Automated Integer Programming Based Separation of Arteries and Veins from Thoracic CT Images

Article *in* Medical image analysis · May 2016

DOI: 10.1016/j.media.2016.05.002

9 authors, including:

Michael Pienn
Ludwig Boltzmann Institute for Lung Vascula…
19 PUBLICATIONS **71** CITATIONS
SEE PROFILE

Zoltán Bálint
Babeş-Bolyai University
61 PUBLICATIONS **564** CITATIONS
SEE PROFILE

Andrea Olschewski
Ludwig Boltzmann Institute for Lung Vascula…
170 PUBLICATIONS **2,127** CITATIONS
SEE PROFILE

Martin Urschler
Ludwig Boltzmann Institute for Clinical-Fore…
87 PUBLICATIONS **555** CITATIONS
SEE PROFILE

Automated Integer Programming Based Separation of Arteries and Veins from Thoracic CT Images

Christian Payer[a,b], Michael Pienn[b], Zoltán Bálint[b], Alexander Shekhovtsov[a], Emina Talakic[c], Eszter Nagy[d], Andrea Olschewski[b,e], Horst Olschewski[b,f], Martin Urschler[a,g,h]

[a]*Institute for Computer Graphics and Vision, Graz University of Technology, Austria*
[b]*Ludwig Boltzmann Institute for Lung Vascular Research, Graz, Austria*
[c]*Division of General Radiology, Department of Radiology, Medical University of Graz, Austria*
[d]*Division of Pediatric Radiology, Department of Radiology, Medical University of Graz, Austria*
[e]*Experimental Anesthesiology, Department of Anesthesia and Intensive Care Medicine, Medical University of Graz, Austria*
[f]*Division of Pulmonology, Department of Internal Medicine, Medical University of Graz, Austria*
[g]*Ludwig Boltzmann Institute for Clinical Forensic Imaging, Graz, Austria*
[h]*BioTechMed Graz, Austria*

Abstract

Automated computer-aided analysis of lung vessels has shown to yield promising results for non-invasive diagnosis of lung diseases. To detect vascular changes which affect pulmonary arteries and veins differently, both compartments need to be identified. We present a novel, fully automatic method that separates arteries and veins in thoracic computed tomography images, by combining local as well as global properties of pulmonary vessels. We split the problem into two parts: the extraction of multiple distinct vessel subtrees, and their subsequent labeling into arteries and veins. Subtree extraction is performed with an integer program (IP), based on local vessel geometry. As naively solving this IP is time-consuming, we show how to drastically reduce computational effort by reformulating it as a Markov Random Field. Afterwards, each subtree is labeled as either arterial or venous by a second IP, using two anatomical properties of pulmonary vessels: the uniform distribution of arteries and veins, and the parallel configuration and close proximity of arteries and bronchi. We evaluate algorithm performance by comparing the results with 25 voxel-based manual reference segmentations. On this dataset, we show good performance of the subtree extraction, consisting of very few non-vascular structures (median value: 0.9%) and merged subtrees (median value: 0.6%). The resulting separation of arteries and veins achieves a median voxel-based overlap of 96.3% with the manual reference segmentations, outperforming a state-of-the-art interactive method. In conclusion, our novel approach provides an opportunity to become an integral part of computer aided pulmonary diagnosis, where artery/vein separation is important.

Keywords: vascular tree reconstruction, artery-vein separation, computed tomography, lung, integer program

1. Introduction

The recent trend towards quantitative pulmonary computer-aided diagnosis (CAD) was enabled by rapid progress in medical imaging modalities. State-of-the-art computed tomography (CT) scanners allow depicting pulmonary anatomical structures like lung vasculature and bronchi on a sub-millimeter level at low radiation doses. Despite not being able to directly visualize the complex capillary network for studying gas exchange, *i.e.* oxygen uptake and carbon dioxide release, this remarkable resolution can help in macroscopic detection of early stage pulmonary nodules (Agam et al., 2005; Murphy et al., 2009) or pulmonary emboli (Masutani et al., 2002), and in CAD of diseases like chronic obstructive pulmonary disease (COPD) (Matsuoka et al., 2010a; Estépar et al., 2013), chronic thromboembolic pulmonary hypertension (CTEPH) (McNeil and Dunning, 2007; Sugiura et al., 2013) or pulmonary arterial/venous hypertension (PAH/PVH) (Linguraru et al., 2010; Helmberger et al., 2014). Quantitative pulmonary CAD requires the analysis of the lung vascular tree. However, the drawback of the high resolution achievable with CT scans is the large amount of generated data, which overwhelms physicians when performing manual analysis. This explains the increased demand for fully automatic extraction of pulmonary tree structures (van Rikxoort and van Ginneken, 2013) and un-

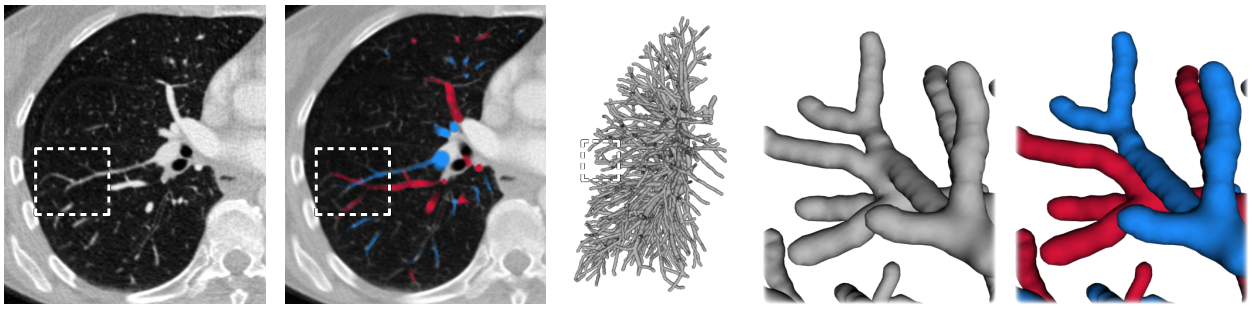

(a) CT slice (b) Labeled CT slice (c) Rendering of a vessel segmentation (d) Zoomed in vessel segmentation (e) Zoomed in A/V segmentation

Figure 1: Thoracic CT image example, where arteries (blue) and veins (red) are close to each other. The highlighted vessel branches in (a) appear as being directly connected, whereas the A/V separation in (b) shows that the branches belong to two distinct vessels. The 3D rendering of the vessel segmentation (c) and its zoom-in (d) visualize these as merged branches in 3D, whereas our proposed A/V separation (e) can tell them apart.

derlines the importance of automated quantitative CAD in the lung (Sluimer et al., 2006). Medical image analysis has therefore produced a large body of research work focusing on automatic enhancement, extraction and segmentation of vascular tree structures from volumetric imaging modality data (Lesage et al., 2009).

While pulmonary vascular tree segmentation has significantly improved and finally matured over the recent years (Rudyanto et al., 2014), a still largely unsolved problem is the fully automated separation of the pulmonary vascular tree into arteries and veins (A/V separation). The main difficulty of this labeling problem is that it is not possible to distinguish arteries and veins based on their image intensity alone. Even in contrast-enhanced scans, the timing for image acquisition can not be done with sufficient accuracy to distinguish arteries and veins robustly. Independent analysis of the arterial and venous pulmonary trees shows high clinical relevance in improving the diagnosis of lung diseases affecting both trees differently. A/V separation allows better understanding of pulmonary structure and function, *e.g.* for lung surgery planning (Stoecker et al., 2013), and improves assessment of pulmonary emboli (Zhou et al., 2007) and CTEPH (where the pressure increase solely occurs in pulmonary arteries). Further, it enables more localized studies of vascular alterations, which is closely related to endothelial dysfunction (Santos et al., 2003) and claimed to be related to emphysema (Kasahara et al., 2001; Matsuoka et al., 2010a,b). Studying the morphology of vascular arterial and venous trees may also yield a macroscopic biomarker for the distinction between PAH and PVH. Here, a non-invasive biomarker derived from thoracic CT would be beneficial, since the current diagnostic gold standard is invasive right-heart catheterization (Hoeper et al., 2013) and, depending on the differential diagnosis of PAH and PVH, different therapeutic decisions are made.

