ORIGINAL ARTICLE

Applying a Convolutional Neural Network Based Iterative Algorithm to Automatically Measure Spinal Curvature on Radiographs for Children with Scoliosis

Solvin Sigurdson[1] · Jason Wong[2] · Marek Reformat[2] · Edmond Lou[2]

Abstract

Purpose Accurate measurement of spinal curvature for adolescent idiopathic scoliosis (AIS) is important because it affects treatment decisions. Currently, the Cobb angle measured on a radiograph is the gold standard for spinal curvature assessment. However, manual measurements introduce inter- and intra-observer reliability challenges, and while fully automatic methods have been developed, performance could be improved. This paper reported a new approach using convolutional neural networks (CNNs) and an iterative vertebra location algorithm to calculate the Cobb angle automatically by segmenting the spinal and vertebral boundaries on posteroanterior radiographs.

Methods Two CNNs for spinal column and vertebra segmentation were trained using 110 and 272 images, respectively. An iterative vertebra location algorithm was developed to localize individual vertebrae in the spinal column for segmentation. To evaluate the accuracy of the automatic Cobb angle measurements calculated from the vertebra segmentations, 100 new radiographs were used. The mean absolute difference (MAD), standard deviation of absolute differences (SD), and percent within clinical acceptance ($\leq 5°$) between manual and automatic measurements were reported as evaluation metrics.

Results The MAD $\pm$ SD was $2.8° \pm 2.8°$ and 88% of the measurements were within 5° of the manual measurements. The result was comparable to other literature, and this method worked for a wide range of curve severities. The average automatic measurement time per image was 90 s, which is clinically acceptable.

Conclusion The automatic measurement method based on CNNs provided a comparable accuracy and speed on spinal curvature measurements on radiographs. It could be a valuable tool for reducing clinical workload and measurement variation.

Keywords Convolutional neural network · Machine learning · Radiograph · Scoliosis · Spinal curvature

1 Introduction

Adolescent idiopathic scoliosis (AIS) is a three-dimensional spinal condition, characterized by lateral curvature and axial vertebral rotation. It affects approximately 3% of adolescents [1]. The Cobb angle, which is the gold standard for quantifying spinal curvature, is measured on posteroanterior (PA) radiographs by identifying the angle between the upper end plate of the superior vertebra and the lower end plate of the inferior vertebra [2], shown in Fig. 1a. To decide which treatment regimen is appropriate, the Cobb angle is measured at a scoliosis clinic. If the measured Cobb angle has increased by more than 6° between two consecutive visits, it indicates that the curve has progressed, and treatment may be recommended to stop the curve progression. Therefore, it is essential to have accurate measurements which ensure that patients receive appropriate treatment in a timely manner. However, manual measurements introduce inter-observer and intra-observer variations and require time and effort to measure in the clinic [3].

Machine learning is a rapidly developing field that has successfully tackled problems that previously challenged traditional computing methods by simulating human learning methods. This is useful because it allows computer programming to solve more general and wide-ranging

✉ Edmond Lou
elou@ualberta.ca

1 Division of Engineering Science, Faculty of Applied Science and Engineering, University of Toronto, Toronto, ON, Canada

