

Sparse-XM: Spine Pose Adjustment with RGB-D Bone Segmentation via Cross-Modality Label Transfer

William R. Warner¹, Indrani Bhattacharya², Linton T. Evans³, Sohail K. Mirza¹, Keith D. Paulsen^{1,2,4}, Xiaoyao Fan^{1*}

¹ Thayer School of Engineering, Dartmouth College, Hanover, NH, USA

² Geisel School of Medicine, Dartmouth College, Hanover, NH, USA

³ Section of Neurosurgery, Dartmouth-Hitchcock Medical Center, Lebanon, NH, USA

⁴ Norris Cotton Cancer Center, Dartmouth-Hitchcock Medical Center, Lebanon, NH, USA
Xiaoyao.Fan@dartmouth.edu

Abstract. In open spine surgery, navigation requires registration between the surgical field and volumetric CT. The spine pose changes between preoperative CT (pCT) and intraoperative patient positioning, and can further change after intraoperative CT (iCT) during surgery, degrading navigation accuracy. In this study, we developed a novel, fully automated end-to-end system for spine pose adjustment driven by intraoperative stereovision imaging (iSV) images. Our system includes three innovative modules. First, we present a method to automatically generate weak bone labels in stereo images via co-registration with iCT images. The automated labeling process addresses the labor and expertise-intensive challenges associated with supervised bone segmentation models that typically require manually segmented labels for training. Second, we train a fully convolutional deep learning method that integrates complementary information from the color (RGB) and depth (D) images to automatically segment bone using the weak labels. Finally, the segmented bone structures are used to perform a pose-adjusted registration. Data collected from 5 porcine cadavers were used for training and validation, and data from 2 porcine cadavers were used for independent testing. Pose-adjusted registration accuracy across all lumbar levels of test specimens was 2.0 ± 1.1 mm, compared to 2.5 ± 1.5 mm using manual segmentation, and 9.1 ± 6.8 mm using a commercially available navigation system. The fully automated pose-adjusting registration framework compensated for spine motion between pCT and intraoperative positioning and overall achieved clinically acceptable accuracy. Our approach was not user or expertise-dependent and holds potential for wider adoptions in open spinal procedures for intraoperative spine motion correction. Code is available at <https://github.com/wRossW/Sparse-XM-Spine-Pose-Adjustment>.

Keywords: Deep Learning Segmentation, Stereovision Images, Surgical Navigation, Spine Motion Adjustment.

1 Introduction

Surgical navigation has become the standard of care in open spinal procedures, providing surgeons with more perception, reducing procedural time, and increasing surgical accuracy [1, 2]. Navigation relies on registration between a volumetric image (e.g., CT) and the surgical field. Preoperative CT (pCT) images are typically acquired in a supine position that differs from the intraoperative surgical position (prone), and the pose changes significantly degrade registration accuracy using pCT. Tracked, mobile intraoperative CT (iCT) units (e.g., O-arm, Medtronic) provide automatic surgical field registration to enable navigation. However, iCT requires large capital expense [3, 4], increases radiation exposure for patients [5] and surgical team, and significantly interrupts workflow by requiring intraoperative time [6] during which the surgical team is required to leave and re-enter sterile field in addition to the imaging time (~15 min). Further pose change may occur during the procedure and degrade registration accuracy.

Compared to fluoroscopy [7] and ultrasound [8, 9] based frameworks that have been proposed to enable intraoperative registration, three-dimensional (3D) vision systems provide radiation-free registration and can account for intraoperative pose change with clinically acceptable accuracy and efficiency [10, 11]. The registration process can be repeated to allow for continual monitoring throughout surgery. We have implemented intraoperative stereovision imaging (iSV) systems to reconstruct the exposed spine surface, and have developed an automated pose-adjusted registration framework (Figure 1) to correct for spinal motion [12, 13]. In prior studies, accurate manual bone segmentation from iSV was required from a user, limiting wider adoption of the approach. Deep learning methods can be leveraged to eliminate user intervention involved in registration driven by iSV. Toward this end, Liebmann et. al [14] and Massalimova et al. [15] segmented bone surface from iSV using color image information, but the study was limited to explanted spines and depth information was not utilized for surface segmentation. Transition to in-situ and in-vivo segmentation can be challenging as deep learning methods typically require labor and expertise to generate labels for supervised training.