Automatic A/V separation is a complex problem, due to the similar intensity values of both vessel trees in CT. Although pulmonary arteries and veins are connected at the capillary level and thus form complex graph structures anatomically, the small capillaries are not visible in CT images. As a result, the vessel graphs appear as multiple disconnected trees, which are intertwined and run in parallel or in close proximity to each other, making their distinction even harder due to the partial volume effect. This problem is illustrated in Fig. 1, where distinction of a branching structure from a close-by running different structure underlines the challenges. Most of the few existing A/V separation algorithms (Park et al., 2006; Saha et al., 2010) start from a vessel segmentation, often using tubularity filters combined with region growing or fast marching methods based on seed points (van Rikxoort and van Ginneken, 2013). Many algorithms require some interactive step to manually correct mislabeled structures (Gao et al., 2012; Park et al., 2013).

In this work we present a novel, fully automatic A/V separation algorithm for thoracic CT images, which requires no manual correction and takes both global and local structural information about vascular trees into account. It is inspired from recently published methods analyzing cerebral vascular structures (Robben et al., 2014; Rempfler et al., 2015). Bottom-up information about the likelihood for being a vessel and the local orientation and proximity to bronchi are taken into account, combined with top-down constraints encouraging a tree structure of vessels with decreasing diameter and the assumption that arterial and venous trees generally have to be of similar size and uniformly dis-

tributed in a given volume since they meet at the capillary level (Tawhai et al., 2011). Based on a lung segmentation and a vessel enhancement step, we formulate the A/V separation problem as two sub-problems. First, the extraction of unlabeled, distinct subtrees from the vessel candidates inside the lung, and second, the labeling of each subtree as either artery or vein. Each subproblem is formulated as an integer program and solved using discrete optimization methods. We evaluate our method on a database of 25 thoracic CT images with manually segmented and labeled A/V trees as a reference and demonstrate the benefits of our method compared to a related method from (Park et al., 2013). The method described in this paper extends a recent conference paper (Payer et al., 2015) by giving a more detailed technical description, proposing a dedicated, more efficient solver for the time-consuming subtree extraction, and finally presenting a more comprehensive experimental evaluation.

The paper is organized as follows: Section 2 places our work in the context of the current related work. In Sec. 3 we explain in detail our novel A/V separation approach which is formalized more compactly and solved more efficiently compared to (Payer et al., 2015). Section 4 defines our experimental setup and the dataset we are evaluating on, while Sec. 5 describes the results of our quantitative evaluation. Finally, we show qualitative results and discuss our findings as well as a few failure cases in Sec. 6, followed by our conclusions in Sec. 7.

2. Related Work

Extraction and segmentation of vascular structures has gained much interest in the past, starting with the seminal works of (Sato et al., 1998; Frangi et al., 1999; Krissian et al., 2000) on model-based detection using the assumption that vessels locally show a tube-like shape. Such methods are based on the eigenanalysis of the Hessian matrix of the image intensity and are referred to as tubularity or vesselness filters. Reviews on this topic can be found in (Kirbas and Quek, 2004; Lesage et al., 2009), while (Rudyanto et al., 2014) presents results of the recent MICCAI Grand Challenge on vessel segmentation (VESSEL12). Another method that enhances tubular structures is from (Law and Chung, 2008), where image regions are enhanced by projecting the image gradient along different axes and measuring the amount of the projected gradient flowing in or out of a local spherical region. Afterwards, an eigenanalysis of this oriented flux determines the direction of the tube, and a multi-scale analysis allows adaptation to different sizes of vessel structures.

A/V separation methods are built on top of vessel enhancement or segmentation approaches. While automated A/V separation of vessels in 2D color images of the retina has recently shown very robust performance (Joshi et al., 2014), this work is not readily extendable to the domain of volumetric 3D images. Early methods dealing with A/V separation were designed for magnetic resonance angiography (MRA) data and relied solely on slight intensity variations or the presence of local gradient information. Applications were restricted to the systemic arteries or the first few generations of the pulmonary vascular tree and user interaction was often required (Lei et al., 2001; Tizon and Smedby, 2002; van Bemmel et al., 2003). Most 3D A/V separation methods for pulmonary CT data use structural information, either for tracking vessels starting at given seed points (Park et al., 2006), or for a voxel-wise distinction of arteries and veins (Saha et al., 2010; Gao et al., 2012; Park et al., 2013; Kitamura et al., 2013), optionally incorporating different anatomical features like proximity of arteries and bronchi (Bülow et al., 2005; Jin et al., 2014). The method of (Park et al., 2013) involves global structural information by constructing a minimum-spanning-tree from weights derived from local vessel geometry measures and an edge cutting step for A/V separation. However, an interactive refinement is necessary to finalize the separation. In (Saha et al., 2010) A/V separation utilizes the close proximity of arteries and veins. By morphological operations with differently sized kernels, equal intensity structures are split and locally separated regions are traced. Gao et al. (2012) extended this method with a GUI enabling efficient refinement. Another promising method is presented in (Kitamura et al., 2013), who formulate an automatic voxel labeling problem based on root detection for both trees. However, it requires a training step and, due to locally restricted image features, it still has problems near the hilum of the lung, *i.e.* where arteries and veins are in close proximity.

Related methods on A/V separation either require contrast-enhancement (Kitamura et al., 2013), involve a significant amount of manual interaction and refinement (Gao et al., 2012; Park et al., 2013) or include a learning step based on tediously annotated data (Kitamura et al., 2013). In our work we explore how well a purely image processing based approach for A/V separation performs, based on the assumption that local and global structural information is used comprehensively in terms of a discrete optimization problem formulation.

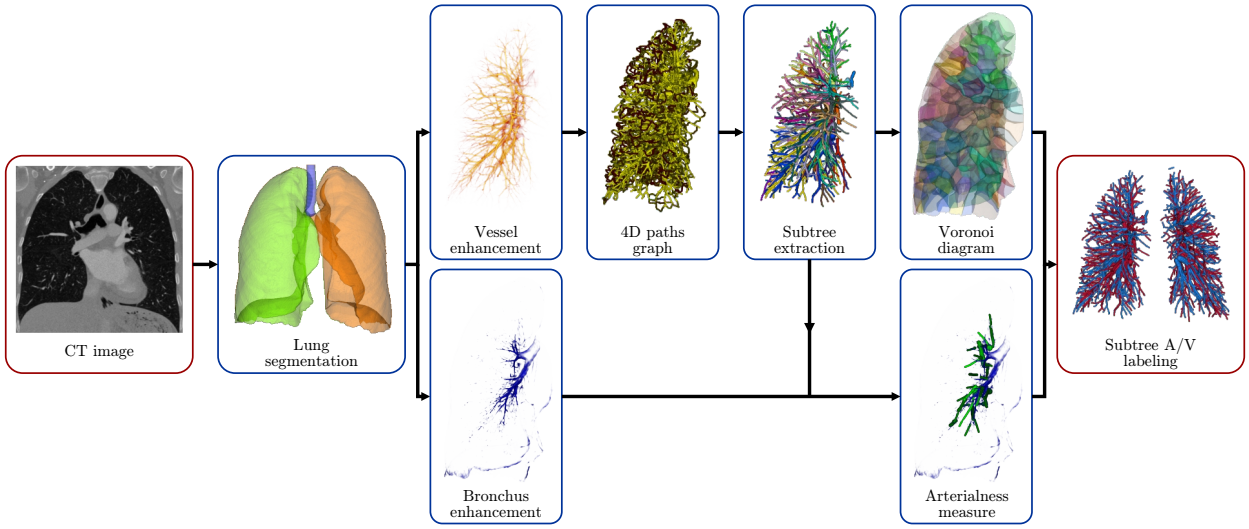

Figure 2: Overview of our proposed A/V separation algorithm. Applied to each segmented lung individually, vessel and bronchus enhancement steps deliver the bottom-up features for both subtree extraction and A/V separation. Top-down knowledge about the geometry of tubular structure paths and the distribution of arteries and veins encoded as a Voronoi diagram are incorporated to increase robustness of the final A/V labeling.

