

YOLOv9 for Fracture Detection in Pediatric Wrist Trauma X-ray Images

Chun-Tse Chien,¹ Rui-Yang Ju,² Kuang-Yi Chou,³ and Jen-Shiun Chiang^{✉, 1}

¹Department of Electrical and Computer Engineering, Tamkang University, New Taipei City, 251301, Taiwan

²Graduate Institute of Networking and Multimedia, National Taiwan University, Taipei City, 106335, Taiwan

³School of Nursing, National Taipei University of Nursing and Health Sciences, Taipei City, 112303, Taiwan

✉Email: jsken.chiang@gmail.com

The introduction of YOLOv9, the latest version of the You Only Look Once (YOLO) series, has led to its widespread adoption across various scenarios. This paper is the first to apply the YOLOv9 algorithm model to the fracture detection task as computer-assisted diagnosis (CAD) to help radiologists and surgeons to interpret X-ray images. Specifically, this paper trained the model on the GRAZPEDWRI-DX dataset and extended the training set using data augmentation techniques to improve the model performance. Experimental results demonstrate that compared to the mAP 50-95 of the current state-of-the-art (SOTA) model, the YOLOv9 model increased the value from 42.16% to 43.73%, with an improvement of 3.7%. The implementation code is publicly available at <https://github.com/RuiyangJu/YOLOv9-Fracture-Detection>.

Introduction: Computer-assisted diagnosis (CAD) helps specialists such as radiologists and surgeons to interpret medical images, including magnetic resonance imaging (MRI), computed tomography (CT), and X-ray images. The application of deep learning techniques to medical images [1–4] has yielded increasingly satisfactory results, making it a popular research focus, especially in fracture detection [5–7].

You Only Look Once (YOLO) series [8–16] are the main neural networks for real-time object detection task, widely employed in fracture detection [17–19]. Wrist fractures in children are more common cases and the GRAZPEDWRI-DX dataset [20] provides 20,327 X-ray images of pediatric wrist trauma that can be used in fracture detection tasks. Research [21] first used the YOLOv8 [16] model for fracture detection on this dataset. Since attention mechanisms [22–25] have excellent results in enhancing the performance of neural network models, Chien *et al.* achieved the state-of-the-art (SOTA) performance by incorporating different attention mechanisms into the YOLOv8 model.

With the presentation of YOLOv9 [26], which achieved remarkable model performance on the MS COCO 2017 [27] benchmark dataset, this paper first trained the YOLOv9 model on the GRAZPEDWRI-DX dataset and obtained the SOTA performance, as shown in Fig 1.

The main contributions of this paper are as follows:

1. This paper is the first to apply YOLOv9 to the fracture detection task, demonstrating that the model not only has the excellent performance in real-time object detection across real-life scenarios, but also has good results in medical image recognition.
2. This paper addresses the issue of information loss in fracture detection on X-ray images by employing YOLOv9 algorithm, and aims to preserve more information during model training on low-features X-ray images, enhancing the performance of the model.
3. The mAP 50-95 of YOLOv9 model trained on the GRAZPEDWRI-DX dataset is significantly improved, achieving the SOTA level.

Related Works: In the field of object detection task, detectors typically employ either one-stage or two-stage algorithms. Compared to two-stage object detectors, the models of YOLO series offer a more balanced combination of accuracy and inference speed, making them suitable for deployment on mobile computing platforms for medical image recognition. Son *et al.* [28] utilized YOLOv4 [9] and U-Net [29] as auxiliary diagnostic tools to assist dentists in identifying mandibular fractures without resorting to cone beam computed tomography (CBCT). Jeon *et al.* [30] employed YOLOv4 [9] to aid surgeons in diagnosing trauma by detecting the fracture and mapping it onto a 3D reconstructed




Fig 1 Comparisons of fracture detection models on the GRAZPEDWRI-DX dataset. In terms of accuracy, our models outperform all previous models at the state-of-the-art level.

