

Robust Registration of Longitudinal Spine CT

Ben Glocker¹, Darko Zikic², and David R. Haynor³

¹ Biomedical Image Analysis Group, Imperial College London, UK

² Microsoft Research, Cambridge, UK

³ University of Washington, Seattle, WA, USA

Abstract. Accurate and reliable registration of longitudinal spine images is essential for assessment of disease progression and surgical outcome. Implementing a fully automatic and robust registration for clinical use, however, is challenging since standard registration techniques often fail due to poor initial alignment. The main causes of registration failure are the small overlap between scans which focus on different parts of the spine and/or substantial change in shape (e.g. after correction of abnormal curvature) and appearance (e.g. due to surgical implants). To overcome these issues we propose a registration approach which incorporates estimates of vertebrae locations obtained from a learning-based classification method. These location priors are used to initialize the registration and to provide semantic information within the optimization process. Quantitative evaluation on a database of 93 patients with a total of 276 registrations on longitudinal spine CT demonstrate that our registration method significantly reduces the number of failure cases.

1 Introduction

Assessment of disease progression and surgical outcome in the context of spinal pathologies is commonly performed using longitudinal imaging [1]. Clinical applications include but are not limited to correction of abnormal curvature, spinal fusion, treatment of fractures, vertebra disc replacement, and quantification of loss of bone mineral density. In order to detect, analyze and quantify changes between structures imaged at different time points it is essential to establish accurate anatomical correspondences which can be obtained by image registration. In the context of spinal imaging, longitudinal data such as pre- and post-operative scans can differ significantly which poses a major challenge to automatic registration methods. Relatively small overlap between pre- and post-operative data is particularly common in trauma cases. While whole-spine images are often acquired for diagnostic purposes when a patient is admitted to the hospital, more restricted scans with focus on the pathological region (e.g. fractures in the cervical part of the spine) are acquired after treatment. The overlap sometimes covers only a few vertebrae that are visible in both images. But even if the overlap of the visible anatomy is comparable, there are often large variations in spinal shape, for example after treatment of scoliosis, and change in appearance due to surgical implants such as metal screws, rods, and cages. These issues are the main cause of registration failure, which we aim to overcome with this work.

Our main contribution is a robust registration method which significantly reduces the number of failure cases in the context of the difficulties mentioned above. Robustness in the registration procedure is achieved by incorporating prior information about approximate locations of vertebrae. This semantic information is automatically extracted from the images by employing a vertebrae classification method. Our experiments demonstrate that the proposed approach yields accurate registrations in challenging cases where alternative methods fail.

The general idea of incorporating additional information into intensity-based registration is not new. Hybrid registration techniques [2], for example, which combine intensity information with automatically detected landmarks are based on such an approach. Also, registration methods which incorporate segmentation priors [3] or other semantic information such as bounding boxes of anatomical structures [4] fall into the same category. The key in such approaches is the integration of effective, application-specific components which can extract useful information from the images prior to registration. However, we are not aware of works that use such additional information to address the robustness in the context of volumetric spine registration. For 2D-3D spine registration robustness is achieved by using fiducial markers [5] which is not applicable in our scenario. Previous works on volumetric spine registration mainly focus on transformation models such as piecewise multi-rigid approaches [6,7] or aim at improving registration accuracy by highlighting spine specific image features [8] and anatomical structures [9]. However, these methods require good initialization which we aim to provide with the proposed registration approach.

2 Robust Registration Using Vertebrae Location Priors

Registration methods commonly rely on iterative optimization procedures which tolerate only a limited amount of initial misalignment. Reasonably accurate pre-alignment is therefore required in order to allow the registration to converge to an optimal solution. Often, application-specific heuristics can be employed in which initial transformations are automatically determined. For example, aligning the centers of intensity masses works often quite well for inter and intra-subject brain registration.

In the context of longitudinal spine registration, such heuristics are more difficult to find. We will later demonstrate that both center of mass alignment and exhaustive search along the main body axis often fail to provide good initialization. The problem of initialization seems particularly challenging for clinical spine data due to the issues discussed earlier such as small overlap, varying field of view and substantial change in shape and appearance. In addition, the presence of repetitive structures such as vertebrae bodies and ribs adds to the difficulty and imposes many local minima for the optimization. In this context, it should be beneficial to incorporate semantic information extracted from the images prior to registration. Recently, significant progress has been made towards automatic labeling and identification of individual vertebrae [10,11,12] and inter-vertebral discs [13,14]. Using such techniques within a registration approach seems a promising direction which can potentially overcome the issues

of initialization. Following this idea, we introduce a registration approach which integrates a recently proposed learning-based vertebrae classification method [12]. This method has been shown to work reasonably well on pathological data which is essential for our purposes. The classification enables us to estimate prior information on approximate locations of vertebrae. Please note that the prior information does not need to be perfect as long as it can be used to start registration within the tolerance regime of the employed method.

