

Predicting Spine Geometry and Scoliosis from DXA Scans

Amir Jamaludin¹, Timor Kadir², Emma Clark³, and Andrew Zisserman¹

¹ Visual Geometry Group, Department of Engineering Science, University of Oxford
² Plexalis

³ Musculoskeletal Research Unit, School of Clinical Sciences, University of Bristol
`{amirj,az}@robots.ox.ac.uk`
`timor.kadir@plexalis.com`
`emma.clark@bristol.ac.uk`

Abstract. Our objective in this paper is to estimate spine curvature in DXA scans. To this end we first train a neural network to predict the middle spine curve in the scan, and then use an integral-based method to determine the curvature along the spine curve. We use the curvature to compare to the standard angle scoliosis measure obtained using the DXA Scoliosis Method (DSM). The performance improves over the prior work of Jamaludin *et al.* 2018. We show that the maximum curvature can be used as a scoring function for ordering the severity of spinal deformation.

1 Introduction

Scoliosis is a disease that appears as an abnormal sideways curvature of the spine often presenting in childhood and affecting up to 3% of children [6]. In its severe form, the disease can cause lifelong disability and pain, however most cases of mild scoliosis present no symptoms and stabilize over time [1,11]. This uncertainty over whether an initial mild curve will stabilize, resolve or progress presents some challenges for effective clinical management. Therefore, our long-term goal is to develop new software tools to assist clinicians in managing such patients, and in particular to predict prognosis.

X-ray imaging is the standard for diagnosing and monitoring scoliotic patients, however, since the disease presents in childhood, the radiation burden is far from ideal, especially if used repeatedly. An alternative approach is to use DXA imaging, which requires a far low dosage of radiation. DXA is more commonly used for measuring bone mineral density (BMD) when osteoporosis is suspected, and occasionally used to detect vertebral fractures ([2][3]) but recently it has been shown to be an accurate method for diagnosis of scoliosis [12]; and techniques to automate the technique of [12] have been proposed [9].

In order to assess whether a patient’s disease is progressing, stabilising or resolving it is necessary to measure the scoliotic curvature accurately. However, [12] focussed on the problem of the binary classification task where patients are identified as having scoliosis or not, measured by as at least one curve having a *scoliotic angle* of greater than 6 degrees. This angle threshold was also adopted

in [9]. The original motivation for introducing the threshold was as a conservative allowance for the fact the DXA scans are taken with the patient lying down, rather than standing, so that the curvature of the spine may be reduced. However, there is a growing interest in so-called micro-curves, that is scoliotic curves that are below the 6 degree threshold but might be very early indications of problematic spines.

In this paper, we extend the work of [9] and make a number of contributions. First, we propose a novel algorithm to accurately predict the *curvature* and angle of the spine curve. This involves a combination of a state-of-the-art deep learning architecture with methods from classical integral geometry. We develop and validate the algorithm on one of the largest DXA databases available comprising 7,645 subjects, and compare the performance to expert defined ground-truth, and to two alternative baseline algorithms. Finally, we show that the resulting curvature prediction can be used to define a score function for ordering severity of scoliosis in DXA scans. An example is shown in Fig. 1.

