

SPINEFM: LEVERAGING FOUNDATION MODELS FOR AUTOMATIC SPINE X-RAY SEGMENTATION

Samuel J. Simons

Department of Engineering Science
University of Oxford
Oxford, UK

Bartłomiej W. Papież

Big Data Institute
University of Oxford
Oxford, UK

ABSTRACT

This paper introduces SpineFM, a novel pipeline that achieves state-of-the-art performance in the automatic segmentation and identification of vertebral bodies in cervical and lumbar spine radiographs. SpineFM leverages the regular geometry of the spine, employing a novel inductive process to sequentially infer the location of each vertebra along the spinal column. Vertebrae are segmented using Medical-SAM-Adaptor, a robust foundation model that diverges from commonly used CNN-based models. We achieved outstanding results on two publicly available spine X-Ray datasets, with successful identification of 97.8% and 99.6% of annotated vertebrae, respectively. Of which, our segmentation reached an average Dice of 0.942 and 0.921, surpassing previous state-of-the-art methods.

Index Terms— Foundation models, Spinal X-ray imaging, Vertebra segmentation

1. INTRODUCTION

In 2020, over 820 million individuals were estimated to suffer from neck and lower back pain due to spinal abnormalities or damage [1, 2]. X-rays are frequently used to diagnose these conditions because of their availability and relatively low cost; however, manual annotation by clinicians remains both time-consuming and technically challenging, often due to patient-specific pathology variations and clinician fatigue [3, 4]. Artificial Intelligence (AI) has the potential to address these issues through automated analysis of images. However, a key foundation of many AI-based methods is accurate vertebra segmentation, making robust segmentation techniques essential for advancing the performance of the diagnostic tools.

Numerous studies have utilized convolutional architectures for vertebra segmentation. For instance, [5] proposed a two-step approach involving vertebra localization followed by segmentation, while [6] introduced a rule-based ensemble model to extract vertebral bodies. Although these methods have shown promising results, they face significant limitations, including limited transferability to new datasets and high demands for data annotations. Recently, foundational models have emerged as a competitive alternative. These models leverage massive training datasets and architectures with hundreds of millions of parameters to develop a comprehensive understanding of the task, achieving competitive zero-shot performance without requiring extensive fine-tuning.

The Segment Anything Model (SAM) [7] has demonstrated strong performance in natural image segmentation. However, this success has not fully translated to the medical imaging domain, where images often exhibit lower contrast and ambiguous boundaries [8]. To address this limitation, the Medical-SAM-Adaptor




Fig. 1: Example X-ray images from NHANES II (left, center) and CSXA (right) datasets, with annotated vertebrae outlined in blue.

(Med-SA) was proposed [9], which incorporates adapter modules within the image encoder, allowing for fine-tuning while keeping the base SAM weights frozen. During development, we found that this model can effectively segment individual vertebrae, making Med-SA an ideal backbone for our spine segmentation pipeline.

Our contributions are as follows. We introduce SpineFM, a novel framework that leverages Med-SA as a general-purpose segmentation foundation model for precise segmentation of challenging spinal radiographs with limited labels. To initialize our pipeline, we first employ R-CNN to obtain rough estimates of vertebrae locations, and the most confident estimates are then input into Med-SA for refined vertebrae segmentation. Following this step, SpineFM applies an inductive approach to locate neighboring vertebrae by exploiting the spine’s ladder-like anatomical structure. The inductive nature of SpineFM, combined with the efficiency of the segmentation foundation model, enables robust performance across a range of anatomical variations, contrasting with traditional U-Net-based or Region-CNN spine segmentation methods, and compensates for missing manual annotations. Consequently, SpineFM achieves state-of-the-art results for both cervical and lumbar spine sections.

2. DATA

Our method was developed and tested using two publicly available datasets: the Second National Health and Nutrition Examination Survey (NHANES II) [10] and the Cervical Spine X-ray Atlas (CSXA) [11]. See Fig. 1 for example images.