2.1 Vertebrae Classification

Vertebrae localization in [12] is posed as a dense classification problem. A voxel-wise classifier based on randomized decision trees is learned using an annotated database of training images. Annotations in this approach are vertebrae centroids, from which dense label maps are generated and used for training.

At test time, the classifier produces probabilistic estimates for each image point. The estimates correspond to the likelihood of a point being part of a particular vertebra. Formally, the output of the classifier is the posterior distribution $p(v|f(x))$ where $v \in \mathcal{V}$ is a vertebra id and $f(x)$ is a feature vector extracted at image point $x \in \Omega_I$ in a test image I . A post-processing step based on mean shift and an outlier removal strategy [12] preserves the confident vertebrae centroid predictions while reducing the number of false positives.

Our approach does not require perfectly accurate prior information, since registration also relies on the rich intensity information. The uncertainty implicitly encoded in the probabilistic estimates will guide the optimization process to rely on the most confident information. The probability estimates generated by the classification approach can be directly integrated into the registration process, both for initialization as well as to guide the subsequent optimization process. To this end, we evaluate the posterior distribution in the spatial domain of the test image in order to produce dense multi-channel probability maps $P_I(v, x) = p(v|f(x))$ with one channel per vertebra v . In Fig. 1, the thresholded maximum-a-posteriori output of the classifier is shown for two images of the same patient. Even though the output is noisy it contains sufficient semantic information for successful registration.

2.2 Registration with Location Priors

The probabilistic vertebrae priors obtained from the classifier are beneficial for two purposes. First, the inferred approximate locations of vertebrae can be used to determine an initial translational offset that is applied to the moving image prior to the actual registration. This offset is determined as the least squares solution which minimizes the distances between the estimated centroids of vertebrae which are detected in both images. Using this as an initialization works well even for image pairs with very small overlap and different field of views. The initial alignment is thus purely based on semantic information extracted from the images, without the need for any additional application-specific heuristics. Initialization results for one example are shown in Fig. 2.

The second purpose of the location priors is to drive the registration towards accurate alignment. This is achieved by augmenting the intensity-based objective function with a matching criterion defined on the probability maps P . The intensity information is suitable for precise alignment of structure boundaries, but only once those structures sufficiently overlap. The prior probabilities on the other hand help to avoid local minima by solving ambiguities between neighboring structures with locally similar appearance. The combination of intensity information and location priors obtained from a classifier makes the registration process very robust. We cast the registration of two images as the following optimization problem

$$\hat{T} = \arg \max_T \psi(I, J, P_I, P_J, T) \quad , \quad (1)$$

where $T : \mathbb{R}^3 \mapsto \mathbb{R}^3$ is the transformation which aligns the moving image I with the fixed image J . The matching criterion ψ has the following form

$$\psi(I, J, P_I, P_J, T) = \rho(T(I), J) + \frac{1}{|\mathcal{V}_{IJ}|} \sum_{v \in \mathcal{V}_{IJ}} \phi(T(P_I(v, \cdot)), P_J(v, \cdot)) \quad , \quad (2)$$

where P_I and P_J are the probability maps for the moving and fixed image, and \mathcal{V}_{IJ} is the set of vertebrae that have been detected in both images. The criterion ψ is thus simply the sum of an intensity-based matching criterion ρ and another criterion ϕ evaluated on the probability maps. In the simplest case, the two criteria can be defined using one of the popular similarity measures. In this work, we use correlation coefficient both on the intensity information and on the prior probability maps and both terms equally contribute to the objective function. It could be beneficial to introduce a weighting factor that balances the two terms, in particular if other similarity measures are considered.

The actual registration procedure with the objective defined in Eq. (2) follows a common hierarchical setup. We first establish linear alignment with six degrees of freedom defining a rigid-body transformation. Subsequently, we run a non-rigid refinement using free-form deformations with three levels using 80, 40, and 20mm control grid spacing. Both the linear and the non-linear components make use of gradient-free optimization methods. For the linear registration we employ the Downhill-simplex optimizer, and the non-rigid part is optimized using discrete optimization [15]. This has the advantage that the objective function in Eq. (2) can be easily integrated into the existing registration code.

