Learning an airway atlas from lung CT using semantic inter-patient deformable registration

Abstract. Pulmonary image analysis for diagnostic and interventions often relies on a canonical geometric representation of lung anatomy across a patient cohort. Bronchoscopy can benefit from simulating an appearance atlas of airway cross-sections, intra-patient deformable image registration could be initialised using a shared lung atlas. The diagnosis of pneumonia, COPD and other respiratory diseases can benefit from a well defined anatomical reference space. Previous work to create lung atlases either relied on tedious and often ambiguous manual landmark correspondences and/or image features to perform deformable inter-patient registration. In this work, we overcome these limitations by guiding the registration with semantic airway features that can be obtained straightforwardly with an nnUNet and dilated training labels. We demonstrate that accurate and robust registration results across patients can be achieved in few seconds leading to high agreement of small airways of later generations. Incorporating the semantic cost function improves segmentation overlap and landmark accuracy.

1 Introduction

Creating lung atlases from inter-subject registration of CT scans has already been studied nearly two decades ago [1]. Yet this work relied on manual one-to-one correspondences that are difficult to obtain and might be ambiguous with respect to the topology of the airway tree. Newer work [2] established an elaborated inter-subject registration pipeline that aims to tackle the challenges of large deformations across anatomical variation based on keypoint alignment and demonstrated discrimination capabilities for COPD (chronic obstructive pulmonary disease) and calcifications.

Establishing airway atlases can be beneficial for large-scale analysis of anatomical topologies like tree structures, that have been shown to be statistically applicable in medical diagnosis [3], and may be applied to further use cases like bronchoscopy guidance.

Deformable inter-patient registration is widely used in medical image analysis to define dense correspondences across different subjects. Recent work on fast 3D medical registration [4] has demonstrated the benefits of directly incorporating deep learning based segmentation features into conventional (discrete) optimisation algorithms. For inter-patient abdominal CT alignment the best supervised deep-learning registration approach [5] was matched in [4] by decoupling semantic feature extraction and optimisation.

In this work we propose a modular technique that is easy to reproduce and yields promising results for creating an airway atlas based on fast large-deformation estimation and learned binary airway features. **Our contribution** is three-fold: First, we adapt the loss for the well-established nnUNet framework [6] using dilated airway labels to improve accuracy for smaller structures. Second, we combine these airway features with handcrafted image features and incorporate them into a GPU-accelerated registration framework and evaluate the accuracy on unseen images using both multi-label structural overlap and landmark errors. Third, we demonstrate that airway keypoints are also beneficial within a purely geometric registration approach.

2 Materials and Methods

Our method aims to incorporate semantic information for inter-patient lung registration to yield an improved atlas. We obtain semantic information using a segmentation network and employ two different frameworks for inter-patient registration. Fig. 1 depicts a schematic representation of our method.

2.1 Dataset and Segmentation

All experiments are performed using a public dataset containing 40 thoracic CT scans from the LIDC-IDRI Dataset [7] as well as 20 scans from the EXACT'09 Challenge Dataset [8]. Annotations for these cases have been provided by [9]. A total of 31 (primary, secondary and tertiary) individual bronchi as well as the trachea are distinguished.

We train a full-resolution 3D U-Net architecture in the nnUNet framework [6] to segment the airway divided into singular bronchi (32 unique labels). To improve the performance on smaller bronchi, we dilate segmentations by 1 mm. Prior to training, images have been resampled to isotropic resolution and cropped around the lung area. In regards to data augmentation, mirroring has been disabled to provide the network with further contextual information to distinguish left and right bronchi. The numerical evaluation of the segmentation accuracy is given in Section 3.

2.2 Inter-patient registration

The quality of building an anatomical CT atlas mainly depends on accurate inter-patient registration. Previous methods for lung atlases have been restricted to either using manual correspondences or solely relying on intensity-based similarity metrics. We incorporate the semantic information of the automatic airway segmentation in two different ways. 1) We sample 64 keypoints from each predicted label class via farthest point sampling, resulting in a total of 2048 keypoints inside the predicted airway. 2) We use one-hot encodings of the predicted labels on unseen test images as structural airway features.

