An Open-Source Workflow for Automated 3D Spine Reconstruction from Tracked Ultrasound: A Pilot Study in Senegal

Pape Mady Thiao[1], Gabriella d'Albenzio[2⋆], Andrew Kim[2], Vivian Nguyen[2], Zhuoxin Guo[2], Idrissa Seck[3], Kyle Sunderland[2], Ousmane Seydina Ba[3], Mamadou Samba Camara[3], Ngor Sine Diagne[1], Tamas Ungi[2], and Gabor Fichtinger[2]

[1] Department of Physical Medicine and Rehabilitation, Centre Hospitalier National Universitaire de Fann (CHNU Fann), Dakar, Senegal

[2] Laboratory for Percutaneous Surgery, School of Computing, Queen's University, Kingston ON, Canada

[3] École Supérieure Polytechnique (ESP), Cheikh Anta Diop University of Dakar (UCAD), Dakar, Senegal

Abstract. Scoliosis monitoring in resource-limited settings remains constrained by limited access to radiography and concerns related to repeated radiation exposure. We present an open-source workflow for automated 3D spine reconstruction from tracked B-mode ultrasound, evaluated on a Senegalese cohort representing an underrepresented population in scoliosis imaging research. The workflow combines optical tracking, PLUS, OpenIGTLink, 3D Slicer, and automated vertebral segmentation using UltraSAM. We compare zero-shot and fine-tuned UltraSAM with nnU-Net and US-Medical Matting on held-out healthy and scoliotic ultrasound sequences. Fine-tuned UltraSAM achieved the best overall performance, with a DSC of 0.655 and HD95 of 4.254, and maintained improved performance in scoliotic subjects despite more complex anatomy. Qualitative reconstruction results showed better preservation of the central vertebral trajectory, supporting its use for downstream curvature assessment. This work provides a reproducible, radiation-free framework for spine assessment in low-resource clinical environments.

Keywords: Scoliosis · Tracked ultrasound · Spine reconstruction · Medical image segmentation · UltraSAM · Resource-limited settings

1 Introduction

In sub-Saharan Africa, the prevalence of scoliosis is estimated to be between 3.3 and 5.5%, affecting millions [1]. Scoliosis is a three-dimensional deformity of the spine that impacts patients' health and quality of life, especially during its development [2,3]. This condition mainly affects children and adolescents; therefore,

⋆ Corresponding author: gabriella.dalbenzio@queensu.ca

early detection is paramount for treatment and management [2]. Treatment is primarily determined by the Cobb angle measured on full-spine X-rays, which are acquired every 4–6 months to monitor disease progression.

Management of scoliosis in sub-Saharan Africa is often hindered by limited healthcare infrastructure and diagnostic resources [4]. Notably, in Senegal, there is severe diagnostic inequality with only a single specialized low-dose radiation X-ray operational in the capital city of Dakar. Due to higher costs and limited accessibility to specialized centers in resource-limited countries, the use of this system is currently discouraged [5,6]. Moreover, public hospitals have long waiting lists for scoliosis radiographs, and many families cannot afford private imaging. As a result, most patients are diagnosed very late in the course of the disease – typically only after developing noticeable deformities, back pain, or respiratory issues. Consequently, many children present to specialists with already severe curvatures (i.e., Cobb angles $> 40^{\circ}$ at first evaluation). This delay in diagnosis leads to a high proportion of advanced cases and highlights the need for more accessible early screening in similar, low-resource contexts.

Since low-dose radiation X-rays are not widely accessible, scoliosis patients in many low-resource settings rely on repeated conventional X-rays for monitoring. Cumulative radiation exposure during childhood has been associated with increased long-term cancer risk in individuals with scoliosis [3,7]. This motivates safer imaging solutions for assessing and monitoring scoliosis. One such approach that has shown promise is ultrasonography coupled with an automatic segmentation algorithm [8,9]. Over the past decade, *Ungi et al.* have demonstrated the viability of using ultrasonography for assessing and monitoring scoliosis patients. Initially, a tracked ultrasound was used to measure spine curvature, and showed that using an ultrasound was a feasible substitute for X-ray [9]. This then evolved into integrating a deep learning model for reconstructing a 3D volume from automatic segmentation [8]. The development of nnU-Net [10] further improved segmentation performance through automatic dataset-specific configuration. The automatic segmentation pipeline became a successful strategy for other tasks using ultrasound segmentation, such as 3D needle guidance for percutaneous nephrostomy [11].