bone image, providing a clear display of the fracture region through the red mask overlaid on the 3D bone image. Hrzić *et al.* [18] employed the YOLOv4 [9] model for fracture detection on the GRAZPEDWRI-DX dataset [20], which was the first to demonstrate that the models of YOLO series can assist radiologists in more accurately predicting wrist injuries in children on X-ray images. Ahmed *et al.* [31] demonstrated the potential of one-stage algorithm models in enhancing the accuracy of diagnosis on pediatric wrist X-ray images by employing YOLOv5 [12], YOLOv6 [13], YOLOv7 [15], and YOLOv8 [16] models for wrist anomaly detection, respectively. Warin *et al.* [32] utilized the YOLOv5 [12] model to detect mandibular fractures in panoramic X-ray images, demonstrating the ability of the YOLOv5 model to recognize mandibular fractures at an expert level. Gaikwad *et al.* [33] applied the YOLOv5 [12] model to detect major and minor fractures of the C1 to C7 vertebrae, achieving an accuracy rate of 89%. Zou *et al.* [34] investigated various fracture morphologies throughout the body, including angle fractures, normal fractures, line fractures, and disoriented angle fractures. They integrated the YOLOv7 [15] model with the attention mechanism [35], achieving the superior performance on the FracAtlas [36] dataset. Samothai *et al.* [17] demonstrated that the YOLOX [10] model exhibits faster convergence speed and higher accuracy than YOLOR [11] by detecting the fracture region through the methods such as detector head decoupling, anchor-free, and augmentation strategies. They also showed that YOLOX can locate fractures even in low-featured X-ray images. Moon *et al.* [37] proposed a computer-aided facial bone fracture diagnosis (CA-FBFD) system based on the YOLOX model, effectively reducing the workload of doctors in diagnosing facial fractures in facial CT scans. While the application of the models of YOLO series to medical image recognition is a hot research topic, to date, no one has utilized YOLOv9 [26] for fracture detection.

Method:

YOLOv9: Neural networks often have the challenge of information loss since the input data undergoes multiple layers of feature extraction and spatial transformation, resulting in the loss of the original information. This issue is particularly pronounced in X-ray images, where the low-features present significant difficult in fracture detection tasks. Specifically, models trained on such low-featured images tend to perform poorly, and addressing the problem of information loss could substantially enhance the accuracy of model predictions. To address this, we utilizing the YOLOv9 algorithm, which leverages the Programmable Gradient Information (PGI) and the Generalized Efficient Layer Aggregation Network (GELAN) to more effectively extract key features.

Programmable Gradient Information: Programmable Gradient Information (PGI) is an auxiliary supervision framework designed to manage the propagation of gradient information across various semantic lev-

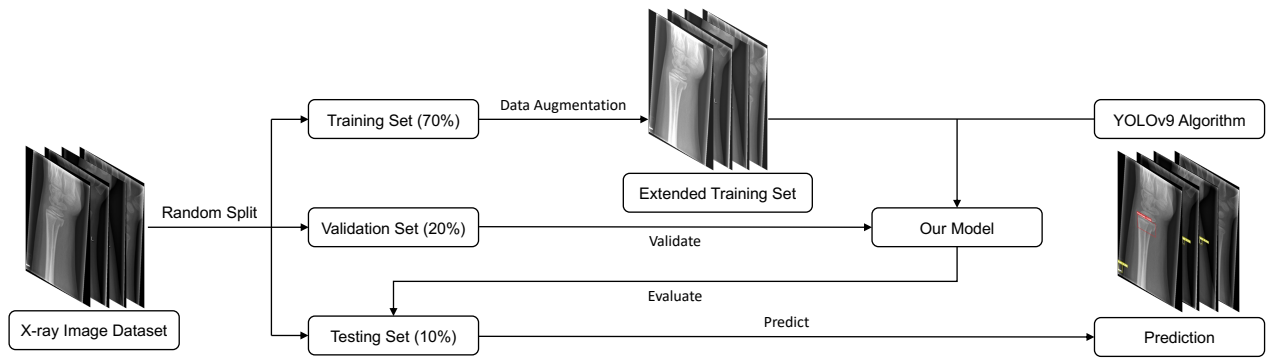


Fig 2 An overall flowchart of the YOLOv9 algorithm model applied to the fracture detection task.

Table 1. Quantitative comparison with other state-of-the-art models for fracture detection on the GRAZPEDWRI-DX datasets when the input image size is 640.