We combine structural airway information with two frameworks for deformable image registration that are able to decouple feature learning and fast optimisation.

The first registration approach we employ is sparse loopy belief propagation (sLBP) [10], which is a geometric method using sampled keypoints instead of a dense voxel grid to calculate displacement fields. It incorporates a large displacement correspondence search that uses the input features (defined for each node) to calculate the cost term for keypoint displacements. The graph-based inference algorithm, loopy belief propagation, is used to obtain smooth displacements by enforcing pairwise constraints. The global solution for all keypoints is jointly found using a message passing scheme and discrete optimisation. Keypoint displacements are then extrapolated to obtain a dense displacement field. By using the proposed sampling strategy, the algorithm focusses directly on the geometrically relevant regions of the airway tree. Besides airway keypoints, we may additionally use Förstner keypoints inside of a lung mask as proposed in [10]. Moreover, we can make use of MIND-SSC features [11] (that work well in inhale-exhale lung registration) and semantic airway features (based on nnUNet predictions) to determine the similarity-based cost function for each keypoint displacement.

The second registration approach uses a regular grid and hence cannot directly benefit from the airway tree geometry for the transformation model. It combines convex optimisation of discrete displacements [12] with Adam optimisation and performed second best in the 2021 Learn2Reg challenge[1]. Possible displacements for each voxel are limited within a quantised set of displacements, for which a cost term is defined using input features for each voxel. The resulting cost minimisation problem is globally solved using coupled convex optimisation. A subsequent refinement step is realized by gradient descent using the well-known Adam optimiser. We calculate the cost function for possible displacements using MIND-SSC features as well as our proposed semantic airway features.

3 Results

All cases (individual patients) have been randomly split into 40 training and 20 test cases. Segmentation results on the unseen test cases are presented in Tab. 1. Segmenting the airway divided into multiple bronchi labels and subsequent binarisation provides comparable results to directly learning a binary segmentation network (Dice: 94.6%).

Tab. 1. Dice score coefficients (and standard deviations across cases) for multilabel nnUNet segmentation, averaged over bronchi of the same depth.

	Trachea	Primary	Secondary	Tertiary	Average	Binary
Dice (%)	94.6 (0.03)	92.9 (0.04)	83.1 (0.13)	73.0 (0.21)	78.0 (0.04)	94.2 (0.02)

To evaluate our contribution, we compared the registration quality in several different setups. We explore the benefits of using the proposed airway derived keypoints compared