In 2025, *Ungi et al.* performed a comparative analysis of X-ray and bone-enhanced 3D ultrasound (BEUS) measurements in 34 pediatric scoliosis patients, where each participant underwent same-day spinal imaging with both modalities [12]. Minimal inter-observer variability was found, and 93% of BEUS measurements fell within a clinically acceptable range of $\pm 5^{\circ}$ of X-ray values [12]. This confirmed that combining ultrasound with 3D reconstruction was a comparable alternative for scoliosis assessment. While segmentation algorithms with ultrasonography have shown promise, data scarcity remains an issue. Good quality paired label data from low-resource areas like Senegal is difficult to obtain, and many Senegalese patients tend to have an under-represented anatomy in available scoliosis datasets, as they tend to be older with mature spines. Recent ultrasound foundation models such as OpenUS [13] and UltraSAM [14] benefit from representation learning with large ultrasound datasets as training data, but do not directly solve the label uncertainty problem for thin, curved, partially visible spine structures. Due to image resolution, artifacts, and operator error, among other factors, anatomical structures are often not clearly visible, if at all, in every ultrasound frame. Particularly for ambiguous or complex targets like the spine, improvements in automatic segmentation provide massive downstream implications for 3D reconstruction.

Contributions. The main contributions of this study are threefold. First, we present a tracked spine ultrasound dataset acquired in Senegal, addressing the limited representation of African and resource-constrained populations in scoliosis imaging research. Second, we develop an open-source workflow for automated 3D spine reconstruction using B-mode ultrasound, optical tracking, PLUS, OpenIGTLink, and 3D Slicer. Third, we benchmark multiple vertebral segmentation approaches and show that fine-tuned UltraSAM provides the most reliable performance for both healthy and scoliotic subjects. The implementation and accompanying usage documentation are publicly available at `https://github.com/dalbenzioG/slicer_spine_us`, supporting reproducibility and reuse by the research community.

2 Materials and Methods

2.1 Open-Source System Overview

We developed an open-source workflow for tracked ultrasound acquisition, automated vertebral segmentation, and three-dimensional spine reconstruction in 3D Slicer. B-mode ultrasound images were acquired using a MicrUs ultrasound system equipped with a C5-2R60S-3 convex probe (Telemed Medical Systems, Milan, Italy), with the imaging depth fixed at 90 mm. A freehand sagittal sweep was performed along the posterior spine. The long axis of the probe was aligned with the cranio–caudal direction to obtain sagittal images, while the acoustic beam was maintained approximately normal to the local skin surface as the probe was translated along the back.

Probe motion was measured using an OptiTrack Duo optical tracking system. A rigid marker attached to the ultrasound probe defined the probe coordinate system, while a second marker attached to the participant established a patient-relative reference coordinate system. Ultrasound images and tracking measurements were acquired and synchronized using the PLUS Toolkit [15] and transmitted to 3D Slicer through OpenIGTLink. SlicerIGT [16] was used for spatial calibration, transform management, real-time visualization, and sequence recording. Each ultrasound frame was stored together with the corresponding probe pose in the patient reference coordinate system. The tracked sequence was subsequently processed frame by frame using an automated segmentation pipeline, and the resulting predictions were compounded into a volumetric reconstruction. Segmentation therefore served as an intermediate step for identifying vertebral structures in the acquired images, while the principal output of the workflow was a patient-referenced 3D representation of the scanned spine.

