3D Shape Analysis of Scoliosis

Emmanuelle Bourigault[1], Amir Jamaludin[1], Emma Clark[2], Jeremy Fairbank[3], Timor Kadir[4], and Andrew Zisserman[1]

Visual Geometry Group, Department of Engineering Science, University of Oxford[1]
Musculoskeletal Research Unit, University of Bristol[2]
NDORMS, University of Oxford[3]
Plexalis[4]
{emmanuelle,amirj,az}@robots.ox.ac.uk
Emma.Clark@bristol.ac.uk
jeremy.fairbank@ndorms.ox.ac.uk
timor.kadir@plexalis.com

Abstract. Scoliosis is typically measured in 2D in the coronal plane, although it is a three-dimensional (3D) condition. Our objective in this work is to analyse the 3D geometry of the spine and its relationship to the vertebral canal. To this end, we make three contributions: first, we extract the 3D space curve of the spine automatically from low-resolution whole-body Dixon MRIs and obtain coronal, sagittal and axial projections for various degrees of scoliosis; second, we also extract the vertebral canal as a 3D curve from the MRIs, and examine the relationship between the two 3D curves; and third, we measure the angle of rotation of the spine and examine the correlation between this 3D measurement and the 2D curvature of the coronal projection. For this study, we use 48,384 MRIs from the UK Biobank.

Keywords: MRI · Spine Geometry · 3D/2D Correspondences

1 Introduction

Scoliosis is defined as a lateral deformation of the spine in the coronal plane, usually manually diagnosed on anteroposterior (AP) X-rays, by measuring the Cobb angle, where an angle over 10 degrees is considered scoliotic [4]. More recently, it has been shown that scoliosis can also be diagnosed from DXA (Dual-energy X-ray Absorptiometry) scans, which are less costly and involve a 10 times lower radiation dose than conventional X-rays [23]. However, both X-rays and DXAs do not capture the complex 3D deformation of the spine [8]. The convenience of using coronal radiographs to measure scoliosis has meant that the axial and sagittal planes have been widely disregarded.

In this work, we explore scoliosis in 3D by analysing the 3D shape of the space curve of the spine, and its relationship to the 3D space curve of the vertebral canal. For this study, we use the Dixon MRIs available in the UK Biobank. We segment both the spine and the vertebral canal in axial slices. These segmentations allow us to extract 3D curves for the spine and canal, as illustrated in Figure 1.

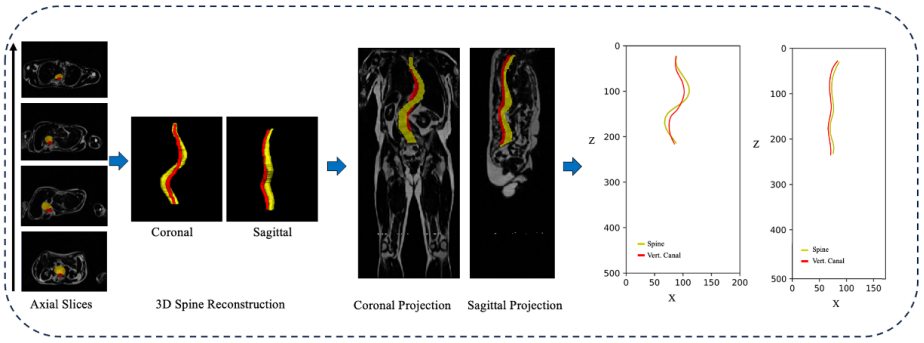

Fig. 1: Overview of the geometry pipeline. The spine (yellow) and the vertebral canal (red) are segmented in each axial slice. The centroids of the spine segments over all axial slices form a 3D space curve (similarly for the canal). The space curve is projected onto the coronal and sagittal planes, and a 2D spline curve fitted to the projected points. Curvature and angles are computed from the spline curve.

Our objective is to study how the 3D spine curve deforms for a scoliotic spine, and also how the vertebral canal adapts to scoliosis. We analyse the 3D spine curve by projecting it onto coronal, sagittal and axial planes, and determine the severity of scoliosis on the coronal plane. It is worth noting that the MRIs from the UK Biobank are uniquely suitable for scoliosis measurement in 3D as there exists an established scoliosis measurement on the paired MRI to 2D DXA for the coronal projection which serves as our point of reference [2, 9].

We then investigate the relationship between the spine and vertebral canal curves on the three planes, and also measure the deviation between the two curves. In addition, we measure the curvature of the coronal projection and the angle of axial rotation of the spine; and investigate their relation.

Section 2 outlines our method for extracting the geometry of the spine and vertebral canal from MRIs, and describes the measures we use for the analysis of the geometry. Then, Section 3 describes the dataset, and presents the results of the analysis, together with several visualizations of the geometry. Finally, Section 4 summarises the findings and the implications of this research.

1.1 Related Work

Research on the relationship between deformations on the sagittal, axial and coronal planes is still in its early phase [11, 14].

The UK Biobank dataset used in this paper is of adults. However, most work on scoliosis focuses on adolescent idiopathic scoliosis (AIS), while scoliosis in adults has been relatively unexplored in past literature. Grown adults can develop degenerative scoliosis as a result of wear and tear on the discs of the spine.

It has been shown that the right thoracic curves are predominant in AIS [13] but this kind of shape analysis of the spine in adult scoliosis is rare.

To date, the vast majority of scoliosis research has focused on 2D shape analysis of the spine, but not in 3D at a large scale. Limitations of 2D spine analysis arise particularly in classifying curve shape. Indeed, deviations are not limited to the coronal plane. They include twisting of the spine in multiple directions [18]. The closest work to ours is by Pasha *et al.* [15,16] in which they look at 3D curves in scoliosis; the main differences between their work and ours are: they used EOS which is quite a niche imaging modality compared to MRIs, they focused on AIS as opposed to degenerative scoliosis, and the number of samples is small (n=103).

