Grenze International Journal of Engineering and Technology, June Issue

Improving YOLO26-OBB for Oriented Fracture Detection in X-Ray Images

Aarthy R[1] and Thirumal P C[2]
[1-2]Kumaraguru College of Technology/Information Technology, Coimbatore, India
Email: raarthy1553@gmail.com, thirumal.pc@kct.ac.in

***Abstract*— In this study, we propose a novel architecture based on YOLO26-OBB for oriented fracture detection in X-ray images. While traditional approaches such as horizontal bounding boxes (HBB) and segmentation are widely used, they often fail to capture the true orientation and precise boundaries of fractures. Oriented bounding boxes (OBB) provide a more accurate representation by aligning with the fracture direction, improving localization quality. To enhance detection, we introduce lightweight architectural modules including a Rotational Feature Enhancement Block (RFEB) for capturing orientation-specific features and an Edge-Aware Attention Module (EAAM) for improving boundary refinement. Additionally, training strategies are applied to handle the limitations of small datasets like FracAtlas. Experimental results show improved mAP@0.5:0.95 and recall with consistent localization under stricter IoU thresholds and minimal impact on inference time.**

***Index Terms*— Oriented Object Detection, YOLO26-OBB, Fracture Detection, X-Ray Imaging, Edge-Aware Attention, Ultralytics YOLO26-OBB.**

I. INTRODUCTION

Accurate detection of bone fractures in X-ray images is an important step in medical diagnosis, as it helps doctors take timely and correct decisions for treatment. In recent years, deep learning-based object detection and segmentation methods have been widely used for this task.

Most object detection models use horizontal bounding boxes (HBB), which are simple to implement but do not capture the actual orientation of fractures. This becomes a limitation when the fracture is tilted or has an irregular shape. Segmentation methods, on the other hand, give detailed pixel-level results but are computationally heavy and not always required in practical settings.

Oriented bounding boxes (OBB) provide a better alternative by adding an angle parameter to the bounding box. This allows the box to align with the direction of the fracture, leading to more accurate localization. With OBB, it is possible to estimate useful information such as fracture direction, approximate length, and exact region of interest. These details can support doctors in better understanding the fracture pattern and making informed decisions.

However, applying OBB detection on medical datasets is challenging due to limited data and variation in fracture appearance. In this work, we improve YOLO26-OBB using lightweight architectural modules and training strategies to better capture fracture orientation and boundary details, enabling reliable performance on small datasets.

Grenze ID: 01.GIJET.12.2.43_24

This paper is organized as follows. Section II discusses the related work in object detection, oriented detection, and medical imaging. Section III explains the proposed methodology and model enhancements. Section IV presents the experimental setup and training details. Section V discusses the results and analysis, followed by conclusions and future scope in Section VI.

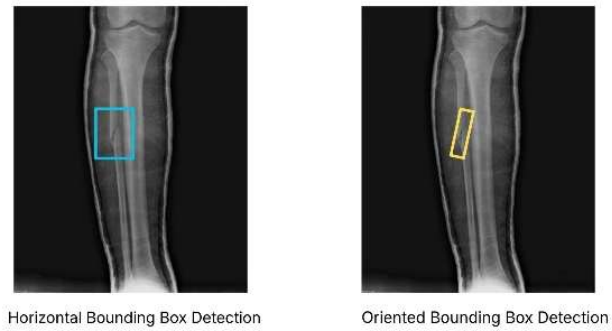

Figure 1. Horizontal vs oriented bounding box detection comparison

II. Related Work

Recent progress in object detection is largely driven by the YOLO family due to its efficiency in real-time applications. In the general computer vision domain, Ranjan Sapkota, Rahul Harsha Cheppally, Ajay Sharda, and Manoj Karkee [1] introduced YOLO26 with NMS-free inference and improved label assignment, enabling better deployment on edge devices. Ranjan Sapkota and Manoj Karkee [4] and Sapthak Mohajon Turjya [5] analysed YOLO evolution and showed that, despite improvements, these models still struggle with small objects and complex geometries. In the medical imaging domain, Zhenhui Cai, Kaiqing Zhou, and Zhouhua Liao [3] highlighted that YOLO-based models face difficulty in detecting low-contrast and fine anatomical structures.

