Journal of Advances in Developmental Research (IJAIDR)

E-ISSN: 0976-4844 ● Website: www.ijaidr.com ● Email: editor@ijaidr.com

AI-Assisted Spine X-Ray Analysis for Early Scoliosis Diagnosis Using YOLOv8-Based Object Detection and Automated Cobb Angle Estimation

Manju darshan M J[1], Supreetha Gowda H D[2]

[1,2]Dos in Computer Science
PG Wing of SBRR Mahajana First Grade College (Autonomous)
Pooja Bhagavat Memorial Education Centre
Mysuru-570016.India

Abstract:
Scoliosis is a musculoskeletal condition involving an abnormal lateral curvature of the spine that, if left undiagnosed, can progress to chronic pain, postural imbalance, and reduced respiratory function. Conventional diagnosis relies on manual radiographic interpretation and Cobb angle measurement performed by radiologists or orthopedic specialists, a process that is reliable but time-intensive and subject to inter-observer variability. This paper presents an automated, deep-learning-driven screening pipeline that detects vertebral structures and spinal abnormalities directly from spine radiographs using a YOLOv8 object-detection model trained on a three-class annotated dataset (Vertebra, Scoliosis Spine, Normal Spine). Detected vertebral bounding-box centroids are used to fit upper and lower spinal reference lines, from which an approximate Cobb angle is computed; the resulting angle is mapped to a four-level severity scale (normal, mild, moderate, severe). The complete pipeline is exposed through a Flask-based web application supporting user authentication, image upload, annotated result visualization, and persistent storage of prediction history in a relational database. Evaluation on held-out validation and test partitions shows strong detection performance, with an overall test mAP@50 of 94.7% and class-wise precision and recall exceeding 0.86 across all three categories. These results indicate that combining real-time object detection with automated angle estimation can provide a practical, reproducible first-pass screening tool, while the system is explicitly positioned as a decision-support aid rather than a replacement for clinical diagnosis.

Keywords: Scoliosis detection; YOLOv8; Cobb angle estimation; spine X-ray; deep learning; medical image analysis; object detection.

1. Introduction

Scoliosis is defined clinically as a lateral curvature of the spine exceeding 10 degrees when measured by the Cobb angle technique. It occurs across all age groups but is most frequently identified during adolescence, and its causes range from congenital and neuromuscular factors to degenerative changes and

idiopathic origins. Without timely intervention, scoliosis can progress to cause postural asymmetry, chronic back pain, restricted lung capacity, and psychosocial distress.

Standard clinical practice for scoliosis assessment depends on visual inspection of spinal radiographs and manual Cobb angle measurement, in which a specialist selects the most tilted vertebrae above and below the curve apex and measures the angle subtended by their endplates. While this approach remains the clinical reference standard, it is labor-intensive, requires specialist availability, and is susceptible to measurement variation between and within observers, particularly in high-volume diagnostic settings.

Advances in deep learning, and in convolutional and single-shot object-detection architectures in particular, have created new opportunities for automating structural analysis of medical images. Models such as YOLO (You Only Look Once) perform localization and classification in a single forward pass, making them well suited to near-real-time clinical screening tools. This work builds on that capability by combining a YOLOv8 detector with a lightweight geometric procedure for Cobb angle approximation, packaged inside an accessible web application.

The remainder of this paper is organized as follows. Section 2 reviews related work on AI-assisted Cobb angle estimation and vertebral detection. Section 3 describes the proposed methodology, including the system architecture, dataset, detection model, angle-estimation procedure, and severity classification scheme. Section 4 reports experimental results and discusses their implications. Section 5 concludes the paper and outlines directions for future work.

2. Related Work

Automated Cobb angle estimation has been approached through several complementary strategies in the recent literature. Segmentation-based methods construct pixel-level masks of vertebral bodies to support subsequent line fitting; an early example used a U-Net-style encoder-decoder to outline vertebral boundaries, demonstrating that learned segmentation could replace manual boundary tracing, albeit at higher computational cost than bounding-box approaches [1].

A second line of work favors landmark or keypoint regression, in which a network predicts the corner coordinates of vertebral endplates directly. Reported results show that keypoint-based models generalize well across exposure conditions and achieve low mean absolute error relative to expert measurements [13]; subsequent studies extended this idea to direct comparison between automated keypoint detection and manual grading, finding high agreement between the two and supporting the case for AI-assisted consistency in scoliosis screening [10].