Though we use MRIs in our work, it is worth noting that most works on spinal MRIs focus on non-scoliosis spinal disorders and as such put more emphasis on segmenting the vertebral bodies and discs individually rather than the spine as a whole [10–12,14,24].

2 The 3D Geometry of the Spine and Vertebral Canal

It is essential for the 3D geometrical analysis of scoliosis that we capture the whole shape of the spine. To this end, we segment the two main structures that can be seen in the axial Dixon MRIs: these are (i) the "spine" itself, which is comprised of the vertebral bodies and the intervertebral discs, and (ii) the "vertebral canal", which is the space occupied by the spinal cord and filled with cerebrospinal fluid. We do this segmentation on a per-slice basis for each axial image in a given scan volume. The centroid (a 2D point) of each segmentation can now be extracted from each slice and stacked vertically according to axial slice numbering, and spaced appropriately with the axial slice thickness, for a given volume.

A spline curve can then be fitted in 3D (or to the 2D projections), to smooth out the noise in the measurements. The full implementation details are given in the appendix, and the process is summarised in Figure 1. The spine segmentation gives us the 3D spine curve and the vertebral canal segmentation gives us the 3D vertebral canal curve. These 3D space curves can be projected to the coronal (X,Z), the sagittal (Y,Z), as well as the axial (X,Y) plane. Figure 2 shows an example of two curves in our dataset rotating in space.

2.1 Measuring the Deviation of the Curves

Now that the 3D curves of the spine and the vertebral canal have been extracted, we can then proceed to the analysis of the curves. For a given normal spine, it can be observed that two curves should overlap when projected onto the coronal plane and should be parallel in the sagittal plane. As such, the simplest measurement that is indicative of how far away from the norm a given pair of curves are is to measure the deviation between these two curves. Simply put, to measure the distance, 'd', between the two curves we can simply compute

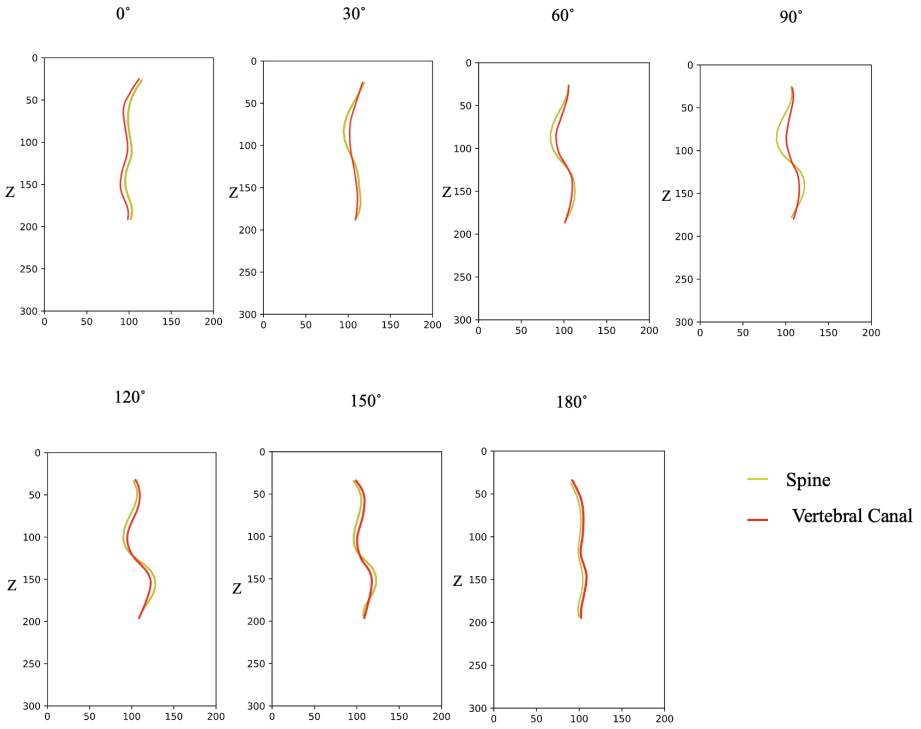

Fig. 2: Spine and Canal 2D Projections for every 30 degree of rotation for a severe 'S shape' spine. This example is in Figure 1. The 0 degree projection corresponds to the sagittal projection, and the 90 projection to the coronal projection.

the vector joining each point of the spine (x_{spine}, y_{spine}) to the vertebral canal curve (x_{canal}, y_{canal}) (see Figure 3). If we project these vectors to the coronal plane, then the sum of their magnitudes measures the 'deviation' between the two curves. For a normal ideal spine, the 'deviation' will be zero in the coronal plane (since the spine and vertebral canal will project on top of one another), and in the sagittal plane, the point-wise difference between the two curves will have a set but constant 'deviation'.

For a given pair of spine and vertebral canal curves, we compute the point-to-point distance in the axial plane:

$$\delta_{spine-canal} = \sum_{i=1}^{N} \sqrt{(x_{ispine} - x_{icanal})^2 + (y_{ispine} - y_{icanal})^2} \tag{1}$$

where i is the slice index and N is the total number of axial slices containing the spine and canal for a given scan.

Then, we can obtain the maximum deviation by taking the maximum of the point-wise distances of the spine-canal deviation (1). For a normal spine, the

maximum deviation will be zero on the coronal plane. For the sagittal plane, a normal spine has inward curvature (lordosis) for the cervical and lumbar sections, and outward spinal curvature (kyphosis) for the thoracic section. Sagittal malalignment is as an exaggeration or deficiency of the normal lordosis or kyphosis curves. Therefore, we measure how parallel the spine and canal curves are on the sagittal plane by taking the standard deviation of the spine-canal deviations.

