A Simple Modality-Agnostic Representation for Scoliosis Phenotyping

Owen Pullen, Amir Jamaludin, and Andrew Zisserman

Visual Geometry Group, University of Oxford, Oxford, UK
{owen,amirj,az}@robots.ox.ac.uk

Abstract. Scoliosis is a spinal disorder characterised by a lateral curvature of the spine, typically diagnosed using X-ray imaging. In this paper we propose a *modality-agnostic* method to predict phenotypes of scoliosis. These phenotypes describe the *curve pattern*, and include the number of (significant) curves, the location and direction of the largest curve, as well as whether the spine in general is scoliotic or not. The method is modality-agnostic in the sense that it can be applied to multiple imaging modalities. The method involves representing the spine curve using the coefficients of a low-dimensional Fourier sine series expansion, and then obtaining the phenotypes from these coefficients using a simple feed-forward neural network. The network is trained on curves extracted from DXA images, but can then be applied 'as-is' (without fine-tuning) to curves extracted from MRIs and X-rays. We evaluate the performance of the method on three datasets, one for each modality, and demonstrate excellent performance.

Keywords: Scoliosis · Spine · Geometry.

1 Introduction

Scoliosis is a condition that causes a lateral curvature in the spine, traditionally diagnosed using X-ray images. Medical image diagnosis and classification neural networks have exhibited excellent performance across many different medical image classification tasks, however, networks are limited to the image modality they are trained on – networks trained to use DXA images will not be able to classify accurately or at all on X-ray images, MRI or CT. Image networks often decline in performance even when exposed to out of domain data (e.g. different imaging protocols) on the same imaging medium. This domain gap presents a key problem to the adoption of neural networks in medical image diagnosis and classification.

Rather than training a neural network for scoliosis phenotyping for each modality (DXA, X-ray, MRI, CT, ...) we instead propose a simple two stage procedure: in the first stage, the spine is segmented (this stage may require a modality specific model, but these are available); in the second stage, the spine is represented solely as a geometric curve (by the mid-points of the segmentation)

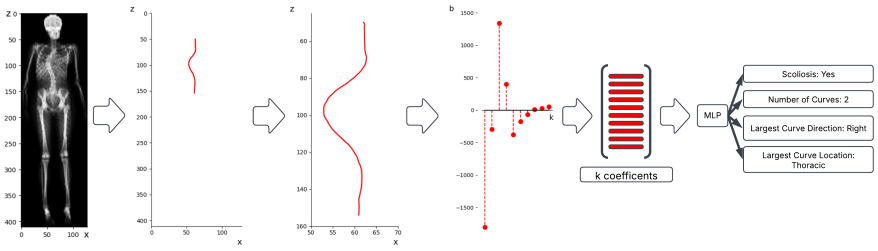

Fig. 1: **Overview of the representation and phenotyping for spinal scoliosis.** The spine in a coronal projection is first represented geometrically, as its mid-point curve (here shown for a DXA image). This curve is then represented by the k coefficients of a Fourier sine expansion (displayed as vertical lines for each value). This low dimensional representation (typically $k = 10$) is sufficient to capture all the geometric information required for the phenotyping. A simple Multi-layer Perceptron (MLP) network is then used to predict the phenotypes, such as scoliosis classification and the location of highest curvature.

and a *modality-agnostic* model is then used to obtain the phenotypes. In this way, an analysis can be easily applied and *standardised* across multiple modalities.

In our model, the spinal curve is represented by the coefficients of a low-dimensional Fourier sine series expansion. The scoliosis phenotypes are then predicted from this low-dimensional representation using a feed forward neural network. The phenotypes are: (i) the presence of scoliosis, (ii) the number of curves, (iii) the direction, and (iv) the location of the largest curve. Since the spine curve can be extracted for any modality, the method is a modality-agnostic (or invariant) approach. The method is illustrated in Fig. 1.

In summary, our contributions include the low-dimensional Fourier representation of the spine curve, and the two-stage procedure for modality independent phenotype prediction. We show that (1) this simple, low-dimensional representation is sufficient to enable phenotypes to be predicted with high accuracy; and (2) that the network can be trained on one modality (DXA in our case), and then applied without change or fine-tuning to other modalities (MRI and X-rays). All code and trained model weights will be publicly released.

Related work. The standard procedure to examine the spine for the presence of scoliosis is using antero-posterior (AP) X-rays and measuring the angle between the most tilted vertebrae [5], this is known as the Cobb angle. Automated Cobb angle measurements are not uncommon; [22, 23] are two such examples which predict vertebral landmarks to predict the Cobb angle. However, X-ray is not the only modality of choice when diagnosing scoliosis, DXA which has now been shown to be a valid alternative [17, 20]. A lower angle threshold of 6° is typically used, compared to the usual criteria of a Cobb angle of 10° as DXA scans are taken supine (instead of standing) which causes spinal curves to reduce in size

compared to a standing position [20]. Automatic scoliosis measurement using DXA has become quite popular in recent times [2, 8, 9, 10]. Some works have also shown the feasibility of using MRIs to measure scoliosis in 3D [1, 3].

