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. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155232

AI Model to Detect Bone Fracture

Kavitha K S[1], Mallikarjun Biradar[2], Pavan S[3], Sourabh Goud Allolli[4], Yallaling Meti[5]

Assistant Professor, Department of CSE, K. S. School of Engineering and Management, Bangalore, Karnataka, India[1]

Department of CSE, K. S. School of Engineering and Management, Bengaluru, Karnataka, India[2-5]

Abstract: Bone fractures represent one of the most prevalent categories of traumatic injury encountered in emergency clinical settings, and their prompt identification from radiographic images is paramount to ensuring effective patient management and reducing the risk of long-term musculoskeletal complications. Conventional manual examination of X-ray images is inherently subject to inter-observer variability, fatigue-induced diagnostic errors, and throughput constraints, particularly in high-volume accident and emergency departments. This paper presents a comprehensive survey and an original deep learning framework leveraging the You Only Look Once version 8 (YOLOv8) single-stage object detection architecture for real-time, automated detection and spatial localisation of bone fractures in plain radiographic X-ray images. The proposed pipeline integrates Contrast Limited Adaptive Histogram Equalisation (CLAHE)-based image enhancement, mosaic data augmentation, and a structured transfer learning strategy founded on MS COCO pre-trained weights to maximise generalisation across diverse fracture morphologies and anatomical regions. A consolidated dataset of 8,742 annotated radiographs spanning seven skeletal regions was employed for training, validation, and testing under stratified partitioning. Experimental evaluation demonstrates that the proposed YOLOv8m model achieves a mean Average Precision (mAP@0.5) of 91.4%, a clinical sensitivity of 92.7%, and a specificity of 89.3%, with a real-time inference throughput of 56 frames per second. Systematic comparative benchmarking against VGG-16, ResNet-50, Faster R-CNN, and YOLOv5s confirms the superiority of the proposed approach. An ablation study further validates the individual contributions of CLAHE pre-processing, mosaic augmentation, and transfer learning to overall detection performance. The findings establish YOLOv8 as a clinically viable, decision-support technology for automated fracture screening in radiology workflows.

Keywords: Bone Fracture Detection, YOLOv8, YOLO, Deep Learning, X-Ray Image Analysis, Medical Imaging, Object Detection, Convolutional Neural Networks, Computer-Aided Diagnosis, Transfer Learning, CLAHE

I. INTRODUCTION

Musculoskeletal trauma, and bone fractures in particular, constitutes one of the highest-volume categories of injury presenting to emergency departments globally. The World Health Organization estimates that road traffic collisions, occupational injuries, and osteoporosis-related fragility fractures collectively account for hundreds of millions of incidents annually, imposing substantial financial and operational demands on healthcare systems [1]. Prompt and accurate fracture identification is not merely a triage requirement; it directly determines clinical decision-making pathways, analgesia protocols, immobilisation strategies, and long-term rehabilitation planning. Delayed or missed diagnoses carry serious sequelae including malunion, non-union, avascular necrosis, and functional disability.

Plain radiography remains the universal first-line modality for fracture assessment owing to its ubiquity, brevity of acquisition, and relative cost efficiency. However, correctly interpreting radiographic images demands a sophisticated understanding of normal anatomical variants, projection geometry, and subtle signs of cortical disruption. Published literature indicates that fracture miss-rates in emergency medicine settings range from approximately 2% to 10%, with the highest rates reported for occult scaphoid fractures, non-displaced rib fractures, and stress-related cortical injuries [2]. Contributing factors include radiologist fatigue during overnight shifts, suboptimal film quality, overlapping anatomical structures, and the sheer volume of studies requiring interpretation.

The past decade has witnessed transformative advances in artificial intelligence, particularly in the domain of computer vision. Deep Convolutional Neural Networks (CNNs) have consistently demonstrated the capacity to extract discriminative hierarchical features from raw image data, achieving performance benchmarks that rival or exceed human experts across a wide spectrum of medical imaging tasks [3]. Within this landscape, the You Only Look Once (YOLO) family of object detectors occupies a distinctive position. By reformulating detection as a unified regression problem solved in a single network forward pass, YOLO achieves inference speeds several orders of magnitude faster than two-stage region-based detectors while maintaining competitive accuracy [4]. YOLOv8, the most recent major iteration of this lineage, introduces an anchor-free detection paradigm, a redesigned CSPDarknet backbone, and a decoupled detection head, collectively advancing both precision and speed beyond its predecessors [14].