NHANES II includes 17,000 spine X-ray images covering cervical and lumbar regions, with 544 images manually annotated. These annotations provide landmark coordinates for the C2-T1 vertebrae in cervical scans and T12-S1 in lumbar scans, which were used to generate ground truth masks. However, some images lack annotations

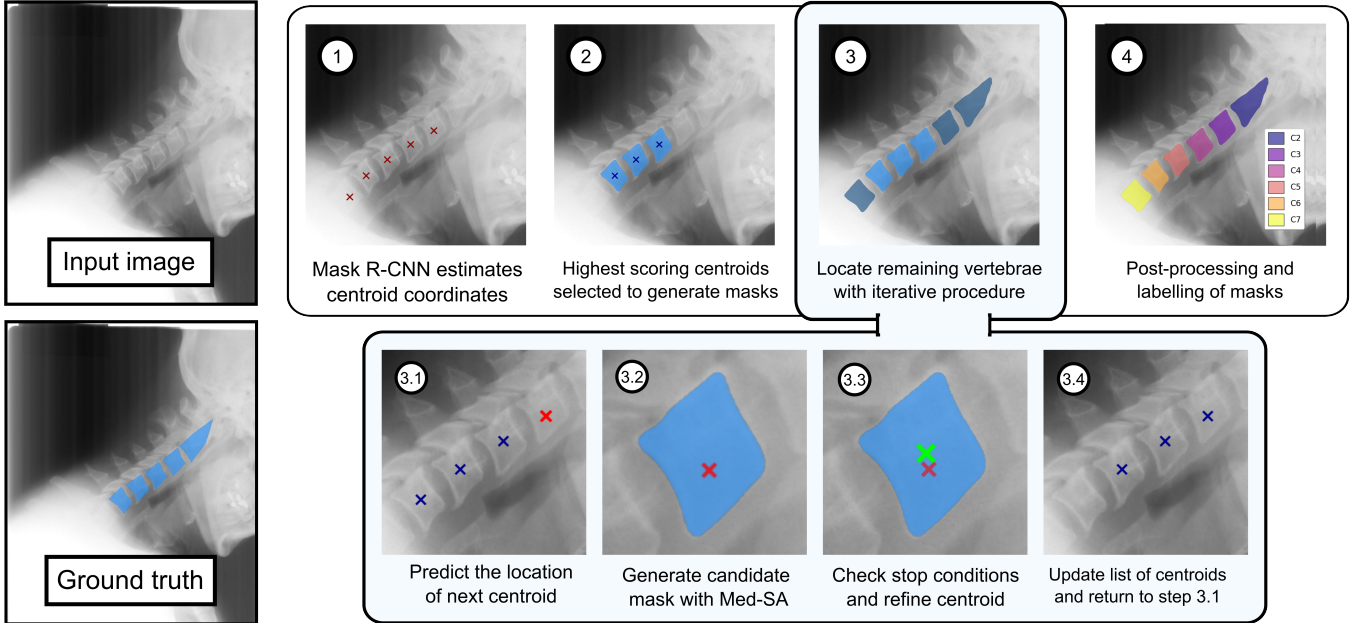


Fig. 2: Diagram of the SpineFM pipeline for automated vertebra segmentation in spine X-rays.

for certain vertebrae due to limited visibility caused by poor contrast or limited field of view (FOV). We split this dataset into 70%, 15%, and 15% for training, testing, and validation, respectively.

CSXA comprises 4,963 cervical spine X-rays, each with complete annotations marking all four corners of the C3-C7 vertebrae. Unlike NHANES II, CSXA images are of higher quality and exhibit consistent FOV. We used a subset of 600 samples from this dataset, split evenly between training, testing and validation.

3. METHODOLOGY

Unlike previous studies [5, 6], which located all vertebral bodies simultaneously, we aim to exploit the repeatable structure of the spine by iteratively predicting the centroid of the next vertebra based on the centroid of previous vertebrae. This inductive approach is similar to methods used for predicting vertebra corners from spinal CT scans [12]. To initialize this inductive process, our pipeline must first estimate the locations of enough vertebrae to infer the remaining ones. Additionally, we incorporate a stopping mechanism to interrupt the process when necessary. To achieve these functionalities, we combine several distinct models into a multi-step pipeline (illustrated in Fig. 2), which can be divided into two main stages: initial predictions and inductive predictions, as outlined next.

3.1. Initial Prediction

The main goal of the initial stage is to predict the centroid coordinates of three successive vertebrae from the original image. The first step involves using a region-based CNN model, Mask R-CNN [13], to generate a list of potential vertebrae masks along with corresponding confidence scores. Masks with confidence scores below 0.6 are immediately discarded. The remaining masks are then used to calculate centroid coordinates, which are sorted based on their projection onto a line of best fit through all points. The three coordinates with the highest average confidence scores are selected as the initial es-

timates. A patch of the X-ray around each coordinate is fed into the Med-SA, which generates the first three masks. These masks are subsequently used to recalculate the estimated centroid coordinates, forming the basis for the inductive stage. Unlike previous approaches [12, 6], our method does not require the detection of a predefined reference vertebral body to initialise segmentation. This characteristic of SpineFM enhances robustness, particularly when these vertebrae are outside the X-ray’s field of view.

