HHS Public Access
Author manuscript
Proc SPIE Int Soc Opt Eng. Author manuscript; available in PMC 2016 June 20.

Published in final edited form as:
Proc SPIE Int Soc Opt Eng. 2016 ; 9786: . doi:10.1117/12.2216913.

Automatic Masking for Robust 3D-2D Image Registration in Image-Guided Spine Surgery

M. D. Ketcha[a], **T. De Silva**[a], **A. Uneri**[b], **G. Kleinszig**[d], **S. Vogt**[d], **J.-P. Wolinsky**[c], and **J. H. Siewerdsen**[a,b,*]

[a]Department of Biomedical Engineering, Johns Hopkins University, Baltimore, MD

[b]Department of Computer Science, Johns Hopkins University, Baltimore, MD

[c]Department of Neurosurgery, Johns Hopkins University, Baltimore, MD

[d]Siemens Healthcare XP Division, Erlangen, Germany

Abstract

During spinal neurosurgery, patient-specific information, planning, and annotation such as vertebral labels can be mapped from preoperative 3D CT to intraoperative 2D radiographs via image-based 3D-2D registration. Such registration has been shown to provide a potentially valuable means of decision support in target localization as well as quality assurance of the surgical product. However, robust registration can be challenged by mismatch in image content between the preoperative CT and intraoperative radiographs, arising, for example, from anatomical deformation or the presence of surgical tools within the radiograph. In this work, we develop and evaluate methods for automatically mitigating the effect of content mismatch by leveraging the surgical planning data to assign greater weight to anatomical regions known to be reliable for registration and vital to the surgical task while removing problematic regions that are highly deformable or often occluded by surgical tools. We investigated two approaches to assigning variable weight (i.e., "masking") to image content and/or the similarity metric: (1) masking the preoperative 3D CT ("volumetric masking"); and (2) masking within the 2D similarity metric calculation ("projection masking"). The accuracy of registration was evaluated in terms of projection distance error (PDE) in 61 cases selected from an IRB-approved clinical study. The best performing of the masking techniques was found to reduce the rate of gross failure (PDE > 20 mm) from 11.48% to 5.57% in this challenging retrospective data set. These approaches provided robustness to content mismatch and eliminated distinct failure modes of registration. Such improvement was gained without additional workflow and has motivated incorporation of the masking methods within a system under development for prospective clinical studies.

1. Introduction

Intraoperative imaging plays a vital role in target localization and verification of the surgical product in spine surgery. For example, intraoperative digital radiographs (DR) are commonly acquired in both open and minimally invasive approaches to assist the surgeon in

* jeff.siewerdsen@jhu.edu; phone (443) 287-6269; fax (410) 955-9826; http://istar.jhu.edu.

localization and guidance. Despite such practice, wrong-level surgery occurs at unacceptable frequency, constituting the second most common form of surgical site error and with up to 50% of neurosurgeons claiming (self-reported) wrong-level surgeries[1]. Accordingly, surgeons go to great lengths to avoid such a "never event," including meticulous (manual) level counting and even preoperative tagging of the surgical target under CT guidance - each costing time, expense, and stress. Recent work has advanced a system, called LevelCheck, to map preoperatively annotated vertebral labels from preoperative 3D CT imaging to intraoperative 2D DR via image-based 3D-2D registration[2], providing a potentially valuable means of decision support[3]. In challenging cases, however, robust registration can be confounded due to content mismatch between the CT and the DR, particularly due the presence of extraneous surgical tools and anatomical deformation[4] as illustrated in Figure 1. To overcome the challenges caused by such mismatch, manually delineated masks have been previously applied to the intraoperative DR to constrain the region of interest and exclude surgical tool gradients[2,4]. However, these masks are time consuming, subject to user variability, and complicate workflow by requiring additional user input. In the work reported below, we develop and evaluate an alternative approach that automatically masks the preoperative CT and/or projection domain similarity calculation by leveraging information already defined in preoperative CT in the course of surgical planning. The method is tested in particularly challenging clinical scenarios drawn from an ongoing clinical study.