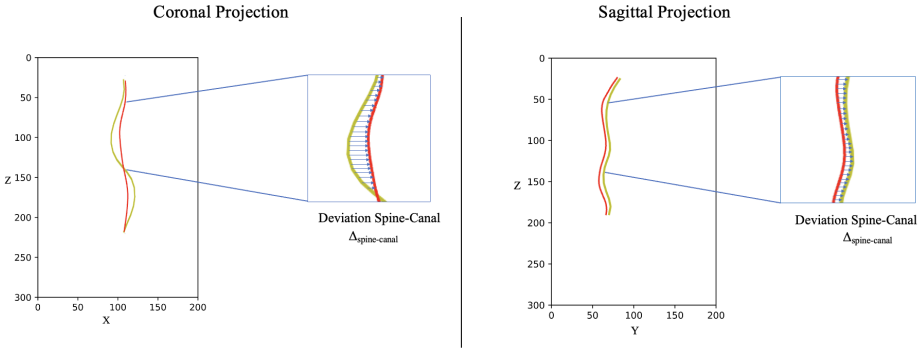

Fig. 3: Measurement of deviation between spine (yellow) and canal (red) curves. Shown is a coronal projection (left) and sagittal projection (right) and a zoom in on the maximum coronal curvature point.

2.2 Curvature of the Spine Curve

The spline is continuous everywhere, as are its first and second derivatives. This is sufficient to determine the curvature κ with the standard mathematical formula:

$$\kappa = \frac{(y'' x' - x'' y')^{\frac{3}{2}}}{(x'^2 + y'^2)} \tag{2}$$

For the results in section 3 the maximum absolute curvature in the coronal plane is used to define three classes of scoliosis severity (normal, mild, severe) according to thresholds obtained on a set of 2K DXA scans annotated for Cobb angles. The threshold for scoliosis is $|\kappa| = 0.083$, mild scoliosis is: $0.083 < |\kappa| \leq 0.118$; and $|\kappa| > 0.208$ is severe scoliosis.

2.3 Angle of Spinal Axial Rotation

Aside from measuring the deviation of the two curves, we can also evaluate the lateral shift of the spine relative to the vertebral canal by measuring the angle of rotation. This is done by using two landmarks: the centroid of the spine and the centroid of the vertebral canal (see Figure 4). The angle between the line through

these centroids and the vertical is the axial rotation (under the assumption that the patient is lying on their back). Note, there are several definitions of the angle of axial rotation. They all rely on measuring the relative positions of anatomical landmarks such as the pedicles, vertebral body, and spinous processes. We use a similar approach to that of [1] and [6], but choose to detect the vertebral canal as a landmark on our axial slices as it is continuous throughout the spine.

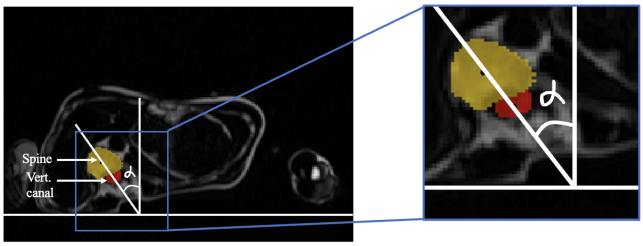

Fig. 4: The angle of axial rotation, α, is the angle between the line through the centroids of the spine (yellow) and vertebral canal (red), and the vertical direction.

3 Results and Discussion

In this section, we compare the 2D projected curvature in relation to the 3D spine. We investigate how the canal curve varies with respect to the spine in subsection 3.2. And in section 3.3, we analyse the coronal and sagittal curvatures and their relation to the angle of maximum axial rotation.

3.1 Dataset

Our dataset 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 [21]. MRIs in the UKBiobank are of much lower resolution than standard clinical scans. Scans are resampled to be isotropic and cropped to a consistent resolution ($501 \times 156 \times 224$). The dataset is split into 80:10:10 for training (38,707), validation (4,838), and testing (4,839) for the segmentation task. 250-200-200 MRI scans are annotated for train-validation-testing for spine and vertebral canal for the baseline segmentation model and checked by an expert clinician. A part of the testing set (1,929) has been annotated by experts for Cobb angles using a modified Ferguson method in whole-body DXA scans as described in [23]. We use this annotated set to define the threshold for scoliosis in our experiment; otherwise this test set is unused in the training of our pipeline. Appendix A.1 gives details of the segmentation.