Paper layout. Section 2 describes the Fourier sine series representation of the spinal mid-curve. Section 3 introduces the three evaluation datasets – one for each modality (DXA, MRI, X-rays) – and their annotations, as well as the DXA training dataset. Then, Section 4 describes the network for predicting the scoliotic phenotypes from the Fourier representation. The modelling choices and performance are evaluated in Section 5.

2 A Fourier sine series curve representation

Our low-dimensional representation of the spinal (mid-point) curve is based on a one-dimensional discrete Fourier sine representation. The displacement x as a function of the vertical position z can be represented by the Fourier Sine Series:

$$x(z) = \sum_{n=0}^{N-1} b_n sin \frac{n\pi z}{N} \tag{1}$$

and $0 \leq z \leq N$. And the coefficients

$$b_n = \frac{2}{N} \sum_{i=0}^{N-1} x_i sin \frac{n\pi i}{N} \tag{2}$$

are obtained from the sequence of spine curve points x_i spaced by $1/N$ over the vertical length of the spine, L.

The expansion maps the curve into its frequency components (real-valued sine coefficients), a sequence of real numbers b_n, that represent the amplitude of each sinusoidal coefficient. Filtering by taking the first k coefficients gives a low-dimensional representation. The original curve can be reconstructed from these coefficients, allowing the accuracy of the representation to be visualised. Fig. 2 shows several curves with varying degrees of scoliosis and their corresponding low-dimensional Fourier coefficient representations. A visual representation of the curve reconstruction varying k is shown in Fig. 3.

In practice, we use $k = 10$. This choice is determined by the performance of the scoliosis prediction model, as described in Section 5. In the implementation the spinal mid-point curve is first isotropically normalized to a fixed length, and rotated so that the end points both have zero x displacement.

There are two reasons for using a Fourier expansion to represent the shape of the spine curve. First, it simplifies the input by providing only concise geometric information to the network (rather than hundreds of point coordinates); Second, fitting the truncated Fourier expansion provides some de-noising of the original curve, allowing for smoothed input and suppressing noise from the imaging and segmentation process. Note, a sine expansion can be used (rather than the more general expansion which also contains cosines) as the curve is geometrically transformed such that the curve starts and ends have zero displacement.

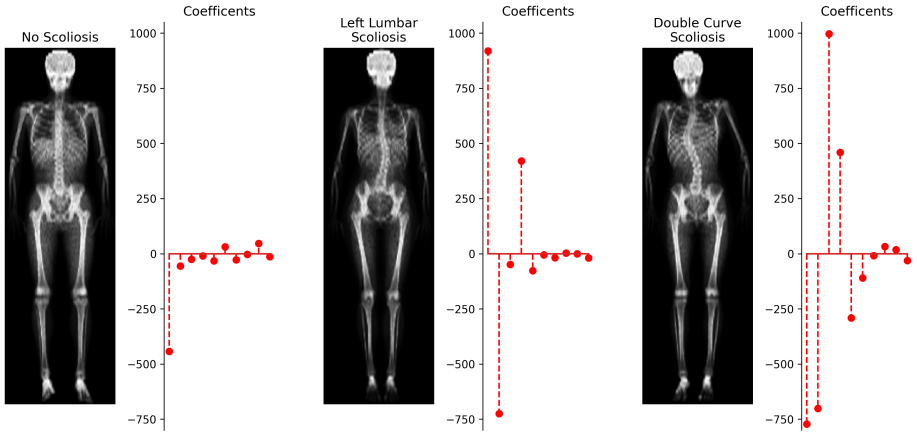

Fig. 2: **Different curve types and their representations.** The type and phenotypes of the curves can be directly 'read off' from the patterns of the Fourier coefficients (with sufficient practice). The coefficients are displayed as a vertical line for each of the k values. The participant without scoliosis (left) shows a relatively small Fourier coefficient magnitude across the first few coefficients (k), except for the first coefficient, than the two participants with scoliosis (centre, right). The participant with left lumbar scoliosis (centre) has larger magnitudes on the first few Fourier coefficients relative to the rest. The participant with "Double Curve" scoliosis has Fourier coefficients with larger magnitudes in both directions (positive and negative), the first few coefficients also have a greater magnitude relative to the latter coefficients.

3 Datasets

In this section we describe the three datasets used, each of a different modality: DXA, MRI, and X-Ray. The DXA dataset is used both for training and evaluation, whereas the MRI and X-Ray datasets are only used for evaluation. Distributions of the unseen test datasets are in Table 1.