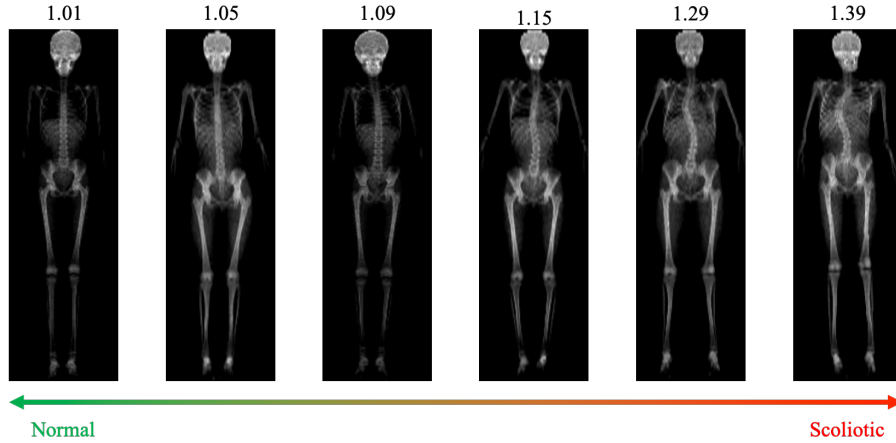


Fig. 1: **Severity of scoliosis:** DXA scan samples ordered by maximum curvature of the spine i.e. the number on top of each scan. Scans with curvature around 1 are at the normal end of the spectrum, while curvatures above 1 and increasing are more scoliotic. In this example, the three samples on the left are normal scans and the three samples on the right have scoliosis.

In the following we describe several possible methods to predict severity of scoliosis from DXA: (i) a baseline method to directly regress the angle using a CNN (Section 2); (ii) predicting the angle by an automated version of the manual method of [12] (Section 3); and (iii) measuring curvature of the spine (Section 4). We predict angles since angle-based measures are widely used by the medical community for scoliosis measurement but as aforementioned we are advocating the use of curvature since that is the underlying symptom of scoliosis. The dataset used is similar to that used in [9] and as such we also compare the

network we used in this work to predict tasks in [9] such as: (i) scoliosis, (ii) body positioning error, and (iii) number of curves. These tasks are used as pre-training for the regression network in Section 2.

2 The Classification & Angle Regression Network

The goal is to predict from a given scan the angle of the largest curve of the spine. Since the number of scans annotated with angle measurements is small, around 927, we pre-train the regression network with several classification tasks that exist in the dataset.

2.1 Classification as Pre-training

We pre-train the network on three different classification tasks: (i) a binary classification of scoliosis vs. non-scoliosis, where an angle of 6° and above is labelled as scoliosis and vice versa, (ii) a binary classification of body positioning error which is dependant on the straightness of the whole body in the DXA scan, and (iii) the number of curves of a scoliotic spine (only on cases with scoliosis). The number of curves is divided into three different classes: no curve (normal spine); one curve, i.e. a “C” shaped spine; and more than one curve, which includes the classical “S” shaped spine with two curves or more. As noted in the introduction, the 6° binary cut-off was suggested in several works using DXA to measure scoliosis namely in [12], [5], and [6]. The network is based on a ResNet-50 but takes in two distinct inputs: (i) the raw DXA scan, and (ii) the soft segmentation mask produced by the segmentation network of [9]. The soft segmentation mask outputs for each row or scanline, the score for which each pixel is most likely to belong to the middle point of one of six body parts. The two input streams are merged after Conv1 via addition. See Fig. 2 for an overview of the network

Classification loss: As in [9], we use a multi-task balanced loss which can be expressed as minimizing a combination of the softmax log-losses of the three classification tasks:

$$\mathcal{L}_t = - \sum_{n=1}^N \left(y_c(x_n) - \log \sum_{j=1}^{C_t} e^{y_j(x_n)} \right) \quad (1)$$

where t corresponds to each classification task, with $t \in \{1 \dots 3\}$, x is the input scan, C_t which corresponds to the number of classes in task t , y_j is the j^{th} component of the classification output, and c is the true class of x_n . The loss for each classification is also balanced with the inverse of the frequency of the class to emphasize the contribution of the minority class e.g. only 8% of the scans have scoliosis.




Fig. 2: **Classification CNN:** The network is a modified ResNet50 [8] but with multiple inputs. The inputs are (1) raw DXA scan, and (2) soft segmentation masks where each mask corresponds to the midcurve of a specific body part. From left to right, the segmented body parts are: (1) head, (2) spine, (3) pelvic cavity, (4) pelvis, (5) right leg, and (6) left leg. The classification outputs are a binary scoliosis prediction (Sco), a binary positioning error prediction (Pos), and a multi-class prediction of the number of curves (NOC).

2.2 Regression

The same network architecture is used for regression of the angle, where now instead of three classification outputs in Fig.2, we have two outputs for the regression – the angle and its uncertainty.