3 Experiments

The data used for experimental evaluation is a subset of our publicly available spine CT database¹. Within this database we identified 93 patients for which at least two longitudinal scans are available. In some cases up to five follow-up scans are present. For each patient, we perform registration between all possible pairs of

¹ <http://research.microsoft.com/spine/>




Fig. 1. Example of a challenging case of longitudinal spine registration where the source and the target image focus on different parts of the spine with only a few common vertebrae. After running our classification algorithm, location priors are obtained which be used to solve the initialization problem. The shown colored labelmaps correspond to the maximum-a-posteriori labels at each image point where each color corresponds to a particular vertebrae

scans if there is at least one common vertebra visible in both images. This results in a total of 276 intra-subject longitudinal registrations. Each dataset comes with manual annotations of vertebrae centroids which allows us to quantify the target registration. To this end we apply the resulting transformations to the manually annotated centroids of the moving image and determine the Euclidean mean distances to the centroids in the fixed image. The vertebrae classifier that is employed in our registration was trained on the remaining data from the original database, for which follow-up scans are not available.

3.1 Pre-processing

All evaluated registration methods operate on clamped intensity images. During initial experiments with different intensity-based methods we found that this yields best registration results when the aim is to establish correspondences between bony structures of the spine. We only consider intensities within an HU range of [100, 1500]. Setting an upper limit has also the effect of suppressing surgical implants and other artifacts caused by metal. This is particularly helpful when registering pre- and post-operative data.

3.2 Baselines

We compare our prior-augmented registration approach to three baselines. Each baseline employs a different strategy for obtaining an initial alignment, followed




Fig. 2. The left image presents initialization results for three different approaches for the moving and fixed image shown in Fig. 1. Green circles are vertebrae centroids of C6 and T3 mapped from the fixed to the moving image space. Red circles are ground truth centroids. Initialization via alignment of ‘Centers of Mass’ fails completely in this case. The ‘1D Exhaustive’ search converges to a local minimum from which subsequent registration fails (not shown). The initialization via location ‘Priors’ yields sufficient pre-alignment for successful registration (not shown). The plot on the right summarizes the alignment errors after initialization for all 276 image pairs grouped with respect to the amount of image overlap, expressed by number of vertebrae contained in both scans. Our prior-based alignment is particularly beneficial when the overlap is small

by a standard registration driven by intensity information only. The first baseline ‘Centers of Mass’ uses the common heuristic of pre-aligning the centers of intensity masses. The second baseline ‘1D Exhaustive’ uses a one-dimensional exhaustive search along the main anatomical axis with a small step size (around 2mm). The idea of this baseline is that it can potentially find reasonable alignments for images with very different a field of view. The third baseline ‘Init with Priors’ uses the proposed location prior-based initialization, but an intensity-only registration subsequently. This allows us to evaluate the impact of using priors for both initialization and within the objective function as it is done in our method. All baselines use exactly the same intensity-based registration components as our method. By using the identical setup for the transformation models and optimization procedures for our method and the baselines, we can isolate the effect of the location priors on the accuracy of the final registration result.

3.3 Results

Fig. 3 summarizes the main quantitative results in terms of error statistics, categorized by the number of overlapping vertebrae. We observe significantly better registration performance for our method compared to the three baselines on all cases. In particular, the difference in performance for the small overlap cases is remarkable. The overall mean registration error after non-rigid refinement on all 276 registration is 12.3mm for our method, 39.8mm for ‘Centers of Mass’, 35.5mm for ‘1D Exhaustive’, and 14.0mm for ‘Init with Priors’. With increasing overlap the registration errors decrease substantially. However, even if more than 10 vertebrae are visible in both images our registration method still outperforms the baselines, including the one that uses prior information for initialization.

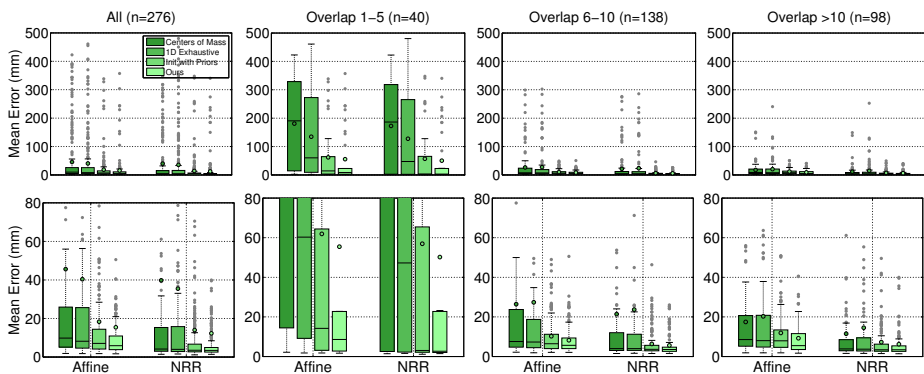


Fig. 3. Statistics of registration errors for two intensity-only baselines, a third baseline that uses prior maps for initialization only, and our registration which uses priors for both initialization and within the objective function. The errors are categorized by the number of overlapping vertebrae. The top row shows the full error range including extreme outliers. The bottom row is a zoom-in on a smaller error range. We observe that our method has significantly lower failure cases, and mean, median and standard deviations of registration errors are consistently lower in particular for the cases with small overlap

This indicates the importance of integrating the location priors also in the objective function. In terms of computational performance, our method is very efficient. Evaluating the classifier to obtain the probability maps takes less than one minute per image. The running time of the registration depends on the size of the images and takes about 2 minutes on average.