ALSPAC DXA Dataset: ALSPAC DXA study provides 33,473 full body DXA scans with multiple scans per person [4]. There are 9,914 useable labelled scans from age groups 9 and 15. These are split into train, validation & test with 7884, 1028 & 1002 images of 632, 80 & 79 positives, respectively. To obtain spine segmentations, we employ the model of [9] to generate soft-segmentation masks. The mid-points are extracted from these soft-segmentations to obtain the spine curve.

Annotations. Details of the annotation process for ALSPAC can be found in [20]. For each curve in the spine, the angle of the curve as well as the direction (left or right) and location (thoracic or lumbar) are annotated. Following [20], we use the largest angle and apply a 6° threshold to label the spine as positive for scoliosis. A lower threshold is used, as DXA scans are taken in the supine coronal plane,

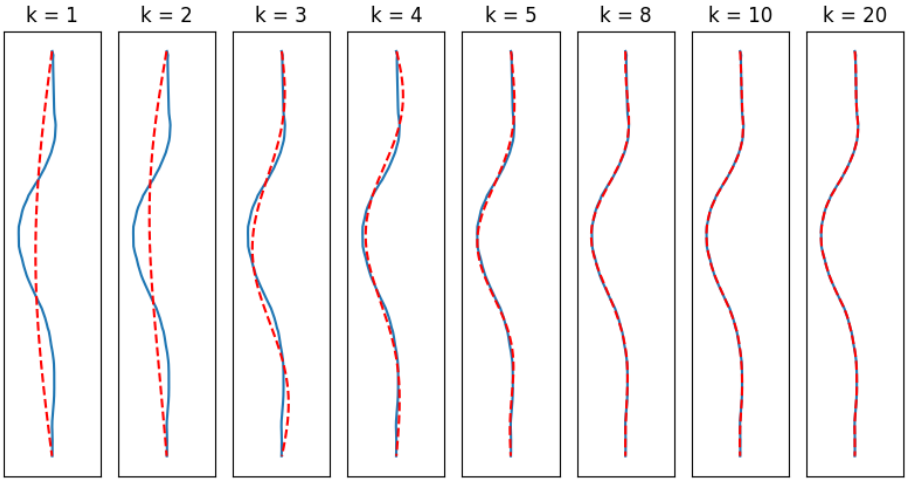

Fig. 3: Reconstructions with a different number k of Fourier sine coefficients for a scoliosis positive sample. The reconstructed curves are shown in red and the original curve in blue. k is varied over $\{1, 2, 3, 4, 5, 8, 10, 20\}$. We find that $k = 10$ is the optimum number of coefficients, where the shape of the curve is maintained and noise is reduced.

and this depresses the scoliotic curvature, thus reducing the angle, compared to X-rays that are usually taken in the standing coronal plane.

Phenotypes. There are 4 classification tasks: (i) Scoliosis Classification: Scoliosis or No Scoliosis, and if a participant has Scoliosis (ii) Number of Curves: 1 curve or 2 (or more), (iii) Largest Curve Direction: left or right, and (iv) Largest Curve Location: thoracic or lumbar. Fig. 4 gives examples of the annotations for these tasks. The prevalence of Adolescent Idiopathic Scoliosis (AIS) worldwide is approximately 0.47–5.20% [13]. The prevalence in our DXA test set in Table 1 is 7.88% which is similar to the a real world prevalence. Additionally, the prevalence of double curves are similar to the typical prevalence 17.72% compared to 15.36% (male and female combined) [13]. Whilst our definitions are slightly different from [13], the prevalences of different curve phenotypes in our test set are similar to real-world prevalences reported. Our dataset is more representative of real-world prevalence and captures more typical curve patterns and borderline cases than other datasets.

UK Biobank MRI Scans: The UK Biobank is comprised of 48,384 whole-body MRIs from the UK Biobank, a large open-access medical dataset with scans from more than 500,000 volunteers [18]. MRIs in the UKBiobank are of much lower resolution than standard clinical scans but the whole of the spine is clearly visible. Scans are resampled to be isotropic and cropped to a consistent resolution ($497 \times 224 \times 174$). In this work, we use a part of the testing set (1,929) that has been annotated by experts using a modified Ferguson DXA Scoliosis Method

Phenotype	Class	MRI	X-Ray	DXA
Number of samples		1646	481	1002
Scoliosis	No scoliosis	1391	18	923
	Scoliosis	255	463	79
Number of curves	1 Curve	232	45	65
	2 or more curves	20	418	14
Largest angle direction	Left	114	-	44
	Right	140	-	35
Largest angle location	Thoracic	100	463	42
	Lumbar	154	0	37

Table 1: **Statistics of the Evaluation Datasets**. The distribution of the number of labels for each of the test evaluation datasets.

(DSM) in whole-body DXA scans as described in [20]; note that only 1,646 MRI scans (255 scoliosis cases) were available for testing. Though these annotations were made on the DXA scans, previous work has shown that, since the MRIs were taken around the same time, the two modalities are largely comparable [21]. To extract the curves, we use the segmentation method described in [6], and calculate the centroid of each intervertebral disc and vertebral body segmented in the axial view along the spine. The stack of centroids is then projected to the coronal plane.