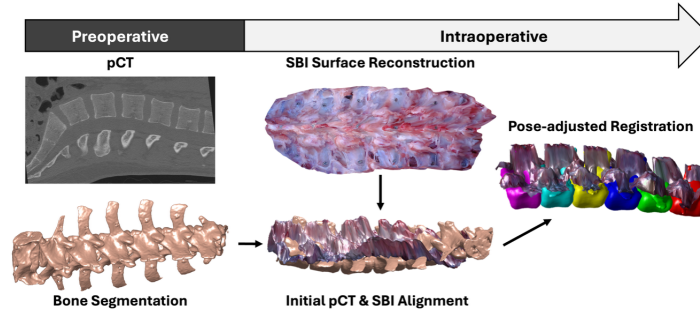


Fig. 1. Pose-adjusted registration process. pCT was acquired in a supine position while intraoperative positioning is prone. iSV images were acquired in the intraoperative prone position and used to drive a registration to compensate for pose changes between each vertebra in pCT (colored) and their counterpart in iSV.

In this paper, we propose a *Sparse-XM approach*: Spine Pose Adjustment with RGB-D Bone Segmentation via Cross-Modality Label Transfer. The contribution of this approach is threefold. First, we developed a cross-modality label transfer scheme for automatically generating weakly-labeled datasets to overcome the challenges associated with labor-intensive labeling process. Second, we adapted and trained the Stanford Prostate Cancer Network (SPCNet) [16], a network based on a modified holistically-nested edge detection (HED) architecture [17], in order to segment bone from color (RGB) and depth (D) images. Third, we completed a fully automated spine pose adjustment framework to account for spine motion between pCT and intraoperative positioning using predictions from the adapted SPCNet. We evaluated the accuracy performance of the Sparse-XM approach in 2 testing specimens, and compared to accuracies achieved with manual segmentation and automatically generated weak labels.

2 Proposed Methods

2.1 Dataset

Data from seven (7) porcine cadavers were collected. For each specimen, a posterior midline incision was performed exposing the posterior surface of the lumbar spine (L1-L6). A total of 18 mini screws (1.5-mm diameter; Stryker) were implanted as fiducial markers in the spine with three at each vertebral level (one at the spinous and one at each transverse process). The specimen was then placed in a supine position and a pCT acquired representing preoperative imaging. A navigation system (StealthStation, Medtronic) was used for optical tracking. The animal was then positioned prone and a reference frame (Spinous Reference Frame, Medtronic) was rigidly attached to the sacrum. An additional iCT was acquired as ground truth. A series of iSV images of the exposed spine surface were acquired in the same surgical prone position using a custom designed imaging system described in [12]. The iSV system was tracked via a rigidly attached tracker (sureTrak active, Medtronic) and calibrated, and the locations of reconstructed surfaces were computed relative to the reference tracker. A point-based rigid registration was performed between the surgical field and iCT with the navigation system using 10 out of the 18 fiducials (maximum number allowed in the system) to achieve co-registration between iSV and iCT. A second point-based rigid registration per individual image (8 ± 2 fiducials/image) was performed. A total of 352 tracked iSV images were acquired across 7 specimens (50 ± 16 images per specimen) and the mean fiducial registration error was 0.6 ± 0.2 mm across images used for model training.

2.2 Cross Modality Label Transfer for Generation of Weak Labels

The spine was segmented from iCT via thresholding, and a 3D point cloud of the spine surface was generated. A 3D point cloud of the exposed spine was reconstructed from each iSV image. The iCT and iSV point clouds were co-registered in the same coordinate system through tracking and iCT registration. Points from iSV within a distance threshold of 3 mm to iCT were classified as bone surface points (Figure 2A) [18]. Briefly, a closest point distance algorithm was used to calculate the distance between

iCT and iSV point clouds, and the distance threshold was applied to select valid bone points from iSV. Since each point in iSV point clouds corresponded to a pixel in the iSV image, the bone surface classification results were transformed into the original pixel space of the rectified stereo image to create labels of the bone surface for all 352 RGB-D images. The resulting labels were weak labels of the bone as the automatic transfer contained various errors and noises from components in the system, such as errors from tracking, 3D reconstruction, calibration, and iCT registration [12].