Following [10], in training we assume a Laplace distribution of the measurements with a regression loss that is the negative log-likelihood of the Laplace distribution:

$$\mathcal{L} = \sum_{n=1}^N \left(-\ln \frac{1}{2\sigma_n^2} \exp \left(-\frac{|y_n - \hat{y}_n|}{\sigma_n^2} \right) \right) \quad (2)$$

where for each sample n we predict both the target angle \hat{y}_n and the uncertainty σ_n^2 ; y_n in the ground truth angle. To ensure positivity of σ_n^2 , we employ a softplus non-linearity to the output of the network.

3 Predicting the Scoliotic Angle via Geometry

The standard way to measure scoliosis in whole body DXA was first introduced in [12] and was called the DXA scoliosis method (DSM). The angle measurement part of DSM is a modified Ferguson method since the standard Cobb method cannot be used due to the low resolution of the DXA scans. We automate this process by following and modifying DSM, as illustrated in Fig. 4.

Normal spine line: The method starts by identifying the normal spine line which acts as a reference line to measure the actual spine curve against. In [12], the normal spine line must cross the centre point of the spine at the level of the first rib and ends at the centre of the spine at the fifth lumbar vertebra (L5). We approximate the spine curve by using the soft segmentation produced in [9]. Each row in the soft segmentation output of the network is a probability map of where the midpoint of a certain body part is; the pixel with the highest score is the predicted midpoint (see Fig. 2). The automated normal spine line is drawn from the midpoints of the 3rd and 97th percentiles of the soft segmentation mask for the spine. Fig. 3 shows the automated normal spine line.

Measuring the angle: The soft segmentation output is used to draw the predicted middle spine curve. For each middle spine curve, the apex of a curve is defined as the point furthest away from the line which lies in between two intersections of the middle spine curve and the normal spine line. The angle is predicted as the inner angle of the apex minus 180°. There can exist multiple curves for a given scan, only the maximum angle is used.

4 Measuring Curvature via an Integral method

Scoliosis is essentially a measure of curvature of the spine and as such we propose to directly measure curvature instead of the angle. Measuring curvature via differential based methods is extremely unreliable especially in our dataset. [7] describes an integral based curvature estimator using digital shapes; we adapt and simplify the method of [7] for curvature measurement here. In short, we compare the two areas for a given shape e.g. a circle centred on the spine curve. An example of the shape on the middle spine curve can be seen in Fig. 5. The estimate of the curvature is defined as:

$$\kappa = Area_{max}/Area_{min} \quad (3)$$

When the spine curve is completely straight, the ratio between the two areas is 1 (see Fig. 6).




Fig. 3: Automated normal spine line obtained from the soft segmentation.

4.1 Mapping the curvature to the angle

Curvature is not directly comparable to the measured angle provided in the ground truth annotation. To map the detected curvature to an angle measurement, we train a simple fully-connected neural network regressor that takes in the max curvature of a given spine and outputs a continuous number. We assume that the max curvature corresponds to the max angle for a given spine measured at the apex of the curve as in the manual method. The regressor was trained via the same loss discussed in Section 2.2 and details of the network are given in Fig. 7.

5 Dataset & Training Details

The dataset is from the Avon Longitudinal Study of Parents and Children (ALSPAC) cohort that recruited pregnant women in the UK. The DXA scans of the subjects were obtained from two different time points; when the subjects were 9 and 15 years of age. This difference in acquisition period and the variation of height between different individuals results in a difference of scan heights. Fig. 8 shows the variation of scan heights in the dataset.