2. Methods

2.1. 3D-2D Registration Framework

To aid the surgeon during intraoperative localization, vertebral levels identified in the preoperative CT image are projected onto the intraoperative DR via 3D-2D registration. During registration, image similarity between the intraoperative radiograph and a digitally reconstructed radiograph (DRR), formed by forward projecting the preoperative CT image, was optimized in a rigid 6 degree of freedom (DOF) transformation space as described in Otake et al[2].

Assuming an imaging geometry with the detector centered at the origin and the x-ray source positioned fixed at (x_s, y_s, z_s), multi-start covariance matrix adaptation-evolution strategy (CMA-ES)[5] was used to optimize the 6 DOF space of the object position, in this case the preoperative CT positon, within the defined imaging geometry. The 6 DOF space consists of 3 translation (x, y, z) and 3 rotation (η, θ, φ) parameters, embedded in the projective transformation matrix denoted by $T_{3\times4}$.

$$T_{3\times4}=\begin{pmatrix} Z_s & 0 & x_z & 0 \\ 0 & Z_s & y_s & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} & R_{3\times3(\eta,\theta,\phi)} & & \begin{matrix} x-x_s \\ y-y_s \\ z-z_s \end{matrix} \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (1)$$

Simultaneous searches of this space were performed after initializing at 50 multi-start locations, distributed within a range of (±75, ±100, ±200) mm and (±15, ±10, ±10) degrees in the 6 DOF space.

Gradient Orientation (GO), as described in De Silva et al[6], was used to evaluate similarity between the DRR and the intraoperative DR. GO calculation examines the pixel-wise correspondence in gradient direction at high gradientmagnitude regions of the two images. Alignment at each pixel, $w'(i)$, is a function of difference in gradient direction, θ_i in radians, between the two images at that pixel location

$$w'(i) = \frac{2 - \ln(|\theta_i| + 1)}{2} \quad (2)$$

To permit fast registration consistent with intraoperative workflow, DRR generation and similarity metric computations were parallelized on GPU. As pre-processing steps, radiographs were downsampled to a pixel dimension of 2 mm in the x and y directions, a soft-tissue threshold of 150 HU was applied to the CT to remove low-density regions, and a rectangular region was defined on the radiographs to remove areas containing collimation and burnt-in text annotations. Image orientations were initialized with the CT image simply translated in the longitudinal direction of the patient to ensure initial overlap between the radiograph and the DRRs (with registration error following basic initialization ~20-200 mm). In this framework, registration performance due to such initialization errors has been shown to be robust up to ±200 mm in the longitudinal direction, De Silva et al[6].

2.2. Automatic Volumetric Mask Creation

To emphasize the anatomical region of interest in registration, a volumetric mask was automatically created centered on the vertebrae centroid locations already defined in preoperative CT images. Note that such definition can be performed automatically (e.g., using Siemens FastSpine Application[7]) and is a planning step consistent with conventional preoperative workflow. To identify an optimal method for applying masks within our registration framework, we investigated multiple approaches to masking and compared their performance. As the initial step for defining the masks, the vertebrae centroids were connected to form a 3D line skeleton – (Figure 2) that is used to compute two different types of volumetric masks with only one input parameter – the mask width. The first is a 3 dimensional Gaussian (Figure 2) centered along the line – with Gaussian width σ:

$$\text{Gaussian} \quad \text{Mask}\,(\mathrm{i}) = e^{\frac{-\mathrm{dist}(\mathrm{i},\mathscr{L})^2}{2\sigma^2}} \quad (3)$$

The second is a binary (0 or 1) mask (Figure 3e,f) centered on – with radius r:

$$\text{Binary} \quad \text{Mask}\,(\mathrm{i}) = \begin{cases} 1 & \mathrm{dist}\,(\mathrm{i},\mathscr{L}) < r \\ 0 & \mathrm{o/w} \end{cases} \quad (4)$$

These two forms of volumetric mask were used in two different masking implementations, detailed below.

2.2.1. Projection Masking—Projection masking aims to provide greater weight to vertebral regions during the similarity metric computation step of the registration. The similarity map consists of pixel-wise contributions to similarity between the DRR and the intraoperative DR [GO(*x,y*)], prior to summation to compute the overall metric value [GO]. By using the same $T_{3\times4}$ that generated the DRR, the volumetric mask can be forward projected to generate a projection mask applied to the similarity map, giving greater weight to the vertebral region (Figure 2). With the two types of volumetric mask defined above, three variations of projection mask were defined: (1) Gaussian projection masking (GaussPM), in which the projection of the Gaussian mask is used as the weight for the similarity map (Figure 3d); (2) scalar projection masking (ScalarPM), in which the binary mask is projected (Figure 2e); and (3) binary projection masking (BinPM), in which the mask weight is set to 1 if the projector passes through the binary volume (Figure 2f).