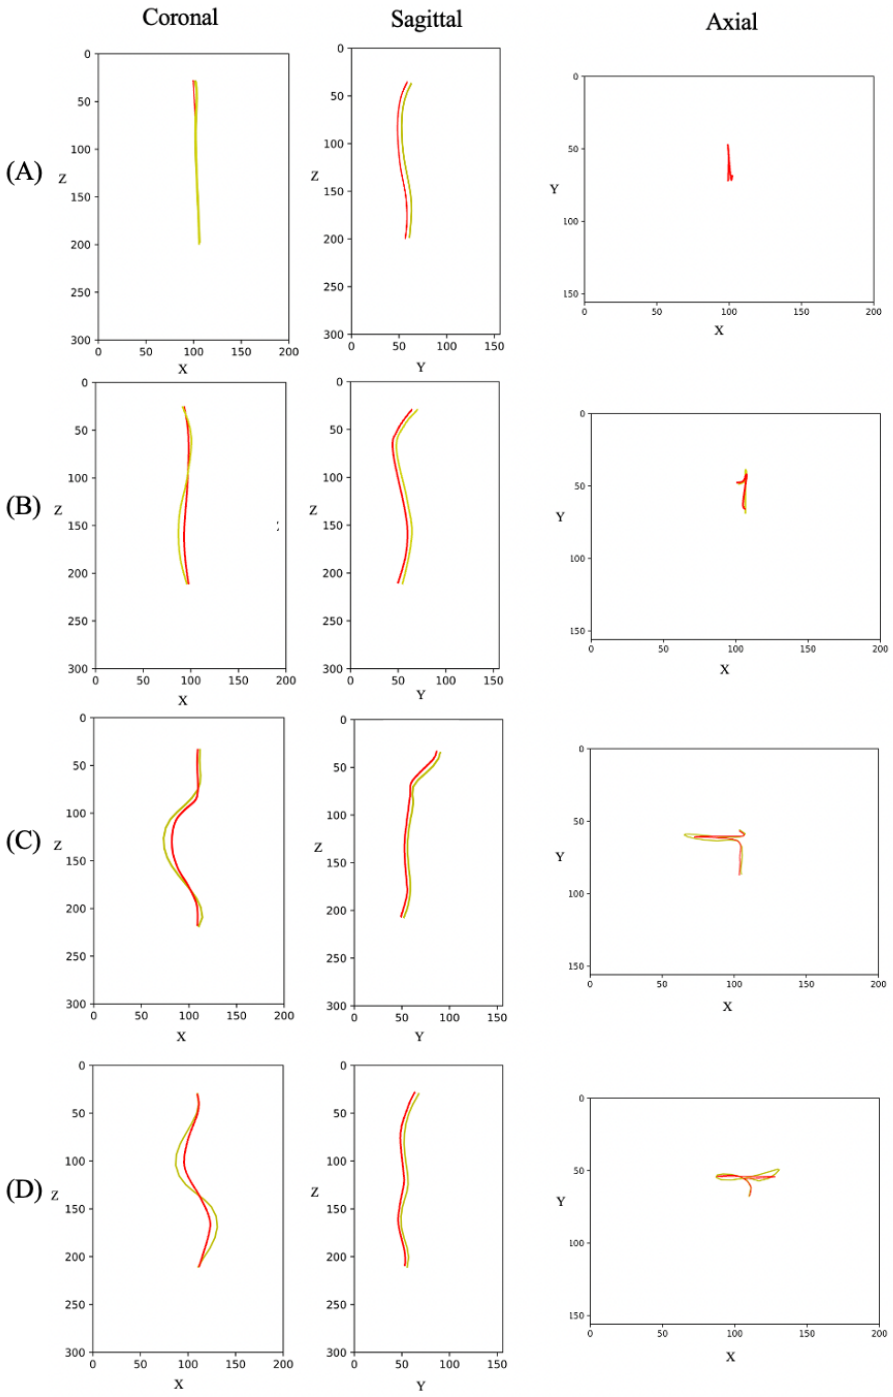

Fig. 5: Comparison of coronal, sagittal and axial 2D projections from 3D curve for normal (A), mild (B) and severe C shape (C) and severe S shape (D) scoliosis cases. Spine curve is in yellow and vertebral canal curve in red.

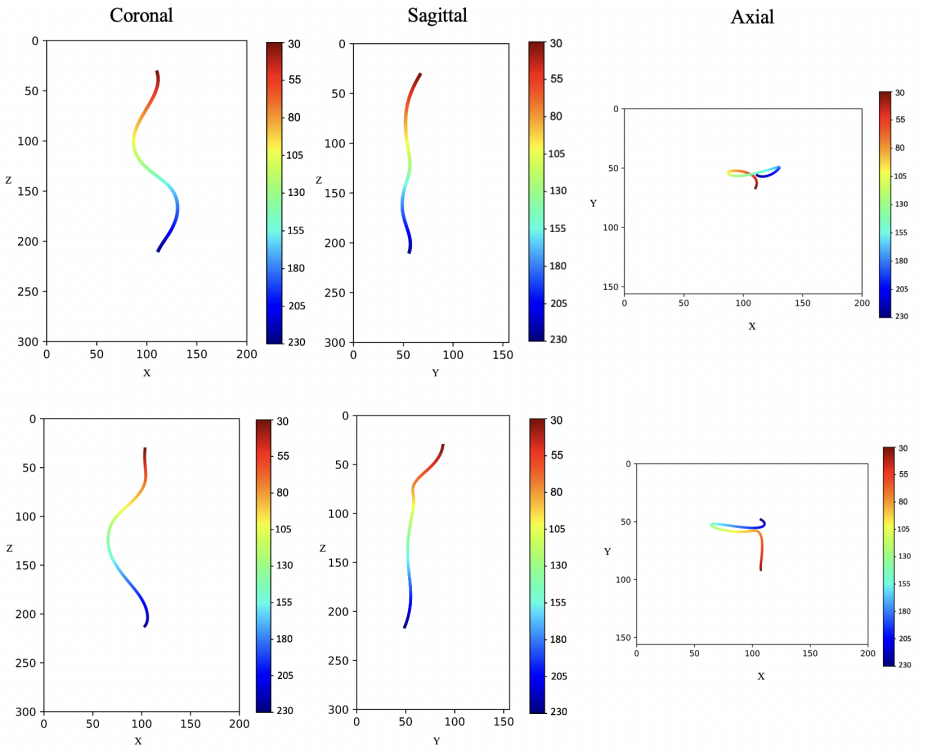

Fig. 6: Comparison of coronal, sagittal and axial 2D projections from 3D curve for a severe S shape (top), and a severe C shape (bottom) scoliosis case. The axial curves (3rd column) are more challenging to interpret. Spines are colour-coded on the z axis to visually indicate the order of the curve in the other projections.

3.2 Geometry of the Spine: Deviation of the Spine and Vertebral Canal

For a normal case, the spine and canal overlap in coronal, and are at a constant separation in sagittal (see Figure 5). By comparing the curves of the spine and canal for normal versus scoliosis cases, we observe that the curves on coronal for scoliosis cases no longer overlap. We also observe that the vertebral canal is less curved than the spine suggesting that it deforms less than the spine. On the sagittal plane, the curves straighten from normal to scoliosis cases (see Figures 5 and 6).

We study how the 3D deviation measurements relate to 2D. The results confirm a strong correlation in deviations between the spine and vertebral canal in 2D coronal and 3D curves (see Figure 7A). The threshold for scoliosis is $|\kappa| = 0.083$. We define mild scoliosis as: $0.083 < |\kappa| \leq 0.118$; and $|\kappa| > 0.208$ for severe scoliosis. Distribution of spine-canal deviations can be discretised according to scoliosis severity (see Figure 7B). This suggests that spine-canal deviations (mm)

can potentially be used as another quantitative measurement of scoliosis. We then investigate how the vertebral canal is varying for different scoliosis severities ranging from normal, mild to severe C and S shape curves.