Several groups have proposed end-to-end convolutional pipelines purpose-built for Cobb angle regression. One representative system used a multi-task feature-fusion network to jointly capture local vertebral detail and global spinal curvature [14]; another combined convolutional feature extraction with attention-based Transformer blocks to model long-range spinal relationships, reporting improved accuracy at the cost of greater computational demand [13]. A modified region-based detector incorporating vertebral corner and centroid detection has also been used to support interpretable, visualized Cobb angle output across radiograph and EOS imaging modalities [3].

Ensemble and multi-stage architectures represent a further direction, combining several model components to improve robustness for adolescent idiopathic scoliosis diagnosis, though typically at increased training and deployment complexity [11]. More recent work has explored morphology-aware networks that fuse centerline, boundary, and region features for improved regression accuracy on

benchmark datasets [8], as well as fully automated systems validated directly against multi-expert manual readings, reporting strong inter-method agreement [6]. Comparative studies evaluating multiple AI algorithms side-by-side have further shown that model architecture choice materially affects measurement consistency [5].

Relative to this body of work, the present study adopts a deliberately lightweight design: rather than dense segmentation masks, fine-grained keypoints, or attention-heavy architectures, the proposed pipeline uses YOLOv8 bounding-box detections of vertebrae as the sole geometric input to a simple line-fitting procedure for Cobb angle approximation. This trades some measurement precision for substantially lower computational overhead, real-time inference, and straightforward integration into a deployable web application — properties that are valuable for preliminary screening, educational use, and resource-limited settings rather than for clinical-grade quantification.

3. Proposed Methodology

The proposed system follows a sequential pipeline that begins with user-submitted spine radiographs and ends with a severity-labeled, database-logged screening result. Figure 1 presents the overall block diagram of the system.

3.1 System Architecture

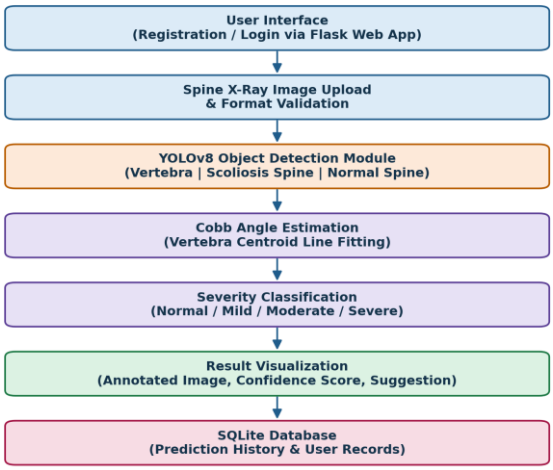

Fig. 1. Block diagram of the proposed AI-assisted scoliosis screening pipeline.

As shown in Figure 1, an authenticated user uploads a spine radiograph through the web interface. The image is first checked against a whitelist of supported formats (JPG, JPEG, PNG, WEBP, BMP); files that fail validation are rejected before reaching the detection stage. Validated images are passed to the YOLOv8 detector, which localizes and classifies vertebral and spine-condition regions. The resulting vertebra bounding boxes feed a geometric Cobb angle estimator, whose output angle is mapped to a severity category. The annotated image, detected classes, confidence scores, estimated angle, and severity label are then rendered to the user and persisted in a relational database for later retrieval.

3.2 Dataset

The detection model is trained on a publicly available, YOLO-format annotated spine X-ray dataset partitioned into training, validation, and test splits, consistent with dataset conventions used in comparable Cobb-angle and vertebra-detection studies [1], [7]. Three object classes are labeled: Vertebra, Scoliosis Spine, and Normal Spine. Because each radiograph contains multiple vertebral bodies but only one condition-level label, the Vertebra class is numerically dominant relative to the two condition classes, producing a degree of class imbalance that is taken into account when interpreting per-class metrics. Standard augmentation operations — rotation, brightness and contrast adjustment, horizontal flipping, zoom, cropping, and additive noise — are applied during training to improve robustness to variation in radiograph acquisition.