Model	Params (M)	FLOPs (G)	F1 (%)	mAP 50 (%)	mAP 50-95 (%)	Speed ^a (ms)
YOLOv8[16]	43.61	164.9	59	62.44	40.32	3.6
YOLOv8+SA[22]	43.64	165.4	62	63.99	41.49	3.9
YOLOv8+ECA[23]	43.64	165.5	61	62.64	40.21	3.6
YOLOv8+GAM[24]	49.29	183.5	60	63.32	40.74	8.7
YOLOv8+ResGAM[38]	49.29	183.5	62	63.97	41.18	9.4
YOLOv8+ResCBAM[25]	53.87	196.2	62	62.95	40.10	4.1
YOLOv9-C (Ours)	51.02	239.0	64	65.31	42.66	5.2
YOLOv9-E (Ours)	69.42	244.9	64	65.46	43.32	6.4

Note: The model size of all YOLOv8 and its variants listed in the table is large.

^a Speed is the total time for preprocessing, inference, and post-processing.

els, to improve the detection capability of the model. PGI comprises three main components: main branch, auxiliary reversible branch, and multi-level auxiliary information. During the inference process, it exclusively employs the main branch, which handles both forward and back propagation. As the network becomes deeper, an information bottleneck may occur, leading to loss functions that fail to produce useful gradients. In such cases, auxiliary reversible branch employs reversible functions to preserve information integrity and mitigate information loss in the main branch. Additionally, multi-level auxiliary information addresses the issue of error accumulation from the deep supervision mechanism, improving the learning capacity of the model through the introduction of supplementary information at different levels. Notably, research [26] highlighted the efficacy of PGI in preserving information during training, particularly in scenarios with limited features. This provides the theoretical basis for the YOLOv9 model to have excellent performance in fracture detection tasks.

Generalized Efficient Layer Aggregation Network: To enhance information integration and propagation efficiency in model training, YOLOv9 introduced a novel lightweight network architecture named Generalized Efficient Layer Aggregation Network (GELAN). GELAN integrates CSPNet [39] and ELAN [40] to efficiently aggregate network information, reducing information loss in propagation and enhancing inter-layer information interaction. This architecture is particularly suitable for fracture detection in environments with limited computing resources due to its lower parameters and computational complexity.

Data Processing and Augmentation: Fig 2 illustrates the flowchart of the experiments conducted in this study. Since the publisher of the GRAZPEDWRI-DX [20] dataset did not provide predefined training, validation, and test sets, we randomly assigned 70% to training set, 20% to validation set, and 10% to test set during the data processing. Moreover, due to the limited brightness diversity of low-featured X-ray images, models trained only on these images may not generalize well to X-ray images in other environments. To enhance the robustness of the

Table 2. Quantitative comparison with other state-of-the-art models for fracture detection on the GRAZPEDWRI-DX datasets when the input image size is 1024.

Model	Params (M)	FLOPs (G)	F1 (%)	mAP 50 (%)	mAP 50-95 (%)	Speed ^a (ms)
YOLOv8[16]	43.61	164.9	62	63.63	40.41	7.7
YOLOv8+SA[22]	43.64	165.4	63	64.25	41.64	8.0
YOLOv8+ECA[23]	43.64	165.5	65	64.26	41.94	7.7
YOLOv8+GAM[24]	49.29	183.5	65	64.26	41.00	12.7
YOLOv8+ResGAM[38]	49.29	183.5	64	64.98	41.75	18.1
YOLOv8+ResCBAM[25]	53.87	196.2	64	65.78	42.16	8.7
YOLOv9-C (Ours)	51.02	239.0	66	65.57	43.70	12.7
YOLOv9-E (Ours)	69.42	244.9	66	65.62	43.73	16.1

Note: The model size of all YOLOv8 and its variants listed in the table is large.

^a Speed is the total time for preprocessing, inference, and post-processing.

model, we employed data augmentation techniques to extend the training set. Specifically, we fine-tuned the contrast and luminance of the X-ray images using the addWeighted function from the OpenCV library.

Experiment:

Dataset: GRAZPEDWRI-DX [20] is a public dataset provided by the Medical University of Graz, which contains 20,327 X-ray images of pediatric wrist trauma. These X-ray images were collected by a team of pediatric radiologists at the University Hospital Graz from 2008 to 2018. The dataset comprises 6,091 patients and 10,643 studies, with a total of 74,459 labeled images, representing 67,771 labeled objects.