3. Method

Our proposed algorithm for A/V separation from thoracic CT images is shown in Fig. 2. It starts with generating a lung segmentation, which is based on (Helmberger et al., 2014). All subsequent steps are performed for left and right lungs independently.

For later vessel extraction, vascular structures are enhanced with a multi-scale tubularity filter (Law and Chung, 2008), *i.e.* the optimally oriented flux (OOF). In order to generate vessel paths, regularly spaced local maxima are identified in the vessel-enhanced images and connected with four-dimensional tubular paths (Benmansour et al., 2013) consisting of spatial locations and the radius estimate. This generates a local maxima graph $G = (V, E)$ (see Fig. 3a,b), where edges $(i, j) \in E$ connect maxima $i \in V$ in a local neighborhood similar to (Türetken et al., 2013). For every edge, a path between its two endpoints is extracted (see Fig. 3c), which minimizes the geodesic distance in the four-dimensional path space, thus penalizing small tubularity values along the path (Benmansour et al., 2013). To drastically prune these edges, a filtering step is performed removing all paths that fully contain any other path. The resulting graph still contains many spurious edges, but also the real arterial and venous vessel paths. A subsequent integer program groups these tubular paths in order to represent distinct vascular subtrees. A main contribution of our work is that we do not need declared root nodes for this subtree calculation, but generate them implicitly by minimizing a constrained objective function. As each of these distinct subtrees is either an artery or a vein, they have to be labeled accordingly. This is done in our second contribution, a novel integer program based A/V labeling step that uses two anatomical properties of pulmonary vessels. The first anatomical property is the uniform distribution of arteries and veins throughout the lungs. We quantify this uniform distribution by maximizing the contact surface of neighboring arterial and venous regions of a Generalized Voronoi Diagram (GVD). The second anatomical property that we use is the proximity of pulmonary arteries and bronchi. It is measured with an "arterialness" value similar to the method described in (Bülow et al., 2005). Solving the latter integer program delivers the final labeling that separates the subtrees into arteries and veins. In the following subsections we explain the two central optimization problems for subtree extraction and A/V labeling in more detail.

3.1. Subtree Extraction

In order to identify anatomically meaningful vessel trees and to prepare the input for A/V separation, we extract a set of directed subtrees from the local maxima graph G. Objectives of this step are: 1) to find the right directions of vessels in G and 2) to enforce that vessels form a tree. Since it is not possible to extract a single tree unambiguously, it is preferable to extract a set of trees. This reduces the risk of merging incorrectly arterial and venous subtrees, which could not be fixed later

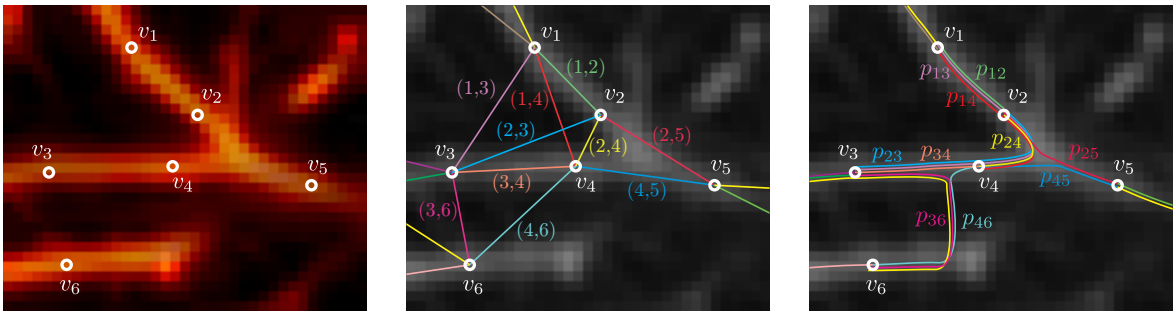

(a) Local maxima (b) Neighbor edges (c) Paths

Figure 3: Example of a local maxima graph extracted from the tubularity image. In (a) the regularly spaced local maxima v_i of a tubularity image are shown, while (b) illustrates the edges (i, j) formed by neighboring vertices. Note that here undirected edges are shown for simplicity, while we extract two directed edges between vertices to form our graph. Finally, (c) shows candidate paths for the subtree extraction.

in the labeling step. The subtree extraction is formulated as an optimization procedure similar to (Türetken et al., 2013) but in contrast to them, we extract multiple trees simultaneously and do not need to declared root nodes explicitly. We evaluate two equivalent formulations: one using the 0-1 integer quadratic program (IQP) similar to (Türetken et al., 2013), and a new one using a Markov random field (MRF) model. In both cases, a configuration of binary variables minimizing a constrained objective function is calculated, involving weights of adjoining vessel paths (see Sec. 3.1.3). The result of the MRF reformulation is a drastic improvement in computation speed. The MRF model is solved using the method of (Shekhovtsov et al., 2015), based on linear programming (LP) relaxation and its fast dual block-coordinate solver (Kolmogorov, 2006). It turns out that for our problem in many instances the method returns a globally optimal solution in less than a second (Sec. 5, Table 1). For some more difficult instances, a good approximate solution was found and a significant part of this solution was determined provably optimal.

3.1.1. IQP Formulation

(Türetken et al., 2013) have proposed the following formulation. Let $t_{ij} \in \{0, 1\}$ be an edge selection variable for each directed edge $(i, j) \in E$. Assignment $t_{ij} = 1$ means that edge (i, j) is a part of an extracted subtree. The optimal subtree problem is formulated as

$$\arg\min_{t} \sum_{(i,j),(j,k)\in E} w_{ijk}\, t_{ij}\, t_{jk} \tag{1}$$

subject to that the subgraph formed by edges with $t_{ij} = 1$ is a tree. They argued that many models for subtree extraction in the literature, if discretized, can be cast as this optimization problem, known as *quadratic arborescense*. Note, that costs w_{ijk} are assigned to pairs of incident edges in the graph and can conveniently model geometrical compatibility of the respective tubular paths.

We first propose a variant adopted to our case, where we search for an optimal *branching*, which is a forest rather than a single tree. For this purpose we introduce a virtual root node, denoted 0. The graph G is extended to a graph $G' = (V \cup 0, E')$ in which node 0 is connected to all other vertices by a directed edge. We then can use the problem formulation (1) for graph G'. It is clear that any tree in G' corresponds to a forest in G (by excluding the virtual root) and vice-versa. In the objective

$$\sum_{(i,j),(j,k)\in E'} w_{ijk}\, t_{ij}\, t_{jk} \tag{2}$$

we distinguish the *root-penalizing* factor:

$$w_{0jk} = \sigma, \tag{3}$$

imposed on pairs originating from the virtual root, and *edge pair compatibility* weights w_{ijk} for $(i, j), (j, k) \in E$.