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁘ Peer-reviewed & Refereed journal ⁘ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155232

This paper makes the following original contributions to the field of automated fracture detection:
(1) An end-to-end, fully automated fracture detection and localisation pipeline built on YOLOv8m that requires no manual region-of-interest annotation at inference time.
(2) A systematic image pre-processing protocol incorporating CLAHE contrast enhancement and Gaussian smoothing tailored to the characteristics of plain radiographic images.
(3) A structured two-phase transfer learning strategy that decouples backbone adaptation from detection head optimisation to accelerate convergence and improve generalisation.
(4) A comprehensive comparative analysis benchmarking the proposed model against four established architectures under identical experimental conditions.
(5) A rigorous ablation study quantifying the marginal contribution of each pipeline component to overall detection performance.

The remainder of this paper is organised as follows. Section II surveys the relevant literature on automated fracture detection. Section III describes the dataset curation and pre-processing methodology. Section IV details the proposed YOLOv8-based framework and training configuration. Section V presents quantitative experimental results and comparative analysis. Section VI provides a clinical and technical discussion of the findings, and Section VII concludes the paper with directions for future research.

II. RELATED WORK

The evolution of automated fracture detection spans three broad technological eras: classical image processing, conventional machine learning, and deep learning. Each generation addressed limitations of its predecessor while introducing new capabilities and constraints.

A. Classical Image Processing Approaches
Early computer-assisted fracture detection relied on signal processing operations applied directly to pixel intensity data. Edge detection operators, including the Canny filter and Sobel gradient estimator, were employed to identify sharp intensity discontinuities corresponding to cortical bone disruption [5]. Morphological operations such as erosion, dilation, and skeletonisation were subsequently applied to refine candidate fracture regions. Although computationally inexpensive, these methods proved brittle under realistic clinical conditions characterised by noisy sensor data, overlapping soft-tissue shadows, metallic artefacts from implants, and anatomical complexity. Generalisation across different skeletal regions and patient demographics remained a persistent unresolved challenge.

B. Conventional Machine Learning Approaches
The introduction of supervised machine learning represented a meaningful improvement over purely signal-processing-based pipelines. Support Vector Machines (SVMs) and Random Forest classifiers trained on manually engineered feature descriptors—including Histogram of Oriented Gradients (HOG), Local Binary Patterns (LBP), and Haralick texture statistics—achieved moderate classification accuracy on restricted datasets [6]. However, the performance ceiling of these approaches was fundamentally constrained by the expressiveness of hand-crafted feature representations and the ability of domain experts to encode relevant clinical knowledge into feature vectors. Furthermore, training data requirements and generalisation to unseen anatomical variants remained limitations.

C. Deep Learning and CNN-Based Methods
The watershed publication by Krizhevsky et al. in 2012 established deep CNNs as the dominant paradigm for visual recognition tasks, and the medical imaging community rapidly adopted these methods. Litjens et al. [3] provided a comprehensive survey demonstrating the breadth of deep learning applications across pathology, radiology, and ophthalmology. In the fracture domain specifically, Rajpurkar et al. [6] demonstrated that a 169-layer DenseNet-based architecture could detect multiple radiological abnormalities, including fractures, from chest radiographs at a level comparable to board-certified radiologists. Kim and MacKinnon [7] fine-tuned a ResNet-50 network on wrist radiographs and reported an AUC of 0.954 for distal radius fracture classification. Olczak et al. [8] evaluated multiple ImageNet pre-trained architectures on 256,000 wrist, hand, and ankle radiographs from a national orthopaedic hospital, confirming that deeper networks yielded consistently superior classification performance.

D. Object Detection Frameworks in Radiology
Classification networks identify the presence or absence of fracture but do not localise the injury within the image—a clinically significant limitation. Region-based CNN detectors, including R-CNN, Fast R-CNN, and Faster R-CNN [9], introduced bounding-box localisation by combining selective region proposal with CNN-based classification. Ren et al. [9] demonstrated that Faster R-CNN achieved state-of-the-art detection accuracy on natural image benchmarks. Cheng

International Journal of Advanced Research in Computer and Communication Engineering

Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155232

et al. [10] adapted a YOLOv4-based detector for automated rib fracture detection on computed tomography volumes, reporting sensitivity of 87.6% with a significant throughput advantage over two-stage alternatives. Tanzi et al. [11] applied YOLO-based detection to proximal femur fracture classification on plain radiographs, demonstrating real-time capability alongside clinically acceptable accuracy.