3.3 YOLOv8-Based Vertebra and Spine-Condition Detection

YOLOv8 performs object localization and classification in a single forward pass, predicting bounding boxes, class identities, and confidence scores directly from the input image without a separate region-proposal stage [2]. This single-stage design gives it a favorable speed-accuracy trade-off for screening applications where near-instant feedback is desirable, in line with reported advantages of single-pass detectors in high-volume clinical triage settings [4]. During training, the network learns discriminative visual features for vertebral bodies and for the broader curvature patterns that distinguish scoliosis spine from normal spine. At inference time, detections below a configurable confidence threshold (set to 0.25 in this work) are discarded, and the remaining detections are used both for direct classification feedback and as geometric input to the Cobb angle estimator.

3.4 Automated Cobb Angle Estimation

Because clinical Cobb angle measurement normally requires precise identification of the most tilted vertebral endplates [9], the proposed system approximates this procedure using the centroids of detected vertebra bounding boxes rather than full endplate geometry. For every vertebra detection retained after confidence filtering, the bounding-box centroid (cx, cy) is computed from its corner coordinates. The collected centroids are sorted along the vertical (cranio-caudal) image axis, and the topmost and bottommost subsets — approximating the upper and lower thirds of the visible spine — are each fitted with a least-squares regression line, following the general centroid-based line-fitting logic used in comparable landmark-driven Cobb angle systems [13]. The acute angle each fitted line makes with the vertical axis is computed, and the absolute difference between the upper and lower line angles is taken as the estimated Cobb angle. If fewer than four vertebra detections are available, the system reports an insufficient-detection message rather than producing an unreliable estimate; if the image is classified as a normal spine with no scoliosis-spine evidence, angle estimation is skipped entirely. Algorithm 1 summarizes this procedure.

Algorithm 1. Cobb Angle Estimation from Detected Vertebra Centroids

1. Input: list of YOLOv8 detections D for an image
2. V ← centroids of all boxes in D with class = Vertebra and confidence ≥ 0.25
3. if image is classified Normal Spine and no Scoliosis Spine detection exists then
4. return "not calculated"

5. if |V| < 4 then
6. return "insufficient vertebra detections"
7. Sort V by vertical coordinate (top to bottom)
8. Vupper ← top third of V; Vlower ← bottom third of V
9. Fit line Lupper through Vupper; fit line Llower through Vlower (least-squares)
10. θupper ← angle of Lupper from vertical; θlower ← angle of Llower from vertical
11. CobbAngle ← |θupper − θlower|, rounded to 2 decimal places
12. return CobbAngle

3.5 Severity Classification

The estimated Cobb angle is converted into a four-level severity category using clinically established angle thresholds, summarized in Table 1. This conversion gives the end user an interpretable result rather than a raw numerical angle alone.

Estimated Cobb Angle	Severity Category
< 10°	Normal
10° – 25°	Mild scoliosis
25° – 40°	Moderate scoliosis
> 40°	Severe scoliosis

Table 1. Cobb angle to severity-category mapping used by the system.

3.6 Web Application and Data Logging

The complete pipeline is exposed through a Flask web application that handles user registration and authentication (with hashed password storage), image upload and validation, model inference, result rendering, and database persistence. Each completed prediction — including the original image path, annotated result path, detected classes, confidence scores, estimated Cobb angle, severity label, and timestamp — is written to a relational (SQLite) store, allowing users to review prediction history and supporting downstream auditing or longitudinal tracking.

4. Results and Discussion

4.1 Detection Performance

Model performance is reported using standard object-detection metrics — precision, recall, mAP@50, and mAP@50-95 — which are more appropriate than classification accuracy for a bounding-box detection task. Table 2 reports validation-set metrics, and Table 3 reports test-set metrics, for each of the three detection classes and for the overall (all-class) average.

Class	Precision	Recall	mAP@50	mAP@50-95
All classes	0.913	0.904	0.947	0.531
Vertebra	0.964	0.868	0.939	0.530

Class	Precision	Recall	mAP@50	mAP@50-95
Scoliosis spine	0.809	0.969	0.955	0.611
Normal spine	0.967	0.876	0.947	0.452

Table 2. YOLOv8 validation-set detection performanc

Class	Precision	Recall	mAP@50	mAP@50-95
All classes	0.908	0.954	0.978	0.574
Vertebra	0.862	0.946	0.951	0.541
Scoliosis spine	0.919	1.000	0.994	0.597
Normal spine	0.941	0.917	0.989	0.584

Table 3. YOLOv8 test-set detection performance

On the test partition, the model achieves an overall mAP@50 of 0.978, with the scoliosis-spine class reaching perfect recall (1.000) and the normal-spine class achieving the highest precision (0.941–0.967) across both splits. The consistently lower mAP@50-95 relative to mAP@50 reflects the stricter IoU thresholds used in that metric, a pattern that is common in medical object-detection tasks where exact bounding-box alignment with anatomical boundaries is inherently difficult to achieve [4].