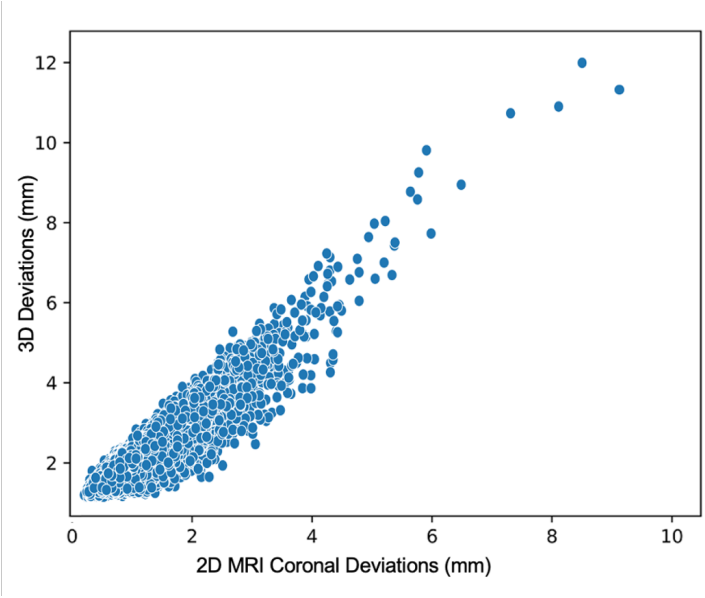

(A)

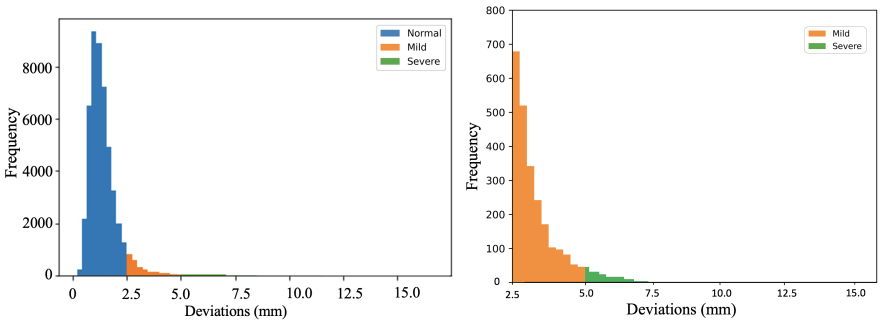

(B) (C)

Fig. 7: (A) Scatter plot of spine-canal point-wise deviations (mm) from 2D coronal projection versus 3D (mm) (Pearson's $\rho = 0.86$, p-value < 0.05, n = 48,384). (B) Histogram with density function displaying the distribution of 2D spine-canal deviation values (n = 48,384) for normal, mild and severe scoliosis cases. (C) Zoom in on mild and severe scoliosis cases from plot in (B). The threshold for scoliosis based on human angles ($> 6°$ in whole body DXA as defined by [23]) in terms of curvature is 0.083. This threshold corresponds to $2.5mm$ of spine-canal deviation.

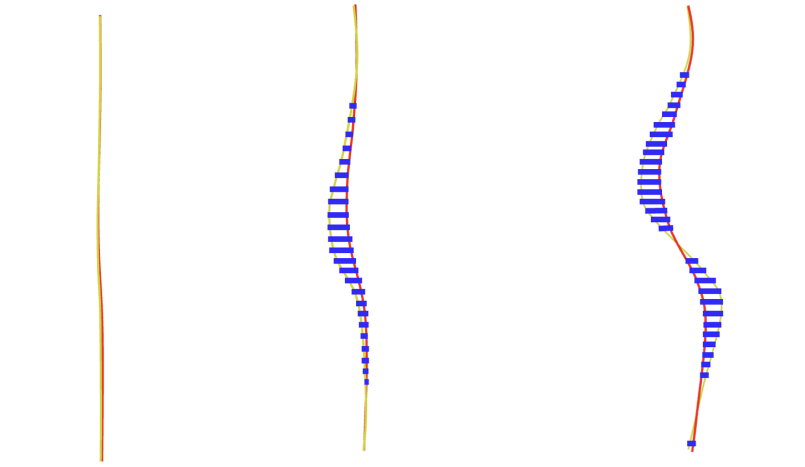

Fig. 8: Visualisation of Spine and Canal Deviations for a normal (A), mild (B) and severe (C) scoliosis cases on the coronal plane. Spine in yellow, vertebral canal in red, and deviations between the spine and vertebral canal in blue.

We can now investigate the properties of the spine that are obtained from the projections of the 3D space curve (see Figure 2 for a severe S shape curve). **Coronal vs Sagittal.** We measured the deviation of the spine and canal at the point of maximum coronal curvature. Comparing the MRI coronal and sagittal spine-canal deviations at point of maximum coronal curvature, we observe an inverse correlation (Pearson's ρ = -0.64, n = 48,384). Curves on the sagittal plane are challenging to accurately assess due to the natural variations of the spine. We notice that severe scoliosis cases tend to have straighter spines in the sagittal plane (see Figure 5). This inverse correlation between coronal and sagittal plane deviations is in accordance with past studies on biplanar radiographs curvature measurements [5]. Moreover, we observe a correspondence between the coronal plane and axial plane. The spine and vertebral canal deviation is greater on the axial projection for severe cases (see Figure 5).

3.3 Curvature measurement in MRI and relation to axial plane

The correlation between maximum coronal curvature and angle of maximum axial rotation is moderately strong (ρ = 0.77, n = 48,384) which may suggest a critical role of the axial plane in relation to curvature on the coronal plane. This is in line with recent research on reconstructed 3D images [7, 11]. Previous work suggested a causal link between axial deformations and onset of coronal deformations due to compensatory mechanical factors [17].