In all, there are 7,645 unique subjects in the dataset, most of which have two scans, which totals to 12,040 scans. Most scans are annotated with several annotations which include labels such as body positioning error, and angle measurement of the curves of the spine. Angle measurements were only annotated for




Fig. 4: Spine angle construction. Left: example of the soft segmentation of the spine overlaid with the predicted normal spine line (yellow) and middle spine curve (red). Middle: the spine line and curve overlaid on top of the actual DXA scan. Right: close up view of the spine line and curve now with the intersections; there are three intersections resulting in two curves. The ground truth angle for this particular case is 35° , while the measured angles are 33.8° and 18.1° .

scans with curves deemed to be close to scoliotic or larger. The label distribution of the different classification tasks is given in Table 1 while the frequency of the angle of the biggest curve of the spine is shown in Table 2. We use a 80:10:10 (train:test:validation) random split to train the CNN both for the pre-training on the classification task and the angle regression, on a per patient basis (about 9.6k:1.2k:1.2k scans). A single subject and all of its scans can only appear in one of the training, validation or test sets. Two different random splits were used and we show the mean and standard deviation of the performance at test time for the two splits.

Pre-processing: The scans are normalized such that both the head and feet are roughly in the same region for all the scans regardless of age and original

	Normal	Abnormal	
Positioning	10147 (84.3%)	1893 (15.7%)	
Scoliosis	9563 (92.0%)	814 (8.0%)	
	0	1	>1
NOC	9435 (91.1%)	768 (7.4%)	159 (1.5 %)

Table 1: **Distribution of classification labels:** There are three different classification tasks: (i) scoliosis, (ii) positioning error, and (iii) number of curves (NOC). There are 12,040 scans but fewer labels, since not all scans have labels for all three tasks.

Angle	0°	1° – 4°	5°	6°	7°	8°	9°	10°	11°	12°	13°	14°	15°	16°+
Frequency	9450	28	85	116	146	144	113	81	60	45	18	16	12	63

Table 2: **Distribution of annotated angle in the dataset:** There are 12,040 scans but only 927 were annotated; any scan annotated as 6° and above are labelled as scoliosis. Cases marked as 0° are unreliable as they may actually have a small curvature, but since they were deemed to be below the scoliosis cut-off, they were marked down as 0° during the annotation process. As such, only scans with > 0° are used in our validation experiments. Only the angle of the largest curve for a given scan is shown here; a scan can have multiple curves. In all, we have 814 scoliotic and 113 non-scoliotic cases.




Fig. 5: Curvature via area of a circle. The circle is centred on the spine curve and the curvature is measured via the ratio of the left and right areas of the circle. The circle is moved from the top part of the spine curve to the end, and curvature is calculated at each point.




Fig. 6: When the circle is centred on a curve, the areas are not equal ($Area_A \neq Area_B$), but if the circle is centred on a straight line the areas are equal ($Area_A = Area_B$).




Fig. 7: Network architecture that maps angle from curvature. Each fully-connected (FC) layer is followed by a ReLU activation and batch normalization.

height of the scans. First, empty spaces on top of the head and below the feet are also removed. Then, the scans are resized and cropped isotropically to prevent distortion and to keep the aspect ratio the same as the original. The dimensions of the scans after normalization is 416×128 pixels while the raw dimensions of the scans vary from 173×128 to 411×128 pixels.

Training Details: Both the classification and regression networks are optimized via Adaptive Moment Estimation (Adam) from scratch. The hyperparameters are; batch size of 128; beta1 0.9; beta2 0.999; initial learning rate is 0.0001 and




Fig. 8: **Height Normalization:** The top row shows examples of scans prior to height normalization for both time points (the first five examples are from 9 year old subjects while the last four are from 15 year old subjects), while the bottom row shows the height normalized scans.

lowered by a factor of 10 as the loss plateaus. The network were trained via PyTorch using an NVIDIA Titan X GPU. We employ several training augmentation strategies: (i) translation of ± 24 pixels in the x-axis, (ii) translation of ± 24 pixels in the y-axis, and (iii) random horizontal flipping. At test time, the final prediction is calculated from the average prediction of an image and its flip.

6 Experiments & Results

We compare all the methods to predict angles from DXA scans which includes the regression network in Section 2, the automated automated DSM in Section 3, and the curvature via integral geometry in Section 4.