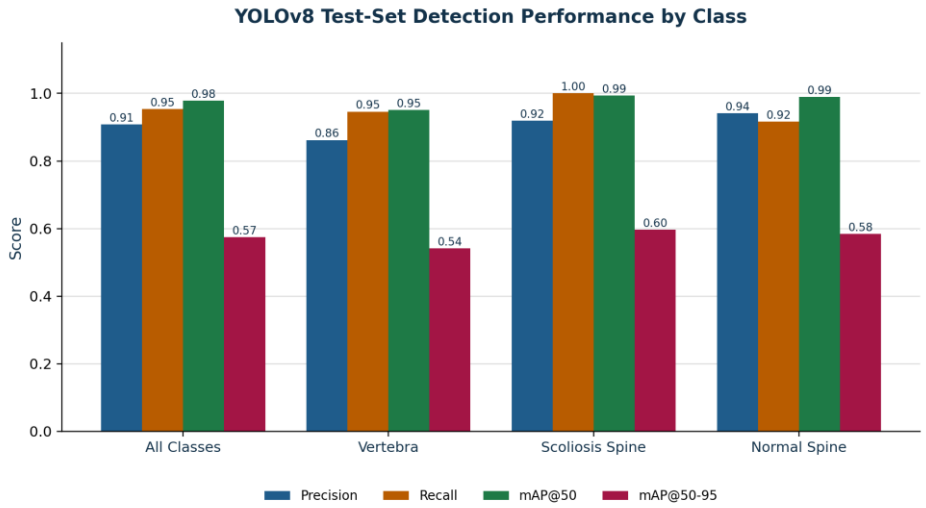

Fig. 2. Test-set precision, recall, mAP@50, and mAP@50-95 by detection class.

Figure 2 makes three analytical points visible that are easy to miss when scanning Table 3 row by row. First, recall consistently outpaces precision for the vertebra and scoliosis-spine classes (0.946 vs. 0.862, and 1.000 vs. 0.919, respectively), while the opposite holds for normal spine (0.917 recall vs. 0.941 precision). This asymmetry suggests the model is biased toward generating a generous number of candidate detections for vertebra and scoliosis-spine regions — finding nearly every true instance at some cost to false positives — while behaving more conservatively for normal spine, where it produces fewer but more reliable detections. For a screening tool, this is a favorable bias: missing a true scoliosis-spine

region (a false negative) is clinically more costly than flagging an extra candidate region for review (a false positive), so recall-leaning behaviour on the disease-relevant classes is the safer failure mode.
Second, the scoliosis-spine class achieves the highest mAP@50 (0.994) of the three classes despite having the lowest precision (0.919) among the condition-level classes. This combination indicates that the bounding boxes the model does produce for scoliosis-spine regions are spatially well-localized and consistently above the 0.50 IoU threshold across nearly the full confidence range, even though a modest number of additional low-confidence false-positive boxes pull precision down slightly. In other words, the precision dip is a confidence-calibration effect rather than a localization-quality problem. Third, the vertebra class is the weakest performer on every metric in Table 3 (0.862 precision, 0.946 recall, 0.951 mAP@50, 0.541 mAP@50-95) — a direct consequence of vertebrae being small, numerous, and tightly packed within each radiograph, which makes both detection and precise localization intrinsically harder than for the single, larger condition-level region per image.