Weights of edge pairs can easily incorporate the geometric relation of adjoining edge paths as well as weights reviewassociated to individual edge paths, in a single combined weight. Only if both edges (i, j) and (j, k) are selected in the solution, the weight w_{ijk} is added to the sum. The detailed modeling of the weights is deferred to Sec. 3.1.3.

The root penalizing factor is the cost for introducing a new subtree (branch) in the solution. By assigning cost σ to every edge pair that originates from the virtual root, it penalizes the number of distinct subtrees in the solution but also the number of branches that start from the same node. Indeed, it is preferable that every subtree

had a root node of degree 1, *i.e.* was starting from an edge rather than a branching point, as the former gives more reliable subtrees.

We further replace the tree constraint with the following simple constraints:

$$\sum_{i\,|\,(i,j)\in E'} t_{ij} = 1 \qquad \forall j \in V, \tag{4a}$$

$$t_{ij} + t_{ji} \leq 1 \qquad \forall (i,j) \in E. \tag{4b}$$

Constraint (4a) ensures that each node (except the virtual root) has exactly one incoming selected edge. Note that an incoming edge from the virtual root node is allowed and is not penalized in our model unless there is an outcoming edge too forming an edge pair (and hence starting a subtree). Constraint (4b) ensures that there are no 2-edge cycles in the solution. Longer cycles are theoretically permitted in the solution, *i.e.* a feasible solution to (4) need not be a forest. Contrary to that, we found experimentally that optimal solutions in this formulation were cycle-free for all 25 tested instances (also during parameter optimization). Given that a single lung contains a lot of sub-trees, we suggest that cycles are not a major limitation of the method. Additionally, in Appendix B we discuss possible strategies to tackle this problem in case cycles would occur (existing approaches and a simple heuristic).

It can be seen that the formulation (2)-(4) is fully equivalent (including the possibility of cycles) to the one we proposed in (Payer et al., 2015). There, an explicit variable indicating a root edge was used. Introducing a virtual root node allowed to write it in a more compact form[1], which is also more convenient for the following MRF reformulation.

3.1.2. MRF Formulation

In this formulation, instead of selecting active edges, we let each node $i \in V$ to select a parent from all its incoming edges in G', including the virtual parent 0 (meaning that node i is a root of a subtree). Since the constraint (4a) ensures that each node has exactly one incoming edge, such encoding of the solution is fully equivalent. Let $\mathcal{L}_i = \{j \in V \cup 0 \,|\, (j,i) \in E'\}$ be the set of tail vertices of the incoming edges to i in the extended graph G'. A *labeling* $x = (x_i \in \mathcal{L}_i \,|\, i \in V)$ defines a spanning subgraph X of G'. We will minimize the following *pairwise energy* function:

$$\min_x \sum_{(i,j)\in E'} f_{ij}(x_i, x_j). \tag{5}$$

Pairwise energy terms f_{ij} are designed to match formulation (2)-(4):

$$f_{ij}(x_i, x_j) = \begin{cases} w_{hij} & \text{if } x_i = h \neq j,\ x_j = i, \\ \infty & \text{if } x_i = j,\ x_j = i, \\ 0, & \text{otherwise.} \end{cases} \tag{6}$$

The infinity cost implements constraint (4b) penalizing 2-edge cycles; the weight w_{hij} is active only if the solution selects (h,i) as an incoming edge for i and (i,j) as incoming for j, *i.e.* same as in (2). Zero penalty corresponds to cases when edge (i,j) is not selected in the solution. Since, by construction, both directed edges (i,j) and (j,i) are present in graph G, the cost of the reverse linkage with $x_i = j$ and $x_j = h \neq i$ is charged to the directed edge (j,i).

In fact, methods (Shekhovtsov et al., 2015) and (Kolmogorov, 2006) will reformulate the energy minimization (5) as an integer linear program (ILP) and solve its LP relaxation, a strategy similar to what an off-the-shelf solver would apply to solve the 0-1 IQP (2)-(4). The experimental evidence confirms that we obtain a better performance with the MRF approach due to: i) a tighter resulting LP relaxation and ii) application of a fast specialized dual solver. The ILP formulation of the MRF model (5) is provided for comparison in Appendix A.

3.1.3. Edge Pair Weight

The choice of the edge pair weight w is a critical part of our subtree extraction model. It provides local evidence distinguishing between physically plausible and implausible vessel paths (see Fig. 4). Similar to (Türetken et al., 2013), we use a combined weight of adjoining edge pairs $(i,j),(j,k)$ to incorporate the geometric relationship of vessel paths.

The edge pair weight of the edges (i,j) and (j,k) with their corresponding paths p_{ij} and p_{jk} is defined as

$$w_{ijk} = \alpha\, w_{ijk}^{distance} + \beta\, w_{ijk}^{direction} + \gamma\, w_{ijk}^{radius} + \delta, \tag{7}$$

with $\alpha, \beta, \gamma, w_{ijk}^{distance}, w_{ijk}^{direction}, w_{ijk}^{radius} \in \mathbb{R}_0^+$, and $\delta \in \mathbb{R}^-$.

As the objective function (2) will be minimized, weights of physically plausible paths need to be negative. Otherwise the trivial solution, *i.e.* no path extracted, would always be a minimizer. The only term of (7) that may be negative is δ, all other terms are positive. The choice of the parameters $\alpha, \beta, \gamma, \delta$ is important and requires a parameter search over a range of

[1] The previous formulation had a bit larger modeling power, *e.g.* the possibility to penalize the number of subtrees and branches starting from the same node differently, which however was not used.

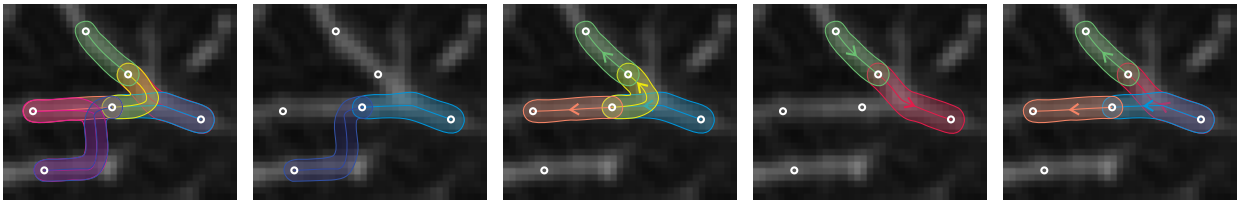

(a) Example path graph (b) Large distance weight (c) Large direction weight (d) Large radius weight (e) Extracted subtree

Figure 4: Example of a small path graph and visualization of the edge pair weight. The edge pair weight should be defined such that minimizing the objective function (2) of graph (a) leads to subtree (e). The images (b), (c) and (d) show example paths pairs with large distance weight, direction weight, and radius weight, respectively. For better visualization, the images show 3D (x, y, r) instead of 4D (x, y, z, r) paths, where the radius coordinate r is visualized as the border of the vessel.

suitable values. The value of δ is constant for all edge pairs and balances weights to be negative when a path is physically plausible and positive if not. The direction weight $w_{ijk}^{direction}$, the radius weight w_{ijk}^{radius}, and the distance weight $w_{ijk}^{distance}$ are normalized to lie between 0 and 1 and penalize physically implausible paths by making their total edge pair weight w_{ijk} positive.