X-Rays: The 2019 MICCAI data challenge dataset from the challenge Accurate Automated Spinal Curvature Estimation, MICCAI 2019 [15]. The dataset consists of 609 X-ray images of the spine taken in the standing coronal plane; we use 481 as the test set (463 scoliosis cases), while the remaining 128 as our calibration set. Three Cobb angles are annotated in each image; we took the maximum Cobb angle and applied the widely used 10° threshold of a Cobb and of $\geq 10°$ to label Scoliosis and $< 10°$ No Scoliosis [13]. The images were also annotated with corner points of the 17 vertebral bodies (4 corner points each, 2 points for each lower and upper endplate). To extract the curve for each image, we used a combination of the midpoint of the upper endplate, the centroid of the vertebral body, and the midpoint of the lower endplate; all calculated using the 4 corner points. These 17 x 3 points results in a curve of 51 points along the midpoint of the spine in each image.

4 Network and Implementation Details

We first introduce the network that predicts the phenotypes, given the Fourier sine representation of the spine curve. We then describe the loss function and training.

4.1 Network architecture

The 2D curve is obtained from the mid-points of the spine segmentation, as described in section 3. The curve is then transformed isotropically by a rotation

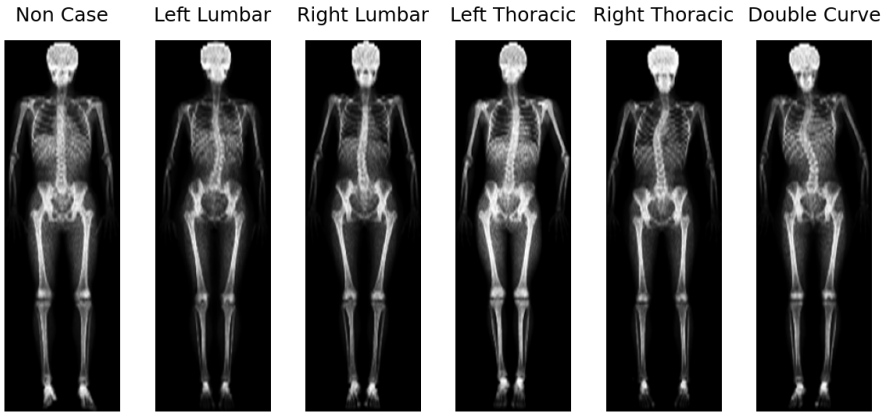

Fig. 4: **Different curve patterns.** Apart from the first image, all are scoliotic. Phenotypes also include the location of the curve (lumbar/thoracic), the number of curves, and the direction. Note, curve direction (left/right) is defined from the viewpoint of the individual, not for the image.

and 2D translation such that its end points both lie at $x = 0$, and it has length L in the vertical direction, as illustrated in figure 1. The curve is then resampled to ensure a standard length and is represented as a Fourier series of sine coefficients that are fed into a neural network.

Model Structure. We design a series of simple Multi-layer perceptron (MLP) feed-forward neural networks to predict the presence of Scoliosis, the number of curves, the location of the largest curve and the direction of the largest curve. The input is a vector of the first k sinusoidal coefficients of the Fourier sine series expansion of the spine curve. The network is comprised of l linear layers with a non-linear ReLU activation function [16] following each layer, width w and has input of a vector length k coefficients, the outputs are logits and the dimensions are fixed to the number of classes in each task, 2 for all the tasks. For the best model, the variable dimensions are $l = 5, w = 64, k = 10$. The actual value of the parameters is determined in section 5.

Input Normalisation. We use per coefficient, k_i, whitening input normalisation, $X_i = \frac{X_i - \bar{X}_i}{\sigma_i}$ for each element of the input before feeding into the network, where $\bar{X}_i$ is the mean and σ_i the standard deviation, for each Fourier coefficient k_i. These values are determined on the training set. This helps to balance the inputs to avoid excessively large values from causing errors in the training, and the propagation of excessively large values within the network.

4.2 Network loss and training

Objective. The loss is defined as

$$\mathcal{L}_{\text{total}} = \sum_{t=1}^{T} w^t \, \mathcal{L}_{\text{CE}}^t$$

where $t \in \{\text{Scoliosis, No. Curves, Direction, Location}\}$ indexes the, $T = 4$, four classification tasks, and L_{CE} is the binary cross entropy loss for each task.

There is significant class imbalance for each task, the population prevalence of adolescent idiopathic scoliosis is 2-3% globally [19] so this is not unexpected. To overcome this we used inverse frequency class weighting. Positive and negative samples were weighted using the inverse class frequency within each task.