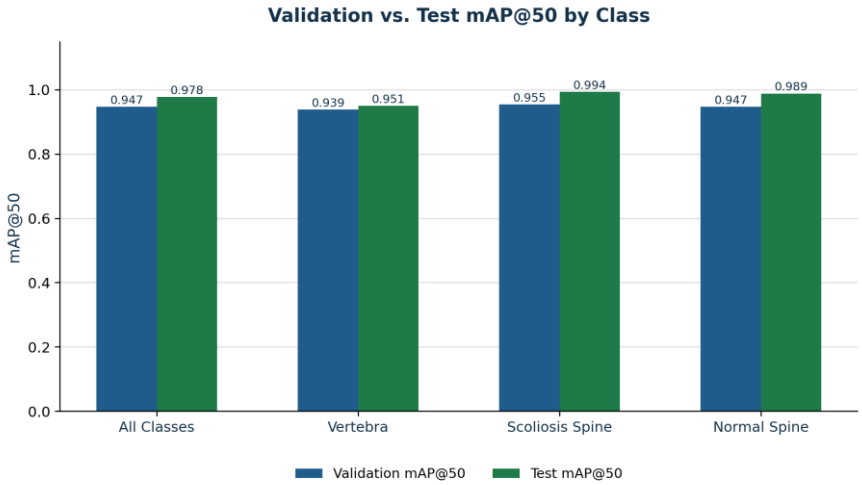

Fig. 3. Validation vs. test mAP@50 by class, illustrating consistency across data splits.

Figure 3 compares validation- and test-set mAP@50 side by side for each class. Two observations follow. First, test-set mAP@50 is equal to or higher than validation-set mAP@50 for every class (by margins of 0.008–0.042), rather than lower, which rules out the most common failure pattern in model evaluation — a model that scores well on validation data used during development but degrades on a genuinely held-out test set. Second, the gap between validation and test performance is smallest for vertebra (0.012) and largest for scoliosis spine (0.039); since scoliosis-spine images are the least numerous condition-level class in the dataset, a larger split-to-split fluctuation for this class is consistent with ordinary sampling variability in a smaller per-class test subset, rather than indicating a systematic generalization problem.

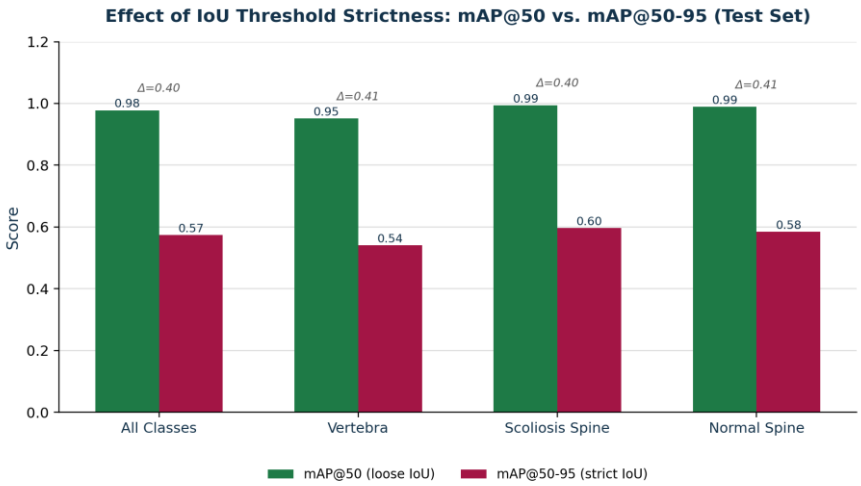

Fig. 4. mAP@50 versus mAP@50-95 by class, isolating the effect of IoU threshold strictness on measured performance.

Figure 4 isolates the effect of evaluation strictness alone by holding the model and class fixed and varying only the IoU threshold convention. The gap between mAP@50 and mAP@50-95 is remarkably uniform across all four rows (Δ = 0.40–0.41), which is itself an analytical finding: if one class had a markedly larger gap than the others, it would suggest that class's bounding boxes are well-classified but poorly localized (correct detections, imprecise boxes); the near-identical gap across all classes instead indicates that the loosening from strict to loose IoU benefits every class by roughly the same amount, consistent with a single shared source of localization imprecision — most plausibly the inherent difficulty of drawing pixel-tight boxes around curved, overlapping anatomical structures — rather than a class-specific weakness in the detector.