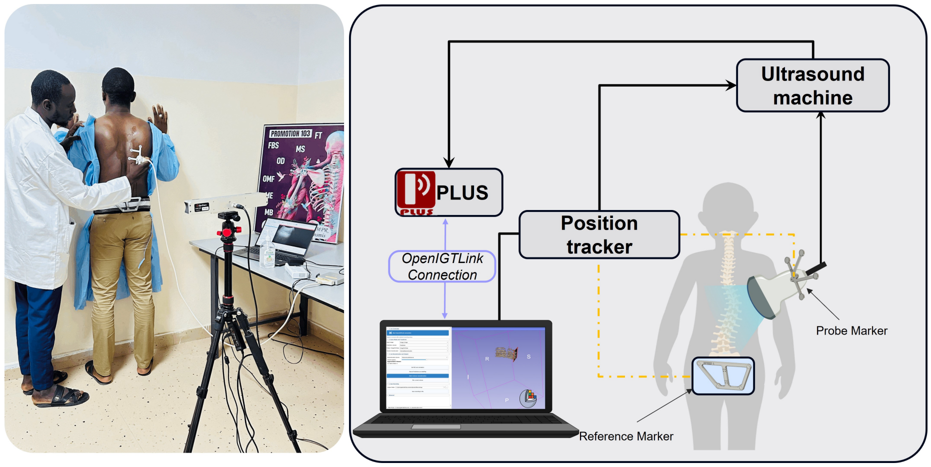

Fig. 1. Tracked ultrasound system for spine reconstruction. Left: clinical acquisition setup with B-mode ultrasound and optical tracking. Right: workflow showing the ultrasound machine, optical tracker, probe marker, and patient reference marker. Image and tracking streams are sent through PLUS to 3D Slicer via OpenIGTLink for synchronized acquisition and 3D reconstruction.

2.2 Frame-Wise Vertebral Segmentation

UltraSAM [14] was used to segment vertebral structures in each ultrasound frame. The model was operated using bounding-box prompts that were generated automatically from preceding predictions, removing the need for manual interaction during sequence processing. For the first frame, the prompt was initialized as a bounding box covering the entire ultrasound image. UltraSAM then produced a binary segmentation at the original image resolution. When the prediction contained foreground pixels, the tightest axis-aligned bounding box enclosing the predicted mask was used as the prompt for the following frame. In this way, the prompt was updated throughout the sequence according to the spatial location of the most recent successful prediction. When UltraSAM returned an empty mask, the prompt was not updated. Instead, the bounding box from the last frame with a valid prediction was retained and used to process the next frame. The empty prediction itself contributed no foreground information to the reconstruction. This strategy prevented an isolated segmentation failure from resetting the prompt or interrupting the remaining sequence. Only the bounding-box location was propagated between frames; the previous mask was not copied, registered, or provided directly to the model. Each frame was therefore segmented independently based on its ultrasound appearance and the automatically propagated spatial prompt. No ground-truth annotations or manually defined prompts were used during inference.

Table 1. Demographic and clinical characteristics of the TRUSD-S study cohort. Values are group means. Units: age, years; height, m; weight, kg; BMI, kg/m^2; Cobb angle, degrees. M, male; F, female. An em dash indicates that the measurement was not applicable.

Group	N	Gender	Age	Height	Weight	BMI	Cobb	Risser
Healthy Controls	21	61.9% M / 38.1% F	22.05	1.72	67.70	22.74	—	—
Scoliotic Patients	9	33.3% M / 66.7% F	25.33	1.68	61.67	21.90	25.86	3.83
Total Cohort	30	53.3% M / 46.7% F	23.69	1.70	64.69	22.32	—	—

2.3 3D Reconstruction

The frame-wise segmentation masks were combined into a volumetric reconstruction using the Volume Reconstruction functionality in 3D Slicer. Each predicted mask was processed together with the tracking transform associated with the corresponding ultrasound frame. Using the calibration and tracking transform chain, the two-dimensional mask was positioned in the patient reference coordinate system according to the location and orientation of the ultrasound probe at the time of acquisition.

The spatially positioned masks were inserted sequentially into a common three-dimensional reconstruction grid. Contributions from overlapping frames were compounded within the same volume, allowing vertebral regions observed from successive probe positions to be combined across the sagittal scanning trajectory. Frames with empty segmentation masks did not add foreground information but were otherwise retained in the recorded sequence.