[1]https://github.com/multimodallearning/convexAdam/

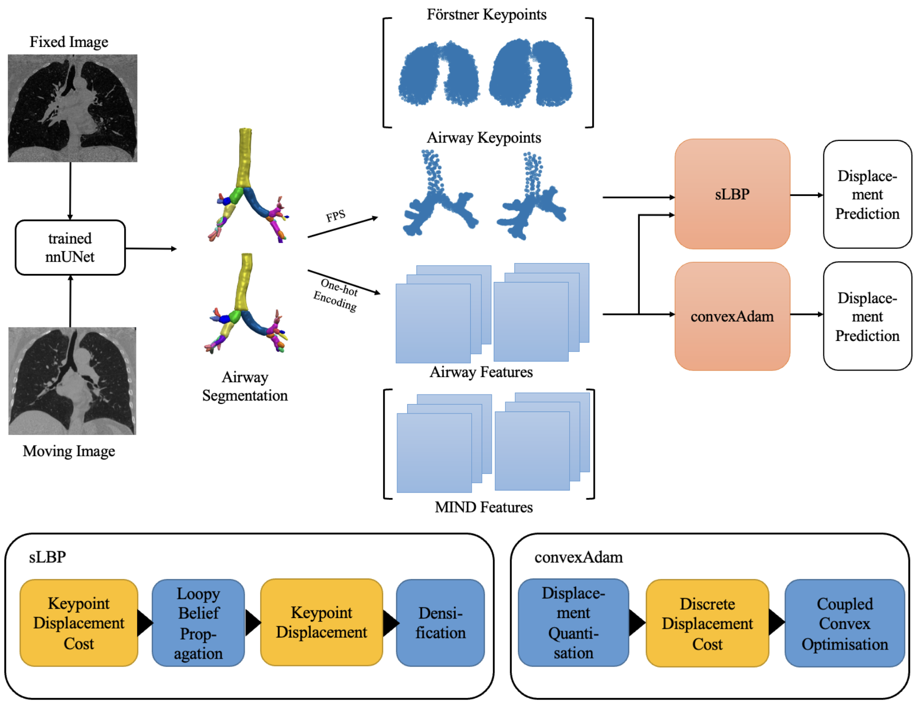

Fig. 1. Out methods consists of extraction of semantic airway information and subsequent image registration. We predict nnUNet airway segmentations of fixed and moving images, which are then used to generate airway keypoints and airway features. In addition to airway keypoints and features we use Förstner keypoints as well as MIND features directly derived from the grey value images. These keypoints and features are applied in registration using a geometric (sLBP) and a dense registration approach (convexAdam). These methods differ regarding how possible displacements are determined as well as optimisation of image similarity and displacement smoothness.

to more generic interest points (Förstner), and we evaluate the impact of including the semantic nnUNet features in comparison to the more general handcrafted features (MIND). Each registration method has been evaluated on a total of 60 pairs of test images (inter-patient). In addition to evaluating the overlap of airway structures, we algorithmically located 34 anatomical landmarks for each scan based on the ground truth airway annotations, each corresponding to the position of a bi- or trifurcation in the airway tree. Note, that we are not using these landmarks to drive the registration (as done in [1]), because it would bias our analysis. We only use these landmarks to calculate target registration error (TRE) in the airway area. Results are presented in Tab. 2.

For geometric registration, the incorporation of airway keypoints has a very positive influence on registration accuracy in the corresponding area. Using semantic features as cost function improves performance even further. Combining MIND-SSC with semantic airway features yields the best results for all employed methods. While dense registration

Tab. 2. Average performance of the proposed methods for 60 inter-patient registrations. Dice score is calculated using dilated ground truth airway segmentations and averaged over all 32 labels. Target registration error (TRE) is averaged over 34 landmarks at bifuractions of airways.

	TRE (px) ↓			Dice (%) ↑		
Initial	17.25			5.7		
Sampling	MIND	Label	MIND+Label	MIND	Label	MIND+Label
Förstner keypoints	11.13	20.00	9.77	13.3	23.1	27.3
Airway keypoints	8.20	6.47	**6.25**	40.8	45.5	50.4
Airway+Förstner keypoints	13.30	23.05	7.02	15.2	2.9	51.0
Dense grid (convexAdam)	9.15	13.27	6.92	37.7	56.1	**62.6**

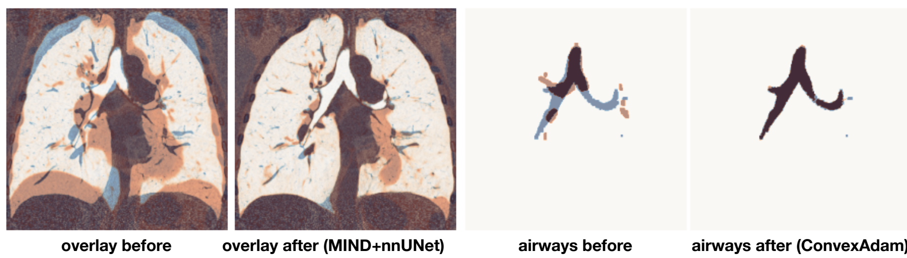

Fig. 2. Visual result of proposed semantic inter-patient registration of lung CT (coronal slice of 3D). The overlay shows fixed in blue and warped/moving scan in orange colours (good alignment appears greyish). The alignment of airways before and after registration confirms the good registration quality of our numerical results.

performs best with regard to overall overlap of the semantic segmentations, sLBP yields better results for the target registration error of specific anatomical bifurcation landmarks. Dense registration of the airway can also be improved with only binary airway features in addition to MIND features (TRE: 7.62px, Dice: 51.4%), though using multilabel features yields better results. Each registration runs in around 1 second at test time on an Nvidia RTX GPU. A visual result of convexAdam registration with MIND and airway features is presented in Fig. 2.