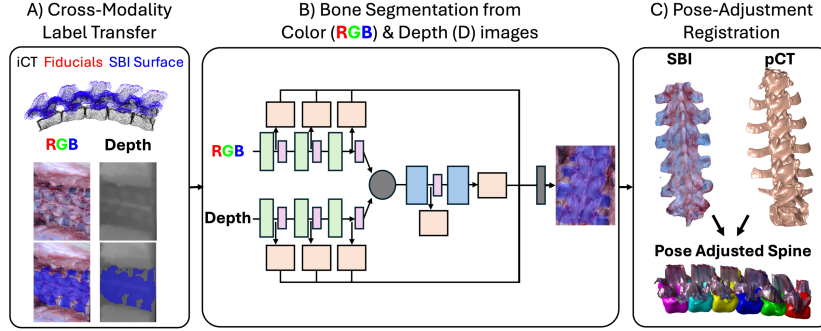


Fig. 2. The Sparse-XM approach. A) Cross Modality Label Transfer relies on coregistration of iSV and iCT surfaces to transfer bone labels to RGB-D images. B) Cross Modality weak labels are leveraged for training bone segmentation network. C) predictions from the network are passed as input to the pose-adjustment registration framework.

2.3 Learning Bone Segmentation from RGB & D Images

A modified SPCNet network architecture (Figure 2B) was adapted for bone segmentation. The adapted-SPCNet architecture leveraged deep supervision that was applied to multiple side outputs as well as the final downstream output of the architecture to learn both multi-level and multi-scale features [17]. These design components were beneficial for aiding the extraction of edges in the bone segmentation task and the reason for SPCNet selection. The original architecture was modified to handle multimodal inputs: two parallel input branches taking the RGB and D images respectively from a single iSV input pair. Each input was passed into a series of convolutional layers and were fused into a single fusion branch where additional convolutions were applied. Upsampled side outputs were extracted along each branch and concatenated with the final fusion output for binary segmentation. The RGB and D pair and its corresponding weak label formed a training triplet. Training involved optimizing a loss function containing binary cross entropy (BCE) component and a Dice Coefficient component (Equation 1). Best loss utilized an equal weighting ($\alpha = \beta = 0.5$ in Equation 1).

$$\mathcal{L}_{total} = \alpha \cdot \mathcal{L}_{BCE} + \beta \cdot \mathcal{L}_{Dice} \quad (1)$$

The adapted-SPCNet was trained using five-fold cross validation with an 80-20 training-validation split. Images from 5 of the porcine specimens (a total of 283 images) were used for training and validation, and two porcine specimens (a total of 69 images)

were reserved for independent testing. Each model in the cross-validation was trained for 100 epochs using early stopping criteria. Augmentation included random horizontal and vertical flipping (RGB and D), and color jitter (RGB only). Z-score normalization was performed per specimen.

2.4 Pose-Adjustment Registration

Figure 2C shows the pose-adjustment registration process. The spine was segmented from pCT via thresholding and a surface point cloud was generated. Bone segmentation was performed by adapted SPCNet on a series of iSV images extending over the entire exposed lumbar spine. The predicted masks were applied to corresponding iSV point clouds, and the point clouds from all images were merged using tracking information into a composite point cloud that represented the exposed spine surface. Both pCT and iSV point clouds were rectified into a neutral spine pose using a nonlinear principle component analysis algorithm [19]. Depth images were then generated, respectively, and registered via a denseSIFT algorithm [20]. The 2D correspondence between depth images were used to compute 3D correspondences between pCT and iSV point clouds [13]. A global rigid registration was performed between pCT and iSV point clouds. Each vertebral level was then segmented using depth image from rectified pCT, and each level of pCT point cloud was rigidly registered with its iSV counterpart [13].

2.5 Evaluation

To quantify performance of a trained model, an average Dice Coefficient was calculated for the validation data set for hyperparameter selection. Dice Coefficient between the predicted segmentations and weak labels was computed for test images. Given labels are weak and noisy, the accuracy of Dice Coefficient does not faithfully portray the accuracy of the bone segmentation algorithm. However, they do provide a means to roughly assess the algorithm for the next module – pose-adjustment registration.