Volume reconstruction was performed using an output voxel spacing of 1.0 mm isotropic and nearest-neighbor interpolation, with maximum-value compounding used to combine overlapping samples. Hole filling was disabled. The reconstructed volume was stored as a scalar 3D Slicer volume node in the patient reference coordinate system. For three-dimensional visualization, the reconstructed volume was displayed directly using GPU volume rendering (MR-Default transfer-function preset).

3 Experiments

3.1 Dataset and Data Splits

The study used a single lumbar spine ultrasound dataset comprising 45 tracked B-mode ultrasound video sequences acquired from 36 participants (17 healthy controls and 19 participants with scoliosis) from the Tracked Ultrasound Spine

Table 2. Composition of the training, validation, and test subsets. The participant column sums to 38 because two participants contributed sequences to both the training and validation sets; the test set is strictly participant-disjoint from all development data.

Subset	Participants (healthy / scoliotic)	Sequences	Annotated frames
Training	24 (11 / 13)	31	7,876
Validation	5 (3 / 2)	5	1,096
Test	9 (5 / 4)	9	3,335
Total	36 (17 / 19)	45	12,307

Dataset (TRUSD-S; Approval No. SN25/110). TRUSD-S is an ongoing data-collection initiative designed to improve the representation of populations that remain underrepresented in medical imaging research. Several participants contributed more than one sequence. Complete demographic and clinical information was available for 30 of the 36 participants and is summarized in Table 1. For six participants with scoliosis, complete clinical characterization was unavailable because radiographic examination could not be obtained due to financial constraints. Their ultrasound acquisitions were nevertheless retained, as the imaging data and segmentation annotations required for model development were complete.

The dataset was divided into fixed training, validation, and test subsets defined at the participant level. Table 2 reports the number of participants, sequences, and annotated frames in each subset. The training set comprised 31 sequences from 24 participants (7,876 frames), the validation set 5 sequences from 5 participants (1,096 frames), and the test set 9 sequences from 9 participants (3,335 frames). The test set was strictly participant-disjoint from all data used for model development: no test participant contributed any sequence to training or validation. Two participants each contributed one sequence to the validation set while their remaining sequences were assigned to training; all other participants appeared in a single subset. The test set included five healthy and four scoliotic participants, enabling group-wise evaluation.

3.2 Evaluation of Segmentation Approaches

We evaluated nnU-Net, US-Medical Matting, and UltraSAM to identify the most suitable approach for frame-wise vertebral segmentation and subsequent 3D reconstruction. All methods were evaluated on the same test sequences and at the original image resolution.

nnU-Net. We trained nnU-Net [10] from scratch on the TRUSD-S training set. To address the small proportion of foreground pixels occupied by the vertebral structures, the cross-entropy component of the Dice–cross-entropy loss was assigned class weights of 0.1 for background and 0.9 for foreground [17]. Dropout was applied with a probability of 0.20.

US-Medical Matting. We adapted the Medical Matting architecture of Wang et al. [18] to investigate whether modeling ambiguous bone boundaries improved segmentation. Because only one annotation was available per frame, the unknown trimap region was generated synthetically by extending the annotated bone surface across locally plausible continuations supported by hyperechoic ridges or acoustic shadowing. The resulting trimap represented assumption-based ambiguity rather than inter-annotator variability. The original two-stage architecture was trained from scratch without structural modification on the TRUSD-S dataset. It generated eight probabilistic segmentation samples and used their aggregate uncertainty to refine a continuous alpha prediction. The training objective and adaptive loss weighting followed the original Medical Matting formulation. At inference, the model received only the ultrasound image, and the predicted alpha map was thresholded at 0.5 for binary evaluation.

UltraSAM. UltraSAM [14] was evaluated under two configurations: (i) using the publicly released pretrained weights without task-specific adaptation, and (ii) after fine-tuning on the training subset of TRUSD-S. Model selection during fine-tuning was performed using the validation subset, while final performance was assessed exclusively on the held-out test set. For both configurations, test sequences were processed in their original acquisition order using the same fully automated prompting strategy employed during 3D reconstruction. Ground-truth annotations and manually defined prompts were not used during inference.