3.2. Inductive Prediction

With the first three centroids estimated, a small fully connected neural network is employed to estimate the coordinates of the next vertebra. A fixed-size patch of the X-ray around this coordinate, approximately twice the size of a vertebra, is fed into the Med-SA model to predict the corresponding mask, which is then used to refine the estimate for the vertebra centroid. Next, we check for overlap between the new mask and the previous vertebra mask by calculating their intersection-over-union (IoU). If the IoU exceeds the threshold of 0.1, the new mask is rejected and the algorithm terminates. If the IoU condition is satisfied, a patch around the new centroid is input into a ResNet classifier, which distinguishes between background, a regular vertebra, or one of the two spine-end vertebrae, S1 and C2. If the background is detected, the new mask is rejected, and the algorithm terminates. If one of the spine-end vertebrae is detected, the new mask is retained and marked as the spine-end vertebra, terminating the algorithm once again. Finally, if a regular vertebra is identified, the new mask is kept and the process repeats with a new set of three initial points: the new centroid and the two previous centroids. This iterative process continues both upward and downward along the spine. Our framework is highly data-efficient, as it operates on a patch level, unlike [12, 6], which utilizes the entire image for prediction. This patch-based approach enables to achieve reliable results even with extremely limited training labeled data.

Cervical	C2	C3	C4	C5	C6	C7	T1	Avg
% identified	100%	100%	100%	100%	100%	96%	62.5%	98.8%
Located DSC	0.912	0.944	0.936	0.934	0.933	0.937	0.897	0.932
Overall DSC	0.912	0.944	0.936	0.934	0.933	0.904	0.561	0.916
Lumbar	T12	L1	L2	L3	L4	L5	S1	Avg
% identified	92%	97%	100%	97%	100%	92%	100%	97.3%
Located DSC	0.953	0.962	0.961	0.967	0.966	0.945	0.892	0.952
Overall DSC	0.880	0.937	0.961	0.942	0.966	0.873	0.892	0.925

Table 1: Results for the NHANES II dataset

Cervical	-	C3	C4	C5	C6	C7	-	Avg
% identified	-	98%	100%	100%	100%	100%	-	99.5%
Located DSC	-	0.922	0.922	0.922	0.923	0.917	-	0.921
Overall DSC	-	0.899	0.922	0.922	0.923	0.917	-	0.917

Table 2: Results for the CSXA dataset

3.3. Post-Processing

Once the final list of masks is generated, each mask is converted from logits to a binary mask using a sigmoid function followed by a threshold of 0.9. Each mask is then smoothed by a Gaussian filter and re-thresholded. The positions of the known spine-end vertebrae are subsequently used to infer the labels of the remaining vertebrae along the spine.

3.4. Models

This section provides brief details on each model used within the pipeline. Further details on specific training parameters for each model can be found in the supplementary code.

3.4.1. Medical-SAM-Adaptor


The Med-SA [9] model is employed to segment individual vertebrae when provided with an initial estimation of the centroid coordinates. During training, samples are generated by selecting a random point within the ground truth mask and creating a patch around that point. For each epoch, the model generates a new patch for each ground truth mask. The large variation in prompts ensures that the model remains robust to inaccuracies in centroid coordinates during inference. It is important to note that the publicly available version of Med-SA includes modifications only to the SAM image encoder, while the prompt embedder and mask decoder remain unchanged.

3.4.2. Vertebra Detection and Classification

Mask R-CNN is employed for the initial centroid predictions. Beginning with the COCO V1 pre-trained weights, we fine-tuned the model independently on each dataset i.e. NHANES II and CSXA. We utilize ResNet-101 [14] to classify predictions as either a spine-end vertebra (S1 or C2), a regular vertebra, or background. We began with weights pre-trained on ImageNet 1K, which were then fine-tuned on our datasets.

3.4.3. Point Predictor

Our point predictor is a shallow fully connected neural network. It consists of six input channels, representing the three pairs of coordinates for the known centroids. The network features a hidden layer with 50 units and ReLU activation, and it produces two output channels for the next centroid coordinates.



