

Pose as Clinical Prior: Learning Dual Representations for Scoliosis Screening

Zirui Zhou¹, Zizhao Peng^{1,2}, Dongyang Jin¹, Chao Fan³,
Fengwei An¹, and Shiqi Yu^{1*}

¹ Southern University of Science and Technology, Shenzhen, China

² The Hong Kong Polytechnic University, Hong Kong, China

³ School of Artificial Intelligence, Shenzhen University, Shenzhen, China

Abstract. Recent AI-based scoliosis screening methods primarily rely on large-scale silhouette datasets, often neglecting clinically relevant postural asymmetries—key indicators in traditional screening. In contrast, pose data provide an intuitive skeletal representation, enhancing clinical interpretability across various medical applications. However, pose-based scoliosis screening remains under-explored due to two main challenges: (1) the scarcity of large-scale, annotated pose datasets; and (2) the discrete and noise-sensitive nature of raw pose coordinates, which hinders the modeling of subtle asymmetries. To address these limitations, we introduce **Scoliosis1K-Pose**, a 2D human pose annotation set that extends the original Scoliosis1K dataset, comprising 447,900 frames of 2D keypoints from 1,050 adolescents. Building on this dataset, we introduce the **Dual Representation Framework (DRF)**, which integrates a continuous *skeleton map* to preserve spatial structure with a discrete *Postural Asymmetry Vector (PAV)* that encodes clinically relevant asymmetry descriptors. A novel *PAV-Guided Attention (PGA)* module further uses the PAV as clinical prior to direct feature extraction from the skeleton map, focusing on clinically meaningful asymmetries. Extensive experiments demonstrate that DRF achieves state-of-the-art performance. Visualizations further confirm that the model leverages clinical asymmetry cues to guide feature extraction and promote synergy between its dual representations. The dataset and code are publicly available at <https://zhouzi180.github.io/Scoliosis1K/>.

Keywords: Scoliosis Screening · Biometrics in Healthcare · Computer Vision.

1 Introduction

Scoliosis is a complex, three-dimensional spinal deformity that affects approximately 3.1% of adolescents worldwide [10]. If untreated, it may lead to chronic pain, respiratory problems, and a diminished quality of life [16]. Early and accurate screening is therefore essential. Traditional screening methods such as visual inspection, the Adam’s forward bend test, and trunk rotation measurements are time-consuming, require specialized clinical expertise, and raise privacy concerns due to the need for exposed torso examinations. These limitations hinder large-scale screening, especially in resource-constrained settings [9].

Artificial intelligence (AI) offers a promising solution to overcome these limitations. Early AI-based methods relied on static trunk imaging using photogrammetry [19,18] or 3D surface topography [2,1] to assess back asymmetry. While reducing the need for expert interpretation, these methods still depend on controlled environments and require exposed torsos. Recent approaches explore dynamic biomarkers in naturalistic settings. Visual gait analysis, among these methods, captures walking-related movements and enhances both privacy and scalability in scoliosis screening. Zhou et al. [21] introduced this paradigm with the Scoliosis1K dataset, comprising human silhouettes extracted from natural walking videos. They also introduced ScoNet-MT, a CNN-based model designed to extract fine-grained body shape features for scoliosis screening. Despite improving data accessibility and reducing acquisition constraints, these methods often overlook clinically relevant asymmetries—such as shoulder imbalance and pelvic tilt [3]—that are critical for accurate and interpretable assessments.

Human pose data, with its explicit skeletal structure, offers a natural way to integrate these clinical evaluation criteria and has demonstrated value in various medical applications [12,14,8,15]. Despite its potential, the use of pose data in scoliosis screening remains underexplored, mainly due to two key challenges: (1) the scarcity of large-scale, annotated pose datasets; and (2) the discrete and noise-sensitive nature of raw pose coordinates, which hinders the modeling of subtle body asymmetries.