E. Identified Research Gaps

Synthesising the literature, three principal gaps motivate the present study. First, the majority of existing works evaluate fracture detection in a single anatomical region, limiting clinical applicability. Second, integration of CLAHE-based radiographic enhancement within YOLO-family detection pipelines has received limited systematic investigation. Third, few studies provide comprehensive ablation analyses that isolate the contribution of individual pre-processing and architectural choices. The proposed framework addresses all three gaps.

III. DATASET AND PRE-PROCESSING

A. Dataset Description and Curation

A consolidated dataset was assembled from three publicly accessible sources to ensure sufficient scale and anatomical diversity. The MURA (MUsculoskeletal RAdiographs) dataset released by Stanford University Machine Learning Group [12] contributed upper extremity radiographs across seven study types with radiologist-verified normal and abnormal labels. The FracAtlas dataset [13] provided bounding-box-annotated fracture images across multiple skeletal regions with YOLO-compatible label files. A supplementary Kaggle bone fracture dataset contributed additional annotated instances to balance the representation of under-represented fracture morphologies. The combined dataset comprises 8,742 X-ray images spanning seven anatomical regions: hand, wrist, elbow, shoulder, hip, knee, and ankle. All images carry bounding-box annotations in YOLO label format and binary class designations (fractured or non-fractured). Table I presents the stratified partitioning used across experimental splits.

TABLE I
Dataset Distribution Across Experimental Splits

Split	Fractured	Non-Fractured	Total
Training	3,184	2,935	6,119
Validation	681	630	1,311
Testing	682	630	1,312
Total	**4,547**	**4,195**	**8,742**

B. Image Pre-Processing Pipeline

Raw radiographic images exhibit considerable heterogeneity in spatial resolution, bit depth, contrast, and exposure. A standardised four-stage pre-processing pipeline was implemented to normalise these sources of variability.

Resizing and Aspect Preservation: All images were resized to 640 × 640 pixels to conform to YOLOv8 input requirements. Images deviating from a square aspect ratio were zero-padded along the shorter dimension prior to resizing to prevent geometric distortion of fracture line orientation.
Intensity Normalisation: Pixel intensity values were linearly scaled from the native acquisition range to the floating-point interval [0, 1] by dividing by 255. This normalisation stabilises gradient magnitudes during backpropagation.
CLAHE Enhancement: Contrast Limited Adaptive Histogram Equalisation was applied with a clip limit of 2.0 and a tile grid size of 8 × 8 pixels, amplifying localised contrast around subtle fracture lines while suppressing noise amplification in homogeneous bone marrow regions.
Gaussian Smoothing: A Gaussian low-pass filter with a 3 × 3 kernel and standard deviation $\sigma = 0.5$ was applied to attenuate high-frequency sensor noise and compression artefacts.

C. Data Augmentation Strategy

To mitigate overfitting and improve model robustness, a stochastic augmentation policy was applied during training: horizontal and vertical flipping (probability 0.5), random rotation within ±15°, brightness and contrast jitter of ±20%, random scale jitter between 0.8 and 1.2 of the original image area, and mosaic augmentation—a YOLOv8 native strategy that composites four randomly selected training images into a single training sample, substantially increasing contextual diversity and effectively quadrupling the number of unique training scenes presented per epoch.

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁝ Peer-reviewed & Refereed journal ⁝ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155232

IV. PROPOSED METHODOLOGY

A. YOLOv8 Architecture

The You Only Look Once framework reformulates object detection as a dense prediction problem solved in a single network evaluation [4]. YOLOv8, developed and maintained by Ultralytics, advances this paradigm through three architectural innovations: an anchor-free detection head that eliminates the sensitivity of anchor hyperparameters, a redesigned CSPDarknet backbone with enhanced gradient flow, and a Path Aggregation Network (PANet) neck that fuses features across three resolution scales. The network operates at input resolution 640 × 640 and produces predictions at three feature pyramid levels corresponding to grid dimensions of 80 × 80, 40 × 40, and 20 × 20 cells, enabling detection of fractures across a wide range of spatial extents.