Quantitative and Qualitative Assessment: Binary predictions were evaluated using the Dice coefficient (DSC), intersection over union (IoU), precision, recall, and the 95th-percentile Hausdorff distance (HD95). The US-Medical Matting alpha predictions were thresholded at 0.5 before computing these metrics. Results were calculated independently for each frame and are reported as mean and standard deviation for the complete test set and separately for healthy and scoliotic participants. External generalization was evaluated on the handheld subset of the CT-referenced lumbar ultrasound dataset of Cavalcanti et al. [19], comprising 2,322 annotated frames from 7 subjects. This evaluation was performed only for the approach showing the strongest and most consistent segmentation performance on our internal test set, i.e., the model selected for integration into the 3D reconstruction workflow. The same fine-tuned model and inference settings were used without retraining, adaptation, or dataset-specific parameter selection.

4 Results and Discussion

Segmentation Performance. The comparison in Table 3 was designed to identify the most suitable segmentation approach for the proposed workflow among realistically available methods. nnU-Net was trained from scratch using its standard self-configuring setup, whereas UltraSAM benefits from large-scale ultrasound pretraining. However, pretraining alone was insufficient, as zero-shot UltraSAM underperformed nnU-Net in DSC; the performance advantage emerged only after fine-tuning on the local dataset. Fine-tuned UltraSAM achieved the best overall performance in HD95, DSC, IoU, and precision, reducing HD95 by 55% and improving DSC, IoU, and precision by 56.7%, 74.1%, and 32.1%, respectively, over the second-best method. Zero-shot UltraSAM achieved the highest recall but substantially lower precision, indicating over-segmentation. Performance decreased for all methods in the scoliotic subgroup, reflecting the greater anatomical variability and less regular ultrasound appearance. Nevertheless, fine-tuned UltraSAM remained the strongest method, reducing HD95 by 42.3% and improving DSC and IoU by 60.1% and 75.4%, respectively. US-Medical Matting showed the largest degradation in scoliotic cases, while fine-tuned UltraSAM provided the best overall balance across segmentation metrics. Comparison with task-specific models pretrained on large-scale ultrasound datasets remains an important direction for future work as such models become publicly available.

Table 3. Segmentation performance of different methods on our local test set. The overall test set (All) comprises 9 participants and 3,335 frames; the healthy subgroup 5 participants and 1,639 frames; the scoliotic subgroup 4 participants and 1,696 frames. Results are reported as mean ± standard deviation over frames. Best results are shown in bold and second-best results are underlined.

Method	**Group**	**HD95 ↓**	**DSC ↑**	**IoU ↑**	**Prec. ↑**	**Rec. ↑**
nnU-Net	**All** (9; 3,335)	19.20 ± 21.41	<u>0.42 ± 0.24</u>	<u>0.29 ± 0.19</u>	<u>0.49 ± 0.25</u>	0.48 ± 0.31
	Healthy (5; 1,639)	16.32 ± 19.39	<u>0.54 ± 0.20</u>	<u>0.39 ± 0.17</u>	<u>0.53 ± 0.21</u>	0.62 ± 0.24
	Scoliotic (4; 1,696)	21.69 ± 22.72	0.31 ± 0.23	0.21 ± 0.17	<u>0.45 ± 0.28</u>	0.35 ± 0.30
US-Medical Matting	**All** (9; 3,335)	25.94 ± 23.79	0.31 ± 0.28	0.22 ± 0.21	0.37 ± 0.32	0.34 ± 0.34
	Healthy (5; 1,639)	17.52 ± 19.58	0.49 ± 0.24	0.35 ± 0.19	0.49 ± 0.25	0.56 ± 0.30
	Scoliotic (4; 1,696)	34.07 ± 24.65	0.14 ± 0.21	0.09 ± 0.15	0.26 ± 0.34	0.12 ± 0.21
UltraSAM zero-shot	**All** (9; 3,335)	<u>9.46 ± 6.79</u>	0.39 ± 0.25	0.28 ± 0.23	0.29 ± 0.23	$\mathbf{0.80 \pm 0.20}$
	Healthy (5; 1,639)	<u>8.69 ± 7.46</u>	0.43 ± 0.28	0.32 ± 0.25	0.32 ± 0.25	$\mathbf{0.81 \pm 0.23}$
	Scoliotic (4; 1,696)	<u>10.20 ± 5.98</u>	<u>0.36 ± 0.21</u>	<u>0.24 ± 0.19</u>	0.25 ± 0.19	$\mathbf{0.78 \pm 0.17}$
UltraSAM finetuned	**All** (9; 3,335)	$\mathbf{4.25 \pm 6.97}$	$\mathbf{0.66 \pm 0.18}$	$\mathbf{0.51 \pm 0.18}$	$\mathbf{0.65 \pm 0.17}$	<u>0.71 ± 0.23</u>
	Healthy (5; 1,639)	$\mathbf{2.57 \pm 5.03}$	$\mathbf{0.74 \pm 0.11}$	$\mathbf{0.60 \pm 0.14}$	$\mathbf{0.72 \pm 0.13}$	<u>0.78 ± 0.15</u>
	Scoliotic (4; 1,696)	$\mathbf{5.89 \pm 8.10}$	$\mathbf{0.57 \pm 0.18}$	$\mathbf{0.42 \pm 0.17}$	$\mathbf{0.57 \pm 0.18}$	<u>0.64 ± 0.26</u>

