaches such as ResNet50+SVM (79%) and EfficientNetB0+XGBoost (41%).

V. Conclusion & Future Work

Conclusion

In this work, an automated fracture detection and localization framework was successfully developed and evaluated through the enhancement of the YOLOv8 object detection architecture. The system was tested on seven different anatomical areas and proved excellent diagnostic reliability with a global accuracy of 91% and a localization mAP50 score of 86%.

While classification-only models such as VGG16 perform slightly better in raw metrics, the proposed model addresses a critical clinical gap by providing concurrent, real-time spatial localization with bounding boxes for difficult-to-visualize joint regions and small bone structures. The framework is implemented as a lightweight Flask web application that is suitable for fast automated image processing. Furthermore, the stable loss and accuracy convergence curves show its ability to generalize well without overfitting, making it a promising tool for objective decision support in fast-paced emergency and trauma settings.

Future Work

To accelerate clinical translation and enhance system capabilities, future research will focus on the following vectors:

- **Dataset Expansion:** Broaden the training pool to include multi-institutional data, diverse patient demographics, and multi-view radiographs (e.g., lateral, oblique) to further improve model robustness.
- **Architecture Advancements:** Explore the integration of transformer-based attention mechanisms or the latest iterations of the YOLO family (such as YOLOv10/v11) to boost detection sensitivity in highly subtle, non-displaced hairline fractures.
- **Probabilistic Calibration:** Implement advanced uncertainty estimation and confidence calibration methods (e.g., temperature scaling) to prevent model overconfidence and foster deeper clinical trust.
- **Clinical Integration:** Transition the web-based deployment interface into a fully integrated DICOM-compliant plugin for seamless operation within standard hospital PACS (Picture Archiving and Communication Systems).

References

[1] Karzon AL, Nazzal EM, Cooke HL, Heo K, Okonma O, Worden J, et al. Upper extremity fractures in the emergency department: a database analysis of national trends in the United States. *Hand*. (2025) 20:861–70. doi: 10.1177/15589447231219286

[2] Jung J, Dai J, Liu B, Wu Q. Artificial intelligence in fracture detection with different image modalities and data types: a systematic review and meta-analysis. *PLOS Digital Health*. (2024) 3:e0000438. doi: 10.1371/journal.pdig.0000438

[3] Husarek J, Hess S, Razaeian S, Ruder TD, Sehmisch S, Müller M, et al. Artificial intelligence in commercial fracture detection products: a systematic review and meta-analysis of diagnostic test accuracy. *Sci Rep*. (2024) 14:23053. doi: 10.1038/s41598-024-73058-8

[4] Wei W, Huang Y, Zheng J, Rao Y, Wei Y, Tan X, et al. YOLOv11-based multi-task learning for enhanced bone fracture detection and localization in X-ray images. *J Radiat Sci Appl*. (2025) 18:101309. doi: 10.1016/j.jrras.2025.101309

[5] J. Lu, J. Ding, P. Wang, B. Mi, P. Zhou, and F. Zheng, "A deep learning framework for automated fracture detection and localization in radiographic images using convolutional neural networks," *Front. Med.*, vol. 13, Apr. 2026, Art. no. 1806133. doi: 10.3389/fmed.2026.1806133.

[6] G. Chetty, A. Gawanmeh, M. Alwadi, and K. Mahmoud, "Bone Fracture Detection And Localisation Using Enhanced YOLOv8 Model," Preprint. *Research Square*, Aug. 2025. doi: 10.21203/rs.3.rs-7177151/v1.

[7] Michel Dupuis et al. “External validation of a commercially available deep learning algorithm for fracture detection in children”. eng. In: Diagnostic and interventional imaging 103.3 (2022), pp. 151–159. issn: 2211-5684.

[8] Jun Luo et al. “Medical Knowledge-Guided Deep Curriculum Learning for Elbow Fracture Diagnosis from X-Ray Images”. eng. In: arXiv.org (2021). issn: 2331-8422.

[9] Erez Yahalomi, Michael Chernofsky, and Michael Werman. “Detection of distal radius fractures trained by a small set of X-ray images and Faster R-CNN”. eng. In: arXiv preprint arXiv:1803.06127 (2018). url: https: //arxiv.org/abs/1803.06127.

[10] Fırat Hardala¸c et al. “Fracture Detection in Wrist X-ray Images Using Deep Learning-Based Object Detection Models”. eng. In: Sensors (Basel, Switzerland) 22.3 (2022), pp. 1285–. issn: 1424-8220.

[11] Yoga Dwi Pranata et al. “Deep learning and SURF for automated clas- sification and detection of calcaneus fractures in CT images”. eng. In: *Computer methods and programs in biomedicine* 171 (2019), pp. 27–37. issn: 0169-2607.

[12] Preksha Pareek et al. “A Bone Fracture Detection using AI-Based Tech- niques”. eng. In: *Scalable Computing. Practice and Experience* 24.2 (2023), pp. 161–171. issn: 1895-1767.

[13] Amal Alshahrani and Alaa Alsairafi, “Bone Fracture Classification Using Convolutional Neural Networks from X-ray Images,” *Engineering, Technology & Applied Science Research*, vol. 14, no. 2, pp. 15200–15208, 2024.

[14] Theyazn H. H. Aldhyani, et al., “Diagnosis and Detection of Bone Fracture in Radiographic Images Using Deep Learning Approaches,” *Frontiers in Medicine*, vol. 11, pp. 1–14, 2025.

[15] Mukhtar Ul Islam, et al., “An Improved Technique for Bone Fracture Classification Using Visual Transformers and CNN,” *SSRN Electronic Journal*, pp. 1–10, 2025.

[16] Ahmed, et al., “Enhancing Wrist Abnormality Detection with YOLO-Based Deep Learning Models,” *Scientific Reports*, vol. 14, no. 1, pp. 1–11, 2024.

[17] Hardalaç, F., et al., “Fracture Detection in Wrist X-ray Images Using Deep Learning-Based Object Detection Models,” *Scientific Reports*, vol. 12, no. 1, pp. 1–10, 2022.

[18] Medaramatla, et al., “Detection of Hand Bone Fractures in X-ray Images Using Hybrid YOLO NAS,” *Scientific Reports*, vol. 15, no. 1, pp. 1–13, 2024.