2.2.1. Binary Volume Masking—In Binary Volume Masking (BinVM), the mask is applied directly to the CT volume. In this approach, the binary mask is applied to the original CT to emphasize the region that is most relevant to the surgery and eliminate sensitivity to anatomical regions distal to the surgical target, such as the pelvis and ribs, which can impair registration (Figure 4). The mask is applied once in the initial step, such that the registration optimization loop is unaltered.

2.3. Experiments

Analysis was performed under an IRB-approved retrospective study to evaluate registration performance in a clinical data set (24 cases, yielding 24 CT images and 61 radiographs) for individuals undergoing thoracolumbar spine surgery. To focus specifically on challenging cases that tend to confound 3D-2D registration ("failure modes" in previous work[4]), 15 radiographs that exhibited registration failure were identified as "challenging" - and were analyzed separately to examine robustness of the proposed approach and then pooled with the 61 radiographs to ensure that the proposed method did not diminish overall performance within the cohort as a whole.

2.3.1. Mask Performance Comparison—Comparisons among the four masks (BinVM, ScalarPM, BinVM, and GaussPM) with respect to "No Masking" were carried out among the challenging subset of 15 radiographs. For each radiograph and masking technique, registration was repeated at each mask width 10 times (noting small, arguably negligible stochasticity in CMA-ES in the current data). Following this analysis, the performance of each mask (at optimal width) was then verified on the entire dataset (61 radiographs). For the experiments throughout, registration accuracy was evaluated using mean projection distance error (PDE) as outlined in Otake et al.,[2] measuring the mean distance between the projected labels and manually identified ("true") vertebral centroids. Gross failure was defined as PDE >20 mm – a threshold reflecting the approximate distance of half of a thoracolumbar vertebral body (at the detector, assuming a conventional magnification factor of 2). To determine statistical significance, *p*-values were computed under the null hypothesis that the binomial parameter (fraction with PDE >20) for a specified mask scenario is greater than or equal to that of the "No Mask" scenario.

3. Results and Breakthrough Work

3.1. Mask Performance Comparison

Figure 5a shows the performance of the four mask types described above, applied to the challenging subset of 15 radiographs. BinVM with r = 50 mm (denoted BinVM-50) provided the most effective mask and improved registration performance by rectifying more than half of the challenging cases (from 44.13% gross error in the "No Mask" case to 19.33% error with BinVM-50, p-value < 0.001).

Further tests consisted of applying projection masking along with BinVM-50 to determine if a combination of the masking results might yield even better registration. However, none of the projection masks at the tested widths yielded a statistically significant improvement in registration from BinVM-50 alone.

To verify the improvement yielded by BinVM-50, performance of the mask was verified on the entire data set of 61 radiographs. Through application of BinVM-50, registration performance improved from a failure rate of 11.48% with no masking to 5.57% with BinVM-50, p-value < 0.001.

4. Conclusion

The ability to quickly and accurately augment intraoperative radiographs with registered vertebral levels offers a potentially valuable means of decision support against wrong-level surgery, and the methods for automatic masking established in this work were shown to successfully mitigate distinct failure modes that can confound registration. This increased robustness in registration comes with no additional manual steps and significantly aided with reliable registration in challenging cases of strong anatomical deformation and instances of high implant density in the intraoperative scene. Automatic masking, particularly BinVM-50 improved registration accuracy and reduced the failure rate in challenging cases without diminishing performance in the general cohort. These results motivate incorporation of automatic masking in the registration system now in translation to prospective clinical studies.

Acknowledgment

Research supported by NIH Grant No. R01-EB-017226 and collaboration with Siemens XP.