2 Department of Electrical and Computer Engineering, 11-263 Donadeo Innovation Centre for Engineering, University of Alberta, 9211-116 St., Edmonton, AB T6G 1H9, Canada
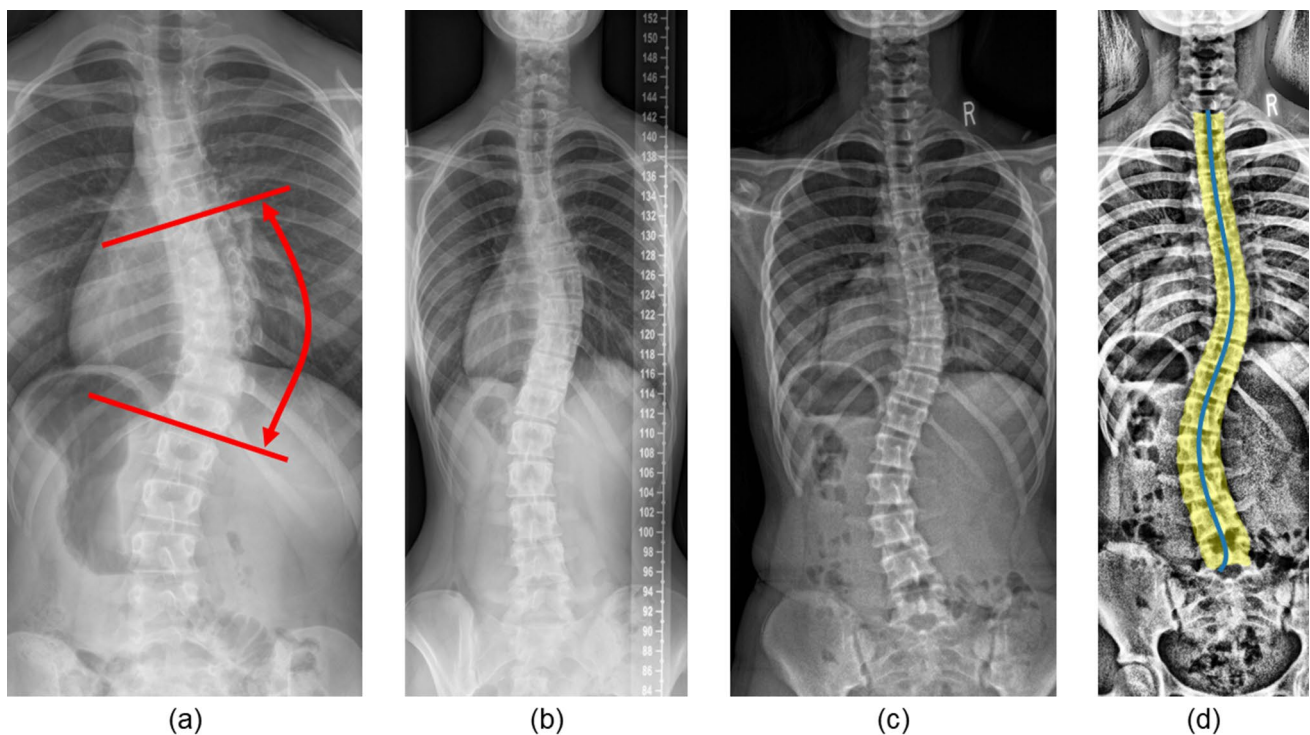
Fig. 1 **a** A Cobb angle measurement on a posteroanterior radiograph; **b** a standard X-ray radiograph, **c** a low-dose radiation EOS radiograph, **d** pre-processed version of the EOS radiograph from **c** with segmented spinal column (yellow) and midline (blue)

problems that do not have a specific set of outputs. One such problem is computer vision and image segmentation, which has been changed by the introduction of the convolutional neural network (CNN). These networks search over an image several times to identify and mark distinguishing features [4]. Of particular success is the U-net from Ronneberger et al., which is a CNN architecture specially designed for segmenting images with small training datasets [5]. This makes the U-net perfect for medical imaging applications with limited training samples.

Recently, Horng et al., developed a fully automatic program that used image processing to isolate the spinal column and a U-net CNN to segment the individual vertebrae. Using the vertebra segmentations, they measured the Cobb angle with an inter-method reliability of 0.94–0.97 between manual and automatic measurements. However, their test set comprised of only 35 images and contained only mild Cobb angles of less than 20°, which is a small subset of overall cases [6]. Fu et al. used a two-network architecture for estimating landmarks and calculating the Cobb angle using these landmarks [7]. They achieved a MAD of 3.15° on 240 radiographs. However, the distribution of Cobb angle severities in their test set is unclear. Other groups have reported on their own automated methods, but do not achieve comparable results to these two papers [8–12].

The objectives of this paper were to report a new quality-based iterative CNN algorithm that automatically measures Cobb angles on PA radiographs with a wide range of curve severities and to determine the accuracy of this method by comparing automatic measurements with clinical records. Developing a fully automatic and accurate Cobb angle measurement algorithm would reduce measurement variation and clinical workload, therefore improving diagnostic accuracy and streamlining clinical workflow. The accuracy of the Cobb angle measurement is important because it affects treatment decisions.

2 Materials and Methods

2.1 Data

From a local scoliosis clinic, over 600 PA radiographs of children with AIS were available for use. Among these radiographs, 300 were acquired using the standard X-ray system (between 2010 and 2015), while the remainder were acquired using the new EOS X-ray imaging system (between 2015 and 2020), which uses an ultra-low dose of radiation. Figure 1b and c compare the image quality of the radiographs obtained from the standard and EOS X-ray systems, respectively. The average size of a PA radiograph obtained

from the standard and EOS X-ray systems is 4820×2335 and 3465×1600 pixels, respectively. Out of the 600 available PA radiographs, 238 were randomly selected and broken down into three groups (120, 18, and 100 images). The first 120 PA images were used for spinal column segmentation, of which 110 were used for training and 10 were used for validation. The 18 PA images were used to generate 282 vertebra images which were used for training (272 images) and validation (10 images) of vertebra segmentation. The last 100 images were used for automatic Cobb angle measurement accuracy testing. Ethics approval (Pro00102044) was granted from the University of Alberta Health Research Ethics Board.