In fracture detection (medical domain), Rui-Yang Ju and Weiming Cai [6] applied YOLOv8 with augmentation to improve detection accuracy, while Junting Zou and Mohd Rizal Arshad [7] improved YOLOv7 using attention and EIoU loss for better localization. Wanmian Wei et al. [20] proposed a multi-task YOLOv11 model combining detection and classification, improving diagnostic performance. Parvathaneni Naga Srinivasu, Gorli L. Aruna Kumari, Sujatha Canavoy Narahari, Shakeel Ahmed, and Abdulaziz Alhumam [22] showed that contrast-based augmentation and tuning enhance fracture detection. However, these works rely on horizontal bounding boxes, which are not suitable for capturing thin and oriented fracture lines.

To address orientation challenges, oriented bounding box (OBB) methods have been explored across multiple domains. In the medical domain, YingLiang Ma, Sandra Howell, Aldo Rinaldi, Tarv Dhanjal, and Kawal S. Rhode [9] used attention-guided rotated detection for localizing wire-like structures in X-ray images. In industrial robotics, Jinshun Dong et al. [10] used deformable convolution to handle dense and rotated objects. In barcode detection (industrial vision), Peidong Luo, Zhixing Ma, Qingyang Wu, Tianhong Zhao, and Xiaole Shen [11] proposed dynamic convolution for accurate angle regression.

In remote sensing, Jian Ding et al. [12] and Xingxing Xie, Gong Cheng, Jiabao Wang, Xiwen Yao, and Junwei Han [13] showed that orientation-aware detection significantly improves performance in aerial images. Kun Wang et al. [14] summarized challenges such as feature misalignment and unstable angle prediction. In medical applications, Sayed Masoud Hosseini et al. [15] and Kwang Hyeon Kim, Hae-Won Koo, and Byung-Jou Lee [16] demonstrated improved localization using OBB, but without specialized orientation-aware feature learning.

Attention mechanisms improve feature representation and boundary precision. Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon [19] proposed CBAM for channel and spatial refinement. In fracture detection, Chun-Tse Chien, Rui-Yang Ju, Kuang-Yi Chou, Enkaer Xieerke, and Jen-Shiun Chiang [18] showed that attention improves feature focus in YOLO models. In medical segmentation, Rashmi R and Girisha S [17] and Gefeng Hu, Wen Zhu, Xinyi Liao, and Qingbo Li [21] introduced edge-aware attention for precise boundary detection. Zhihao Su, Afzan Adam, Mohammad Faidzul Nasrudin, Masri Ayob, and Gauthamen Punganan [8] emphasized the need for better localization and interpretability.

OBB remains underexplored in fracture detection, where orientation is a clinically important feature. While YOLO26 [1] provides an efficient architecture suitable for edge deployment, existing methods still rely on horizontal bounding boxes, which cannot capture fracture orientation. To the best of our knowledge, oriented bounding box-based fracture localization in musculoskeletal radiographs has not been sufficiently addressed.

To bridge this gap, this work proposes an enhanced detection framework that integrates rotational feature learning and edge-aware attention to improve orientation sensitivity and boundary precision in fracture localization.

III. Proposed Methodology

A. Dataset Preparation and Preprocessing

The study utilizes the FracAtlas dataset [2], which is a clinically curated dataset consisting of radiographic X-ray images specifically annotated for fracture detection tasks. The dataset contains a total of 4,083 images, among which 717 images include fracture instances, accounting for 922 annotated fractures. These annotations are provided in the form of segmentation masks along with bounding box information and have been verified by medical experts, ensuring high annotation reliability.

Given that fracture patterns are inherently irregular and often oriented arbitrarily with respect to the image axes, the use of segmentation masks enables precise geometric transformation into oriented bounding box representations. The dataset is systematically divided into training, validation, and testing subsets, maintaining class distribution consistency to avoid sampling bias and to ensure that the model generalizes well across unseen data.

The preprocessing pipeline begins with resizing all images to a uniform spatial resolution of 640 × 640 pixels. This ensures compatibility with the detection architecture while maintaining sufficient spatial detail for small fracture regions. Normalization is implicitly handled within the training framework, ensuring stable gradient propagation during optimization.

To address the limited size of the dataset and improve model robustness, extensive data augmentation is applied. Color space augmentations introduce variations in hue, saturation, and brightness to simulate differences in imaging conditions. Geometric transformations such as rotation, translation, and scaling are applied to increase spatial diversity.