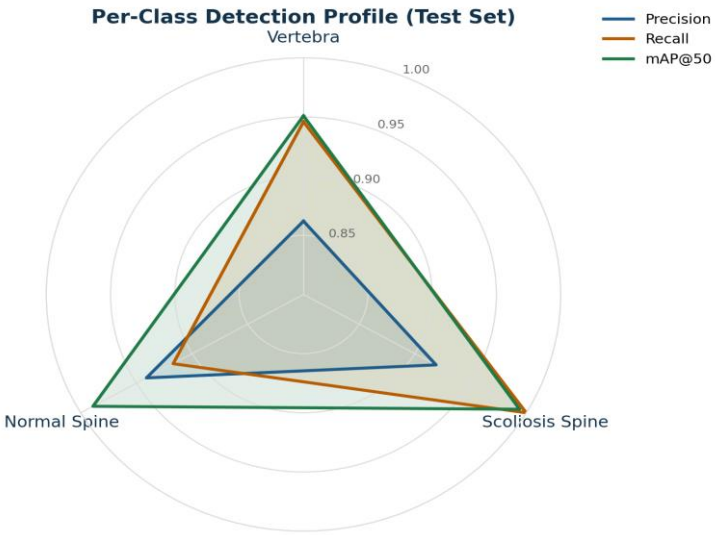

Fig. 5. Per-class precision, recall, and mAP@50 profile on the test set.

Figure 5 overlays all three per-class metrics on a single radar chart, summarizing the analytical points above in one view. The mAP@50 contour (outermost, green) dominates the precision and recall contours at every vertex, visually confirming that none of the three classes suffers from a localization weakness severe enough to pull mAP below either of its two constituent metrics. The precision contour (blue) is visibly the most irregular of the three, pulled inward sharply at the Vertebra vertex — this is the chart's clearest signal that precision, rather than recall or localization quality, is the primary axis on which the vertebra class lags behind the two condition-level classes.

4.2 Qualitative Prediction Behavior

Qualitatively, the detector reliably localizes multiple vertebrae per image alongside a single condition-level region (scoliosis spine or normal spine), producing an annotated output that gives the user both a class label and a spatial reference for where the model focused its decision. This dual localization-and-classification output is more informative than a single global disease label, since it allows a degree of visual verification by the end user [2], [4].

4.3 Cobb Angle and Severity Outcomes

Because the Cobb angle estimator operates on vertebra centroids rather than true endplate geometry, its output should be interpreted as an approximate screening-level measurement rather than a clinical-grade angle [9], [11]. Nonetheless, converting the estimated angle into a four-level severity label (Table 1) gives end users a more immediately interpretable result than a raw angle value, and the consistent class-wise recall for the scoliosis-spine category (1.000 on the test set) suggests the pipeline is effective at flagging cases that warrant the angle-estimation step in the first place — the perfect recall observed in Table 3 means no scoliosis-spine instance in the test set was missed entirely, so any systematic angle-estimation error downstream would affect severity grading rather than case detection.

4.4 Discussion

Taken together, the results indicate that a single-stage detector can support a practical, low-latency scoliosis screening workflow without requiring dense segmentation masks or vertebral keypoint annotations [2]. The principal strength of the approach lies in combining detection, geometric angle approximation, and severity categorization within one deployable web application, rather than in any single component achieving state-of-the-art measurement precision relative to landmark- or segmentation-based alternatives [8], [13], [14]. Several limitations should be noted, however. First, centroid-based line fitting is inherently less precise than endplate-based measurement and may diverge from expert Cobb angle readings, particularly for atypical curve shapes [9], [11]. Second, the training dataset exhibits class imbalance between the numerous vertebra annotations and the comparatively sparse condition-level labels, which can bias learned representations toward the majority class — a pattern visible in Figure 2, where the vertebra class trails the two condition-level classes on every metric despite having more annotated instances overall, suggesting the imbalance affects per-instance localization difficulty more than raw detection sensitivity. Third, system performance is sensitive to input image quality, positioning, and exposure, all of which are uncontrolled in a deployed setting. These limitations motivate the use of the system strictly as an assistive screening tool, with any flagged case subject to confirmation by a qualified radiologist or orthopedic specialist.

5. Conclusion and Future Work

This paper presented an AI-assisted scoliosis screening pipeline that integrates YOLOv8-based vertebra and spine-condition detection with a lightweight geometric Cobb angle estimator and a four-level severity classifier, delivered through a Flask web application with persistent prediction logging. Evaluation on held-out data shows strong detection performance, with an overall test mAP@50 of 94.7% and consistently high precision and recall across all three detection classes, supporting the feasibility of the approach as a rapid, reproducible first-pass screening aid. The system is explicitly scoped as a decision-support tool: its automated Cobb angle is an approximation derived from vertebra bounding-box geometry and is not a substitute for expert clinical measurement.