To bridge this gap, we introduce **Scoliosis1K-Pose**, a 2D human pose annotation set that augments the original Scoliosis1K dataset [21] for scoliosis screening. It provides 2D keypoints extracted from natural walking videos of adolescents. Each frame includes 17 keypoints in the MS-COCO format [11], extracted using ViTPose [17]. Building on this dataset, we propose the **Dual Representation Framework (DRF)**, a pose-based method for scoliosis screening. DRF integrates two complementary pose representations: a continuous *skeleton map* that preserves the global spatial structure of the pose, and a discrete *Postural Asymmetry Vector (PAV)* that explicitly encodes clinically relevant asymmetry descriptors (e.g., vertical, midline, and angular deviations) between paired keypoints. To synergize these representations, our novel *PAV-Guided Attention (PGA)* employs the PAV as a clinical prior to guide the network’s feature extraction, directing attention to the most clinically relevant asymmetries within the skeleton map. This enables DRF to capture subtle yet clinically meaningful postural patterns for scoliosis assessment.

Our main contributions are threefold:

- **Scoliosis1K-Pose:** A new 2D human pose annotation set introduced in this work, augmenting the original Scoliosis1K dataset with detailed keypoint annotations for scoliosis screening.
- **Dual Representation Framework (DRF):** A novel, clinically-informed framework that synergistically combines continuous skeleton maps with discrete, clinically-derived PAV. This is implemented through a PGA module that embeds clinical prior into the feature learning process.
- **State-of-the-Art Performance:** Extensive experiments demonstrate that DRF can achieve superior performance in scoliosis screening. Visualizations further confirm that the model leverages clinical asymmetry cues to guide feature extraction and promote synergy between its dual representations.




Fig. 1: Examples from the Scoliosis1K dataset. Left: silhouettes representing (a) positive, (b) neutral, and (c) negative cases. Right: corresponding 2D skeletal keypoints estimated by ViTPose [17].

2 Dataset

Our research is built upon the Scoliosis1K dataset [21], which originally provided only silhouette data from 1,493 walking videos of 1,050 adolescents. To enable fine-grained postural analysis, we introduce *Scoliosis1K-Pose*, a large-scale 2D human pose annotation set that extends the original Scoliosis1K.

Scoliosis1K-Pose was constructed by processing 447,900 frames using ViTPose [17] to extract 17 anatomical keypoints per frame, following the MS-COCO format [11]. The extracted keypoints provide detailed skeletal representations of facial features (eyes, ears, nose), upper limbs (shoulders, elbows, wrists), and lower limbs (hips, knees, ankles). This augmentation transforms Scoliosis1K into a comprehensive multi-modal resource that integrates silhouette and 2D pose data, as illustrated in Figure 1. Participants are categorized as positive (scoliosis), neutral (borderline), or negative (non-scoliosis) according to established screening protocols. For detailed information on the original data collection and its characteristics, please refer to [21].

3 Method

We propose the **Dual Representation Framework (DRF)**, a novel approach to scoliosis screening that explicitly integrates clinical priors into a deep learning pipeline. The core idea is to emulate how clinicians assess scoliosis by identifying postural asymmetries. As shown in Figure 2(a), the framework first transforms raw pose data into two complementary representations: a continuous *skeleton map* and a discrete *Postural Asymmetry Vector (PAV)*. The skeleton map is processed by a feature encoder. Concurrently, the PAV, serving as a clinical prior, guides a *PAV-Guided Attention (PGA)* module to generate attention weights that modulate the extracted features. Finally, the refined features are used for scoliosis assessment.

The following sections describe the construction of the dual pose representations (Section 3.1) and the PAV-guided feature learning and assessment (Section 3.2).




Fig. 2: Overview of the proposed Dual Representation Framework (DRF). (a) The overall pipeline, from raw pose data to the final scoliosis assessment. (b) The construction process of the Postural Asymmetry Vector (PAV). (c) The architecture of the PAV-Guided Attention (PGA) module, which generates channel and spatial attention weights.

3.1 Dual Pose Representations

Given a sequence of 2D pose keypoints, we first normalize the data to ensure robustness to variations in subject position and scale. This involves aligning the pelvis by translating the midpoint of the hip joints to the origin, followed by height normalization to a standard height of 128 pixels. From the normalized pose, we construct two complementary representations.