To quantify the accuracy of the pose-adjustment registration, a target registration error (TRE) was calculated for each fiducial and each vertebral level. Specifically, the distances between corresponding fiducial locations from ground truth iCT and the pose-adjusted pCT were calculated. Average TRE of each vertebral level and across levels were reported for both test specimens. Additionally, the distances between fiducials in pCT and their counterparts in the surgical field were quantified to measure the preoperative-to-intraoperative motion.

2.6 Experiments & Ablation Studies

While the composite accuracy of the co-registered iCT-iSV dataset was used to select distance threshold for weak label generation, various thresholds ranging from 3 to 5.5 mm were utilized for classification of bone surface via label transfer. Evaluation and selection of the final threshold as 3 mm was guided by an expert user of the surgical navigation system (with >15 years of experience) based on conformity of the resulting bone masks to the transverse process surfaces when displayed in the RGB images and

minimization of the inclusion of surrounding soft tissue, as well as error sources contributing to RGB-D data, e.g. calibration, reconstruction, and tracking [12]. Experimentation with morphological operations was performed and hole filling applied to all weak labels.

Ablation experiments for the bone segmentation model included training with (a) RGB only, D only, and both RGB and D as inputs to the model, and (b) different combinations of binary cross-entropy and/or Dice losses. The model with optimal validation set performance was chosen as our bone segmentation model.

Various bone segmentation masks were used to run the pose-adjustment registration. A series of 5-6 iSV images per test specimen were chosen as inputs to the registration framework that provided full coverage of the lumbar spine. Three sets of bone segmentation masks were applied to this series of iSV images: (1) the weak labels corresponding to each image, (2) manually contoured bone surfaces by an expert user, and (3) the adapted-SPCNet predictions from the selected model. These masks, when applied to the iSV images, created a surface reconstruction of lumbar bone surface that was passed to the registration process. The TREs using each set of bone segmentation masks were computed to compare the overall performance of each segmentation.

3 Experimental Results

3.1 Qualitative Assessment of Cross Modality Label Transfer

The evaluation of the Sparse-XM method was first performed by qualitatively assessing the labels generated by the cross-modality label transfer and the resulting bone segmentation by the adapted-SPCNet. Representative weak labels and predicted labels are shown in Figure 3. The weak labels (green) contained irregular contours due to noises and errors, whereas predicted segmentation masks (red) produced regular shapes.




Fig. 3. Representative examples of the weak label generated via the cross-modality label transfer process (green). Sample predictions are shown from Sparse-XM (red).

3.2 Bone Segmentation Results for Model Selection

The model with RGB and depth images as inputs to the independent branches of our adapted SPCNet model yielded the best performance, over using either RGB or depth images alone as inputs. We posit the preferred performance using RGB+D was due to complementary information in the images captured by the model. The performance of BCE, Dice, and weighted average of BCE and Dice losses were comparable on the test set ($85.4 \pm 8.8\%$, $85.8 \pm 7.1\%$, and $86.4 \pm 6.9\%$ respectively). While the average dice coefficient was calculated, it was a secondary endpoint of the Sparse-XM performance.

3.3 Pose-Adjustment Registration Accuracy

The output of the registration framework was a pose-adjusted model of the pCT such that each vertebral level was aligned with the reconstructed bone surface from iSV imaging. The distance between the fiducial locations from the adjusted pCT model and the locations measured on iCT was reported as a target registration error (TRE) per vertebral level as the primary endpoint of the Sparse-XM performance. Table 1 shows TREs with three types of segmentation mask applied to the series of iSV images: weak label, manual, and prediction by adapted-SPCNet. The average TRE across 12 levels was significantly lower ($p < 0.05$) when driving the pose-adjustment registration framework with predicted masks, compared to weak labels and manual contours (5 and 6 images for test samples 1 and 2 respectively). Normality of the paired-difference was assessed and either a paired t- or Wilcoxon signed-rank test applied as appropriate for determination of statistical significance.