(a) Low-contrast NHANES II image (b) Poor-quality ground truth masks with reduced vertebrae clarity (e.g. squares) from CSXA dataset

Fig. 3: Examples of limitations related to each dataset, showcasing specific challenges encountered in spine X-ray segmentation

4. RESULTS

Tab. 1 and Tab. 2 present our method’s results on each dataset. When evaluating the performance of our method, we individually assessed each vertebra mask prediction. Then, for each vertebra level, we calculated ‘% identified’; the percentage of vertebrae from this level that were detected by our model, ‘Located DSC’; the average dice similarity coefficient (DSC) for the vertebrae detected by our model, and ‘Overall DSC’; the average DSC across all annotated vertebrae, defaulting to zero for each unidentified vertebra. We also included the average of these metrics across each section, accounting for the frequency of occurrence of each vertebra.

5. DISCUSSION

In this paper, we introduced SpineFM, a novel method for spine X-ray segmentation. This method employs an inductive approach to localize each vertebra and utilizes a foundational model to produce high-quality masks. SpineFM significantly outperforms the existing state-of-the-art method, VertXNet, on the NHANES II dataset, achieving overall DSCs of 0.916 and 0.925 for the cervical and lumbar sections, respectively, compared to VertXNet’s DSCs of 0.880 and 0.863 [6]. Mask R-CNN and nnUNet, two widely used models, have also been assessed for spine segmentation [6]. On cervical and lumbar samples from an in-house dataset, Mask R-CNN achieved average DSCs of 0.730 and 0.649, while nnUNet achieved average scores of 0.903 and 0.829. Both are significantly surpassed by SpineFM’s performance for NHANES II. Finally, MDR2U-Net, another state-of-the-art approach that segments the L1-L5 vertebrae, achieves an average DSC of 0.929 on an in-house dataset [5]. In contrast, SpineFM achieves an average DSC of 0.936 for the same set of vertebrae on the NHANES II dataset. Unfortunately, since we are the first study to use the CSXA dataset, we do not have other examples to compare our results to.

Tab 1 demonstrates that the robustness of our pipeline, achieving an average DSC of over 0.93 across most vertebra levels, with some levels even surpassing 0.96. Achieving such high DSCs is challenging due to inherent ambiguities in defining the exact boundaries within X-ray images. Furthermore, since each ground truth mask is constructed as a polygon from just seven landmark vertices, and so it may not perfectly represent the true shape of the corresponding

vertebra (see Fig. 3(b)).

The main limitation of our method in the NHANES II test arose from imperfect vertebra identification, made apparent in the discrepancy between the Located DSC and Overall DSC. This issue occurs when the patch classifier misclassifies a new vertebra mask, leading to its rejection and the termination of the iterative process. Vertebrae at the ends of the spinal sections were particularly affected, especially T1 in the cervical section and T12 and L1 in the lumbar section. This misclassification is primarily due to inconsistent saturation levels between the X-rays, which can result in very poor contrast (see Fig. 3(a)). It is important to note that the NHANES II X-rays are digital scans of physical X-ray films. In contrast, digital X-rays offer superior clarity and contrast, which would mitigate this issue in a clinical setting. Our results on the CSXA dataset, which contains digital X-ray images, support this conclusion, as we successfully identified 99.5% of all labeled vertebrae.

While all CSXA X-ray images are of high quality and have consistent FOV, we observed a similar overall DSC to the NHANES II dataset, along with a decrease in the average DSC for the located vertebrae. We believe that this decrease is not due to a decline in segmentation performance but rather a consequence of the ground truth masks, which inadequately evaluate segmentation quality. Since the dataset only provides annotations for the vertebra corners, each mask is represented as a simple quadrilateral that does not accurately reflect the true shape of the underlying vertebrae. This limitation is evident in Fig.3(b). Consequently, the DSCs calculated from these masks do not truly represent the quality of our predictions.

6. CONCLUSION

This paper presents SpineFM, a novel method for vertebral body segmentation that employs an innovative inductive approach for locating each vertebra and utilizes a powerful foundational model for its effective segmentation. We demonstrate that SpineFM achieves state-of-the-art performance on two separate datasets, NHANES II and CSXA, which include both cervical and lumbar spine sections. Future work may involve employing an agent-based mechanism to model our inductive process of vertebrae identification as state evolution in a dynamic process [15] to further automate our pipeline.