2.2 Proposed Method

The automatic Cobb angle measurement algorithm was broken down into two steps: (a) segmentation of the spinal column and (b) segmentation of the individual vertebrae. The spinal column segmentation aimed to identify the region starting from the top thoracic vertebra (T1) to the bottom lumbar vertebra (L5) on a PA radiograph. Using the spinal column as a guide, the individual vertebrae were located with an iterative algorithm and subsequently segmented. The CNNs used for segmentation were developed using Python, coupled with Tensorflow libraries, and were trained on a Linux virtual machine hosted on the Industry Sandbox & AI Computing (ISAIC) supercomputer, using an NVIDIA Tesla V100 16 GB GPU and an Intel Xeon Gold 6138 dual processor. All images related to CNN development were labelled manually by the primary author using a MATLAB segmentation graphical user interface. The primary author, a novice scoliosis researcher, used 10 extra images to practice and confirm with a senior author who had 20 years of scoliosis research experience prior to labelling the training images.

2.2.1 Spinal Column Pre-processing

To reduce extraneous information from being passed to the CNN, a region of interest (ROI) centered around the spinal column was identified and cropped from the PA radiographs. First, the head was identified, as it was positioned directly above the T1 vertebra. The image intensity of the head was higher than its immediate surrounding area, so the maximum intensity vertical projection of the top 50 rows of the radiograph was calculated. The ROI was determined by setting the region of highest intensity in this projection as the ROI's center, then extending it to 500 pixels on either side, creating a total width of 1000 pixels. Contrast limited adaptive histogram equalization (CLAHE) was applied to the ROI to accentuate the spinal column boundaries. The cropped and histogram equalized images were then resized to 256×128 pixels to standardize the input image size for the CNN. These pre-processed images, along with their labels, were used to train the CNN for spinal column segmentation. Labelling the spinal column involved annotating a ground truth semantic segmentation, similar to the yellow highlighted area in Fig. 1d, where the whole spinal column from T1–L5 was labelled.

2.2.2 Spinal Column Segmentation

A CNN with a U-net architecture was used to segment the spinal column from the pre-processed images [5]. The architecture, which is similar to that of Horng et al. [6], is illustrated in Fig. 2. The spinal column CNN was trained for 400 epochs using a binary cross entropy loss function and an Adam optimizer with a learning rate of 10^{-4} [13]. A dropout probability of 0.1 was used after each downsampling and upsampling layer, and a batch size of 1 was used.

A total of 120 PA radiographs were manually labelled and split into a 110-image training set and 10-image validation

Fig. 2 Network architecture for the spinal column segmentation network. The vertebra segmentation network varies only in terms of input image size
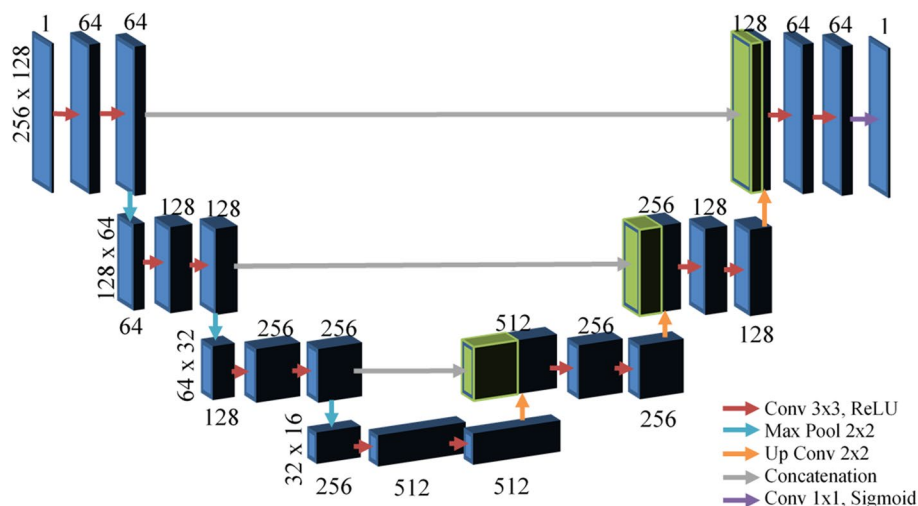
set. The training set consisted of half regular and half EOS radiographs, while the validation set contained only EOS radiographs. Data augmentation methods of random horizontal flips and rotations ($\leq 10°$) were applied to increase the effective training set size, and therefore the robustness of the CNN.

The segmentations underwent post-processing to optimize performance in the upcoming iterative vertebra location algorithm. This consisted of keeping only the largest connected component of the segmentation and estimating the spinal column shape by fitting a tenth-degree polynomial to the segmentation. This polynomial was denoted as the "midline" and is illustrated in Fig. 1d, along with the pre-processed radiograph and spinal column segmentation.

2.2.3 Vertebra Pre-processing

Square images of the vertebrae were manually cropped from the PA radiographs and labelled to train a CNN for automatic vertebra segmentation. Like the spinal column images, CLAHE was applied to these images to highlight the vertebral boundaries. The pre-processed images were then resized to 128×128 pixels. Figure 3a illustrates a pre-processed cropped vertebra image. These pre-processed images and their respective labels were used to train the CNN for vertebra segmentation. Labelling the vertebrae involved annotating a ground truth semantic segmentation of each vertebral body, similar to the yellow highlighted area in Fig. 3b.