The first weight of (7), the distance weight $w_{ijk}^{distance}$ penalizes paths with a large geodesic distance, as visualized in Fig. 4b, and is defined as

$$w_{ijk}^{distance} = \max\left(l\left(p_{ij}\right), l\left(p_{jk}\right)\right), \tag{8}$$

where $l\left(p_{ij}\right)$ and $l\left(p_{jk}\right)$ are the geodesic lengths of paths p_{ij} and p_{jk} that were calculated by the path extraction algorithm. These values represent a minimal geodesic distance on the 4D tubularity image between the two end-points of a path. The distance weight of the edge pair is the maximum value of the geodesic lengths of the two adjoining paths.

As the calculated geodesic distance does not incorporate the direction of paths, path pairs with unusual directions or changes in direction, as shown in Fig. 4c, do not get larger weights than path pairs with physically plausible directions. Therefore, $w_{ijk}^{direction}$ penalizes such path pairs. It is defined as

$$w_{ijk}^{direction} = 1 - \frac{1}{(n-1)(m-1)} \sum_{i=1}^{n-1} \sum_{j=1}^{m-1} \mathbf{u}_i \cdot \mathbf{v}_j, \tag{9}$$

where $\mathbf{u}_i$ is the set of $n-1$ normalized direction vectors between the n points of path p_{ij} and $\mathbf{v}_j$ is the set of $m-1$ normalized direction vectors between the m points of path p_{jk}. Eq. (9) quantifies the orientation similarity of two paths by calculating the mean of all dot products of directions of the first path with the second path. In the best case, *i.e.* when both paths are collinear straight lines, the mean of all dot products is equal to 1. The more differences in direction are present, the lower the mean of its dot products will be. As the direction weight should be 0 in the best case, the final weight is calculated by subtracting this mean from 1.

The last weight of (7), the radius weight w_{ijk}^{radius}, helps detecting the correct direction of a path pair. In the lungs, vessel radii are usually decreasing with distance from the heart. Therefore, the radius weight penalizes path pairs with increasing radius, as visualized in Fig. 4d. It is defined as

$$w_{ijk}^{radius} = \max\left(1 - \frac{r^{start}(p_{ij})}{r^{end}(p_{jk})}, 0\right), \tag{10}$$

where $r^{start}(p_{ij})$ is the radius of the start-point of path p_{ij} and $r^{end}(p_{jk})$ is the radius of the end point of path p_{jk}. When the radius is increasing from start to end, their ratio is greater than 1, otherwise it is less than 1. To give higher weights to path pairs with increasing radii, this ratio is subtracted from 1. As we do not want to prefer path pairs that have decreasing radii, which would be negative after this subtraction, we set the minimum of w_{ijk}^{radius} to 0. In the best case, *i.e.* when path pairs have equal or decreasing radii, the radius weight becomes 0, while in any other case the weight becomes positive.

3.2. Subtree A/V Labeling

After extracting the set of subtrees T, every subtree $s_i \in T$ has to be labeled either as an artery or as a vein. This labeling problem is again modeled as a quadratic integer program, its objective function is explained in Sec. 3.2.1. At its core, the objective makes use of two anatomical properties of the lung vessels.

First, we utilize that arteries and veins are roughly uniformly distributed in the lung, *i.e.* there are no large regions, where just arteries and no veins, or vice versa, are present. This property holds for all subjects, irrespective of possible lung diseases, as blood that flows

from the heart into the pulmonary arteries has to flow back again oxygenated via the venous network to the heart (Tawhai et al., 2011).

We calculate a measure of uniformity with the help of a generalized Voronoi diagram (GVD) (Aurenhammer, 1991). The GVD determines the nearest subtree for every voxel inside the lung segmentation. By maximizing the number of voxels on the contact surface between artery and vein regions, a uniform distribution of those regions is encouraged.

The second anatomical property we use is that arteries run in parallel and in close proximity to bronchi, as previously also noticed and utilized by (Bülow et al., 2005) and (Jin et al., 2014). To quantify this property that encodes an "arterialness" measure, we search for bronchus points near the extracted vessels and calculate their distance and co-orientation.

3.2.1. Objective Function

The quadratic objective function used for A/V labeling incorporates two binary variables for every previously extracted subtree s_i. The first variable, a_i, is set to 1, if and only if the subtree s_i is an artery, whereas the second variable, v_i, is set to 1, if and only if the subtree s_i is a vein. The optimization problem reads:

$$\begin{aligned} \arg\max_{a,v} \quad & \sum_{s_i,s_j \in T} a_i \, v_j \, w_{ij}^{border} + \lambda \sum_{s_i \in T} a_i \, w_i^{artery} \\ \text{subject to} \quad & a_i + v_i = 1 \qquad \forall i \in \{1,\dots,|T|\}, \end{aligned} \tag{11}$$

where $a_i, v_i \in \{0,1\}$, and $\lambda, w_{ij}^{border}, w_i^{artery} \in \mathbb{R}_0^+$. Note that subtrees selected by a and v form a disjoint partition of T and problem (11) is equivalent to the well-known NP-hard Max-Cut problem (Karp, 1972). However, we do not investigate a dedicated solver for this problem since (11) can also be solved optimally very quickly with a general-purpose ILP solver, due to the low number of variables (proportional to the number of subtrees). The first term of (11) counts the number of voxels w_{ij}^{border} on the contact surface between artery and vein regions modeled by a GVD. If the trees s_i and s_j have different labels ($a_i = 1$ and $v_j = 1$) and additionally those two regions share a common border ($w_{ij}^{border} > 0$), their border weight is added to the sum. The labeling that maximizes the number of voxels on the contact surface encourages a uniform distribution of arteries and veins throughout the lung, see Sec. 3.2.2 for details.

The second term of (11) sums up the arterialness values w_i^{artery} for every tree s_i labeled as an artery ($a_i = 1$), thus performing the distinction between arteries and veins. Arterialness values of vessels labeled as arteries increase the sum, whereas those of veins do not change the sum. Therefore, the labeling that maximizes the arterialness sum of arteries encourages vessels with high arterialness values to be labeled as arteries. The arterialness measure is explained in more detail in Sec. 3.2.3.

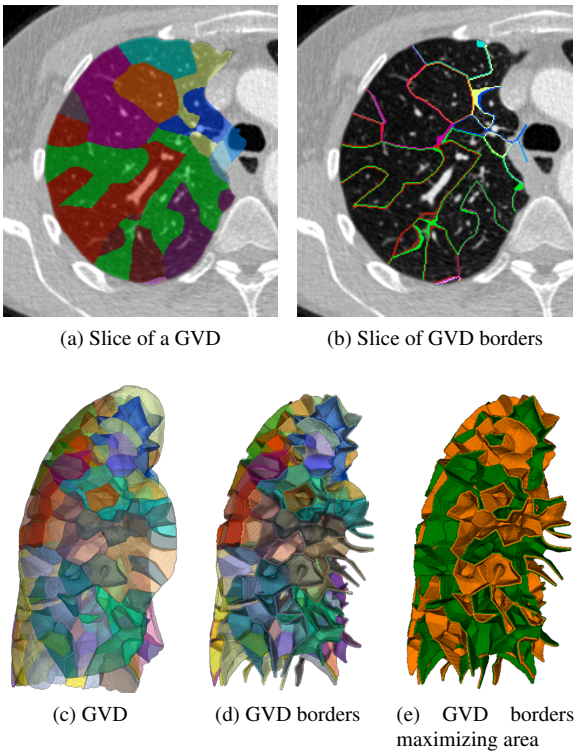

(a) Slice of a GVD (b) Slice of GVD borders

(c) GVD (d) GVD borders (e) GVD borders maximizing area

Figure 5: The generalized Voronoi diagram (GVD) for the subtrees. A single CT slice overlayed with the GVD is shown in (a), the borders of its regions are shown in (b). A 3D rendering of the GVD is shown in (c), of the borders of its regions in (d) and of the binary labeling that maximizes surface area in (e).