Rotation augmentation, in particular, plays a critical role in this work, as it directly enhances the model's ability to learn orientation-invariant representations, which is essential for oriented object detection tasks. Advanced augmentation strategies such as Mosaic and MixUp are also incorporated, enabling the model to learn from composite image contexts and improving generalization.

Figure 2. Polygon annotation with bounding box overlay

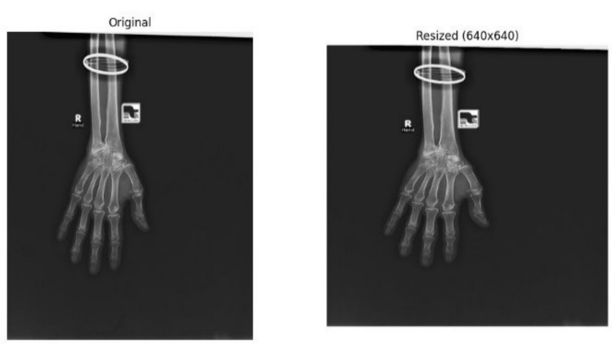

Figure 3. Original and resized X-ray preprocessing comparison

B. Oriented Bounding Box Annotation

The original segmentation masks are converted into oriented bounding box annotations using contour-based geometric processing. Each segmentation mask is first converted into grayscale format, after which contour extraction is performed to identify connected fracture regions. To eliminate noise and irrelevant artifacts, contours with an area smaller than a predefined threshold are discarded.

For each valid contour, a minimum area rotated rectangle is computed, which provides a compact representation of the fracture region. This operation yields the center coordinates, dimensions, and rotation angle of the bounding box. The rotated rectangle is represented using its four corner coordinates as in (1):

$$(x_1, y_1), (x_2, y_2), (x_3, y_3), (x_4, y_4) \tag{1}$$

These coordinates are normalized with respect to image width and height to ensure scale invariance as in (2):

$$x_i' = \frac{x_i}{W}, y_i' = \frac{y_i}{H} \tag{2}$$

The final annotation format used during training consists of the class identifier followed by the normalized coordinates of the four vertices. This representation allows the bounding box to align precisely with the fracture orientation, thereby reducing localization error compared to conventional axis-aligned representations.

C. Model Architecture

The proposed approach is built upon the YOLO26-OBB architecture implemented within the Ultralytics framework. The model follows a standard detection pipeline comprising a backbone for feature extraction, a neck for feature aggregation, and a detection head for prediction. However, to address the specific challenges of fracture detection, two key modules are introduced: the Rotational Feature Enhancement Block (RFEB) and the Edge-Aware Attention Module (EAAM). The oriented bounding box representation used in this work is polygon-based, defined by four vertices instead of the conventional center-width-height parameterization. This allows the model to represent rotated objects explicitly, improving alignment with fracture structures that exhibit arbitrary orientations. The RFEB module is designed to enhance orientation-sensitive feature extraction. Traditional convolution operations are inherently limited in capturing directional variations, especially when dealing with elongated and rotated structures. To address this limitation, RFEB employs multiple depthwise convolution branches, each capturing features at different receptive orientations. The outputs from these branches are combined using learnable attention weights, as shown in (3).

$$F_{\mathrm{RFEB}} = \sum_{i=1}^{n} \alpha_i F_i \tag{3}$$

where F_i represents the feature maps from each branch and α_i denotes the corresponding attention weights. This enables the network to adaptively prioritize orientation-specific features.

The EAAM module is introduced to enhance edge and boundary information, which is crucial for accurately identifying fracture regions. Edge information is extracted using gradient-based operators, where horizontal and vertical gradients are computed as in (4).

$$G_x = \frac{\partial I}{\partial x}, G_y = \frac{\partial I}{\partial y} \tag{4}$$