2.2.4 Vertebra Segmentation

Vertebra segmentation was accomplished using the same U-net architecture shown in Fig. 2, except the input size of the images was 128×128 pixels. It was trained for 200 epochs using a binary cross entropy loss function and an Adam optimizer with a learning rate of 10^{-4}. No dropout was performed and a batch size of 1 was used.

A total of 282 vertebra images were manually labelled and split into a 272-image training set and 10-image validation set. Both the training and validation set comprised of half regular and half EOS radiographs. To increase the effective training set size, random horizontal flipping, rotations ($\leq 45°$), zooms (80–120%), and horizontal and vertical shifts ($\leq 10\%$) were employed as data augmentation methods.

Like the spinal column, only the largest connected component of the vertebra segmentation was kept. A minimum bounding box was placed around each vertebra segmentation to determine its tilt angle relative to horizontal. Figure 3b and c display a segmented vertebra and that segmented vertebra with a bounding box, respectively. Further methods of ensuring high quality vertebra segmentations are detailed in the iterative vertebra location sub-section.

2.2.5 Iterative Vertebra Location

To find all the individual vertebrae within a spine, an iterative algorithm was employed. To identify the first vertebra, five images were cropped around two-thirds of the way down the spinal column segmentation, where the clearest vertebra (T12) is normally present. The images were separated vertically by a quarter of the width of the spinal column segmentation, to be certain that at least one centered image of the vertebra was obtained. Each of these images was segmented by the CNN and the segmentation with the best quality was selected. A new region was then cropped above the initial segmentation at a distance equal to the height of the initial segmentation, with a width of about the spinal column segmentation at the new height. If the quality of this segmentation was good, this step was repeated, resulting in segmentation of the next vertebra upwards. If the quality of the segmentation was poor, the cropping window was
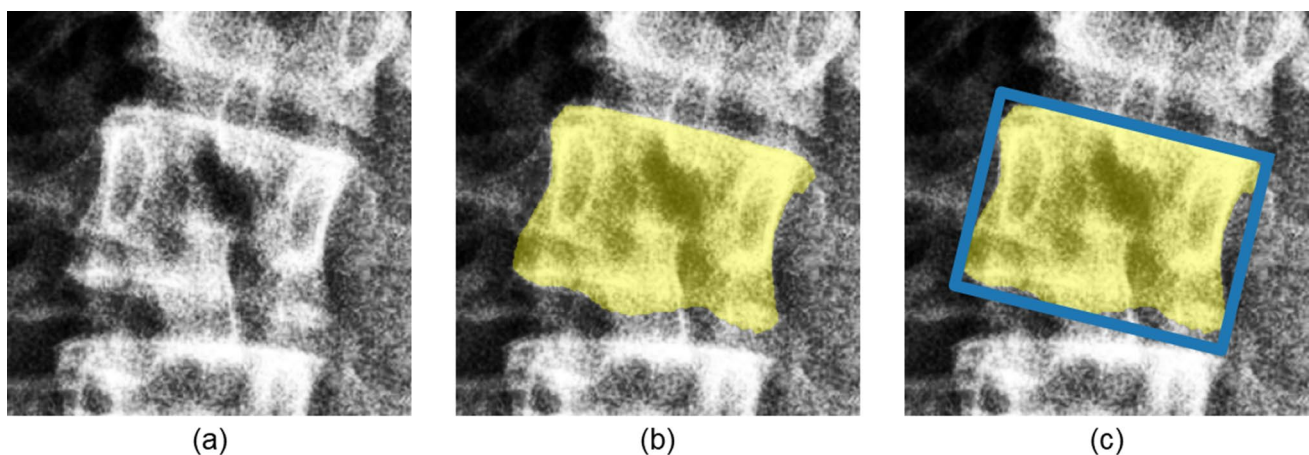

Fig. 3 **a** Cropped region of vertebra, **b** segmented vertebra, **c** segmented vertebra with bounding box

increased in height or width, and the vertebra segmentation was attempted again. Quality was deemed poor if the quality variance or mean (Eqs. 1 and 2) exceeded a value of 100 or 10, respectively. In this manner, vertebra segmentation continued upwards until it reached the top of the spinal column segmentation. Finally, segmentation was performed from the initially segmented vertebra downwards until L5.