The constraint of (11) ensures that both labels are not active at the same time for the same tree s_i. Since w_{ij}^{border} and w_i^{artery} are scaled differently, they need to be normalized between 0 and 1, which is accomplished by dividing by their maximally occurring values. Thus, the two summation terms have a similar range of values and are weighted with a parameter λ balancing the influence of the two anatomical properties.

3.2.2. Generalized Voronoi Diagram of Subtrees

To enforce a uniform distribution of arteries and veins throughout the lung, we use a generalized Voronoi diagram (GVD). The GVD partitions a volume into regions, based on the distance to their nearest seed points. Every voxel in the same connected region has the same nearest seed points (Aurenhammer, 1991). In our case, every extracted subtree acts as a seed, leading to a GVD

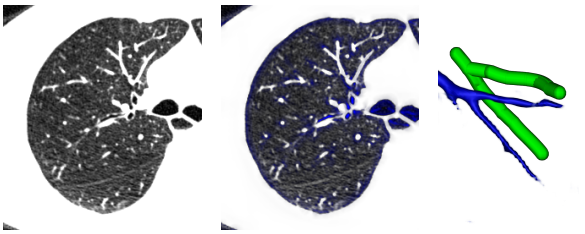

(a) CT slice clamped below -400 HU (b) Bronchus enhancement slice (c) Close artery and bronchus

Figure 6: The steps of calculating the arterialness measure. (a) A slice of the preprocessed input image. (b) The result of the subsequent bronchus enhancement for this slice, where the bronchus response is overlayed in blue. (c) A 3D rendering of two artery segments (green) and a nearby bronchus (blue).

where every voxel inside the lung segmentation is assigned the label of its nearest subtree. GVD calculation is illustrated in Fig. 5. Computation of the GVD involves distance transforms for every subtree, where the region inside the subtree has a distance of 0.

The binary labeling that maximizes the number of voxels on the contact surfaces of different regions, *i.e.* their common border (see Fig 5b), encourages a uniform distribution of both labels throughout the lungs. The number of voxels w_{ij}^{border} on the contact surface of the two Voronoi regions i and j of subtrees s_i and s_j is calculated by counting all voxels inside region i having a voxel of region j as direct neighbor or vice-versa.

Note that by just maximizing the contact surfaces, at this point it is still undefined which of the binary labels corresponds to arteries and veins, respectively. Therefore, we additionally use the arterialness measure as explained in the following subsection.

3.2.3. Arterialness Measure of Vessel Segments

Similar to (Bülow et al., 2005), we propose a measure that gives a high value for arteries. It uses the anatomical property that arteries, in contrast to veins, are usually running parallel and in close proximity to bronchi.

For our proposed arterialness measure, we require a bronchus enhanced image (see Fig. 6) and bronchus direction estimates. In contrast to (Bülow et al., 2005) we compute the OOF (Law and Chung, 2008), but configured to enhance dark-on-bright instead of bright-on-dark structures. The input CT images are first clamped to have a maximal value of -400 Hounsfield Units (HU) to further enhance bronchus regions and suppress responses outside the HU range of interest.

The arterialness measure is calculated independently for each vessel segment, *i.e.* the part of the vessel between branching- or end-points. At each voxel along a segment, we search for similarly oriented bronchial structures giving high tubularity response in a plane orthogonal to the vessel direction. Every point inside this plane with a value below a certain threshold t_{HU} in the input image, a value above t_{OOF} in the bronchus enhanced image, and a maximal angle t_θ relative to vessel direction, is added to a list of bronchus candidate points.

The bronchus near a vessel segment is approximated by a straight line. To connect the extracted probable bronchus points of a whole vessel segment, a line l is fitted to them. The line fitting algorithm needs to be robust to outliers, as some non-bronchus points may be in the list of probable bronchus points as well, *i.e.* points adjacent to a vessel, which may also get high tubularity values and have directions similar to the vessel. If the number of bronchus candidate points, which lie in close proximity to l, normalized to the length of the vessel segment is below a threshold t_ρ, the fitted line does not represent a bronchus branch. In this case, we set the arterialness measure to 0. Otherwise, the final arterialness value is the mean of the dot products of the vessel direction and the direction of the fitted line divided by their point distance. This leads to higher values for arteries running closer and in parallel to bronchi, while veins, typically more distant and deviating stronger from the bronchus direction, will receive lower values. Formally, the arterialness value of a vessel path p is defined as

$$arterialness(p) = \sum_{k=1}^{n-1} \frac{c(\mathbf{u}_k, \mathbf{v})}{d(\mathbf{p}_k, l)}, \tag{12}$$

with $c(\mathbf{u}_k, \mathbf{v}) = |\mathbf{u}_k \cdot \mathbf{v}|$ being the co-orientation of the vectors $\mathbf{u}_k$ and $\mathbf{v}$, and $d(\mathbf{p}_k, l)$ being the minimal Euclidean distance between path point $\mathbf{p}_k$ and line l. The set $\mathbf{u}_k$ is composed of $n-1$ normalized direction vectors between the n points of path p, whereas $\mathbf{v}$ is the normalized direction vector of line l.

The final arterialness weight w_i^{artery} of a tree s_i is the sum of all arterialness values of its vessel segments.

4. Experimental Setup

To evaluate our proposed A/V separation algorithm, we collected 25 datasets from patients with different degrees of lung vascular disease who underwent thoracic contrast-enhanced CT examinations. To show that the proposed method is also applicable to non contrast-enhanced images, we additionally evaluated one dataset of a subject, where both contrast-enhanced as well as native CT data was available. For each of our 25 datasets, a manual reference A/V segmentation was

created. Furthermore, we re-implemented the semi-automatic A/V separation method of (Park et al., 2013), which also extracts and labels vascular subtrees.

4.1. Data Description

The 25 CT datasets (13 female / 12 male) used for evaluation were acquired with two scanners, a Siemens Somatom Definition Flash (D30f kernel) or a Siemens Somatom Force (Qr40d reconstruction kernel). The native dataset was acquired with a Siemens Somatom Definition Flash (B30f kernel). Ten of the 25 CT datasets are the same as already used in (Payer et al., 2015). CT volumes were on average $512 \times 512 \times 475$ voxels with a nearly isotropic phyiscal resolution of 0.6 mm. The average dose of the CT scans was 3.2 ± 1.7 mSv.

For each patient dataset, a voxel-based manual reference labeling was created using ITK-SNAP [2]. This tedious procedure was performed by radiologists and experts in lung vascular analysis. At least two experts were involved in the creation of each manual reference, one doing the vessel segmentation/labeling and the other(s) doing a quality check and corrections. The labelings were restricted to include only voxels above -400 HU. To reduce image noise and simplify manual processing, the input volume was denoised with an edge-preserving denoising method based on total generalized variation (Bredies et al., 2010). The generated manual references cover the main pulmonary artery and the left atrium as well as pulmonary vessels with a diameter down to a few mm. The approximate average time for generating each reference was ten hours, involving both the segmentation/labeling and quality check by one or more additional expert(s).

For one dataset, a native CT scan of the same patient was investigated in addition to the contrast-enhanced scan. To prevent the necessity of creating a separate manual reference A/V segmentation for this dataset, we nonlinearly registered both data sets and warped the reference labeling of the contrast-enhanced scan to the native scan. We obtained the transformation parameters by performing a deformable B-spline registration (Rueckert et al., 1999) based on the mean squared error cost function as provided by the 3DSlicer software [3].