Skeleton Map: Continuous Gaussian Heatmaps. To convert sparse 2D keypoints into a dense, silhouette-like representation suitable for convolutional networks, we construct a two-channel skeleton map. This approach, inspired by related works [6,4], provides a richer input than raw coordinates by capturing both local joint details and global skeletal structure. It consists of two channels:

- **Keypoint Map:** Encodes each joint’s location and confidence with a Gaussian heatmap:

$$J(i, j) = \sum_k \exp \left(-\frac{(i - x_k)^2 + (j - y_k)^2}{2\sigma^2} \right) \cdot c_k, \quad (1)$$

where (x_k, y_k) and σ are the coordinates and confidence score of the k -th joint, respectively.

- **Limb Map:** Represents skeletal connections as line-like heatmaps, preserving structural integrity:

$$L(i, j) = \sum_n \exp \left(-\frac{d_{L2}((i, j), S[n^-, n^+])^2}{2\sigma^2} \right) \cdot \min(c_{n^-}, c_{n^+}), \quad (2)$$

where $S[n^-, n^+]$ denotes the limb segment between joints n^- and n^+ , and $d_{L2}(\cdot)$ is the Euclidean distance from pixel (i, j) to this segment.

The resulting skeleton map yields a robust, continuous representation that mitigates pose estimation noise and serves as the primary input for feature extraction.

PAV: Discrete Clinical Prior. To explicitly incorporate clinical priors into our framework, we introduce the Postural Asymmetry Vector (PAV). Instead of relying on a few hand-picked landmarks, which are susceptible to noise and may miss subtle compensatory patterns, the PAV is designed to capture a holistic asymmetry profile. The PAV is a matrix $v \in \mathbb{R}^{P \times M}$, where P denotes the number of anatomically symmetric keypoint pairs (e.g., eyes, ears, shoulders, elbows, wrists, hips, knees, and ankles in COCO17 format), with the p -th pair represented as (k_p^L, k_p^R) , and M denotes the number of asymmetry metrics.

The PAV is constructed in two stages, as illustrated in Fig. 2(b):

1. Asymmetry Metric Computation. For each frame, a raw matrix $\tilde{v} \in \mathbb{R}^{P \times M}$ is computed using the following three metrics for each keypoint pair:

- Vertical Deviation (L_{VD}): $|y_p^L - y_p^R|$, quantifies height differences, directly corresponding to clinical signs like shoulder and pelvic imbalance.

- Midline Deviation (L_{MD}): $\left| \frac{x_p^L + x_p^R}{2} - x_{\text{midline}} \right|$, measures lateral drift from the body's central axis (x_{midline} is defined by the hip center), reflecting spinal curvature.

- Angular Deviation (α_{AD}): $\left| \arctan \left(\frac{y_p^L - y_p^R}{x_p^L - x_p^R} \right) \right|$, captures the tilt of the segment connecting a keypoint pair, indicative of rotational asymmetries in the trunk and limbs.

2. Sequence-Level PAV Refinement. The raw matrix $\tilde{v}$ is refined across the sequence to obtain a stable PAV v using the following steps:

- (1) Outlier Removal: Statistical filtering is applied independently to each element in $\tilde{v}$ using the interquartile range (IQR) method.

- (2) Temporal Aggregation: For each metric, the mean of valid measurements across all frames is computed, summarizing temporal variation into a stable sequence-level descriptor.

- (3) Normalization: Min-Max scaling is applied across the dataset for all $P \times M$ dimensions, standardizing values to the $[0, 1]$ range and ensuring consistent feature scaling.

This two-stage design enables robust quantification of global postural asymmetry and allows the model to capture complex, clinically relevant inter-joint relationships. The resulting PAV acts as a clinical prior to guide the subsequent feature learning process.

3.2 PAV-Guided Feature Learning and Assessment

Our DRF uses the clinical prior encoded in the PAV to guide visual feature learning from the skeleton map. This is accomplished using a feature encoder followed by our novel PAV-Guided Attention (PGA) module.

Feature Encoder. First, the feature encoder processes the input skeleton map sequence to extract high-level features. For a fair comparison, we adopt the encoder architecture from ScoNet-MT [21]. It uses a ResNet-based backbone followed by temporal and horizontal pooling to generate a feature map $F_{\text{enc}} \in \mathbb{R}^{C \times H}$, where $C = 256$ is the number of channels and H is the number of horizontal body segments. This feature map F_{enc} captures general postural patterns from the data.