4 Discussion

We performed inter-patient lung registration incorporating learned semantic airway information. By decoupling feature learning from fast optimisation of the registration cost function, we obtain state-of-the-art performance without complicated multi-level or cascaded registration architectures. Our results show that high-quality automatic airway segmentation (up to tertiary bronchi) with a simple adaptation (dilation of ground truth segmentations) works well using the 3D nnUNet.

Registration results demonstrate the importance of combining learned semantic information with intensity-based image features for robust and accurate registration. Clear improvement for semantic sampling of keypoints in graph-based registration

(+3.5px better TRE, +23% points better Dice) as well as large improvement for using airway features in dense registration (+2.2px better TRE, +25% points better Dice) compared to current state-of-the-art in CT lung registration has been made. We have also shown that large improvements are possible when using semantic features derived from binary labels. This can substantially ease the manual annotation task and make it less dependent on the distribution of certain topological variants of airway trees in the training population.

Geometric registration performance surpassing dense image registration in regards to landmark error suggests further employing point cloud-based registration approaches for regular image registration tasks and confirms the relevance of keypoint information for generating anatomical atlases.

Inter-patient registration incorporating learned airway features depicts a first step in generating an airway atlas. In order to create a shared lung atlas further work, especially regarding statistical analysis and modelling, is required.

References

1. Li B, Christensen GE, Hoffman EA, McLennan G, Reinhardt JM. Establishing a normative atlas of the human lung: intersubject warping and registration of volumetric CT images. Academic radiology 10.3 (2003), pp. 255–265.
2. Xu K, Gao R, Khan MS, Bao S, Tang Y, Deppen SA et al. Development and characterization of a chest CT atlas. SPIE. Vol. 2021. NIH Public Access. 2021.
3. Feragen A, Owen M, Petersen J, Wille MM, Thomsen LH, Dirksen A et al. Tree-space statistics and approximations for large-scale analysis of anatomical trees. International Conference on Information Processing in Medical Imaging. Springer. 2013, pp. 74–85.
4. Hansen L, Heinrich MP. Revisiting Iterative Highly Efficient Optimisation Schemes in Medical Image Registration. MICCAI. Springer. 2021, pp. 203–212.
5. Mok TC, Chung AC. Large deformation diffeomorphic image registration with Laplacian pyramid networks. MICCAI. Springer. 2020, pp. 211–221.
6. Isensee F, Jaeger PF, Kohl SA, Petersen J, Maier-Hein KH. nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods 18.2 (2021), pp. 203–211.
7. Armato III SG, McLennan G, Bidaut L, McNitt-Gray MF, Meyer CR, Reeves AP et al. The lung image database consortium (LIDC) and image database resource initiative (IDRI): a completed reference database of lung nodules on CT scans. Medical physics 38.2 (2011), pp. 915–931.
8. Lo P, Van Ginneken B, Reinhardt JM, Yavarna T, De Jong PA, Irving B et al. Extraction of airways from CT (EXACT'09). IEEE T-MI 31.11 (2012), pp. 2093–2107.
9. Tan Z, Feng J, Zhou J. SGNet: Structure-Aware Graph-Based Network for Airway Semantic Segmentation. MICCAI. Springer. 2021, pp. 153–163.
10. Hansen L, Heinrich MP. Deep learning based geometric registration for medical images: How accurate can we get without visual features? IPMI. Springer. 2021, pp. 18–30.
11. Heinrich MP, Jenkinson M, Papież BW, Brady M, Schnabel JA. Towards realtime multimodal fusion for image-guided interventions using self-similarities. MICCAI. Springer. 2013, pp. 187–194.
12. Heinrich MP, Papież BW, Schnabel JA, Handels H. Non-parametric discrete registration with convex optimisation. WBIR. Springer. 2014, pp. 51–61.