References

[1]. Longo UG, Loppini M, Romeo G, Maffulli N, Denaro V. Errors of level in spinal surgery An evidencebased systematic review. J. Bone Joint Surg. Br. 2012; 94(11):1546–1550. [PubMed: 23109637]

[2]. Otake Y, Wang AS, Stayman JW, Uneri A, Kleinszig G, Vogt S, Khanna AJ, Gokaslan ZL, Siewerdsen JH. Robust 3D–2D image registration: application to spine interventions and vertebral labeling in the presence of anatomical deformation. Phys. Med. Biol. 2013; 58(23): 8535. [PubMed: 24246386]

[3]. De Silva T, Lo S-F, Aygun N, Aghion DM, Boah A, Petteys R, Uneri A, Ketcha MD, Yi T, Vogt S, Kleinszig G, Wei W, Weiton M, Ye X, Bydon A, Sciubba DM, Witham TF, Wolinsky J-P,

Siewerdsen JH. Utility of the LevelCheck algorithm for decision support in vertebral localization. Spine in review. 2016

[4]. Lo S-F, Otake Y, Puvanesarajah V, Wang AS, Uneri A, De Silva T, Vogt S, Kleinszig G, Elder BD, Goodwin CR, Kosztowski TA. Automatic Localization of Target Vertebrae in Spine Surgery: Clinical Evaluation of the LevelCheck Registration Algorithm. Spine. 2015; 40(8):476–483.

[5]. Hansen N, Kern S. Evaluating the CMA evolution strategy on multimodal test functions. PPSN VIII. 2004:282–291.

[6]. De Silva T, Uneri A, Ketcha MD, Reaungamornrat S, Kleinszig G, Vogt S, Aygun N, Lo S-F, Wolinsky J-P, Siewerdsen JH. 3D-2D Image Registration for Target Localization in Spine Surgery: Investigation of Similarity Metrics Providing Robustness to Content Mismatch. Phys. Med. Biol. 2016 In press.

[7]. Scholtz JE, Wichmann JL, Kaup M, Fischer S, Kerl JM, Lehnert T, Vogl TJ, Bauer RW. First performance evaluation of software for automatic segmentation, labeling and reformation of anatomical aligned axial images of the thoracolumbar spine at CT. Eur. J. Radiol. 2015; 84(3): 437–442. [PubMed: 25554009]

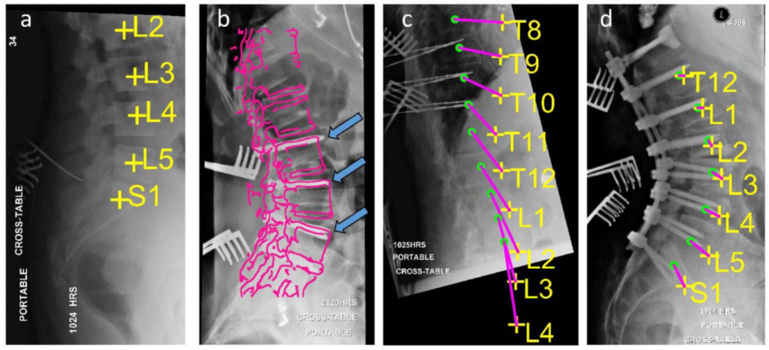

Figure 1.

LevelCheck registration examples. (a) A successful registration with labels correctly placed. (b) An example of a failed registration caused by the presence of surgical instrumentation. The (magenta) gradient overlay of the projection image illustrates that the superior edges of the vertebral bodies erroneously aligned with the edges of surgical tools. (c) Example failed registration due to alignment with non-corresponding anatomy. The pink line segments join the estimated (yellow) and correct (green) label positions and are indicative of PDE. (d) Poor alignment due spinal deformation between the CT and radiograph.

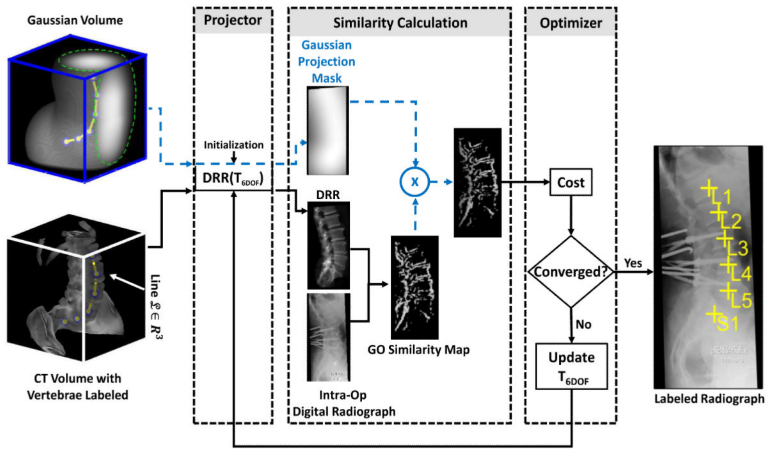

Figure 2.