We show the scatter plot between the MRI axial angle of rotation at the point of maximum coronal curvature and the maximum of the MRI spine coronal

projection in Figure 9A, for all 48,384 scans in the UK Biobank. The correlation is relatively good (Pearson's $\rho = 0.79$) between coronal curvature and axial rotation at point of maximal curvature. This confirms the findings of section 3.2,at a large scale.

As a qualitative example, we compare the spine and vertebral canal masks for a mild scoliosis case (max. abs. curvature = 0.18, brown circle in Figure 9A), and a more severe scoliosis case (max. abs. curvature = 0.29, yellow circle in Figure 9A). Axial slices corresponding to these two cases are shown in Figure 9B and 9C.

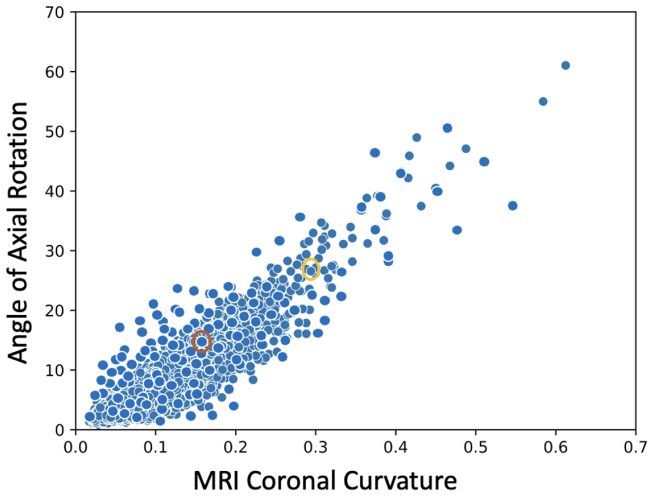

(A)

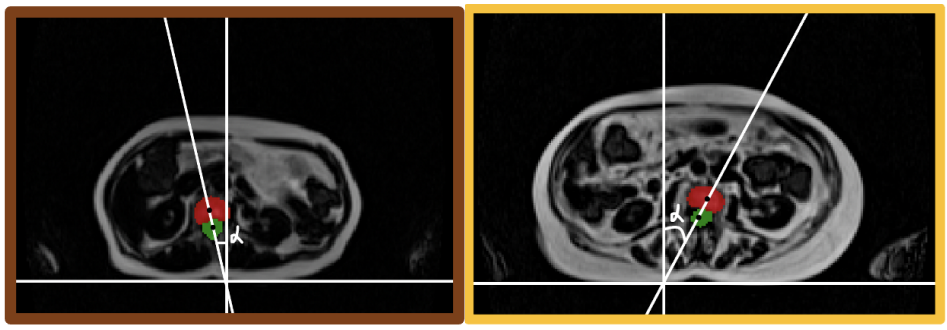

(B) (C)

Fig. 9: (A) Scatter plot of angle of axial rotation vs MRI coronal maximum absolute curvature (Pearson's $\rho = 0.79$, n = 48,384). Angle is given in degrees. (B) and (C) Axial slices corresponding to point of max MRI coronal curvature (yellow and brown circles in (A)). Spine (red) lateral deviation is more prominent on (C) for severe scoliosis case than (B) for mild scoliosis.

4 Conclusion

In this work, we investigated the geometry of scoliosis in 3D, while most prior work has focused on 2D deformations. We measured the curvature of the spine on one of the largest datasets of MRIs. One of the most remarkable outcomes of the visualizations is to see how the vertebral canal arranges itself to have less severe curvature than the spine itself. We also show that the axial plane is quite relevant for the assessment of scoliosis as suggested by the relatively high correlation between the angle of axial rotation and coronal curvature. By considering the spine as a 3D curve, we compared the projected 2D curves of the spine and canal on the coronal and sagittal plane. This efficient method could be used to measure the severity of the spine's deformation.

Ultimately, the goal of this research is to provide an accurate and consistent interpretation of spinal deformations in order to support clinicians in their decision-making process. Prior to the work in this paper, the link between coronal and sagittal curves was not well defined. Also, the role of the axial plane in relation to the coronal and sagittal planes was not yet known. However, one possible future analysis could be to use the relationship between the coronal, sagittal and axial curves as a 3D classification method, without the need to explicitly model the spine in 3D, thus facilitating its adoption in clinics.

Acknowledgements This work was supported by the Centre for Doctoral Training in Sustainable Approaches to Biomedical Science: Responsible and Reproducible Research (SABS: R^3), University of Oxford (EP/S024093/1), and by the EPSRC Programme Grant Visual AI (EP/T025872/1). We are also grateful for the support from the Novartis-BDI Collaboration for AI in Medicine.