The model is trained using the AdamW optimiser [14], this helps to stabilise training and decouple weight decay, increasing the rate of convergence compared to Adam [12]. The AdamW optimiser also provides better adjustment of learning rate during training than Adam. We used an initial learning rate $1e-5$, batch size of 128. The models were trained for an average of 1,560 epochs on an NVIDIA A30 GPU.

Augmentations. We trialled multiple 1-dimensional augmentations when training the model. These included a local Gaussian noise augmentation, and a horizontal flip augmentation. In the Gaussian augmentation, noise was added to each point of the curve, before converting into the Fourier domain. A horizontal flip augmentation flipped the curve direction (from left to right or right to left) and changed the largest angle direction label explicitly. All the augmentations were applied with a probability of 0.5.

5 Results

First determine the optimum parameters for the curve representation and network. We then evaluate the method with these parameters over the datasets for the three modalities: DXA, MRI and X-rays. In all cases the spinal curves are first extracted, as described in section 3, and the curve is then represented by the first k coefficients of the Fourier series (as described in section 2). We then feed the normalised coefficients into the network to classify scoliosis, number of curves, location of largest curve and direction of largest curve.

Evaluation measures. To evaluate the network's performance we use the area under the ROC curve (AUC). This is calculated from the predicted probability scores and evaluates the model's ability to distinguish between two classes at all possible threshold values.

5.1 Model parameter selection

The number of layers l, width w and input dimension k were varied on the DXA training set to determine the best performing network on the test set,

l	w	k	Test Set AUC (%)				
			Scoliosis	N Curves	Direction	Location	Mean
5	64	2	85.73	80.22	88.38	68.40	80.72
5	64	5	95.21	83.74	92.21	84.23	86.73
5	64	10	95.79	88.13	91.56	85.26	90.19
5	64	15	94.82	86.59	92.40	82.75	89.14
5	64	20	92.18	80.11	90.26	81.98	86.13
5	32	10	95.07	83.95	91.17	82.05	88.06
5	64	10	95.79	88.13	91.56	85.26	90.19
5	128	10	95.86	87.36	90.65	84.62	89.62
1	64	10	96.10	85.16	89.29	86.04	89.15
5	64	10	95.79	88.13	91.56	85.26	90.19
10	64	10	93.00	76.48	91.49	77.02	84.50

Table 2: Performance on the DXA test set. k specifies the number of Fourier coefficients used as network input. l specifies the number of layers, and w is the width of each layer in the MLP network.

shown in Table 2. Varying k shows that 10 coefficients is the optimum value. Ablations of $k = \{2, 5, 10, 15, 20\}$ were trained. $k = 10$ has the highest mean AUC value across all tasks, except for direction. With more than 10 coefficients the model may be receiving image and segmentation noise rather than only geometric information relating to the shape of the spine. The best performing network is the configuration with 5 linear layers and hidden layer width 64.

In Table 3 we assess the repeatability and variance of the predictions, we trained ten models with the same hyper parameters, without augmentation, and achieve (AUC $\pm$ standard deviation): scoliosis: 95.40 $\pm$ 0.46; number of curves 86.15 $\pm$ 2.20; direction of largest angle 90.85 $\pm$ 0.70; location of largest angle 85.42 $\pm$ 1.73; mean of all tasks 89.46 $\pm$ 0.78. The higher variance reflects the low number of training samples. The horizontal flip augmentation increased Scoliosis classification AUC to 96.16 $\pm$ 0.29 with minimal effect on most tasks. A local Gaussian noise augmentation, where noise was added to each point of the curve increased a Scoliosis classification AUC slightly (95.56 $\pm$ 0.42) with a positive effect on most tasks. Combining these augmentations produced the best results with increases in classification AUC across most tasks: Scoliosis: 96.20 $\pm$ 0.20; number of curves: 85.63 $\pm$ 2.08; direction of largest curve 91.90 $\pm$ 0.80; location of largest curve: 86.11 $\pm$ 1.74; the mean of all tasks was 89.96 $\pm$ 0.50.

5.2 Performance across modalities

The quantitative results of ROC-AUC for the unseen DXA, MRI and X-Ray Test datasets are compared in Table 4, and the ROC curves are plotted in Fig. 5. Qualitative results on the three datasets are shown in Fig. 6–8. We obtain good results on MRIs and X-rays, despite not training the model on these modalities.
Performance on DXA. We obtain good results for the direction and location of the largest curve. This is important information for clinicians, as certain

Augmentation	Test Set ROC–AUC			
	Scoliosis	No. of Curves	Direction	Location
No Augmentation	95.40 ± 0.46	86.15 ± 2.20	90.85 ± 0.70	85.42 ± 1.73
Horizontal Flip	96.16 ± 0.29	85.43 ± 1.57	91.02 ± 0.96	85.33 ± 1.66
Local Gaussian	95.56 ± 0.42	86.84 ± 1.96	91.23 ± 0.75	85.19 ± 1.76
Local Gaussian & Horizontal flip	96.20 ± 0.20	85.63 ± 2.08	91.90 ± 0.80	86.11 ± 1.74

Table 3: Augmentations during training. Augmentations are applied to curves prior to conversion into the Fourier domain.