4.2. Re-implementation of (Park et al., 2013)

For comparing to related work from the literature, we created a re-implementation of the method from (Park et al., 2013). Thus, we were able to compare both our automated approach as well as the re-implemented method to our manual reference A/V segmentations. As the method of (Park et al., 2013) needs a full vessel segmentation as input, we binarized the result of the vessel-enhancement filter OOF (Law and Chung, 2008) with a threshold of 150 on the filter response. Qualitatively, this gave improved vessel segmentation results compared to what was originally proposed in (Park et al., 2013). Afterwards, the shortest path trees from the center of the input image were calculated for both lungs separately. As proposed in (Park et al., 2013), each of those trees were manually cut to remove the critical and erroneous central regions of the lungs. The regions for the manual cuts were created using a segmentation in ITK-SNAP, where the border of the segmentation served as the cut in the tree. The final A/V labeling of the subtrees was done manually by looking for connections to the heart for every subtree. We carefully tried to optimize the re-implementation of this method to perform as well as possible on our datasets.

4.3. Quantitative Evaluation Experiments

Experiment 1: In the first experiment, the total volume-based overlap between the labeled A/V segmentation results and the manual reference is calculated, in the same manner as presented in (Payer et al., 2015). The manual reference may include correctly labeled voxels, that are not present in the calculated A/V labelings, especially in the main pulmonary artery and left atrium, containing on average 70% of all voxels of the manual reference. Additionally, the calculated A/V segmentations include many correctly labeled voxels, that are not present in the manual reference, as the calculated segmentations usually include many more small vessels, resulting in an average of 40% of automatically labeled voxels that are not included in the manual reference. Therefore, we compared only those voxels that are present in both the calculated A/V segmentation and the manual reference. By only comparing the union of the segmentation and the manual reference, hardly any voxels of the manual reference inside the lungs are ignored, since the algorithm delivers vessels with smaller diameters than available in the manual ground truth. As a drawback of comparing only the union, falsely detected non-vascular structures cannot be quantified. We deal with these non-vascular structures in Experiment 3.

The voxel-based agreement is defined as the ratio of voxels with equal label to the voxels that are present in both segmentations, thus focusing on correct A/V label-

[2] ITK-SNAP version 3.0 or 3.2 from http://www.itksnap.org/
[3] http://www.slicer.org, version 4.5.0

ing. Formally, we define

$$\text{agreement}_{voxel} = \frac{\left|A_{ref} \cap A_{test}\right| + \left|V_{ref} \cap V_{test}\right|}{\left|(A_{ref} \cup V_{ref}) \cap (A_{test} \cup V_{test})\right|}, \quad (13)$$

where A_{ref} and V_{ref} are the sets of arteries and veins of the manual reference, respectively, and A_{test} and V_{test} are the corresponding results of the A/V segmentation.

Experiment 2: Vessels with larger diameter have a higher influence than vessels with smaller diameter, giving a slight bias in Experiment 1. In order to overcome this limitation, the second experiment quantifies the centerline-based agreement with the manual reference, as it is independent from the vessel's diameter. Similarly to the first experiment, we can only compare the label of the centerline that is inside the manual reference. The centerline-based agreement is here defined as the ratio of the length of the centerline with correct label to the length of the total centerline that is included in the manual reference. The centerline used for calculating the agreement is obtained by the subtree extraction as described in Sec. 3.1. We define

$$\text{agreement}_{centerline} = \frac{\left|l_A \cap A_{ref}\right| + \left|l_V \cap V_{ref}\right|}{\left|(l_A \cup l_V) \cap (A_{ref} \cup V_{ref})\right|}, \quad (14)$$

where A_{ref} and V_{ref} are the sets of arteries and veins of the manual reference A/V labeling, l_A and l_V are the centerlines of the calculated arterial and venous subtrees, and lengths are measured by cardinalities of sets. As obtaining the centerline for segmentations from (Park et al., 2013) is not straightforward and we do not expect significantly differing results compared to Experiment 1 due to the interaction step, we perform this experiment only for our proposed method, where the resulting separation is represented as 4D centerlines.

Experiment 3: In this experiment, we evaluate the subtree extraction. We quantify the number of subtrees, that are composed of non-vascular structures, and the amount of subtrees, that are not well separated, *i.e.* consist of merged arteries and veins. In order to run this experiment, we skip the A/V labeling part of the algorithm and perform an automatic labeling by explicitly using the manual references. Every subtree, that has a voxel-based overlap with the manual reference of less than 10% of its total volume, is considered a non-vascular subtree. Every subtree, that is merged with both artery and vein label of the manual reference of at least 10% of its total volume is considered as a merged subtree. We quantify the amount of non-vascular and merged subtrees by dividing their volume by the total volume of the automatic segmentation. Again, as this is not straightforward for the method of (Park et al., 2013), we perform this experiment only for the proposed method.

Experiment 4: In the last experiment, we compare the results and efficiency of the IQP and MRF formulations (Sec. 3.1). We calculate the relative optimality gap, which is defined as

$$\text{gap} = \frac{E - LB}{|E|}, \quad (15)$$

where LB is the lower bound and E is the best objective value found by the respective solver[4]. We additionally compare the runtime of both formulations.

4.4. Implementation Details

The path extraction was configured to extract vessel paths up to 10 mm diameter. The parameters used for the subtree extraction part were obtained by performing a grid-search, where the final parameters used for evaluation provide a good trade-off between a low number of detected subtrees and a high voxel-based overlap with the manual reference. The range of the variables for the grid search was $[0.5, 2.0]$ for α, β and γ, $[-0.1, -0.5]$ for δ, and $[0.1, 0.5]$ for σ, resulting in a total of 1600 combinations. Following the grid search result, the values are set to $\alpha = 2.0, \beta = 0.5, \gamma = 0.5, \delta = -0.2$, and $\sigma = 0.2$. The parameters for the subtree A/V labeling were also obtained by grid search, where the final parameters provide the highest mean voxel-based agreement with the manual references. Here, the grid search was performed on the parameters $\lambda \in [1, 10]$, $t_{HU} \in [-750, -650]$, $t_\theta \in [0.6, 0.8]$, and $t_\rho \in [0.4, 0.8]$, resulting in a total of 450 combinations. The weighting parameter of the objective function is set to $\lambda = 9$, whereas the parameters for the arterialness measure are set to $t_{HU} = -750$, $t_{OOF} = 100$, $t_\theta = 0.8^{\text{rad}}$, and $t_\rho = 0.7$.

The development and testing platform was an Intel Core i7-4820K @ 3.70 GHz with 32 GB RAM under Ubuntu 15.04. The algorithm was developed in C++ with ITK. For 4D path extraction, publicly available code from (Benmansour et al., 2013) was taken. Gurobi Optimizer[5] was used as a solver for integer programs.

5. Results

The automatic algorithm generated an average of 1982 individual vessel segments, composed of 1010 arteries and 972 veins, with diameters ranging from 2

[4]Same optimality measure as used in Gurobi.

[5]Gurobi Optimizer Version 6.5 with academic license from http://www.gurobi.com/

(a) Reference (b) Automatic (c) Overlap

Figure 7: Example segmentation of a dataset with a high voxel-based agreement of 99.4% (Patient # 6). Visualization of reference segmentation (left), automatic segmentation (center) and their overlap (right). Only voxels, that are set in both segmentations, are visualized in the overlap image. Red: veins, blue: arteries, yellow: disagreement between segmentations.

to 10 mm. An example A/V segmentation result and its comparison with the manual reference is shown in Fig. 7. The approximate runtime needed for automatically computing a single A/V segmentation with the proposed algorithm was 50 minutes (lung segmentation 5 min, airway and bronchus enhancement 10 min, generation of graph G and subtree extraction 30 min, A/V labeling 5 min). The re-implementation of (Park et al., 2013) needed 30 minutes computation time with additional 2 hours of manual interaction per dataset.