The quality of a predicted vertebra segmentation was determined automatically by comparing it to several standard vertebra masks from the training images and then calculating an original quality coefficient inspired by Al Arif et al. and their shape-aware parameter [14]. First, the standard masks were rotated to match the angle of the predicted segmentation. Then, the outlines of the predicted mask and each standard mask were compared by finding the distribution of the minimum distances between each point on the predicted mask contour and the standard mask contour. Figure 4a shows an example of the distribution of minimum distances of a predicted mask with a poorly matched standard mask and Fig. 4b shows that distribution with a well-matched mask. The segmentation quality was determined by both the mean and variance of this distribution—the lower these two values, the higher the quality of the segmentation. The lowest mean and variance with one of the standard masks was taken as the quality

metric for the predicted segmentation. Several standard masks were used during comparison because the contours of thoracic vertebrae are different from those of lumbar vertebrae. Equations 1 and 2 detail the formulae for the quality measurements:

$$Quality\ Variance = \min_i\{\mathrm{var}(x_i)\} \qquad (1)$$

$$Quality\ Mean = \min_i\{\mathrm{mean}(x_i)\} \qquad (2)$$

where x_i is the distribution of minimum distances between each point on the predicted mask and the ith standard mask.

2.2.6 Cobb Angle Measurement

When measuring the Cobb angle, any vertebrae with angles that were significantly different from the spinal column midline angle at the vertebra location were discarded, and then angles significantly different from the angles of their neighboring vertebrae were discarded. The steepest remaining pair of opposing angles was then used to measure the Cobb angle for each curve containing one apex. An apex occurred wherever the centroids of the vertebra segmentations reached an extremum. The pseudocode for the algorithm is displayed in Table 1.
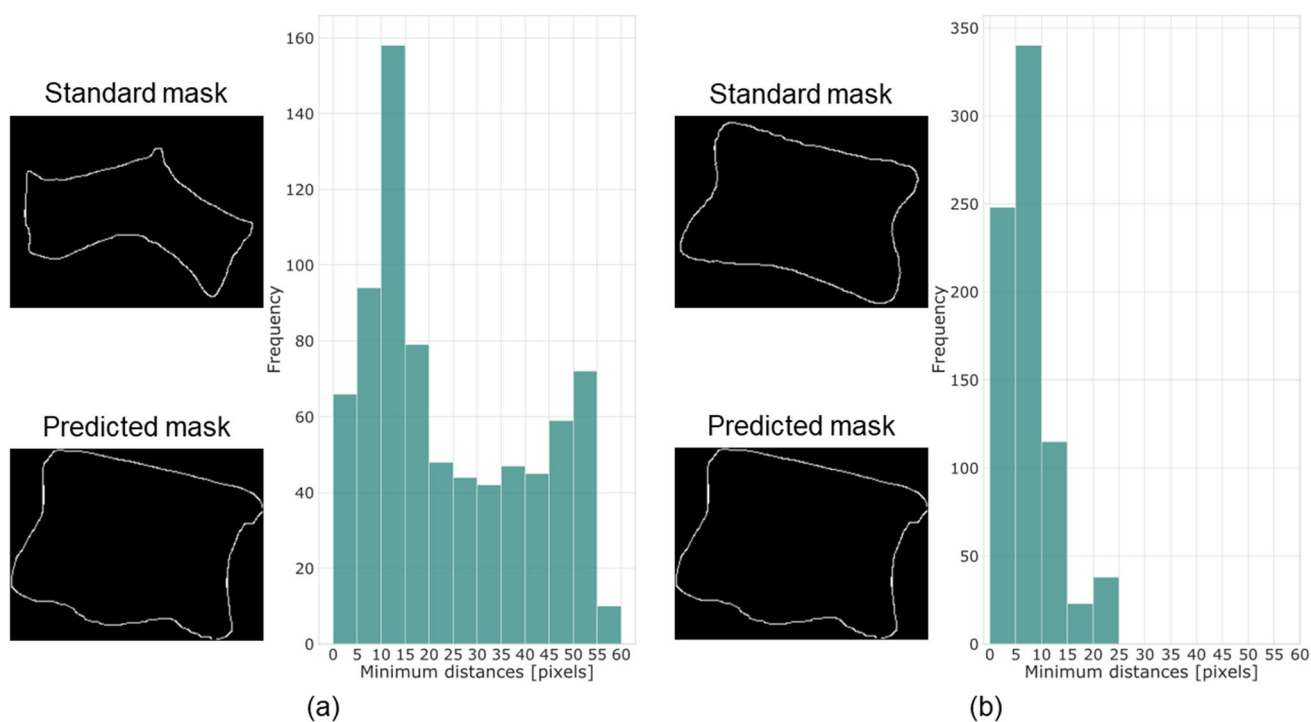

Fig. 4 Distribution of minimum distances between a predicted and standard mask for: **a** a poor match and **b** a good match

Table 1 Pseudocode for the complete iterative algorithm

Segment the spinal column from the PA radiograph

Crop 5 images 2/3 of the way down the spinal column segmentation

For each cropped image:

   Calculate quality coefficients

Set the best quality segmentation as the current vertebra

While the current vertebra has not reached the top of the spinal column segmentation:

   Estimate the location of the next vertebra (above) using the spinal column segmentation

   (*) Crop an image around the estimated location

   Segment a vertebra from the cropped image

   Calculate the quality coefficient of the vertebra segmentation

   If the quality of the next vertebra segmentation is poor:

      Change the cropping parameters (height and width)

      Repeat from (*)

   Set the next vertebra as the current vertebra

Set the lowest segmented vertebra as the current vertebra

While the number of segmented vertebrae is less than 17:

   Estimate the location of the next vertebra (below) using the spinal column segmentation

   (**) Crop an image around the estimated location

   Segment a vertebra from the cropped image

   Calculate the quality coefficient of the vertebra segmentation

   If the quality of the next vertebra segmentation is poor:

      Change the cropping parameters (height and width)

      Repeat from (**)

   Set the next vertebra as the current vertebra

For each vertebra:

   Eliminate the vertebra if its angle is too different from the midline angle at that location

   Eliminate the vertebra if its angle is too different from the neighboring vertebra angles

Calculate the apexes of the spine by finding the extreme vertebrae

For each apex:

   Calculate the greatest angle between a pair of vertebra angles above and below the apex

   Set the greatest angle as the Cobb angle for that apex

2.3 Validation

2.3.1 Network Segmentation

For both spinal column and vertebra segmentation networks, the 10-image validation sets were used to evaluate their performance. The mean accuracy and mean Dice coefficient were used as metrics. The ratio of validation to training images was small, but success of the networks was more strongly evaluated on how well they automatically measured Cobb angles. The spinal column training and validation sets had average Cobb angles of $24.5° \pm 13.2°$ (range 6°–97°) and $25.6° \pm 10.8°$ (range 8°–43°), respectively.

2.3.2 Cobb Angle Measurement

The automatic CNN-based algorithm was run on a new test set of 100 radiographs, and the difference between each automatic and manual Cobb angle measurement was found.

All the manual measurements were performed by a clinician with over 20 years of experience and the variation was within 4°. In the 100 test radiographs, 177 Cobb angles were manually measured with an average Cobb angle of $24.8° \pm 10.1°$ (range 9°–52°). The mean absolute difference (MAD), standard deviation of absolute differences (SD), Pearson correlation coefficient (r), and percentage of differences within clinical acceptance ($\leq 5°$) were reported.

3 Results

Figure 5a shows a test image with all vertebrae segmented, and Fig. 5b shows the final result of the automatic algorithm. In this example, the algorithm detected 3 curves: T3–T5, T5–T12, and T12–L4. However, the clinician's measurement only had two curves: T5–T12 and T12–L4 as the slope of T3 was small ($< 1°$) and therefore not considered as part of a curve.
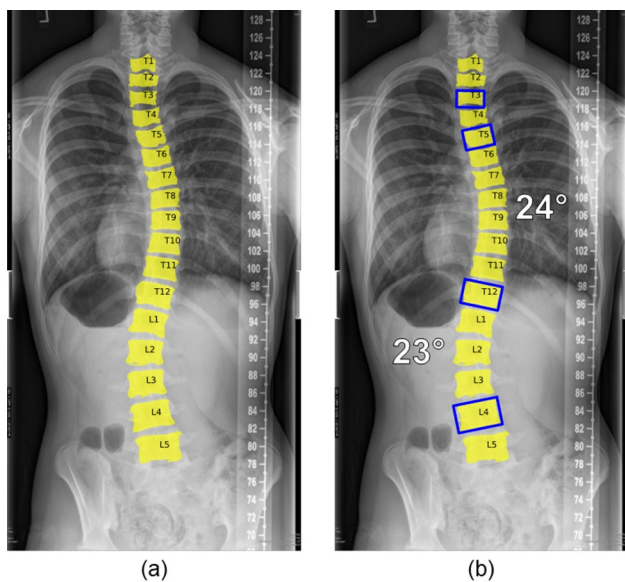
Fig. 5 **a** Results of the iterative algorithm to find individual vertebrae on a test image, **b** bounding boxes on the vertebrae used for Cobb angle measurement
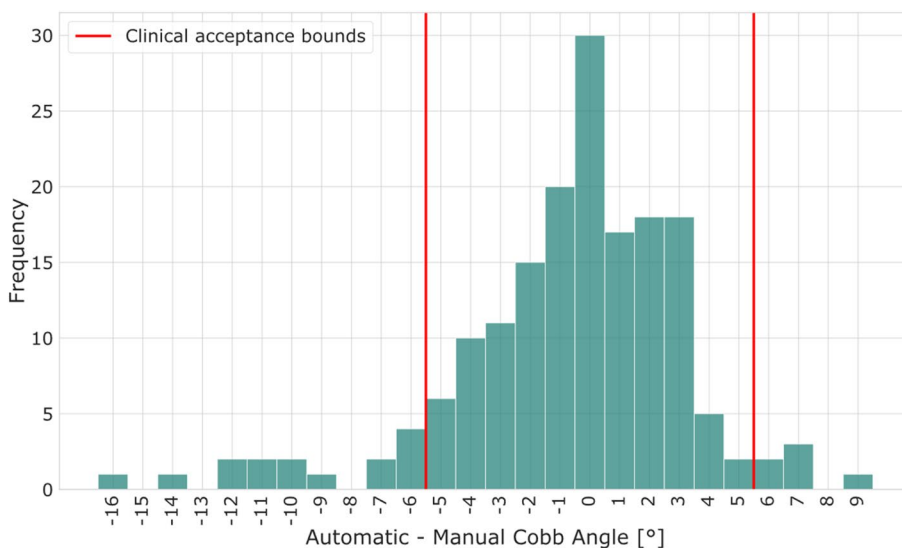
3.1 Network Segmentation