Prior-Guided Attention (PGA) Module. To ensure these features are clinically relevant, we introduce the PGA Module. Unlike standard attention that relies solely on data-driven correlations, our module explicitly incorporates a clinical prior to generate attention weights. In our framework, the prior is the PAV, which guides the model to focus on features critical for scoliosis screening.

As shown in Fig. 2(c), the module receives two inputs: the encoded features F_{enc} and the PAV v . It then generates both channel-wise (w_c) and spatial-wise (w_s) attention weights using the PAV:

$$w_c = \sigma(W_c \cdot v), \quad w_s = \sigma(W_s * v),$$

where σ denotes the sigmoid function, W_c is a learnable linear layer, and W_s is a learnable 1D convolution. The resulting attention weights, $w_c \in \mathbb{R}^{C \times 1}$ and $w_s \in \mathbb{R}^{1 \times H}$, are used to recalibrate the encoded features via element-wise multiplication:

$$F_{\text{out}} = F_{\text{enc}} \odot w_c \odot w_s,$$

where $\odot$ denotes element-wise multiplication with broadcasting. This produces a refined feature map F_{out} that integrates visual patterns with clinical priors, yielding a more informative representation for downstream assessment.

Scoliosis Assessment and Training Objective. The refined features F_{out} are passed to a classification head with fully connected layers to produce the final logits for scoliosis assessment. The entire framework is trained end-to-end with a combined loss function, following the approach of ScoNet-MT [21]:

$$L_{\text{total}} = L_{\text{ce}} + L_{\text{triplet}}.$$

4 Experiments

4.1 Setup

Evaluation Protocol. We use the standard dataset split from [21] (745 training, 748 test sequences) with real-world class distribution (positive:neutral:negative = 1:1:8). To address class imbalance and enable fair comparison, we report overall accuracy along with macro-averaged precision, recall, and F1-score.

Implementation Details. Our method is implemented using PyTorch [13] and OpenGait [5]. Following ScoNet-MT [21], we retain the same network architecture and training protocol, with two key modifications: (1) the Conv0 layer is adapted to process skeleton maps instead of silhouettes, and (2) the proposed PGA module is integrated to guide feature extraction.

Table 1: Comparison with state-of-the-arts on Scoliosis1K dataset. The best results are in bold. ScoNet-MT^{ske} refers to ScoNet-MT [21] adapted for skeleton map input.

Input	Method	Total Accuracy	Macro-average		
			Prec	Rec	F1
Silhouette	ScoNet-MT [21]	82.0	84.0	75.9	75.4
Pose	OF-DDNet [12]	74.1	68.6	68.4	67.0
Coordinates	GPGait [7]	74.9	69.2	68.9	67.5
Skeleton Map	ScoNet-MT ^{ske}	82.5	81.4	74.3	76.6
	DRF (ours)	86.0	84.1	79.2	80.8




Fig. 3: Visualization of feature response heatmaps. (a) ScoNet-MT. (b) ScoNet-MT^{ske}. (c) Our DRF.

4.2 Comparison with State-of-the-Art Methods

We perform comprehensive comparisons with state-of-the-art methods using various input representations. For silhouette-based methods, we compare our model with ScoNet-MT [21]. For pose coordinate-based approaches, we include OF-DDNet [12] and GP-Gait [7]. OF-DDNet is adapted to process 2D pose inputs to ensure fair comparison.

Quantitative Results. As shown in Table 1, our proposed DRF achieves state-of-the-art performance across all evaluation metrics. Pose coordinate-based methods, such as GP-Gait and OF-DDNet, yield the lowest performance, with F1-scores of 67.5% and 67.0%, respectively. Image-based representations perform better overall. The silhouette-based ScoNet-MT model significantly outperforms pose coordinate-based methods. When adapted to use our skeleton map input (ScoNet-MT^{ske}), the model achieves an additional 0.5% improvement in accuracy and a 1.2% gain in F1-score. These results confirm that the skeleton map captures rich structural information. Importantly, integrating the PGA module with PAV enables our final DRF model to significantly outperform the ScoNet-MT^{ske} baseline. Specifically, accuracy increases by 3.5%, and the F1-score improves by 4.2%. This significant improvement demonstrates that explicitly guiding the model with clinical priors on postural asymmetry is highly effective for scoliosis screening.