References

1. Aaro, S., Dahlborn, M., Svensson, L.: Estimation of vertebral rotation in structural scoliosis by computer tomography. Acta Radiologica **19**, 990 – 992 (1978)
2. Bourigault, E., Jamaludin, A., Kadir, T., Zisserman, A.: Scoliosis measurement on DXA scans using a combined deep learning and spinal geometry approach. In: Medical Imaging with Deep Learning (2022)
3. Chen, B., Jiang, J., Wang, X., Wan, P., Wang, J., Long, M.: Debiased self-training for semi-supervised learning (2022). https://doi.org/10.48550/ARXIV.2202.07136, https://arxiv.org/abs/2202.07136
4. Cobb, J.: Outline for the study of scoliosis. Instr Course Lect AAOS **5**, 261–275 (1948)
5. Galbusera, F., Bassani, T., Panico, M., Sconfienza, L.M., Cina, A.: A fresh look at spinal alignment and deformities: Automated analysis of a large database of 9832 biplanar radiographs. Frontiers in Bioengineering and Biotechnology **10** (2022)
6. Ho, E.K., Upadhyay, S.S., Chan, F.L., Hsu, L.C.S., Leong, J.C.Y.: New methods of measuring vertebral rotation from computed tomographic scans. an intraobserver and interobserver study on girls with scoliosis. Spine **18 9**, 1173–7 (1993)
7. Illés, T.S., Lavaste, F., Dubousset, J.: The third dimension of scoliosis: The forgotten axial plane. Orthopaedics & traumatology, surgery & research : OTSR **105 2**, 351–359 (2019)
8. Illés, T.S., Tunyogi-Csapó, M., Somoskeöy, S.: Breakthrough in three-dimensional scoliosis diagnosis: significance of horizontal plane view and vertebra vectors. European Spine Journal **20**, 135–143 (2010)
9. Jamaludin, A., Kadir, T., Clark, E., Zisserman, A.: Predicting scoliosis in DXA scans using intermediate representations. In: MICCAI Workshop: MSKI (2018)
10. Jamaludin, A., Lootus, M., Kadir, T., Zisserman, A.: Automatic intervertebral discs localization and segmentation: A vertebral approach. In: Computational Methods and Clinical Applications for Spine Imaging. pp. 97–103. Springer International Publishing, Cham (2016)
11. Karam, M., Ghanem, I., Vergari, C., Khalil, N., Saade, M., Chaaya, C., Rteil, A., Ayoub, E., Saad, E., Kharrat, K.E., Skalli, W., Assi, A.: Global malalignment in adolescent idiopathic scoliosis: the axial deformity is the main driver. European Spine Journal pp. 1–13 (2022)
12. Khalil, Y.A., Becherucci, E.A., Kirschke, J.S., Karampinos, D.C., Breeuwer, M.M., Baum, T., Sollmann, N.: Multi-scanner and multi-modal lumbar vertebral body and intervertebral disc segmentation database. Scientific Data **9** (2022)
13. Konieczny, M.R., Senyurt, H., Krauspe, R.: Epidemiology of adolescent idiopathic scoliosis. Journal of Children's Orthopaedics **7**, 3–9 (2013)
14. Ma, Q., Wang, L., Zhao, L., Wang, Y., Chen, M., Wang, S., Lv, Z., Luo, Y.: Coronal balance vs. sagittal profile in adolescent idiopathic scoliosis, are they correlated? Frontiers in Pediatrics **7** (2020)
15. Pasha, S.: Data-driven classification of the 3d spinal curve in adolescent idiopathic scoliosis with an applications in surgical outcome prediction. Scientific Reports (11 2018). https://doi.org/10.1038/s41598-018-34261-6
16. Pasha, S., Ecker, M., Ho, V., Hassanzadeh, P.: A hierarchical classification of adolescent idiopathic scoliosis: Identifying the distinguishing features in 3d spinal deformities. PLoS ONE (03 2019). https://doi.org/10.1371/journal.pone.0213406
17. Roaf, R.: Rotation movements of the spine with special reference to scoliosis. The Journal of bone and joint surgery. British volume **40-B 2**, 312–32 (1958)

18. Rockenfeller, R., Müller, A.: Augmenting the cobb angle: Three-dimensional analysis of whole spine shapes using bézier curves. Computer Methods and Programs in Biomedicine **225**, 107075 (08 2022). https://doi.org/10.1016/j.cmpb.2022.107075
19. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. International Conference on Medical image computing and computer-assisted intervention pp. 234–241 (2015)
20. Ruppert, D., Wand, M.P., Carroll, R.J.: Semiparametric regression. No. 12, Cambridge university press (2003)
21. Sudlow, C.L.M., Gallacher, J.E., Allen, N.E., Beral, V., Burton, P., Danesh, J., Downey, P., Elliott, P., Green, J., Landray, M.J., Liu, B.C., Matthews, P.M., Ong, G., Pell, J.P., Silman, A.J., Young, A., Sprosen, T., Peakman, T.C., Collins, R.: 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)
22. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Jorge Cardoso, M.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support. pp. 240–248. Springer International Publishing, Cham (2017)
23. Taylor, H., Harding, I., Hutchinson, J., Nelson, I., Blom, A., Tobias, J., Clark, E.: 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**, 539–547 (2013)
24. Windsor, R., Jamaludin, A., Kadir, T., Zisserman, A.: A convolutional approach to vertebrae detection and labelling in whole spine mri. In: MICCAI (2020)
25. Yi-de, M., Qing, L., Zhi-bai, Q.: Automated image segmentation using improved pcnn model based on cross-entropy. Proceedings of 2004 International Symposium on Intelligent Multimedia, Video and Speech Processing, 2004. pp. 743–746 (2004)
26. Zhang, Z., Liu, Q., Wang, Y.: Road extraction by deep residual u-net. IEEE Geoscience and Remote Sensing Letters **15**, 749–753 (2017)
27. Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support : 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, held in conjunction with MICCAI 2018, Granada, Spain, S... **11045**, 3–11 (2018)

A Segmentation

There are four separate aligned sequences in the MRI Dixon scans used here. These are in-phase, opposed-phase, fat-only and water-only. The fat-only and water-only sequences are best suited to our task, see Figure 10. Note, the MRI scans in the UK Biobank have a lower resolution compared to typical clinical spine scans. We segment the spine using axial slices as they have higher resolution, and also support larger receptive fields for training the deep network.