Table 1. Target Registration Error of Pose-Adjusted Registration using Weak Labels, Manual Contours, and predictions from the Sparse-XM approach

Units of mm	Porcine Test Sample 1			Porcine Test Sample 2		
Level	Weak Label	Manual	Sparse-XM	Weak Label	Manual	Sparse-XM
1	2.6±0.8	2.2±0.8	2.2±1.1	3.8±0.6	2.9±1.3	2.6±1.4
2	2.0±0.0	1.8±0.3	1.0±0.1	2.0±0.4	0.9±0.3	1.0±0.5
3	1.4±0.1	2.5±0.4	1.8±0.4	2.1±0.8	2.2±0.6	2.0±1.5
4	2.3	1.8	1.0	1.7±0.3	2.7±0.6	2.2±1.4
5	3.0±1.2	1.2±0.5	1.4±0.3	2.1±0.6	5.5±3.1	2.3±0.2
6	3.0±0.3	3.5±0.4	3.2±1.6	2.5±1.5	2.8±1.0	2.7±1.6
Mean	2.4±0.6	2.2±0.8	1.8±0.9	2.4±0.8	2.8±1.5	2.1±0.6

To quantify the amount of motion corrected by the pose-adjusted registration, the TREs were reported using the pCT registration and compared to the pose-adjusted registrations using Sparse-XM (Figure 4). The average TRE using pCT was 9.1 ± 6.8 mm across all fiducials. The higher TREs at L1 and L6 (up to 25 mm) indicate the amount of preoperative-to-intraoperative motion that has occurred. In contrast, TREs using Sparse-XM were evenly distributed across levels.




Fig. 4. Comparison of pose-adjusted vs. pCT-based registration averaged over each level for each test specimen. Higher registration error is observed at L1 and L6 of the uncorrected registration. L2 results for Test Samples 1 and 2 are equivalent for the Sparse-XM approach.

4 Discussion

In this study, we present Sparse-XM, a fully automated level-wise registration for pose adjustment in image-guided open spine surgery, through a series of advances. First, the labeling process was fully automated for adoption of deep learning methods for segmentation to address the significant human effort and time required to label training data for deep learning models. In medical imaging segmentation, labeling typically involves manually drawing contours in each image, which requires expertise (e.g., a physician), is subject to inter-/intra-user variability, is time-consuming and resource-intensive, and slows technical development. We leveraged co-registered datasets and generated bone labels automatically in a challenging image modality (iSV) through co-registration with images from a modality where corresponding labels were easily produced (CT). Although iCT images were used to generate weak labels for training, they are not required at inference, making Sparse-XM clinically relevant. Second, bone segmentation from iSV was fully automated by incorporating an adapted-SPCNet that integrates complementary information from RGB and Depth channels of iSV. In prior studies, manual bone segmentation in iSV images was provided by a user, which added time, required technical and medical expertise, and was subject to user error.

The accuracy of Sparse-XM was assessed in terms of fiducial TREs. The pose adjustment in test cases was successful. Registration using pCT showed higher TREs (up to 25 mm) towards the ends of exposure (L1 and L6) due to pose change, whereas the Sparse-XM TREs were evenly distributed across levels (Figure 4) and the overall TREs meet clinically acceptable range [21]. The overall TRE measurements were affected by all components in the registration framework, including iSV reconstruction (~ 1.2 mm), bone segmentation, and tracking accuracy (< 1 mm), and level-wise registration [12]. In addition, the co-registration error in label transfer was subject to similar reconstruction and tracking errors, and affected the labeling accuracy (e.g., Figure 3). We selected 3 mm as a distance threshold to provide sufficient labels while not exceeding the clinically acceptable accuracy [21]. Nevertheless, the adapted SPCNet architecture was robust to noises in weak labels and learned sufficient features. The accuracy of bone segmentation in the test cases was not assessed using Dice Coefficient, as Dice Coefficient did not directly correlate with the overall TREs. For example, a high Dice Coefficient mask with artifacts (false positive) can produce higher TREs than a low Dice Coefficient mask with only true positive segmentation, as segmentation inaccuracy can be propagated through the pipeline and bias downstream point cloud registration. Therefore, Dice Coefficient was used as a secondary measure. Indeed, the TREs using predicted mask outperformed their counterparts using weak labels (2.0 ± 1.1 mm vs. 2.4 ± 0.9 mm). More importantly, the Sparse-XM TREs outperformed their counterparts with manual segmentation (2.0 ± 1.1 mm vs. 2.5 ± 1.5 mm) likely due to the incorporation of additional depth information, not provided to the user for manual segmentation. Clinically, this 0.5 mm improvement corresponds to 10% of a 5 mm lumbar pedicle screw safety margin [22]. Clinical data acquisition is underway for future validation studies.