Qualitative Results. As shown in Figure 3, we visualize the intermediate feature responses using the visualization technique from [20]. ScoNet-MT with silhouette input (Fig.3 (a)) focuses on isolated regions (e.g., head and shoulders), missing holistic structural relationships crucial for scoliosis screening. With skeleton map as input, ScoNet-MT^{ske} (Fig.3 (b)) preserves structural information but mainly attends to one body side. By incorporating our PGA module guided by the PAV, our DRF (Fig. 3 (c)) more effectively attends to key regions on both sides of the body (e.g., imbalanced shoulders and pelvic), which strongly corresponds to our clinical prior integration objective. These visualizations confirm that our PGA module, when guided by PAV, effectively guides feature extraction using clinical asymmetry cues, creating the synergistic relationship between our dual representations.

Table 2: Ablation studies on Scoliosis1K dataset. The best results are in bold.

Channel Spatial	Total Macro-average				Guidance Source	Total Macro-average			
	Acc	Prec	Rec	F1		Acc	Prec	Rec	F1
	82.5	81.4	74.3	76.6	Self-attention	81.4	84.5	73.9	75.3
✓	79.7	80.9	75.7	74.9	All-Ones	82.2	80.8	74.3	74.7
	81.7	77.1	80.8	78.1	Random	75.4	82.2	62.7	57.9
✓	86.0	84.1	79.2	80.8	Learnable	69.4	84.2	53.2	56.4
					PAV (ours)	86.0	84.1	79.2	80.8

(a) Component Analysis of PGA Module. All variants of the PGA module are guided by our PAV.

(b) Comparison of PAV and alternative guidance sources for the PGA Module.

4.3 Ablation Studies

To validate the design of our framework, we conduct two sets of ablation studies. First, we analyze the effectiveness of the PGA module and its components. Second, we verify that the proposed PAV provides superior guidance compared to alternative strategies.

Effectiveness of the PGA Module. We first evaluate the overall contribution of the PGA module, followed by an analysis of its channel-wise and spatial-wise attention branches. We use ScoNet-MT^{ske} as our baseline, which operates without any attention guidance. As shown in Table 2a, integrating the full PGA module (guided by PAV) into the baseline significantly improves performance, with a 3.5% increase in accuracy and a 4.2% gain in F1-score. Furthermore, both attention branches are essential. Disabling either the channel or spatial branch results in a notable performance decline compared to the full model. This highlights the synergistic interaction between channel and spatial attention, which is essential for effectively leveraging the clinical cues provided by the PAV.

Effectiveness of PAV as a Clinical Prior. After confirming the effectiveness of the PGA module’s architecture, we next examine the importance of the guidance source. We compare the PAV with several alternative guidance vectors within the same PGA module architecture. The evaluated alternatives include:

(1) *No External Prior (Self-attention)*: Guidance is generated directly from feature maps, following a standard self-attention mechanism. This setup evaluates whether the performance gain arises solely from the attention mechanism rather than from clinical prior.

(2) *Fixed, Uninformative Priors*: Fixed vectors with the same dimensionality as the PAV are used: (a) All-Ones, representing uniform attention, and (b) Random, representing arbitrary, unstructured guidance. These variants test whether a static structural bias could achieve similar benefits.

(3) *Learnable Prior*: The PAV is replaced by a learnable parameter vector, enabling the network to derive a data-driven guidance signal without explicit clinical constraints.

As shown in Table 2b, using PAV to guide the PGA module significantly outperforms all alternatives. The self-attention approach results in a 5.5% drop in F1-score, indicating that the explicit prior offers greater value than internally derived correlations.

The uninformative priors (All-Ones, Random) and the learnable prior also perform significantly worse, with F1-score reductions ranging from 6.1% to 22.9%. These results confirm that the performance improvement arises not merely from using an attention module, but crucially from the rich, clinically relevant information encoded in the PAV. This underscores the value of incorporating explicit clinical prior for robust scoliosis screening.