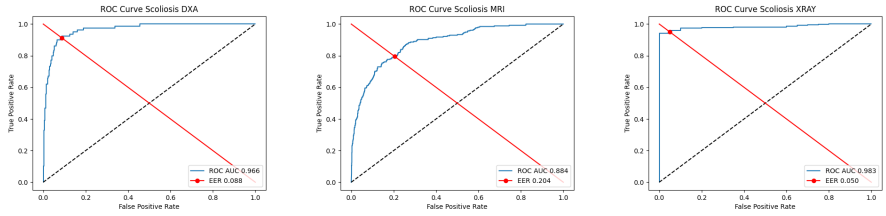

Fig. 5: ROC curves for the DXA, MRI and X-Ray evaluation sets. In each case, we also report the Equal Error Rate (EER). Best DXA model shown, trained with combined local Gaussian & horizontal flip augmentation. Note that the AUCs in Tables 3 & 4 are an average of 10 models, while this is just from a single model.

curves types are more likely to cause complications than others. For example severe Scoliotic curvature in the left thoracic region can cause heart failure and breathing difficulties [11].

Table 4 shows our results across modalities and compares our DXA results to other methods. We achieve comparable results to the method of Jamaludin *et al.* [7], that obtains 97.91% ROC-AUC in Scoliosis classification on our DXA test set. The model was trained using k-fold validation on the ALSPAC DXA dataset, and only the test predictions from each fold were used for the computation on our test split.

Note that our network is much lighter than that of [7], which requires both the DXA image and the spine segmentation mask as input, and then uses a modified ResNet50 image classification network to predict the scoliosis class; 25M vs 18k parameters. Also, unlike [7], we can also provide relevant clinical information on the scoliosis phenotypes of location, direction, and number of curves,

Performance on X-ray & MRI. Note that the model has not been trained on MRI or X-ray (both datasets are unseen). On the X-ray modality, the model shows excellent results on scoliosis classification 98.43% ± 0.16 ROC-AUC and

Modality	ROC–AUC			
	Scoliosis	No. of Curves	Direction	Location
DXA [7]	97.91	–	–	–
DXA	96.20 ± 0.20	85.63 ± 2.08	91.90 ± 0.80	86.11 ± 1.74
MRI	88.12 ± 0.59	76.68 ± 3.65	86.92 ± 0.69	85.30 ± 0.68
X-Ray	98.43 ± 0.16	79.11 ± 4.95	–	–

Table 4: ROC-AUC results for test evaluation on the DXA, X-Ray & MRI datasets. The scoliosis classification method [7] in the first row is an image classification model which requires the raw DXA scan and segmentation masks as input. The method only predicts the scoliosis class, not the other phenotypes.

good results predicting number of curves 79.11% ± 4.95. Direction was unlabelled on this dataset.

Good results were shown on the MRI medium: 88.12% ± 0.59 Scoliosis, 76.68% ± 3.65 No. Curves, 86.92% ± 0.69 Direction, 85.30% ± 0.68 Location. The slightly inferior performance, compared to DXA and X-Rays, can be attributed to the sparser samples obtained from the segmentation of the 3-d MRI, that then have to be projected onto a 2-d plane.

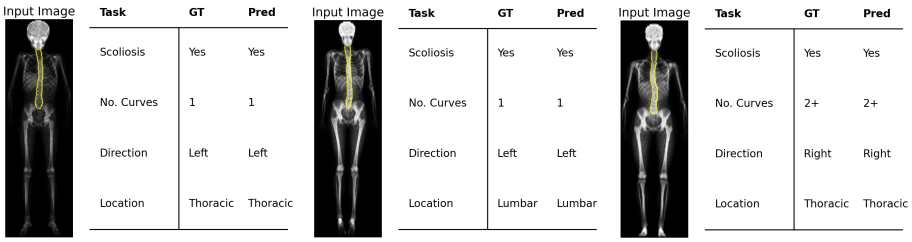

Fig. 6: DXA predictions (Pred) compared to the Ground Truth (GT) labels for three samples from the DXA evaluation set. Curves were extracted from DXA segmentation masks. Single curve scoliosis cases are shown left and middle, and a severe double curve case of scoliosis is shown right.

Limitations. A key limitation of the method is that it is susceptible to positioning error and noise affecting the segmentation network, as this information is lost when representing the spine as a Fourier domain curve. However, image based models are susceptible to this as well.