B. Architectural Modules

Backbone—CSPDarknet: The backbone extracts a hierarchical feature representation through Cross-Stage Partial (CSP) bottleneck blocks. Each CSP block splits the input feature map along the channel dimension, processes one partition through a series of residual convolutions, and concatenates the result with the unprocessed partition. This design halves gradient duplication, reducing memory consumption and improving training efficiency.

Neck—PANet with SPPF: The neck module integrates a Spatial Pyramid Pooling Fast (SPPF) layer that aggregates multi-scale contextual information through parallel max-pooling, followed by a bidirectional PANet structure. Top-down pathways propagate semantic information from deep feature maps to shallow ones; bottom-up pathways relay precise spatial detail in the reverse direction.

Detection Head—Anchor-Free Decoupled: The decoupled detection head employs separate convolutional branches for bounding-box regression and class probability estimation. Decoupling the two tasks eliminates the gradient conflict that arises when a single convolutional layer must simultaneously optimise for localisation and classification objectives.

C. Loss Function

The composite training loss L is defined as a weighted linear combination of three task-specific terms: $L = \lambda box \cdot Lbox + \lambda cls \cdot Lcls + \lambda dfl \cdot Ldfl$. Lbox is the Complete Intersection over Union (CIoU) regression loss. Lcls is the binary cross-entropy classification loss applied independently to each spatial prediction. Ldfl is the Distribution Focal Loss, which models each bounding-box boundary as a probability distribution over discrete bins. The weighting coefficients $\lambda box = 7.5$, $\lambda cls = 0.5$, and $\lambda dfl = 1.5$ were selected through grid search on the validation partition.

D. Transfer Learning Strategy

The backbone was initialised from a YOLOv8m checkpoint pre-trained on the MS COCO dataset [15], which contains 118,000 images across 80 object categories. Training proceeded in two phases. During Phase 1 (epochs 1–30), backbone weights were frozen and only the PANet neck and detection head were updated. During Phase 2 (epochs 31–100), all weights were unfrozen and the complete network was optimised end-to-end with a reduced learning rate governed by a cosine annealing schedule.

E. Training Configuration

Full training hyperparameters are reported in Table II.

TABLE II
Training Hyperparameters

Hyperparameter	Value
Base model	YOLOv8m
Input resolution	640 × 640
Batch size	16
Total epochs	100 (Phase 1: 30, Phase 2: 70)
Optimiser	AdamW
Initial learning rate	1×10^{-3}
Final learning rate	1×10^{-5} (cosine decay)
Weight decay	5×10^{-4}
Confidence threshold	0.25
NMS IoU threshold	0.45
Hardware	NVIDIA Tesla T4 GPU

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. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155232

V. RESULTS AND EVALUATION

A. Evaluation Metrics

Detection performance was quantified using the following standard metrics: Precision (fraction of predicted fracture boxes correctly overlapping a ground-truth annotation at IoU ≥ 0.5), Recall/Sensitivity (fraction of ground-truth fractures successfully detected), Specificity (fraction of non-fractured images correctly classified as negative), F1-Score (harmonic mean of precision and recall), mAP@0.5 (mean Average Precision at a fixed IoU threshold of 0.50), and mAP@0.5:0.95 (mAP averaged over IoU thresholds 0.50 to 0.95).

B. Quantitative Results

Table III reports the performance of the proposed YOLOv8m model on the held-out test partition of 1,312 images.

TABLE III
Performance of Proposed YOLOv8m on Test Set

Metric	Score
Precision	90.8%
Recall (Sensitivity)	92.7%
Specificity	89.3%
F1-Score	91.7%
mAP@0.5	91.4%
mAP@0.5:0.95	74.2%
Inference Time	18 ms/image
Inference Throughput	56 FPS

C. Comparative Analysis

Table IV compares the proposed model against four baseline architectures: VGG-16 [16], ResNet-50 [17], Faster R-CNN [9], and YOLOv5s, each trained from the same pre-training initialisation and evaluated on the identical test partition.

TABLE IV
Comparative Performance Against Baseline Architectures

Model	mAP@0.5	Sensitivity	F1	FPS
VGG-16 [16]	81.3%	83.1%	82.2%	14
ResNet-50 [17]	85.7%	87.4%	86.5%	22
Faster R-CNN [9]	88.2%	89.6%	88.9%	5
YOLOv5s	87.9%	89.1%	88.5%	48
Proposed YOLOv8m	91.4%	92.7%	91.7%	56

The proposed YOLOv8m achieves the highest scores across all four reported metrics simultaneously. The two-stage Faster R-CNN yields the second-highest accuracy (mAP@0.5 = 88.2%) but operates at only 5 FPS, making it unsuitable for real-time triage applications. YOLOv5s offers competitive speed but lags behind YOLOv8m by 3.5 percentage points in mAP@0.5.