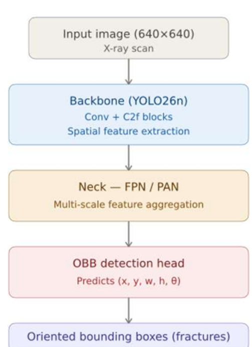

Figure 4. YOLO-OBB26 baseline architecture for fracture detection

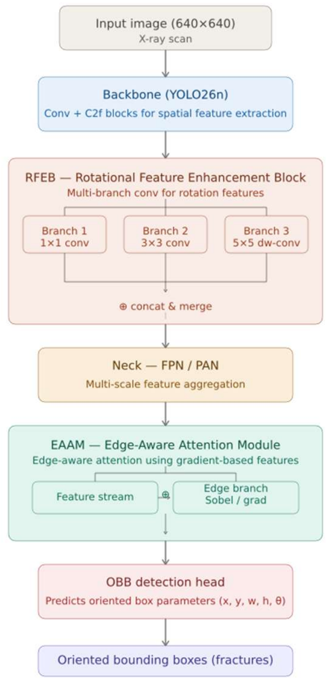

Figure 5. Proposed architecture with RFEB and EAAM modules

The overall edge magnitude is then given by (5):

$$G = \sqrt{G_x^2 + G_y^2} \tag{5}$$

This edge representation is fused with the original feature maps, followed by channel-wise attention to emphasize boundary-relevant features. The refined feature map is expressed as in (6):

$$F_{\text{EAAM}} = \sigma(W \cdot [F \oplus G]) \tag{6}$$

where σdenotes the activation function and $\oplus$ represents feature fusion. The integration of RFEB and EAAM ensures that the model effectively captures both orientation and structural boundary information, leading to improved detection performance.

D. Training Configuration

The model is trained using the Ultralytics training pipeline on a GPU-enabled environment to ensure efficient computation and convergence. The training process is conducted over 100 epochs with a batch size of 16, using an input resolution of 640 × 640. The learning rate is initialized at 0.001, and optimization is carried out using the default optimizer provided by the framework.

The loss function used during training is a combination of bounding box regression loss, classification loss, and Distribution Focal Loss (DFL). The bounding box loss ensures accurate localization, while the classification loss optimizes category prediction.

TABLE I: VALIDATION PERFORMANCE OF THE PROPOSED MODEL

Metric	Value
Precision (P)	0.75
Recall (R)	0.61
mAP@0.5	0.71
mAP@0.5:0.95	0.42

The DFL component improves bounding box precision by modeling the probability distribution of predicted box coordinates, which is particularly beneficial for oriented bounding box regression. Validation is performed periodically during training to monitor performance trends and ensure that the model does not overfit. The final model is selected based on its validation performance.

E. Evaluation Methodology and Results

The model is evaluated using standard detection metrics: Precision measures the fraction of correct detections among all predictions, Recall measures the fraction of actual fractures successfully detected, and mean Average Precision (mAP) summarizes detection quality across confidence thresholds. Two IoU-based metrics are reported: mAP@0.5 for moderate overlap and mAP@0.5:0.95 averaged across stricter thresholds.
The validation performance of the model is summarized in Table I.
Validation results are summarized in Table I. The model achieves a Precision of 0.75, Recall of 0.61, mAP@0.5 of 0.71, and mAP@0.5:0.95 of 0.42. The strong mAP@0.5 confirms reliable fracture detection under standard conditions, while the mAP@0.5:0.95 reflects the inherent difficulty of predicting precise rotated geometries. The RFEB and EAAM modules contribute to improved orientation handling and boundary sensitivity, supporting accurate localization in clinical radiographs.

IV. Experimental Setup and Results

The models are evaluated on a held out test set from the FracAtlas dataset, ensuring a fair and unbiased assessment. A comparison of the experimented models, including YOLO11n OBB, YOLO26s OBB, YOLO26n OBB, and the proposed YOLO26n OBB, is presented in Table II.
The proposed model shows clear improvement in both detection and localization, achieving higher mAP scores compared to the baseline models. It also improves recall, indicating better identification of fracture instances. In several test cases, the model produced confidence scores above 0.80, showing reliable predictions.
Importantly, even after adding the RFEB and EAAM modules, the inference time remains close to the baseline, indicating that the performance gains are achieved without a significant increase in computational cost.