4 Conclusion

We demonstrated how volumetric spine registration can benefit from prior information about vertebrae locations yielding a fully automatic, robust registration tool that could be used in clinical routine on challenging data. Our approach immediately benefits from future improvements on learning-based vertebrae classification which could further reduce the number of failure cases. Establishing correspondences between longitudinal data for a large cohort of patients allows disease-specific modeling and a better understanding of underlying pathophysiological processes. Statistical analysis on such data could potentially be used to predict surgical outcome, which seems an interesting direction of research for future work.

References

1. Jarvik, J.J., Hollingworth, W., Heagerty, P., Haynor, D.R., Deyo, R.A.: The longitudinal assessment of imaging and disability of the back (LAIDBack) study: baseline data. *Spine* 26(10), 1158–1166 (2001)

2. Johnson, H.J., Christensen, G.E.: Consistent landmark and intensity-based image registration. *TMI* 21(5), 450–461 (2002)
3. Studholme, C., Hill, D., Hawkes, D.: Incorporating connected region labelling into automated image registration using mutual information. In: *Mathematical Methods in Biomedical Image Analysis*, pp. 23–31 (1996)
4. Konukoglu, E., Criminisi, A., Pathak, S., Robertson, D., White, S., Haynor, D., Siddiqui, K.: Robust linear registration of CT images using random regression forests. In: *SPIE Medical Imaging* (2011)
5. Russakoff, D.B., Rohlfing, T., Adler Jr., J.R., Maurer Jr., C.R.: Intensity-based 2D-3D spine image registration incorporating a single fiducial marker. *Academic Radiology* 12(1), 37–50 (2005)
6. Hu, Y., Haynor, D.R.: Multirigid registration of MR and CT images of the cervical spine. In: *SPIE Medical Imaging* (2004)
7. Cech, P., Andronache, A., Wang, L., Szekely, G., Cattin, P.: Piecewise rigid multi-modal spine registration. In: Handels, H., Ehrhardt, J., Horsch, A., Meinzer, H.P., Tolxdorff, T. (eds.) *Bildverarbeitung fuer die Medizin 2006*, pp. 211–215. Springer, Heidelberg (2006)
8. Li, W., Sode, M., Saeed, I., Lang, T.: Automated registration of hip and spine for longitudinal QCT studies: integration with 3D densitometric and structural analysis. *Bone* 38(2), 273–279 (2006)
9. Steger, S., Wesarg, S.: Automated Skeleton Based Multi-modal Deformable Registration of Head&Neck Datasets. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) *MICCAI 2012, Part II*. LNCS, vol. 7511, pp. 66–73. Springer, Heidelberg (2012)
10. Klinder, T., Ostermann, J., Ehm, M., Franz, A., Kneser, R., Lorenz, C.: Automated model-based vertebra detection, identification, and segmentation in CT images. *MedIA* 13(3), 471–482 (2009)
11. Huang, S.H., Chu, Y.H., Lai, S.H., Novak, C.L.: Learning-based vertebra detection and iterative normalized-cut segmentation for spinal MRI. *TMI* 28(10), 1595–1605 (2009)
12. Glocker, B., Zikic, D., Konukoglu, E., Haynor, D., Criminisi, A.: Vertebrae Localization in Pathological Spine CT via Dense Classification from Sparse Annotations. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) *MICCAI 2013, Part II*. LNCS, vol. 8150, pp. 262–270. Springer, Heidelberg (2013)
13. Schmidt, S., Kappes, J., Bergtholdt, M., Pekar, V., Dries, S., Bystrov, D., Schnörr, C.: Spine detection and labeling using a parts-based graphical model. In: Karssemeijer, N., Lelieveldt, B. (eds.) *IPMI 2007*. LNCS, vol. 4584, pp. 122–133. Springer, Heidelberg (2007)
14. Oktay, A., Akgul, Y.: Localization of the Lumbar discs using machine learning and exact probabilistic inference. In: Fichtinger, G., Martel, A., Peters, T. (eds.) *MICCAI 2011, Part III*. LNCS, vol. 6893, pp. 158–165. Springer, Heidelberg (2011)
15. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through MRFs and efficient linear programming. *MedIA* 12(6), 731–741 (2008)