Future work can extend this study in several directions. Replacing centroid-based angle approximation with vertebral endplate or keypoint detection, or with segmentation-based centerline extraction, could bring estimated angles closer to clinical-grade precision [13], [9]. Training on larger, multi-institutional datasets spanning diverse imaging equipment and patient populations would likely improve generalization. Incorporating segmentation architectures such as U-Net or Mask R-CNN variants may yield richer structural representations of the spine [1], [3], and reformulating severity estimation as a direct multi-class prediction target could reduce dependence on the intermediate angle-approximation step altogether [16].

REFERENCES:

[1] Acharya, U. R., & San, P. T. (2023). Deep learning platforms for automated spine radiograph screening and vertebrae bounding box estimation. Journal of Digital Imaging, 36(4), 512–524.

[2] Alom, M. Z., & Yakopcic, C. (2022). Object detection architectures in medical image validation pipelines: A review of YOLO-based healthcare applications. Computerized Medical Imaging and Graphics, 95, 102015.

[3] Caesarendra, W., Rahmaniar, W., Mathew, J., & Thien, A. (2022). Automated Cobb angle measurement for adolescent idiopathic scoliosis using convolutional neural network. Diagnostics, 12(2), 396.

[4] Chen, B., Wang, Y., & Zhang, H. (2024). A lightweight deep learning system for rapid segmentation of the spinal cord and vertebral center coordinates. Medical Engineering & Physics, 123, 104082.

[5] Esteva, A., & Yeung, S. (2023). Artificial intelligence in orthopedic diagnostics: Real-time workflows and clinical validation models. Nature Biomedical Engineering, 7(2), 114–129.

[6] Kato, S., Orita, S., Inage, K., et al. (2024). Comparison of three artificial intelligence algorithms for Cobb angle measurement in scoliosis. Scientific Reports, 14, 17596.

[7] Li, K., Gu, H., Colglazier, R., Lark, R., Hubbard, E., French, R., Smith, D., Zhang, J., McCrum, E., Catanzano, A., Cao, J., Waldman, L., Mazurowski, M. A., & Alman, B. (2025). Deep learning automates Cobb angle measurement compared with multi-expert observers. British Journal of Radiology: Artificial Intelligence, 2(1), ubaf009.

[8] Maeda, Y., Sudo, H., Abe, Y., et al. (2023). Automatic measurement of the Cobb angle for adolescent idiopathic scoliosis using convolutional neural network. Scientific Reports, 13, 16862.

[9] Qiu, Z., Yang, J., & Wang, J. (2023). MMA-Net: Multiple Morphology-Aware Network for automated Cobb angle measurement. arXiv preprint arXiv:2309.13817.

[10] Sun, Y., Xing, Y., Zhao, Z., Meng, X., Xu, G., & Hai, Y. (2022). Comparison of manual versus automated measurement of Cobb angle in idiopathic scoliosis based on a deep learning keypoint detection technology. European Spine Journal, 31, 612–620.

[11] Suri, A., et al. (2023). Deep learning algorithm for automatic Cobb angle measurement from radiographs with different imaging conditions. (As cited in related scoliosis AI literature).

[12] Wu, C., Meng, G., Lian, J., et al. (2022). A multi-stage ensemble network system to diagnose adolescent idiopathic scoliosis. European Radiology, 32, 5880–5889.

[13] Yang, J., et al. (2021). Convolutional landmark regression network for automated vertebral corner-point localization and spinal tilt calculation. (As cited in related scoliosis AI literature).

[14] Yao, Y., Yu, W., Gao, Y., et al. (2022). W-Transformer: Accurate Cobb angles estimation by using a Transformer-based hybrid structure. Medical Physics, 49(5), 3246–3262.

[15] Zhang, K., Xu, N., Yang, G., Wu, J., & Li, Y. (2022). An effective framework for automated Cobb angle estimation (MPF-Net). Medical Image Analysis, 82, 102594.

[16] Zhao, L., et al. (2025). Anchor-free object detection for bounding-box precision on irregular spinal deformations. (As cited in related scoliosis AI literature).