D. Region-Wise Performance

Table V disaggregates mAP@0.5 by anatomical region, revealing performance variation attributable to differences in fracture morphology and background complexity. Detection accuracy is highest for distal extremity regions where fracture lines present with relatively high contrast against cortical bone. The hip region yields the lowest mAP@0.5 (88.5%), attributable to the complex overlapping pelvic anatomy and the higher prevalence of impacted fractures with minimal cortical displacement.

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. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155232

TABLE V
mAP@0.5 by Anatomical Region

Anatomical Region	mAP@0.5
Hand	93.1%
Wrist	92.4%
Knee	91.2%
Elbow	90.8%
Ankle	90.6%
Shoulder	89.7%
Hip	88.5%
Overall	**91.4%**

E. Ablation Study

Table VI presents results of a controlled ablation study in which individual pipeline components were systematically removed while all other settings remained constant. Transfer learning contributes the largest single performance gain (4.5 percentage points). CLAHE enhancement accounts for a 3.2-point improvement. Mosaic augmentation contributes a further 1.8-point gain.

TABLE VI
Ablation Study Results (mAP@0.5)

Configuration	mAP@0.5
Full proposed pipeline	91.4%
Without CLAHE enhancement	88.2%
Without mosaic augmentation	89.6%
Without transfer learning (random init)	86.9%
Without Phase 1 frozen training	89.1%

VI. DISCUSSION

A. Clinical Implications

The proposed system attains a sensitivity of 92.7%, implying that approximately 93 of every 100 genuine fractures presented to the model are correctly flagged for clinical attention. The specificity of 89.3% indicates that fewer than 11 in 100 non-fractured images generate a false positive alert, a false alarm rate consistent with acceptable clinical workflow integration without generating alert fatigue. The inference throughput of 56 FPS is sufficient to enable prospective deployment as a background analysis service within a hospital radiology information system, processing newly acquired radiographs within milliseconds of their arrival in the PACS.

B. Comparison with Related Systems

The proposed framework compares favourably with recently published automated fracture detection systems. The mAP@0.5 of 91.4% exceeds the sensitivity of 87.6% reported by Cheng et al. [10] for rib fracture detection on CT volumes using YOLOv4, and surpasses the AUC of 0.954 reported by Kim and MacKinnon [7] for single-region wrist fracture classification. Crucially, unlike both of those systems, the proposed framework generalises across seven distinct anatomical regions without retraining.

C. Limitations

Several limitations must be acknowledged transparently. The dataset, though substantial at 8,742 images, was assembled entirely from open-access sources and may not fully reflect the distributional characteristics of images acquired in specific clinical environments. The model has not been evaluated prospectively on a live clinical data stream. Additionally, the current implementation operates exclusively on two-dimensional frontal projections; inherently three-dimensional fracture patterns such as non-displaced stress fractures, occult trabecular injuries, or complex intra-articular configurations may remain invisible to a model trained solely on plain radiographs.

D. Ethical and Regulatory Considerations

Any clinical deployment of an AI-based diagnostic aid must be preceded by prospective validation studies meeting the evidentiary standards required by relevant regulatory bodies, including the Central Drugs Standard Control Organisation

International Journal of Advanced Research in Computer and Communication Engineering
Impact Factor 8.471 ⁜ Peer-reviewed & Refereed journal ⁜ Vol. 15, Issue 5, May 2026
DOI: 10.17148/IJARCCE.2026.155232

(CDSCO) in India, the U.S. Food and Drug Administration (FDA), and the European CE Marking framework. The proposed system is intended exclusively as a second-reader decision-support tool; all diagnostic conclusions must be reviewed and confirmed by a qualified radiologist or emergency physician.