Examples of classification failures are shown in Fig. 9. Failure cases were observed in DXA scans with increased imaging and/ or segmentation noise. Imaging noise is a common cause of segmentation noise. Failure cases were also observed in scans with positioning error, which our method is particularly susceptible to as it relies on solely geometric information. Positioning error can be reduced by

Input Image

Task	GT	Pred
Scoliosis	Yes	Yes
No. Curves	1	1
Direction	Left	Left
Location	Lumbar	Lumbar

Input Image

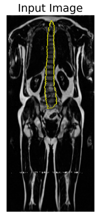

Task	GT	Pred
Scoliosis	Yes	Yes
No. Curves	1	1
Direction	Right	Right
Location	Lumbar	Lumbar

Input Image

Task	GT	Pred
Scoliosis	Yes	Yes
No. Curves	2+	2+
Direction	Left	Left
Location	Lumbar	Lumbar

Fig. 7: MRI predictions (Pred) compared to the Ground Truth (GT) labels for three samples from the MRI evaluation set. Curves were extracted from a 3-d MRI segmentation projected into 2-d. Then the curve is processed as described in Section 2. Single curve scoliosis cases are shown left & middle, and a double curve case of scoliosis is shown right.

Input Image

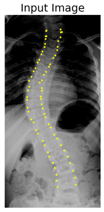

Task	GT	Pred
Scoliosis	Yes	Yes
No. Curves	2+	2+
Direction	Not Labelled	
Location	Thoracic	Thoracic

Input Image

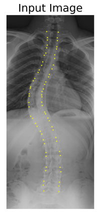

Task	GT	Pred
Scoliosis	Yes	Yes
No. Curves	2+	2+
Direction	Not Labelled	
Location	Thoracic	Thoracic

Input Image

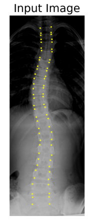

Task	GT	Pred
Scoliosis	Yes	Yes
No. Curves	2+	2+
Direction	Not Labelled	
Location	Thoracic	Thoracic

Fig. 8: X-Ray predictions (Pred) compared to the Ground Truth (GT) labels for three samples from the X-Ray evaluation set. Curves were extracted from the segmentation landmarks. Three severe double curve scoliosis cases are shown left to right.

good scanning technique. Failure also occurred in borderline cases, where the angle is close to the case boundary (≥ 6). Many of these factors can occur in a single scan and cause classification failure, borderline cases may be more affected by noise and errors than non-borderline cases, as a small error can alter the representation relative to the learned boundaries. We did not detect any signs of failure due to out of distribution severe cases on the DXA test set.

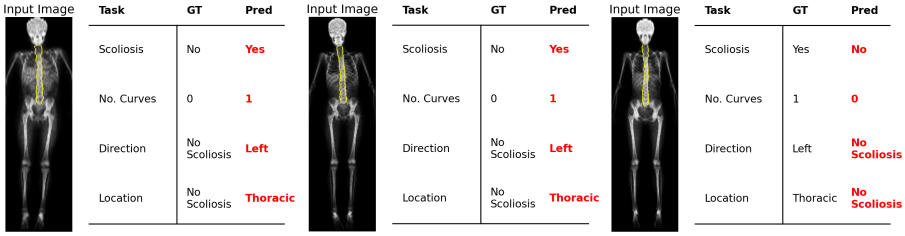

Fig. 9: Scoliosis-negative cases are shown in the left and middle examples, while the right example depicts a scoliosis-positive case. The failure in the left example is likely due to imaging noise affecting segmentation and, subsequently, curve extraction. The classification failure in the middle example is likely due to a positioning error affecting the curve extracted from the segmentation mask. Lastly, in the right example, the model failed to classify a scoliosis-positive participant with a 6° (GT) DSM angle [20]. Borderline cases such as this are among the most challenging for the model to classify.

6 Conclusions

We have shown that a simple Fourier sine representation is sufficient for scoliosis phenotyping, and we demonstrate that a model trained on this representation generalises across modalities; a model trained on representations extracted from DXA scans can be effectively applied to those extracted from MRIs and X-rays. This is important as population studies such as the UK Biobank and ALSPAC have invested heavily in DXA and MRI collections, and now can be analysed with the same geometric model only trained on DXA. Although our method shows high agreement with the labels in the respective datasets, further analysis is needed in future work to adequately assess the correlation with standard clinical measurements for diagnosing scoliosis, such as Cobb angle.

Bibliography

[1] E. Bourigault, A. Jamaludin, E. Clark, J. Fairbank, T. Kadir, and A. Zisserman. 3d shape analysis of scoliosis. In *MICCAI Workshop on Shape in Medical Imaging*, 2023.

[2] E. Bourigault, A. Jamaludin, T. Kadir, and A. Zisserman. Scoliosis measurement on DXA scans using a combined deep learning and spinal geometry approach. In *MIDL*, 2022.

[3] E. Bourigault, A. Jamaludin, and A. Zisserman. 3d spine shape estimation from single 2d dxa. In *MICCAI*, 2024.

[4] A. Boyd, J. Golding, J. Macleod, D. A. Lawlor, A. Fraser, J. Henderson, L. Molloy, A. Ness, S. Ring, and G. D. Smith. Cohort Profile: The 'Children of the 90s'—the index offspring of the Avon Longitudinal Study of Parents and Children. *International Journal of Epidemiology*, 42(1):111–127, Feb. 2013. Publisher: Oxford Academic.