6.1 Classification

First, since we pre-trained the regression network in Section 2 to classify certain tasks, we compare the performance to existing literature. We report better performance owing to the deeper and modern network used, i.e. the ResNet50 [8], against the VGG-M [4] style network used in [9]. See Table 3. The biggest improvement is in the prediction of the number of curves $72.6\% \rightarrow 77.3\%$ which might be correlated with the improvement of scoliosis prediction $90.5\% \rightarrow 94.2\%$ as these two tasks are closely linked i.e. more pronounced scoliosis or curvature of the spine normally appear with more than one curves.

	[9]	ResNet50
Scoliosis	90.5 ± 1.5	94.2 ± 2.1
Body Positioning	80.5 ± 0.3	81.5 ± 1.1
Number of Curves	72.6 ± 1.2	77.3 ± 1.4

Table 3: **Average Per-class Accuracy (mean \pm std %):** We compare against [9].

6.2 Comparing Angle Measurements

Fig. 9 shows scatter plots of the ground truth against the proposed methods. Surprisingly, the correlation between the ground truth annotation and predictions is similar for three different methods. The correlation is: 0.79 for the regression network; 0.82 for the automated DSM; and 0.82 for the curvature method. However, it can be seen that the regression network (Section 2) struggles to predict extreme scoliosis, $> 40^\circ$, and constantly over-predicts 0° ; possibly due to the abundance of cases with micro-curves that were graded as “normal” during pre-training. The other methods, namely automated DSM and curvature, do not suffer from this problem.




Fig. 9: Scatter plot of all the angle annotated versus the three methods.

We also compare the angle prediction of each method against the ground truth angle in Fig. 10.

The average error of all the methods are quite similar with: the curvature-based method (mapping to the angle via a neural network) having an average error of just 1.9° ; the regression network having 2.1° ; and the automated DSM having 2.4° . However, looking at proportion of error we can see a clear difference where 95% of the data fall below 3.6° error for the curvature-based method while the error is $> 5.5^\circ$ for the other methods.

For comparison, a manual method using ultrasound imaging to measure coronal curvature in subjects with scoliosis reported intra-rater correlation coeffi-




Fig. 10: Comparing the three methods against ground truth on the test set. The x-axis is the error in terms of angle between the prediction and the ground truth, while the y-axis is the proportion of the test set. The curvature method works best with 95% of the data having at most 3.6° error.

cients ranging from 0.84 to 0.93 with the standard error ranging from 1.6° to 2.8° [13].

6.3 Qualitative Evaluation

Curvature Heatmaps. Unlike a CNN as in [9] that can produce evidence hotspots, we instead produce heatmaps of curvature by mapping the values to the segmentation mask on a per scanline basis. This is done since there is no CNN involved in producing the curvature of the line. The segmentation is the same soft segmentation used in Section 2; examples can be seen in Fig. 2. We show examples of scoliotic scans alongside their heatmaps in Fig.11. As expected, the heatmaps are brighter on regions with high curvature and vice versa. In Fig.11, we can see that these heatmaps highlight specific regions according to the type of scoliosis i.e. a spine with thoracic scoliosis is brighter around the thoracic region (upper spine) and similarly a spine with lumbar scoliosis is brighter around the lumbar region (lower spine).




Fig. 11: **Scoliosis Heatmaps:** The left pair shows thoracic scoliosis, while the right shows lumbar scoliosis. Each pair includes the input image and the image with the heatmap overlaid.

Severity of Scoliosis. The maximum curvature can be interpreted as a form of scoliosis score; since the curvature directly relates to the shape of the spine, the larger the curvature the more likely that the scan has scoliosis. Fig.1 shows scans on the test set alongside their curvature. This curvature can be used to monitor disease progression of patients with scoliosis, where an increase of the curvature across a period of time, i.e. a longitudinal study of the subject, would mean the scoliosis is getting worse.

7 Conclusion

We have shown that there is a correlation between the curvature of the spine measured automatically in DXA scans and the angles measured by clinical experts. We have also shown that measuring curvature is slightly more beneficial in terms of regression performance in terms of angle and that we can reliably use these detected curvature values to represent how scoliotic a spine is in DXA.