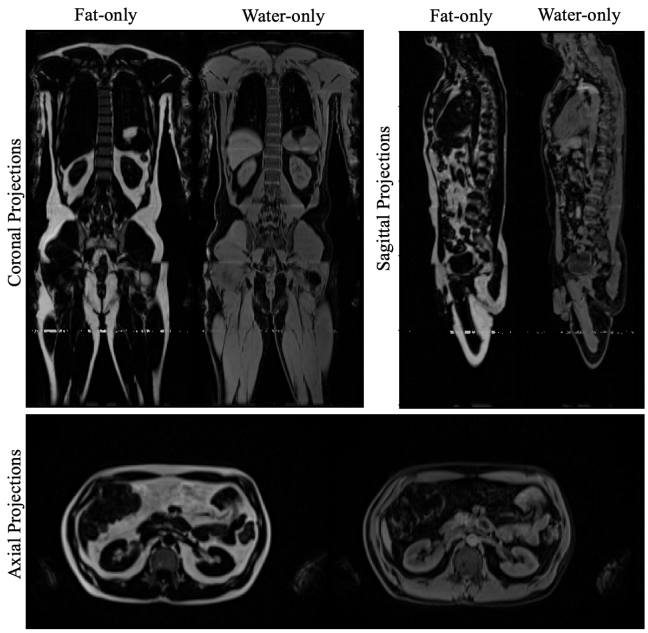

Fig. 10: Coronal, sagittal and axial projections for fat-only and water-only Dixon MRI sequences.

A.1 Segmentation network

A U-Net based network architecture is used for the segmentation task [19, 26]. We use a U-Net++ [27] network with a ResNet-34 encoder. The input is 224 $\times$ 160 $\times$ 6, where we stack three adjacent MRI image slices of the spine region for the two MRI sequences (fat-only and water-only). To avoid partial volume effects, and also to benefit from more context, we ingest three adjacent slices, with the middle slice as output. The output has size 224 $\times$ 160 $\times$2, where 2 refers to the segmentation maps for the spine and vertebral canal.

For training, the loss function is a weighted sum of categorical cross-entropy loss [25] and dice loss [22] computed over a foreground/background/uncertain tri-map to mitigate potentially noisy boundaries in our labels which we define

as ± 2px from the foreground boundary. Networks are trained for a maximum of 500 epochs with early stopping when the validation Dice does not increase by e^{-4}. We use self-training to leverage the whole training set i.e n = 38,707. Inspired from the recent work on confirmation bias reduction in self-training [3], we use an independent head for pseudo-label generation to prevent potentially inaccurate pseudo-label backpropagation.

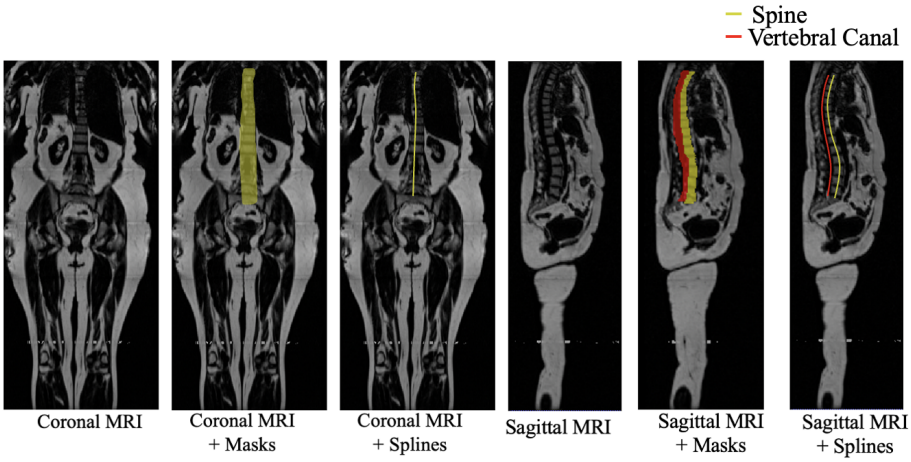

Fig. 11: Visualisation of spine and vertebral canal segmentation masks and mid-point curves on the coronal and sagittal plane.

B Spline Fitting

B.1 2D Spline Fitting

The 2D projected points (in the coronal or sagittal planes) are approximated by a piecewise cubic spline to smooth out any noise due to sampling. For this fitting, we use the method described in [2].

Using a parametrised curve, we construct polynomial piecewise cubic curves. A single cubic curve has only one inflection point, but scoliosis curves may have one or more. A solution could be to add extra control points and using higher order polynomials. However, higher order polynomials are known to be very sensitive to the locations of the control points. A common alternative in computer vision is to construct cubic curves pieced together with a greater number of inflection points. Each pair of control points form one segment of the curve, where each curve segment is a cubic with its own coefficients.

$$f_i(x) = a_i + b_i x + c_i x^2 + d_i x^3 \tag{3}$$

where f is the function representing the curve between control points i and i + 1.

We ensure C^0, C^1, C^2 continuity conditions.

- C^0: Each segment is required to pass through its control points. That is, $f_i(x_i) = y_i$, and $f_i(x_{i+1}) = y_{i+1}$
- C^1: Each curve segment has the same slope at each junction, $f_i^{'}(x_{i+1}) = f_{i+1}^{'}(x_{i+1})$
- Each curve segment has the same curvature at each junction, $f_i^{''}(x_{i+1}) = f_{i+1}^{''}(x_{i+1})$

We improve the method in [2] by changing the uniform placement of a fixed number of knots by automatic knot selection using penalised regression splines [20]. The spline curve is composed of $n-1$ piecewise cubic polynomials where n is the total number of knots. The number of knots is selected in the range from 2 to 10.

n is optimised using a penalty to balance goodness-of-fit and smoothness. The selection of knots is such that the model chooses from a bigger selection of functions. As the number of knots increases, the model overfits the data. Too few knots on the other hand gives a more restrictive function.

B.2 3D Spline Fitting

We now extend the 2D spline fitting to three-dimensional space. We have two systems of linear equations for x and y: $M_x\mathbf{b}_x = \mathbf{x}$ and $M_y\mathbf{b}_y = \mathbf{y}$, where $\mathbf{b}$ is the vector of curve coefficients, $\mathbf{y}$ is the vector of constants, and M is a matrix of continuity conditions ie. C^0, C^1, and C^2. Each system is solved similarly as in 2D section above, except that we are solving two linear systems instead of one.