External Validation Assessment. The external evaluation supports the transferability of the segmentation stage: without any retraining or adaptation, the fine-tuned model retained nearly the same region overlap on the external cohort as on our internal test set (DSC 0.64 versus 0.66), despite differences in ultrasound system, probe, population, and annotation protocol. Boundary errors were larger and more variable externally driven by a small fraction of frames with multiple weakly visible bone structures. These results pertain to the segmentation stage only; the 3D reconstruction workflow was developed and assessed on TRUSD-S data.

Qualitative Reconstruction Assessment. Qualitative inspection of the 3D reconstructions showed that UltraSAM fine-tuned produced the most anatom-

Table 4. External evaluation of the fine-tuned UltraSAM model on the handheld subset of Cavalcanti et al. [19] (7 subjects, 2,322 annotated frames). The internal test-set result (All group, Table 3) is repeated for reference.

Cohort	**HD95** $\downarrow$	**DSC** $\uparrow$	**IoU** $\uparrow$	**Prec.** $\uparrow$	**Rec.** $\uparrow$
Internal (9; 3,335)	4.25 ± 6.97	0.66 ± 0.18	0.51 ± 0.18	0.65 ± 0.17	0.71 ± 0.23
External (7; 2,322)	6.0 ± 4.36	0.64 ± 0.14	0.48 ± 0.14	0.59 ± 0.09	0.73 ± 0.24

ically coherent representation of the spine (Fig. 2). Compared with nnU-Net and US-Medical Matting, the reconstructed vertebral structures appeared more continuous and better organized along the central spinal trajectory. This is clinically relevant because scoliosis assessment depends on a stable representation of the spinal curve, from which ultrasound-derived curvature measurements can be estimated. The nnU-Net reconstruction preserved parts of the posterior spinal anatomy but showed more scattered false-positive components and reduced continuity. US-Medical Matting produced a visibly sparser and less complete reconstruction, with weaker preservation of the central vertebral line. In contrast, UltraSAM fine-tuned better maintained the midline structure corresponding to the deformity visible on the reference radiograph.