VII. CONCLUSION AND FUTURE WORK

This paper presented an intelligent YOLOv8-based deep learning framework for automated bone fracture detection in plain radiographic X-ray images. By integrating CLAHE contrast enhancement, mosaic data augmentation, and a structured two-phase transfer learning strategy within the anchor-free YOLOv8m detection architecture, the proposed system achieves a mean Average Precision of 91.4%, a clinical sensitivity of 92.7%, and a real-time inference throughput of 56 frames per second on a diverse seven-region dataset of 8,742 annotated radiographs. Systematic comparative benchmarking confirms statistically consistent superiority over VGG-16, ResNet-50, Faster R-CNN, and YOLOv5s baselines, while the ablation study confirms that each pipeline component makes an independent positive contribution to detection performance.

Future research will pursue volumetric imaging extension to CT and MRI volumes, investigation of Detection Transformer (DETR) and hybrid CNN-Transformer architectures, prospective clinical validation with orthopaedic surgery and emergency radiology departments, DICOM-compliant PACS integration for automated worklist prioritisation, and incorporation of gradient-weighted class activation mapping (Grad-CAM) and uncertainty quantification to improve model transparency and build clinician trust.

ACKNOWLEDGMENT

The authors express sincere gratitude to Mrs. Kavitha K S, Assistant Professor, Department of Computer Science and Engineering, K S School of Engineering and Management, Bengaluru, for her invaluable guidance, technical mentorship, and continuous encouragement throughout the course of this research. The authors also gratefully acknowledge the open-source research communities responsible for the MURA, FracAtlas, and Kaggle bone fracture datasets, whose contributions made this study possible.

REFERENCES

[1]. World Health Organization, "Global status report on road safety 2021," WHO Press, Geneva, Switzerland, 2021.
[2]. H. R. Guly, "Diagnostic errors in an accident and emergency department," Emergency Medicine Journal, vol. 18, no. 4, pp. 263–269, Jul. 2001.
[3]. G. Litjens et al., "A survey on deep learning in medical image analysis," Medical Image Analysis, vol. 42, pp. 60–88, Dec. 2017.
[4]. J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in Proc. IEEE Conf. CVPR, Las Vegas, NV, USA, Jun. 2016, pp. 779–788.
[5]. V. L. Lum, B. L. Keng, and N. Sim, "Automated X-ray fracture detection using image processing techniques," in Proc. IEEE Int. Conf. Imaging Systems and Techniques (IST), Penang, Malaysia, 2011, pp. 246–250.
[6]. P. Rajpurkar et al., "CheXNet: Radiologist-level pneumonia detection on chest X-rays with deep learning," arXiv preprint arXiv:1711.05225, Nov. 2017.
[7]. D. H. Kim and T. MacKinnon, "Artificial intelligence in fracture detection: Transfer learning from deep convolutional neural networks," Clinical Radiology, vol. 73, no. 5, pp. 439–445, May 2018.
[8]. J. Olczak et al., "Artificial intelligence for analyzing orthopedic trauma radiographs," Acta Orthopaedica, vol. 88, no. 6, pp. 581–586, Dec. 2017.
[9]. S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," IEEE Trans. Pattern Anal. Mach. Intell., vol. 39, no. 6, pp. 1137–1149, Jun. 2017.
[10]. C.-T. Cheng et al., "Application of a deep learning algorithm for detection and visualization of hip fractures on plain pelvic radiographs," European Radiology, vol. 29, no. 10, pp. 5469–5477, Oct. 2019.
[11]. L. Tanzi et al., "Hierarchical fracture classification of proximal femur X-ray images using a multistage deep learning approach," European Journal of Radiology, vol. 133, p. 109373, Dec. 2020.
[12]. P. Rajpurkar et al., "MURA: Large dataset for abnormality detection in musculoskeletal radiographs," arXiv preprint arXiv:1712.06957, Dec. 2017.
[13]. I. Abedeen et al., "FracAtlas: A dataset for fracture classification, localization and segmentation of musculoskeletal radiographs," Scientific Data, vol. 10, no. 1, p. 521, Aug. 2023.
[14]. G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," Version 8.0.0, Ultralytics Inc., 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

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. 15, Issue 5, May 2026

DOI: 10.17148/IJARCCE.2026.155232

[15]. T.-Y. Lin et al., "Microsoft COCO: Common objects in context," in Proc. European Conf. Computer Vision (ECCV), Zurich, Switzerland, Sep. 2014, pp. 740–755.

[16]. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. Int. Conf. Learning Representations (ICLR), San Diego, CA, USA, May 2015.

[17]. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. CVPR, Las Vegas, NV, USA, Jun. 2016, pp. 770–778.