Training the spinal column and vertebra segmentation networks took 51 min and 9 min, respectively. The spinal column CNN achieved 0.989 accuracy and 0.954 DICE, and the vertebra CNN achieved 0.945 accuracy and 0.910 DICE on their respective 10-image validation sets.

3.2 Cobb Angle Measurement

The automatic method reported 173/177 matching measurements, with 88% within clinical acceptance. The 4 missing

curves were a result of vertebra segmentation failure, which was related to image quality or curve ambiguity. Among the 173 comparisons, 87% of the upper and lower vertebrae and 89% of the apical vertebrae from the curves were within ± 1 vertebral level of the manual measurements. This minor variation is usually accepted in clinical assessment. Figure 6 shows a histogram of the signed differences between automatic and manual Cobb angle measurements, with a maximum absolute difference of 16°. The MAD $\pm$ SD of the 173 comparisons was 2.8° $\pm$ 2.8°, and the manual and automatic measurements were strongly correlated ($r = 0.989$). Table 2 summarizes all relevant segmentation and angle measurement results. The automatic algorithm took 90 ± 41 s per image, which is comparable to manual measurements (90 s per image).

4 Discussion

4.1 Spinal Column and Vertebra Segmentation

Overall, both the vertebra and spinal column segmentation networks performed well given the relatively small number of training samples and the complexity of the task, although there are still a few challenges. The vertebra segmentation network appears to have learned to mark the segmentation edge wherever there are lines or high contrast regions. This is not always desirable, as vertebrae occasionally have extraneous bright lines or regions in the middle of the vertebra or around the vertebral edge, leading to false edge detection. Consequently, the bounding box could be incorrect and result in a poor Cobb angle measurement. Furthermore, since the algorithm is iterative, a poor segmentation on the current vertebra may induce improper segmentation of the following vertebrae.

Fig. 6 Distribution of automatic and manual Cobb angle measurement differences

Table 2 Summary of results for segmentation and automatic Cobb angle measurement

Network segmentation	
Spinal column	Accuracy: 0.989
	DICE: 0.954
Vertebra	Accuracy: 0.945
	DICE: 0.910
Cobb angle measurement	
Angles within $\pm 5°$	88%
Vertebrae within ± 1	Measurement: 87%
	Apex: 89%
MAD $\pm$ SD [°]	$2.8° \pm 2.8°$
Pearson correlation coefficient	$r = 0.989$
	$P < 0.00001$

To improve vertebra segmentation, there are several factors that can be changed. First, different preprocessing can be performed. While CLAHE is standard to improve contrast and feature extraction, another preprocessing algorithm may help reveal appropriate features, particularly by highlighting the vertebral edges without simply increasing contrast everywhere. Secondly, a custom loss function can be developed to train the network more efficiently, such as by including a shape-aware parameter [14] or incorporating a Hausdorff distance metric [15].

4.2 Cobb Angle Measurement

Compared to other literature, the developed algorithm performed admirably. Horng et al. tested their algorithm on 35 images and compared to two sets of manual measurements by an expert rater [6]. They had 91% within clinical acceptance on one set and 97% within clinical acceptance on the other. This is higher than the 88% demonstrated by our algorithm. However, all of their curves were under 20°. If our test set is limited to curves under 20°, 93% of the automatic measurements are within 5° of the manual measurements, indicating comparable performance in terms of clinical acceptability. Calculating the MAD $\pm$ SD of the provided measurements from Horng et al., their algorithm returned $3.0° \pm 2.0°$ and $2.5° \pm 1.7°$ between their automatic and two sets of manual measurements. Our results for curves under 20° improved upon this with a MAD of $2.3° \pm 2.1°$. The most significant improvement of the developed method is that it can be accurately applied to a wide range of curve severities. This is important since curves greater than 25° and 45° are recommended respectively for brace treatment and surgery.

Although the results are good, there is still room for improvement. As mentioned before, a poor vertebra segmentation could affect the Cobb angle measurement. The current algorithm only uses the angles of the spinal column midline and neighboring vertebrae as references to eliminate any outlying vertebral tilt angles from poor segmentations; however, this method does not always succeed. Further development can focus on a more robust algorithm that removes erroneous tilt angles every time; this can be achieved by considering the angle between the tops of the pedicles as an alternative to the bounding box angle. Clinicians sometimes use this instead of the vertebral endplates if the vertebral boundaries are unclear. If the pedicles are segmented, their angle can be used as additional information when calculating vertebral tilt for the Cobb angle.