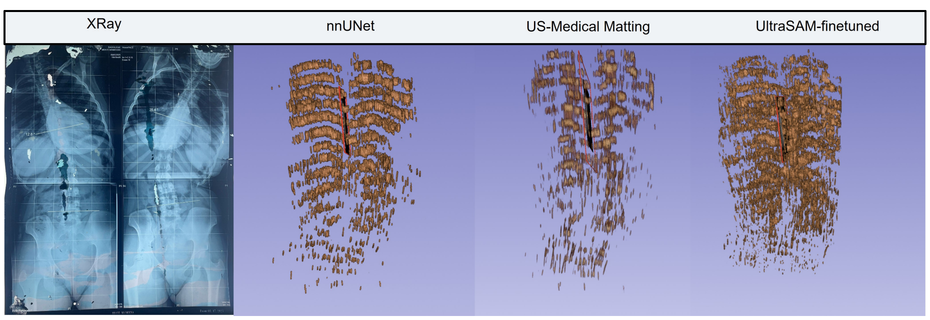

Fig. 2. Qualitative comparison of lumbar spine reconstruction methods for a representative participant with scoliosis. From left to right: the reference radiograph and 3D ultrasound reconstructions obtained using nnU-Net, US-Medical Matting, and fine-tuned UltraSAM. The reconstructed vertebral structures are shown in brown.

5 Conclusions

This pilot study demonstrates the feasibility of a fully open-source, automated 3D spine reconstruction workflow from tracked ultrasound, designed for resource-constrained environments such as Senegal. By integrating a fine-tuned UltraSAM model, the pipeline achieved a substantial improvement in segmentation accuracy over the nnU-Net and US-Medical Matting baselines, maintaining its performance under the complex anatomical variability of scoliotic participants; the fine-tuned model provided the best balance of precision and recall. Relying exclusively on accessible, zero-cost infrastructure—3D Slicer, OpenIGTLink, and the PLUS Toolkit—the workflow constitutes a step toward a radiation-free complement to radiographic spinal assessment in settings where X-ray access is limited, such as sub-Saharan Africa. Its clinical utility as an alternative to radiography remains to be established: this will require validating ultrasound-derived curvature measurements against clinically established radiographic Cobb angles in a larger cohort, which is the objective of the ongoing TRUSD-S data-collection initiative.

Acknowledgments. Gabor Fichtinger is a Canada Research Chair in Computer Integrated Surgery, Tier 1.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Waterkeyn, F., Ikwuegbuenyi, C.A., Woodfield, J., Sommer, F., Magogo, J., Cheserem, B., Schupper, A.J., Shabani, H.K., Hussain, I., Ahmad, A.A., Balsano, M., Mangat, H., Härtl, R.: Evaluating the Feasibility and Outcomes of a Scoliosis Surgical Camp in a Resource-Limited Setting in Sub-Saharan Africa. World Neurosurgery **180**, e550–e559 (2023)
2. Aulisa, A.G., Giordano, M., Guzzanti, V., Falciglia, F., Pizzetti, P., Toniolo, R.M.: Effectiveness of school scoliosis screening and the importance of this method in measures to reduce morbidity in an italian territory. Journal of Pediatric Orthopaedics B **28**(3), 271–277 (2019)
3. Luan, F.J., Wan, Y., Mak, K.C., Ma, C.J., Wang, H.Q.: Cancer and mortality risks of patients with scoliosis from radiation exposure: A systematic review and meta-analysis. European Spine Journal **29**(12), 3123–3134 (2020). https://doi.org/10.1007/s00586-020-06573-7
4. Tiaho, Y., Seri, S.L., Bombo, J., Traoré, S., Faton, D., Akadjé, D., Kinta, H., Tano, M., Kouakou, J., Alloh, A., et al.: Étude rétrospective sur le parcours thérapeutique des patients atteints de scoliose idiopathique Vus en service de médecine physique et de réadaptation à Abidjan [Retrospective study on the therapeutic pathway of patients with idiopathic scoliosis seen in physical medicine and rehabilitation service in Abidjan]. Kinésithérapie, la Revue **21**(236–237), 36–39 (2021), (French)
5. McKenna, C., Wade, R., Faria, R., Yang, H., Stirk, L., Gummerson, N., Sculpher, M., Woolacott, N.: EOS 2D/3D X-ray imaging system: A systematic review and economic evaluation. Health Technology Assessment **16**(14) (2012)
6. Rose, L.D., Williams, R., Ajayi, B., Abdalla, M., Bernard, J., Bishop, T., Papadakos, N., Lui, D.F.: Reducing radiation exposure and cancer risk for children with scoliosis: EOS the new gold standard. Spine Deformity **11**(4), 847–851 (2023)
7. Doody, M.M., Lonstein, J.E., Stovall, M., Hacker, D.G., Luckyanov, N., Land, C.E.: Breast Cancer Mortality After Diagnostic Radiography: Findings From the U.S. Scoliosis Cohort Study. Spine **25**(16), 2052–2063 (2000)