7. ACKNOWLEDGMENTS

This research study was conducted retrospectively using human subject data made available in open access [11, 10]. Ethical approval was not required as confirmed by the license attached with the open-access data. The project was funded by Oxford Big Data Summer Internship Programme 2024. The model weights and code can be accessed here: <https://github.com/sjsimons/SpineFM>

8. REFERENCES

- [1] Wu, Ai-Min, et al., “Global, regional, and national burden of neck pain, 1990–2020, and projections to 2050: a systematic analysis of the Global Burden of Disease Study 2021,” *The Lancet Rheumatology*, vol. 6, no. 3, pp. e142–e155, Mar. 2024.
- [2] Ferreira, Manuela L, et al., “Global, regional, and national burden of low back pain, 1990–2020, its attributable risk factors, and projections to 2050: a systematic analysis of the Global Burden of Disease Study 2021,” *The Lancet Rheumatology*, vol. 5, no. 6, pp. e316–e329, June 2023.
- [3] Dinggang Shen, Guorong Wu, and Heung-II Suk, “Deep Learning in Medical Image Analysis,” *Annual Review of Biomedical Engineering*, vol. 19, no. Volume 19, 2017, pp. 221–248, June 2017.
- [4] Soumya Snigdha Kundu, Yuanhan Mo, Nicharee Srikiaksemwat, and Bartłomiej W Papież, “Spinal osteophyte detection via robust patch extraction on minimally annotated x-rays,” in *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2024, pp. 1–5.
- [5] Dong Hyun Kim, Jin Gyo Jeong, Young Jae Kim, Kwang Gi Kim, and Ji Young Jeon, “Automated Vertebral Segmentation and Measurement of Vertebral Compression Ratio Based on Deep Learning in X-Ray Images,” *Journal of Digital Imaging*, vol. 34, no. 4, pp. 853–861, Aug. 2021.
- [6] Yao Chen, Yuanhan Mo, Aimee Readie, Gregory Ligozio, Indrajeet Mandal, Faiz Jabbar, Thibaud Coroller, and Bartłomiej W. Papież, “VertXNet: an ensemble method for vertebral body segmentation and identification from cervical and lumbar spinal X-rays,” *Scientific Reports*, vol. 14, no. 1, pp. 3341, Feb. 2024.
- [7] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick, “Segment Anything,” Apr. 2023.
- [8] Yuhao Huang, Xin Yang, Lian Liu, Han Zhou, Ao Chang, Xinrui Zhou, Rusi Chen, Junxuan Yu, Jiongquan Chen, Chaoyu Chen, Sijing Liu, Haozhe Chi, Xindi Hu, Kejuan Yue, Lei Li, Vicente Grau, Deng-Ping Fan, Fajin Dong, and Dong Ni, “Segment Anything Model for Medical Images?,” Jan. 2024, arXiv:2304.14660.
- [9] Junde Wu, Wei Ji, Yuanpei Liu, Huazhu Fu, Min Xu, Yanwu Xu, and Yueming Jin, “Medical SAM Adapter: Adapting Segment Anything Model for Medical Image Segmentation,” Dec. 2023.
- [10] National Center for Health Statistics, “National Health and Nutrition Examination Survey II, 1976–1980,” 1980.
- [11] Yu Ran, Wanli Qin, Changlong Qin, Xiaobin Li, Yixing Liu, Lin Xu, Xiaohong Mu, Li Yan, Bei Wang, Yuxiang Dai, Jiang Chen, and Dongran Han, “A high-quality dataset featuring classified and annotated cervical spine X-ray atlas,” *Scientific Data*, vol. 11, no. 1, pp. 625, June 2024.
- [12] Rhydian Windsor and Amir Jamaludin, “The ladder algorithm: finding repetitive structures in medical images by induction,” in *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*. IEEE, 2020, pp. 1729–1733.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick, “Mask R-CNN,” Mar. 2017.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” *CoRR*, vol. abs/1512.03385, 2015.
- [15] Yuanhan Mo, Fangde Liu, Guang Yang, Shuo Wang, Jianqing Zheng, Fuping Wu, Bartłomiej W Papież, Douglas McIlwraith, Taigang He, and Yike Guo, “Labelling with dynamics: A data-efficient learning paradigm for medical image segmentation,” *Medical Image Analysis*, vol. 95, pp. 103196, 2024.