5 Conclusion

We have developed and evaluated a fully automated pose adjustment registration framework to compensate for alignment change in open spine surgery. Cross-modality labels were automatically generated for adoption of a deep learning approach, and bone surface was automatically segmented from an intraoperative vision system using both RGB-D information and registered with preoperative CT for each vertebral level. Quantitative error assessments showed an improved accuracy in two test cases compared to counterparts with manual segmentation. By eliminating user dependency and user error, our Sparse-XM approach can be adopted more broadly for clinical applications.

Acknowledgments. Funding from the National Institutes of Health grant R01EB025747-01, R01CA159324-10, S10OD032308-01, and support from Medtronic Navigation (Medtronic, Louisville, Colorado) are acknowledged

Disclosure of Interests. KP is Co-founder of Cairn Surgical, Inc. and Co-Founder & CTO of InSight Surgical Technologies. XF is an employee of InSight Surgical Technologies. SK is founder of PEER Technologies PLLC and Smart Cavity Creator LLC. Each is a respective inventor of patents assigned to the Trustees of Dartmouth College that have been licensed to each corresponding entity.

References

1. Holly, L.T., Foley, K.T.: Image guidance in spine surgery. *Orthop Clin North Am.* 38, 451–461; abstract viii (2007). <https://doi.org/10.1016/j.ocl.2007.04.001>.
2. Holly, L.T., Foley, K.T.: Intraoperative spinal navigation. *Spine (Phila Pa 1976)*. 28, S54–61 (2003). <https://doi.org/10.1097/01.BRS.0000076899.78522.D9>.
3. Malham, G.M., Wells-Quinn, T.: What should my hospital buy next?-Guidelines for the acquisition and application of imaging, navigation, and robotics for spine surgery. *J Spine Surg.* 5, 155–165 (2019). <https://doi.org/10.21037/jss.2019.02.04>.
4. Virk, S., Qureshi, S.: Navigation in minimally invasive spine surgery. *J Spine Surg.* 5, S25–S30 (2019). <https://doi.org/10.21037/jss.2019.04.23>.
5. Mendelsohn, D., Strelzow, J., Dea, N., Ford, N.L., Batke, J., Pennington, A., Yang, K., Ailon, T., Boyd, M., Dvorak, M., Kwon, B., Paquette, S., Fisher, C., Street, J.: Patient and surgeon radiation exposure during spinal instrumentation using intraoperative computed tomography-based navigation. *Spine J.* 16, 343–354 (2016). <https://doi.org/10.1016/j.spinee.2015.11.020>.
6. Scheufler, K.-M., Franke, J., Eckardt, A., Dohmen, H.: Accuracy of image-guided pedicle screw placement using intraoperative computed tomography-based navigation with automated referencing, part I: cervicothoracic spine. *Neurosurgery.* 69, 782–795; discussion 795 (2011). <https://doi.org/10.1227/NEU.0b013e318222ae16>.
7. Gopalakrishnan, V., Dey, N., Golland, P.: Intraoperative 2D/3D Image Registration via Differentiable X-ray Rendering. Presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024).