Table II: Performance Comparison of Yolo based OBB Models on the FracAtlas Test Set

Model	Precision	Recall	mAP@0.5	mAP@0.5:0.95	Inference Time (ms)
YOLO11n-OBB	0.661	0.503	0.629	0.300	3.8
YOLO26s-OBB	0.716	0.528	0.587	0.305	12.8
YOLO26n-OBB	0.842	0.551	0.672	0.368	6.4
YOLO26n-OBB (Proposed)	0.742	0.634	0.703	0.415	6.8

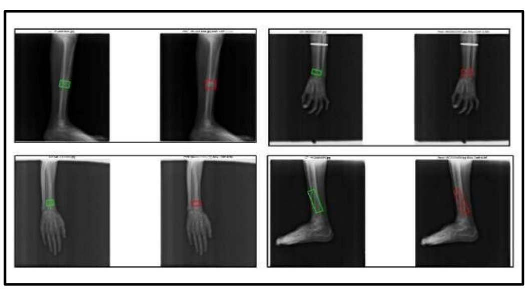

Figure 6. Ground truth versus predicted fracture detections

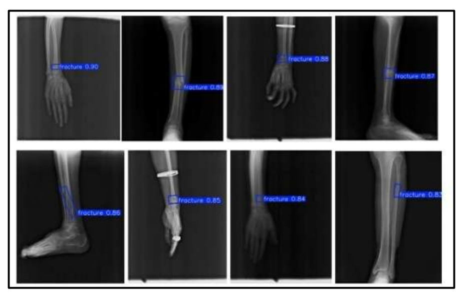

Figure 7. OBB detection outputs with confidence scores

Figure 8. Predicted fracture orientation angles output visualization

V. Conclusions

In this work, we explored oriented object detection for fracture identification using the YOLO26 OBB framework. Unlike horizontal bounding boxes, OBB captures fracture orientation effectively while remaining computationally efficient compared to segmentation methods. The proposed architectural enhancements improve feature representation and achieve better detection performance on the FracAtlas dataset.

The relatively lower improvement in some metrics is mainly due to the limited dataset size and fewer fracture instances. With larger and more diverse data, the model can achieve stronger performance. In future, this work can be extended by integrating segmentation for a complete pipeline, improving OBB specific losses, and validating on larger datasets for better generalization.

References

[1] Sapkota, R., Cheppally, R. H., Sharda, A., & Karkee, M. (2025). YOLO26: Key Architectural Enhancements and Performance Benchmarking for Real-Time Object Detection. ArXiv. https://arxiv.org/abs/2509.25164

[2] Abedeen, I., Rahman, M. A., Prottyasha, F. Z., Ahmed, T., Chowdhury, T. M., & Shatabda, S. (2023). FracAtlas: A Dataset for Fracture Classification, Localization and Segmentation of Musculoskeletal Radiographs. Scientific Data, 10(1), 521. https://doi.org/10.1038/s41597-023-02432-4

[3] Cai, Z., Zhou, K., & Liao, Z. (2025). A Systematic Review of YOLO-Based Object Detection in Medical Imaging: Advances, Challenges, and Future Directions. Computers, Materials and Continua, 85(2), 2255–2303. https://doi.org/10.32604/cmc.2025.067994

[4] Sapkota, R., & Karkee, M. (2025). Ultralytics YOLO Evolution: An Overview of YOLO26, YOLO11, YOLOv8 and YOLOv5 Object Detectors for Computer Vision and Pattern Recognition. ArXiv. https://arxiv.org/abs/2510.09653

[5] Turjya, S. M. (2026). A Comprehensive Review Of Yolov26 And Existing Yolo Variants: Architecture, Training Paradigms, And Applications In Medical And General Computer Vision. https://doi.org/10.2139/ssrn.6159430

[6] Ju, R. Y., & Cai, W. (2023). Fracture detection in pediatric wrist trauma X-ray images using YOLOv8 algorithm. Scientific Reports, 13(1), 20077. https://doi.org/10.1038/s41598-023-47460-7

[7] Zou, J., & Arshad, M. R. (2024). Detection of whole body bone fractures based on improved YOLOv7. Biomedical Signal Processing and Control, 91, 105995. https://doi.org/10.1016/j.bspc.2024.105995

[8] Su, Z., Adam, A., Nasrudin, M. F., Ayob, M., & Punganan, G. (2023). Skeletal Fracture Detection with Deep Learning: A Comprehensive Review. Diagnostics, 13(20). https://doi.org/10.3390/diagnostics13203245