5 Conclusion

This study proposes a novel, clinically informed approach for pose-based scoliosis screening. The proposed Dual Representation Framework (DRF), combined with the Scoliosis1K-Pose dataset, integrates continuous skeletal structures and clinically informed asymmetry descriptors to improve model accuracy and interpretability, enabling accessible and privacy-preserving scoliosis screening. We hope that the work can inspire more research on scoliosis screening using massive video data.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (Grant 62476120) and the Scientific Foundation for Youth Scholars of Shenzhen University (Grant 868-000001033383).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Adankon, M.M., Chihab, N., Dansereau, J., Labelle, H., Cheriet, F.: Scoliosis follow-up using noninvasive trunk surface acquisition. *IEEE Transactions on Biomedical Engineering* **60**(8), 2262–2270 (2013)
2. Adankon, M.M., Dansereau, J., Labelle, H., Cheriet, F.: Non invasive classification system of scoliosis curve types using least-squares support vector machines. *Artificial intelligence in medicine* **56**(2), 99–107 (2012)
3. of the Chinese Orthopaedic Association, S.S.G.: Clinical practice guideline and pathway for screening adolescent idiopathic scoliosis in china. *Chinese Journal of Orthopaedics* **40**(23), 1574–1582 (2020)
4. Duan, H., Zhao, Y., Chen, K., Lin, D., Dai, B.: Revisiting skeleton-based action recognition. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 2969–2978 (2022)
5. Fan, C., Liang, J., Shen, C., Hou, S., Huang, Y., Yu, S.: Opengait: Revisiting gait recognition towards better practicality. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9707–9716 (2023)
6. Fan, C., Ma, J., Jin, D., Shen, C., Yu, S.: Skeletongait: Gait recognition using skeleton maps. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 38, pp. 1662–1669 (2024)
7. Fu, Y., Meng, S., Hou, S., Hu, X., Huang, Y.: Gpgait: Generalized pose-based gait recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 19595–19604 (2023)

8. Islam, M.S., Rahman, W., Abdelkader, A., Lee, S., Yang, P.T., Purks, J.L., Adams, J.L., Schneider, R.B., Dorsey, E.R., Hoque, E.: Using ai to measure parkinson's disease severity at home. *NPJ digital medicine* **6**(1), 156 (2023)
9. Kadhim, M., Lucak, T., Schexnayder, S.: Current status of scoliosis school screening: targeted screening of underserved populations may be the solution. *Public Health* **178**, 72–77 (2020)
10. Li, M., Nie, Q., Liu, J., Jiang, Z.: Prevalence of scoliosis in children and adolescents: a systematic review and meta-analysis. *Front Pediatr* **12**, 1399049 (2024)
11. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: 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)
12. Lu, M., Poston, K., Pfefferbaum, A., Sullivan, E.V., Fei-Fei, L., Pohl, K.M., Niebles, J.C., Adeli, E.: Vision-based estimation of mds-updrs gait scores for assessing parkinson's disease motor severity. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III* 23. pp. 637–647. Springer (2020)
13. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
14. Quan, Y., Zhang, C., Guo, R., Qian, X.: Causality-informed fusion network for automated assessment of parkinsonian body bradykinesia. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 78–88. Springer (2024)
15. Wang, D., Yuan, K., Muller, C., Blanc, F., Padoy, N., Seo, H.: Enhancing gait video analysis in neurodegenerative diseases by knowledge augmentation in vision language model. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 251–261. Springer (2024)
16. Weinstein, S.L., Dolan, L.A., Cheng, J.C., et al.: Adolescent idiopathic scoliosis. *Lancet* **371**(9623), 1527–1537 (2008)
17. Xu, Y., Zhang, J., Zhang, Q., Tao, D.: Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in Neural Information Processing Systems* **35**, 38571–38584 (2022)
18. Yang, J., Zhang, K., Fan, H., Huang, Z., Xiang, Y., Yang, J., He, L., Zhang, L., Yang, Y., Li, R., et al.: Development and validation of deep learning algorithms for scoliosis screening using back images. *Communications biology* **2**(1), 390 (2019)
19. Zhang, T., Zhu, C., Zhao, Y., Zhao, M., Wang, Z., Song, R., Meng, N., Sial, A., Diwan, A., Liu, J., et al.: Deep learning model to classify and monitor idiopathic scoliosis in adolescents using a single smartphone photograph. *JAMA Network Open* **6**(8), e2330617–e2330617 (2023)
20. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2921–2929 (2016)
21. Zhou, Z., Liang, J., Peng, Z., Fan, C., An, F., Yu, S.: Gait patterns as biomarkers: A video-based approach for classifying scoliosis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 284–294. Springer (2024)