8. Ungi, T., Greer, H., Sunderland, K.R., Wu, V., Baum, Z.M.C., Schlenger, C., Oetgen, M., Cleary, K., Aylward, S.R., Fichtinger, G.: Automatic Spine Ultrasound Segmentation for Scoliosis Visualization and Measurement. IEEE Transactions on Biomedical Engineering **67**(11), 3234–3241 (2020)
9. Ungi, T., King, F., Kempston, M., Keri, Z., Lasso, A., Mousavi, P., Rudan, J., Borschneck, D.P., Fichtinger, G.: Spinal Curvature Measurement by Tracked Ultrasound Snapshots. Ultrasound in Medicine & Biology **40**(2), 447–454 (2014)
10. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: NNU-net: A self-configuring method for deep learning-based biomedical image segmentation. Nature Methods **18**(2), 203–211 (2020)
11. Kim, A., Yeung, C., Szabo, R., Sunderland, K., Hisey, R., Morton, D., Kikinis, R., Diao, B., Mousavi, P., Ungi, T., Fichtinger, G.: Percutaneous nephrostomy needle guidance using real-time 3D anatomical visualization with live ultrasound segmentation. In: Medical Imaging 2024: Image-Guided Procedures, Robotic Interventions, and Modeling. vol. 12928, p. 129280S. International Society for Optics and Photonics, SPIE, San Diego, United States of Ameria (2024), `https://doi.org/10.1117/12.3006533`
12. Ungi, T., Wu, C.O., Dance, S.J., Lawall, C., Nagel, C., Khodorov, G., Ford, R.W., Cleary, K., Fichtinger, G., Oetgen, M.E., Borschneck, D.P.: Bone-enhanced 3D ultrasound: A non-ionizing alternative for pediatric scoliosis assessment. European Spine Journal **34**(7), 2662–2668 (2025)
13. Zheng, X., Chen, X., Rauf, A., Fu, Q., Monosi, B., Rivellese, F., Lewis, M.J., Gong, S., Slabaugh, G.: Openus: A fully open-source foundation model for ultrasound image analysis via self-adaptive masked contrastive learning (2025), `https://arxiv.org/abs/2511.11510`
14. Meyer, A., Murali, A., Zarin, F., Mutter, D., Padoy, N.: Ultrasam: A Foundation model for ultrasound using large open-access segmentation datasets. International Journal of Computer Assisted Radiology and Surgery **21**(1), 93–102 (2025)
15. Lasso, A., Heffter, T., Rankin, A., Pinter, C., Ungi, T., Fichtinger, G.: Plus: open-source toolkit for ultrasound-guided intervention systems. IEEE transactions on biomedical engineering **61**(10), 2527–2537 (2014)
16. Ungi, T., Lasso, A., Fichtinger, G.: Open-source platforms for navigated image-guided interventions. Medical image analysis **33**, 181–186 (2016)
17. Hashtrudi-Zaad, K., Ungi, T., Yeung, C., Baum, Z., Cernelev, P.D., Hage, A.N., Schlenger, C., Fichtinger, G.: Expert-guided optimization of ultrasound segmentation models for 3d spine imaging. In: 2024 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE). pp. 680–685. IEEE (2024)
18. Wang, L., Ye, X., Ju, L., He, W., Zhang, D., Wang, X., Huang, Y., Feng, W., Song, K., Ge, Z.: Medical matting: Medical image segmentation with uncertainty from the matting perspective. Computers in Biology and Medicine **158**, 106714 (2023)
19. Cavalcanti, N.A., Li, R., Arango, L., Davoodi, A., Van Assche, K., Ao, Y., Massalimova, A., Salehi, M., Zingg, L., Götschi, T., et al.: A large, paired dataset of robotic and handheld lumbar spine ultrasound with ground-truth ct benchmarking. Scientific Data **12**(1), 1766 (2025)