Acknowledgements. We are extremely grateful to all the families who took part in this study, the midwives for their help in recruiting them, and the whole ALSPAC team, which includes interviewers, computer and laboratory technicians, clerical workers, research scientists, volunteers, managers, receptionists and nurses. The UK Medical Research Council and the Wellcome Trust (Grant ref: 102215/2/13/2) and the University of Bristol provide core support for ALSPAC. This publication is the work of the authors and Amir Jamaludin will serve as guarantor for the contents of this paper.

This research was specifically funded by the British Scoliosis Research Foundation, and the DXA scans were funded through the Wellcome Trust (grants 084632 and 079960).

References

1. Asher, M.A., Burton, D.C.: Adolescent idiopathic scoliosis: natural history and long term treatment effects. *Scoliosis* 1(1), 2 (Mar 2006)
2. Bromiley, P.A., Adams, J.E., Cootes, T.F.: Localisation of Vertebrae on DXA Images Using Constrained Local Models with Random Forest Regression Voting, pp. 159–171. Springer International Publishing, Cham (2015), https://doi.org/10.1007/978-3-319-14148-0_14
3. Bromiley, P.A., Adams, J.E., Cootes, T.F.: Automatic localisation of vertebrae in dxa images using random forest regression voting. In: Vrtovec, T., Yao, J., Glocker, B., Klinder, T., Frangi, A., Zheng, G., Li, S. (eds.) *Computational Methods and Clinical Applications for Spine Imaging*. pp. 38–51. Springer International Publishing, Cham (2016)
4. Chatfield, K., Simonyan, K., Vedaldi, A., Zisserman, A.: Return of the devil in the details: Delving deep into convolutional nets. In: *Proc. BMVC.* (2014)
5. Clark, E.M., Taylor, H.J., Harding, I., Hutchinson, J., Nelson, I., Deanfield, J.E., Ness, A.R., Tobias, J.H.: Association between components of body composition and scoliosis: a prospective cohort study reporting differences identifiable before the onset of scoliosis. *J. Bone Miner. Res.* 29(8), 1729–1736 (Aug 2014)
6. Clark, E.M., Tobias, J.H., Fairbank, J.: The Impact of Small Spinal Curves in Adolescents Who Have Not Presented to Secondary Care: A Population-Based Cohort Study. *Spine* 41(10), E611–617 (May 2016)
7. Coeurjolly, D., Lachaud, J.O., Levallois, J.: Integral based curvature estimators in digital geometry. In: R. Gonzalez-Diaz, M.J. Jimenez, B.M. (ed.) *17th International Conference on Discrete Geometry for Computer Imagery (DGCI 2013)*. pp. 215–227. *Lecture Notes in Computer Science*, Springer Verlag (March 2013), <http://liris.cnrs.fr/publis/?id=5866>
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385* (2015)
9. Jamaludin, A., Kadir, T., Clark, E., Zisserman, A.: Predicting scoliosis in DXA scans using intermediate representations. In: *MICCAI Workshop: Computational Methods and Clinical Applications for Spine Imaging* (2018)
10. Novotny, D., Larlus, D., Vedaldi, A.: Capturing the geometry of object categories from video supervision. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018)
11. Pehrsson, K., Bake, B., Larsson, S., Nachemson, A.: Lung function in adult idiopathic scoliosis: a 20 year follow up. *Thorax* 46(7), 474–478 (Jul 1991)
12. Taylor, H.J., Harding, I., Hutchinson, J., Nelson, I., Blom, A., Tobias, J.H., Clark, E.M.: Identifying scoliosis in population-based cohorts: development and validation of a novel method based on total-body dual-energy X-ray absorptiometric scans. *Calcif. Tissue Int.* 92(6), 539–547 (Jun 2013)
13. Zheng, R., Chan, A.C., Chen, W., Hill, D.L., Le, L.H., Hedden, D., Moreau, M., Mahood, J., Southon, S., Lou, E.: Intra- and Inter-rater Reliability of Coronal Curvature Measurement for Adolescent Idiopathic Scoliosis Using Ultrasonic Imaging Method-A Pilot Study. *Spine Deform* 3(2), 151–158 (Mar 2015)