8. Ma, X., Zhang, X., Wang, Y., Nycz, C., Sungarian, A., Ji, S., Zhang, H.K.: Feasibility of Pointcloud-based Ultrasound-CT Registration towards Automated, Robot-Assisted Image-Guidance in Spine Surgery*. In: 2024 International Symposium on Medical Robotics (ISMR). pp. 1–7 (2024). <https://doi.org/10.1109/ISMR63436.2024.10586071>.
9. Ma, X., Zhang, X., Wang, Y., Nycz, C.J., Sungarian, A., Ji, S., Huang, X., Zhang, H.K.: Cross-Modality Registration using Bone Surface Pointcloud for Robotic Ultrasound-Guided Spine Surgery. *Journal of Medical Robotics Research*. 10, 2540004 (2025). <https://doi.org/10.1142/S2424905X25400045>.
10. Comstock, C.P., Wait, E.: Novel Machine Vision Image Guidance System Significantly Reduces Procedural Time and Radiation Exposure Compared With 2-dimensional Fluoroscopy-based Guidance in Pediatric Deformity Surgery. *J Pediatr Orthop*. 43, e331–e336 (2023). <https://doi.org/10.1097/BPO.0000000000002377>.
11. Abdelrahman, A., Bangash, O.K., Bala, A.: Percutaneous posterior lumbar interbody fusion using optical topographic navigation: Operative technique. *Interdisciplinary Neurosurgery*. 29, 101561 (2022). <https://doi.org/10.1016/j.inat.2022.101561>.
12. Fan, X., Durtschi, M.S., Li, C., Evans, L.T., Ji, S., Mirza, S.K., Paulsen, K.D.: Hand-Held Stereovision System for Image Updating in Open Spine Surgery. *Oper Neurosurg (Hagerstown)*. 19, 461–470 (2020). <https://doi.org/10.1093/ons/opaa057>.
13. Cai, Y., Wu, S., Fan, X., Olson, J., Evans, L., Lollis, S., Mirza, S.K., Paulsen, K.D., Ji, S.: A level-wise spine registration framework to account for large pose changes. *Int J Comput Assist Radiol Surg*. 16, 943–953 (2021). <https://doi.org/10.1007/s11548-021-02395-0>.
14. Liebmann, F., von Atzigen, M., Stütz, D., Wolf, J., Zingg, L., Suter, D., Cavalcanti, N.A., Leoty, L., Esfandiari, H., Snedeker, J.G., Oswald, M.R., Pollefeys, M., Farshad, M., Fürnstahl, P.: Automatic registration with continuous pose updates for marker-less surgical navigation in spine surgery. *Medical Image Analysis*. 91, 103027 (2024). <https://doi.org/10.1016/j.media.2023.103027>.
15. Massalimova, A., Liebmann, F., Jecklin, S., Carrillo, F., Mazda, F., Fürnstahl, P.: SurgPointTransformer: Vertebrae Shape Completion with RGB-D Data, <https://arxiv.org/abs/2410.01443>, (2024).
16. Seetharaman, A., Bhattacharya, I., Chen, L.C., Kunder, C.A., Shao, W., Soerensen, S.J.C., Wang, J.B., Teslovich, N.C., Fan, R.E., Ghanouni, P., Brooks, J.D., Too, K.J., Sonn, G.A., Rusu, M.: Automated detection of aggressive and indolent prostate cancer on magnetic resonance imaging. *Med Phys*. 48, 2960–2972 (2021). <https://doi.org/10.1002/mp.14855>.
17. Xie, S., Tu, Z.: Holistically-Nested Edge Detection, <http://arxiv.org/abs/1504.06375>, (2015). <https://doi.org/10.48550/arXiv.1504.06375>.
18. Warner, W.R., Fan, X., Duke, R.B., Chen, K.L., Stoner, H., Li, C., Wu, S., Ji, S., Mirza, S.K., Paulsen, K.D.: Towards accounting for intraoperative spine motion:

- a simulation study of registration between stereovision surfaces. In: *Medical Imaging 2023: Image-Guided Procedures, Robotic Interventions, and Modeling*. pp. 350–357. SPIE (2023). <https://doi.org/10.1117/12.2654396>.
19. Scholz, M., Fraunholz, M., Selbig, J.: Nonlinear Principal Component Analysis: Neural Network Models and Applications. In: Gorban, A.N., Kégl, B., Wunsch, D.C., and Zinovyev, A.Y. (eds.) *Principal Manifolds for Data Visualization and Dimension Reduction*. pp. 44–67. Springer, Berlin, Heidelberg (2008). https://doi.org/10.1007/978-3-540-73750-6_2.
 20. Liu, C., Yuen, J., Torralba, A.: SIFT flow: dense correspondence across scenes and its applications. *IEEE Trans Pattern Anal Mach Intell.* 33, 978–994 (2011). <https://doi.org/10.1109/TPAMI.2010.147>.
 21. Lien, S.-B., Liou, N.-H., Wu, S.-S.: Analysis of anatomic morphometry of the pedicles and the safe zone for through-pedicle procedures in the thoracic and lumbar spine. *Eur Spine J.* 16, 1215–1222 (2007). <https://doi.org/10.1007/s00586-006-0245-2>.
 22. Chua, M.J., Siddiqui, S., Yu, C.S., Nolan, C.P., Oh, J.Y.-L.: The Optimal Screw Length of Lumbar Pedicle Screws during Minimally Invasive Surgery Fixation: A Computed Tomography-Guided Evaluation of 771 Screws. *Asian Spine J.* 13, 936–941 (2019). <https://doi.org/10.31616/asj.2018.0276>.