Experiment setup: The experiments in this paper utilized one single NVIDIA GeForce RTX 3090 GPU, employing Python with the PyTorch framework. Before training our model, we employed the YOLOv9 model weights pretrained on the MS COCO 2017 [27] dataset. In the training process, we trained the model using the SGD [41] optimizer, with a weight decay rate set to 5e-4 and a momentum of 0.937. We followed the research [21] to set the initial learning rate to 1e-2, the number of epochs to 100. Due to resource limitations (24GB memory) imposed by a single GPU, a batch size of 16 was employed for training the model.

Experimental Results: To evaluate the performance of YOLOv9 and other SOTA models in real diagnostic scenarios, this study compares model size (parameters and floating-point operations per second), accuracy (F1 score, mean average precision at 50% (mAP 50), and mean average precision from 50% to 95% (mAP 50-95)), and inference time. It is widely recognized that using larger input image sizes improves prediction accuracy but also requires more computational resources. Therefore, we conducted two experiments with input image sizes of 640 and 1024 for various scenarios, and the results are presented in Tables 1 and 2. With the input size of 640, both YOLOv9-C (Compact) and YOLOv9-E (Extended) demonstrate significantly improved mAP, while maintaining a reasonable inference speed. Specifically, YOLOv9-E achieves mAP

50-95 of 43.32%, which is 4.4% higher than 41.49% achieved by the current SOTA model YOLOv8+SA. When the input image size is 1024, the mAP 50-95 of YOLOv9-E reaches 43.73%, which also obtains the SOTA performance. However, due to the increased inference time, it is more suitable for deployment on devices with high computing resources.

Conclusion: The models of YOLO series can serve as CAD to assist radiologists and surgeons in interpreting X-ray images. However, the predictions from the previous models are often unsatisfactory due to the low features of X-ray images. This paper first introduces the application of YOLOv9 to fracture detection, addressing the issue of information loss during model training by employing the newly proposed PGI and GELAN. Experimental results indicate that the YOLOv9 model achieves SOTA performance on the GRAZPEDWRI-DX dataset, proving the effectiveness of this method.

Acknowledgments: This research is supported by National Science and Technology Council of Taiwan, under Grant Number: NSTC 112-2221-E-032-037-MY2.

© 2024 The Authors. *Electronics Letters* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology

References

- Chung, S.W., et al.: Automated detection and classification of the proximal humerus fracture by using deep learning algorithm. *Acta orthopaedica* 89(4), 468–473 (2018)
- Choi, J.W., et al.: Using a dual-input convolutional neural network for automated detection of pediatric supracondylar fracture on conventional radiography. *Investigative radiology* 55(2), 101–110 (2020)
- Tanzi, L., et al.: Hierarchical fracture classification of proximal femur x-ray images using a multistage deep learning approach. *European journal of radiology* 133, 109373 (2020)
- Adams, S.J., et al.: Artificial intelligence solutions for analysis of x-ray images. *Canadian Association of Radiologists Journal* 72(1), 60–72 (2021)
- Gan, K., et al.: Artificial intelligence detection of distal radius fractures: a comparison between the convolutional neural network and professional assessments. *Acta orthopaedica* 90(4), 394–400 (2019)
- Yahalomi, E., Chernofsky, M., Werman, M.: Detection of distal radius fractures trained by a small set of x-ray images and faster r-cnn. In: *Intelligent Computing: Proceedings of the 2019 Computing Conference*, Volume 1, pp. 971–981. Springer (2019)
- Blüthgen, C., et al.: Detection and localization of distal radius fractures: Deep learning system versus radiologists. *European journal of radiology* 126, 108925 (2020)
- Redmon, J., et al.: You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 779–788. (2016)
- Bochkovskiy, A., Wang, C.Y., Liao, H.Y.M.: Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934* (2020)
- Ge, Z., et al.: Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430* (2021)
- Wang, C.Y., Yeh, I.H., Liao, H.Y.M.: You only learn one representation: Unified network for multiple tasks. *arXiv preprint arXiv:2105.04206* (2021)
- Glenn, J.: Ultralytics yolov5. GitHub. <https://github.com/ultralytics/yolov5> (2022)
- Li, C., et al.: Yolov6: A single-stage object detection framework for industrial applications. *arXiv preprint arXiv:2209.02976* (2022)
- Ju, R.Y., et al.: Resolution enhancement processing on low quality images using swin transformer based on interval dense connection strategy. *Multimedia Tools and Applications* pp. 1–17. (2023)
- Wang, C.Y., Bochkovskiy, A., Liao, H.Y.M.: Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7464–7475. (2023)
- Glenn, J.: Ultralytics yolov8. GitHub. <https://github.com/ultralytics/yolov8> (2023)
- Samothai, P., et al.: The evaluation of bone fracture detection of yolo series. In: *2022 37th International Technical Conference on Circuits/Systems, Computers and Communications (ITC-CSCC)*, pp. 1054–1057. IEEE (2022)
- Hrzić, F., et al.: Fracture recognition in paediatric wrist radiographs: An object detection approach. *Mathematics* 10(16), 2939 (2022)
- Su, Z., et al.: Skeletal fracture detection with deep learning: A comprehensive review. *Diagnostics* 13(20), 3245 (2023)
- Nagy, E., et al.: A pediatric wrist trauma x-ray dataset (grazpedwri-dx) for machine learning. *Scientific Data* 9(1), 222 (2022)
- Ju, R.Y., Cai, W.: Fracture detection in pediatric wrist trauma x-ray images using yolov8 algorithm. *arXiv preprint arXiv:2304.05071* (2023)
- Zhang, Q.L., Yang, Y.B.: Sa-net: Shuffle attention for deep convolutional neural networks. In: *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2235–2239. IEEE (2021)
- Wang, Q., et al.: Eca-net: Efficient channel attention for deep convolutional neural networks. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11534–11542. (2020)
- Liu, Y., Shao, Z., Hoffmann, N.: Global attention mechanism: Retain information to enhance channel-spatial interactions. *arXiv preprint arXiv:2112.05561* (2021)
- Woo, S., et al.: Cbam: Convolutional block attention module. In: *Proceedings of the European conference on computer vision (ECCV)*, pp. 3–19. (2018)
- Wang, C.Y., Yeh, I.H., Liao, H.Y.M.: Yolov9: Learning what you want to learn using programmable gradient information. *arXiv preprint arXiv:2402.13616* (2024)
- Lin, T.Y., et al.: Microsoft coco: Common objects in context. In: *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pp. 740–755. Springer (2014)
- Son, D.M., et al.: Combined deep learning techniques for mandibular fracture diagnosis assistance. *Life* 12(11), 1711 (2022)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*, pp. 234–241. Springer (2015)
- Jeon, Y.D., et al.: Deep learning model based on you only look once algorithm for detection and visualization of fracture areas in three-dimensional skeletal images. *Diagnostics* 14(1), 11 (2023)
- Ahmed, A., et al.: Enhancing wrist abnormality detection with yolo: Analysis of state-of-the-art single-stage detection models. *Biomedical Signal Processing and Control* 93, 106144 (2024)
- Warin, K., et al.: Assessment of deep convolutional neural network models for mandibular fracture detection in panoramic radiographs. *International Journal of Oral and Maxillofacial Surgery* 51(11), 1488–1494 (2022)
- Gaikwad, D., et al.: Identification of cervical spine fracture using deep learning. *Australian Journal of Multi-Disciplinary Engineering* pp. 1–9. (2024)
- Zou, J., Arshad, M.R.: Detection of whole body bone fractures based on improved yolov7. *Biomedical Signal Processing and Control* 91, 105995 (2024)
- Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141. (2018)
- Abdeen, I., et al.: Fracatlas: A dataset for fracture classification, localization and segmentation of musculoskeletal radiographs. *Scientific Data* 10(1), 521 (2023)
- Moon, G., et al.: Computer aided facial bone fracture diagnosis (cafbfd) system based on object detection model. *IEEE Access* 10, 79061–79070 (2022)
- Chien, C.T., et al.: Yolov8-am: Yolov8 with attention mechanisms for pediatric wrist fracture detection. *arXiv preprint arXiv:2402.09329* (2024)
- Wang, C.Y., et al.: Cspnet: A new backbone that can enhance learning capability of cnn. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pp. 390–391. (2020)
- Wang, C.Y., Liao, H.Y.M., Yeh, I.H.: Designing network design strategies through gradient path analysis. *arXiv preprint arXiv:2211.04800* (2022)
- Ruder, S.: An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747* (2016)