[5] J. Cobb. Outline for the study of scoliosis. *Instr Course Lect AAOS*, 5:261–275, 1948.

[6] R. Graf, P.-S. Platzek, E. O. Riedel, C. Ramschütz, S. Starck, H. K. Möller, M. Atad, H. Völzke, R. Bülow, C. O. Schmidt, et al. Totalvibesegmentator: full body mri segmentation for the nako and uk biobank. *arXiv preprint arXiv:2406.00125*, 2024.

[7] A. Jamaludin, J. Fairbank, I. Harding, T. Kadir, T. J. Peters, A. Zisserman, and E. M. Clark. Identifying Scoliosis in Population-Based Cohorts: Automation of a Validated Method Based on Total Body Dual Energy X-ray Absorptiometry Scans. *Calcified Tissue International*, 106(4):378–385, Apr. 2020. Publisher: Springer.

[8] A. Jamaludin, J. Fairbank, I. Harding, T. Kadir, A. Zisserman, and E. M. Clark. Automated measurement of size of spinal curve in population-based cohorts: Validation of a method based on total body dual energy X-ray absorptiometry scans. *Bone*, 172:116775, July 2023. Publisher: Elsevier.

[9] A. Jamaludin, T. Kadir, E. Clark, and A. Zisserman. Predicting Scoliosis in DXA Scans Using Intermediate Representations. In *MICCAI Workshop: Computational Methods and Clinical Applications for Spine Imaging*, 2018.

[10] A. Jamaludin, T. Kadir, E. Clark, and A. Zisserman. Predicting spine geometry and scoliosis from dxa scans. In *In MICCAI Workshop: Computational Methods and Clinical Applications for Spine Imaging*, 2019.

[11] J. A. Janicki and B. Alman. Scoliosis: Review of diagnosis and treatment. *Paediatrics & Child Health*, 12(9):771–776, Nov. 2007. ISBN: 30/10/20072:2 Publisher: Oxford Academic.

[12] D. P. Kingma and J. L. Ba. Adam: A Method for Stochastic Optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*, Dec. 2014. Publisher: International Conference on Learning Representations, ICLR.

[13] M. R. Konieczny, H. Senyurt, and R. Krauspe. Epidemiology of adolescent idiopathic scoliosis. *Journal of Children's Orthopaedics*, 7(1):3–9, Feb. 2013. Publisher: Springer Verlag.

[14] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

[15] MICCAI. AASCE - MICCAI 2019 Challenge: Accurate Automated Spinal Curvature Estimation, 2019.

[16] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, page 807–814, Madison, WI, USA, 2010. Omnipress.

[17] P. T. T. Ng, L. Straker, K. Tucker, M. Izatt, and A. Claus. Advancing use of dexa scans to quantitatively and qualitatively evaluate lateral spinal curves, for preliminary identification of adolescent idiopathic scoliosis. *Calcified tissue international*, 112, 03 2023.

[18] C. L. M. Sudlow, J. E. Gallacher, N. E. Allen, V. Beral, P. Burton, J. Danesh, P. Downey, P. Elliott, J. Green, M. J. Landray, B. C. Liu, P. M. Matthews, G. Ong, J. P. Pell, A. J. Silman, A. Young, T. Sprosen, T. C. Peakman, and R. Collins. Uk biobank: An open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Medicine*, 12, 2015.

[19] S. Sung, H. W. Chae, H. S. Lee, S. Kim, J. W. Kwon, S. B. Lee, S. H. Moon, H. M. Lee, and B. H. Lee. Incidence and Surgery Rate of Idiopathic Scoliosis: A Nationwide Database Study. *International Journal of Environmental Research and Public Health*, 18(15):8152, Aug. 2021. Publisher: Multidisciplinary Digital Publishing Institute (MDPI).

[20] H. J. Taylor, I. Harding, J. Hutchinson, I. Nelson, A. Blom, J. H. Tobias, and E. M. Clark. Identifying scoliosis in population-based cohorts: development and validation of a novel method based on total-body dual-energy x-ray absorptiometric scans. *Calcified tissue international*, 92:539–547, 2013.

[21] R. Windsor, A. Jamaludin, T. Kadir, and A. Zisserman. Self-supervised multi-modal alignment for whole body medical imaging. In *MICCAI*, 2021.

[22] H. Wu, C. Bailey, P. Rasoulinejad, and S. Li. Automatic landmark estimation for adolescent idiopathic scoliosis assessment using boostnet. In *MICCAI*, 2017.

[23] C. Zhang, J. Wang, J. He, P. Gao, and G. Xie. Automated vertebral landmarks and spinal curvature estimation using non-directional part affinity fields. *Neurocomputing*, 438:280–289, May 2021.