[9] Ma, Y., Howell, S., Rinaldi, A., Dhanjal, T., & Rhode, K. S. (2024). Real-Time Device Detection with Rotated Bounding Boxes and Its Clinical Application. Lecture Notes in Computer Science, 83–93. https://doi.org/10.1007/978-3-031-73083-2_9

[10] Dong, J., Deng, L., Wan, D., Liu, C., Yin, J., Guo, M., Zhang, H., Lin, S., Liu, H., & Liu, L. (2025). Oriented bounding box detection algorithm for dense scenarios of robotic arm operation. Expert Systems with Applications, 298, 129678. https://doi.org/10.1016/j.eswa.2025.129678

[11] Luo, P., Ma, Z., Wu, Q., Zhao, T., & Shen, X. (2026). Efficient adaptive rotated object detection for 1D and QR barcodes. Scientific Reports, 16(1), 4689. https://doi.org/10.1038/s41598-025-34854-y

[12] Ding, J., Xue, N., Xia, G. S., Bai, X., Yang, W., Yang, M. Y., Belongie, S., Luo, J., Datcu, M., Pelillo, M., & Zhang, L. (2021). Object Detection in Aerial Images: A Large-Scale Benchmark and Challenges. ArXiv. https://doi.org/10.1109/TPAMI.2021.3117983

[13] Xie, X., Cheng, G., Wang, J., Yao, X., & Han, J. (2021). Oriented R-CNN for Object Detection. ArXiv. https://arxiv.org/abs/2108.05699

[14] Wang, K., Wang, Z., Li, Z., Su, A., Teng, X., Pan, E., Liu, M., & Yu, Q. (2025). Oriented object detection in optical remote sensing images using deep learning: a survey. Artificial Intelligence Review, 58(11). https://doi.org/10.1007/s10462-025-11256-0

[15] Hosseini, S. M., Seyed Ali Mohtarami, Shahin Shadnia, Rahimi, M., Peyman, Babak Mostafazadeh, Azadeh Memarian, & Heidarli, E. (2025). Detection of Body Packs in Abdominal CT scans Through Artificial Intelligence; Developing a Machine Learning-based Model. Archives of Academic Emergency Medicine, 13(1), e23–e23. https://doi.org/10.22037/aaemj.v13i1.2479

[16] Kim, K. H., Koo, H. W., & Lee, B. J. (2024). Deep Learning-Based Localization and Orientation Estimation of Pedicle Screws in Spinal Fusion Surgery. Korean Journal of Neurotrauma, 20(2), 90. https://doi.org/10.13004/kjnt.2024.20.e17

[17] R, R., & S, G. (2024). ETiSeg-Net: edge-aware self attention to enhance tissue segmentation in histopathological images. Multimedia Tools and Applications. https://doi.org/10.1007/s11042-024-20542-y

[18] Chien, C. T., Ju, R. Y., Chou, K. Y., Xieerke, E., & Chiang, J. S. (2024). YOLOv8-AM: YOLOv8 Based on Effective Attention Mechanisms for Pediatric Wrist Fracture Detection. ArXiv. https://doi.org/10.1109/ACCESS.2025.3549839

[19] Woo, S., Park, J., Lee, J.-Y., & Kweon, I. S. (2018). CBAM: Convolutional Block Attention Module. Computer Vision – ECCV 2018, 3–19. https://doi.org/10.1007/978-3-030-01234-2_1

[20] Wei, W., Huang, Y., Zheng, J., Rao, Y., Wei, Y., Tan, X., & OuYang, H. (2025). YOLOv11-based multi-task learning for enhanced bone fracture detection and classification in X-ray images. Journal of Radiation Research and Applied Sciences, 18(1), 101309. https://doi.org/10.1016/j.jrras.2025.101309

[21] Hu, G., Zhu, W., Liao, X., & Li, Q. (2025). WA-NET: enhanced boundary-aware segmentation of skin lesions via frequency-spatial feature fusion and attention-guided edge refinement. Scientific Reports, 15(1). https://doi.org/10.1038/s41598-025-25583-3

[22] Srinivasu, P. N., Kumari, G. L., Narahari, S. C., Ahmed, S., & Alhumam, A. (2025). Exploring the impact of hyperparameter and data augmentation in YOLO V10 for accurate bone fracture detection from X-ray images. Scientific Reports, 15(1), 9828. https://doi.org/10.1038/s41598-025-93505-4