Flowchart for 3D-2D registration with Gaussian Projection Masking (GaussPM) with modifications to original registration workflow highlighted in blue. Note that line –, demarcating the centroid line through vertebral bodies, can be determined automatically and is used to create the volumetric mask. ScalarPM and BinPM have analogous frameworks.

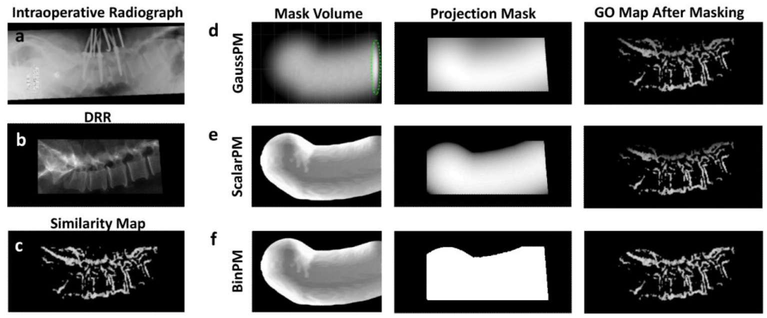

Figure 3.

Projection Masking. (a) Intraoperative radiograph. (b) DRR of the CT at solution. (c) GO similarity map at solution. (e, f, g) Volumetric masks, their corresponding projection masks, and the resulting GO map following (e) GaussPM, (f) ScalarPM, and (g) BinPM.

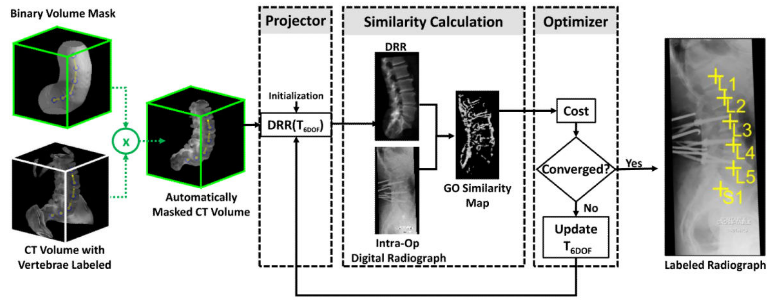

Figure 4.
Flowchart for 3D-2D registration with the Binary Volume Masking (BinVM) with modifications to original registration workflow high-lighted in green.

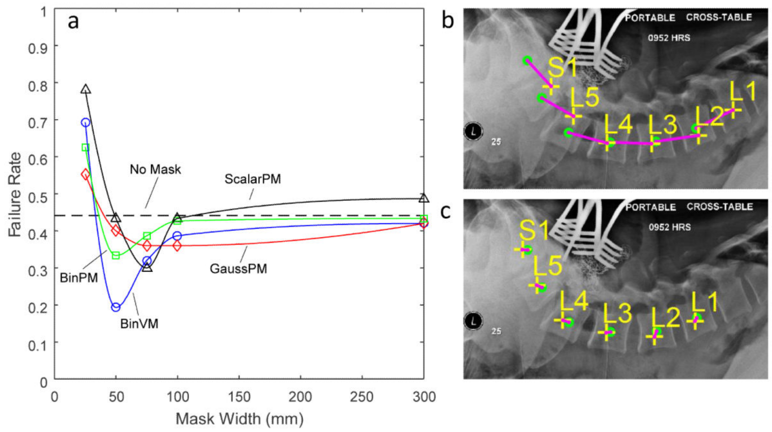

Figure 5.

Evaluation of mask performance in the "challenging" subset of clinical studies that tend to confound registration. (a) Individual masks at various widths. (b) Example of a failure case when no masking technique is applied, with the pink lines joining the estimated (yellow) and correct (green) label positions. (c) The accurate registration result when BinVM-50 is applied.