5 Conclusion

This paper reported a robust iterative method based on convolutional neural networks to automatically measure the Cobb angles on PA radiographs for children with AIS. The MAD $\pm$ SD between manual and automatic measurements was $2.8° \pm 2.8°$ and 88% of automatic measurements were within 5°, which is clinically acceptable error. Comparable accuracy and timing to manual measurements was achieved automatically, and thus the proposed method could be a valuable tool for reducing measurement variation and clinical workload. Further improvements can be performed by increasing the accuracy of the vertebra segmentation, especially on the vertebrae which contribute to a Cobb angle measurement.

Declarations

Conflict of interest The authors have no conflicts of interest to declare that are relevant to the content of this article.

Ethical Approval Ethics approval was granted from the University of Alberta Health Research Ethics Board (Pro00102044—Chart Review).

References

1. Weinstein, S., Dolan, L., Cheng, J., Danielsson, A., & Morcuende, J. (2008). Adolescent idiopathic scoliosis. *The Lancet, 371*(9623), 1527–1537. https://doi.org/10.1016/S0140-6736(08)60658-3
2. Kotwicki, T. (2008). Evaluation of scoliosis today: Examination, X-rays and beyond. *Disability and Rehabilitation, 30*(10), 742–751. https://doi.org/10.1080/09638280801889519

3. Gstoettner, M., Sekyra, K., Walochnik, N., Winter, P., Wachter, R., & Bach, C. M. (2007). Inter- and intraobserver reliability assessment of the Cobb angle: Manual versus digital measurement tools. *European Spine Journal, 16*(10), 1587–1592. https://doi.org/10.1007/s00586-007-0401-3

4. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems—Volume 1* (pp. 1097–1105).

5. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2015* (pp. 234–241). Munich, Germany.

6. Horng, M. H., Kuok, C. P., Fu, M. J., Lin, C. J., & Sun, Y. N. (2019). Cobb angle measurement of spine from X-Ray images using convolutional neural network. *Computational and Mathematical Methods in Medicine*. https://doi.org/10.1155/2019/6357171

7. Fu, X., Yang, G., Zhang, K., Xu, N., & Wu, J. (2020). An automated estimator for Cobb angle measurement using multi-task networks. *Neural Computing and Applications*. https://doi.org/10.1007/s00521-020-05533-y

8. Sardjono, T. A., Wilkinson, M. H. F., Veldhuizen, A. G., van Ooijen, P. M. A., Purnama, K. E., & Verkerke, G. J. (2013). Automatic Cobb angle determination from radiographic images. *Spine, 38*(20), E1256-1262. https://doi.org/10.1097/BRS.0b013e3182a0c7c3

9. Kim, K. C., Yun, H. S., Kim, S., & Seo, J. K. (2020). Automation of spine curve assessment in frontal radiographs using deep learning of vertebral-tilt vector. *IEEE Access, 8*, 84618–84630. https://doi.org/10.1109/ACCESS.2020.2992081

10. Pan, Y., Chen, Q., Chen, T., Wang, H., Zhu, X., Fang, Z., & Lu, Y. (2019). Evaluation of a computer-aided method for measuring the Cobb angle on chest X-rays. *European Spine Journal, 28*(12), 3035–3043. https://doi.org/10.1007/s00586-019-06115-w

11. Wu, H., Bailey, C., Rasoulinejad, P., & Li, S. (2018). Automated comprehensive adolescent idiopathic scoliosis assessment using MVC-Net. *Medical Image Analysis, 48*, 1–11. https://doi.org/10.1016/j.media.2018.05.005

12. Galbusera, F., Niemeyer, F., Wilke, H. J., Bassani, T., Casaroli, G., Anania, C., Costa, F., Brayda-Bruno, M., & Sconfienza, L. M. (2019). Fully automated radiological analysis of spinal disorders and deformities: A deep learning approach. *European Spine Journal, 28*(5), 951–960. https://doi.org/10.1007/s00586-019-05944-z

13. Kingma, D., & Ba, J. (2015). Adam: A method for stochastic optimization. In *3rd International conference for learning representations*.

14. Al Arif, S. M. M. R., Knapp, K., & Slabaugh, G. (2018). Shape-Aware deep convolutional neural network for vertebrae segmentation. In B. Glocker, J. Yao, T. Vrtovec, A. Frangi, & G. Zheng (Eds.), *Computational methods and clinical applications in musculoskeletal imaging* (pp. 12–24). Springer.

15. Karimi, D., & Salcudean, S. E. (2020). Reducing the hausdorff distance in medical image segmentation with convolutional neural networks. *IEEE Transactions on Medical Imaging, 39*(2), 499–513. https://doi.org/10.1109/TMI.2019.2930068