The quantitative results of the volume-based agreement (Experiment 1) are shown in Fig. 8. The proposed method achieves a mean voxel-based agreement with all manual references of 91.1% (median: 96.3%, range: 56.5% to 99.4%), the method of (Park et al., 2013) 91.2% (median: 91.4%, range: 82.4% to 95.3%), see Fig. 8a and 8b. We have also performed a separate evaluation of arteries and veins and compared it with the reference. Our method achieves a mean voxel-based agreement of 91.2% for arteries (median: 96.1%, range: 56.3% to 99.3%) and 90.5% for veins (median: 95.2%, range: 56.6% to 99.8%), respectively, indicating that there is no difference in performance between both vascular structures. Additionally, to compare the results of the proposed algorithm with the results reported in our previous work (Payer et al., 2015), we evaluated the algorithm just on the same ten datasets used there. Here, the proposed method achieves a mean voxel-based agreement of 95.2% (median: 96.0%, range: 88.0% to 99.4%), as compared to the previously reported 94.1% (median: 95.0%, range: 85.0% to 98.7%), see Fig. 8c and 8d. The centerline-based agreement (Experiment 2) with the manual references showed similar results as the voxel-based overlap with a

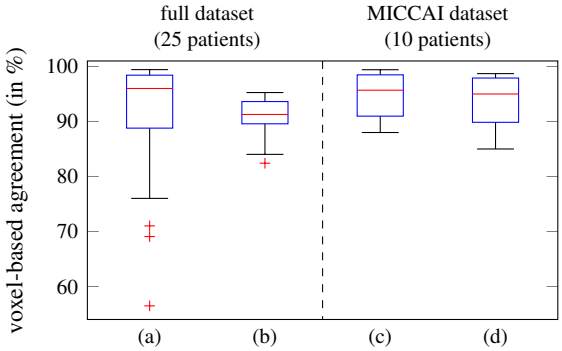

Figure 8: Box-whisker plots of the volume-based agreement (Experiment 1) with the manual reference. The results on all 25 datasets are shown in (a) for the proposed automatic method and in (b) for the re-implementation of the interactive (Park et al., 2013) method. (c) and (d) show the results on the same 10 datasets that have been used in (Payer et al., 2015), where (c) shows the results with the currently proposed method and (d) the results as presented previously.

mean value of 90.9% (median: 95.5%, range: 55.1% to 99.6%). The mean volume ratio of erroneously detected non-vascular structures compared to the total volume of the segmentation is 1.6% (median: 0.9%, range: 0.0% to 10.8%), whereas the mean volume ratio of merged subtrees is 1.8% (median: 0.6%, range: 0.0% to 8.3%) (Experiment 3). The individual quantitative results for all datasets for Experiments 1-3 are shown in Table 2, while Fig. 10 and 11 show more qualitative results.

The comparison of the solution quality of IQP and MRF (Experiment 4) is shown in Table 1. Note that the solver for the IQP was set to terminate, when the last found improvement of the objective value was more than 10 minutes ago, as this leads to a good tradeoff be-

tween optimality gap and runtime. Additionally, Fig. 9 shows which part of a solution was certified as globally optimal (persistent) by the MRF method.

Table 1: Comparison of IQP and MRF formulation. The values show the mean ± standard deviation over all 25 datasets.

IQP		MRF	
gap	runtime	gap	runtime
3.2% ± 4.3%	1122s ± 482s	0.1% ± 0.6%	0.38s ± 0.43s

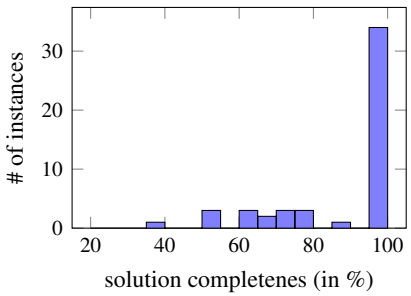

Figure 9: Identified part of globally optimal solution in the subtree extraction problem. Out of 50 instances (25 left and right lungs, respectively) 33 were solved completely (100%), for the remaining instances a significant part of the solution was proven to be optimal.

To show the applicability of the method on non contrast-enhanced images, we repeated the experiments for one dataset, where a CT examination with and without contrast agent was available (Patient # 7). The resulting A/V segmentations for both images are obtained with the same parameters that have been used for the whole dataset of 25 patients. As compared to the image with contrast agent, on the native image less vessels are detected, 2852 compared to 1599. The results of the experiments for the native dataset are as follows: The voxel-based agreement is 94.6% (compared to 98.9% with contrast agent), the centerline-based agreement is 94.6% (99.0%), the amount of erroneously detected non-vascular structures is 0.7% (0.8%), and the amount of merged subtrees is 5.2% (0.5%).

6. Discussion

Our novel fully automatic algorithm for separating arteries and veins in thoracic CT images achieves a higher median and a similar mean overlap with the manual reference A/V segmentations compared to the method of (Park et al., 2013), which requires time-consuming explicit manual correction. These results were obtained on a challenging, clinically relevant dataset of 25 patients with different degrees of lung vascular disease.

Table 2: Results for all datasets. The results are shown for every dataset individually. Mean (μ), standard deviation (σ) and median are calculated with all datasets. All values are shown in %.

	Experiment 1		Experiment 2	Experiment 3	
Patient	voxel-based		centerline-	non-	
#	*Proposed*	*Park*	based	vascular	merged
1	**96.7**	84.0	97.0	3.3	2.2
2	**95.7**	90.2	94.9	1.7	4.0
3	**84.4**	82.4	85.1	0.4	0.6
4	**98.5**	90.5	98.1	0.0	1.4
5	**98.6**	93.9	98.4	0.2	3.0
6	**99.4**	88.0	99.1	1.3	1.2
7	**98.9**	91.2	99.0	0.8	0.5
8	**94.6**	90.4	94.9	5.6	1.5
9	89.9	**91.4**	89.3	0.6	0.1
10	56.5	**87.6**	55.1	0.0	7.9
11	**96.3**	91.9	95.7	2.6	8.3
12	**98.3**	94.3	98.2	1.6	2.2
13	**89.6**	88.0	90.7	0.3	4.2
14	**98.6**	93.4	97.8	0.2	0.2
15	69.1	**92.2**	73.6	1.4	0.2
16	71.0	**90.9**	74.6	0.3	0.1
17	93.5	**94.0**	91.5	1.4	0.5
18	**96.8**	93.8	95.5	10.8	4.3
19	76.0	**88.9**	71.4	0.9	0.1
20	**98.0**	94.8	98.0	0.7	0.4
21	**98.8**	93.2	98.4	0.0	0.0
22	88.0	**90.6**	87.2	0.3	0.0
23	**99.4**	94.7	99.6	1.6	0.0
24	**98.3**	95.3	97.6	2.7	3.0
25	92.4	**93.3**	91.2	1.4	0.0
μ	91.1	**91.2**	90.9	1.6	1.8
σ	11.3	**3.3**	11.1	2.3	2.4
median	**96.3**	91.4	95.5	0.9	0.6

We assume that our integer program is better modeling the anatomical constraints involved in A/V separation compared to the minimum-spanning-tree of (Park et al., 2013), by exploiting the uniform distribution of arteries and veins throughout the lungs as well as proximity and similar orientation of arteries and bronchi. Compared to our previous work in (Payer et al., 2015), we have significantly improved computational efficiency of our method from several hours to less than one hour by switching to our novel MRF based solver. Now the majority of computational effort goes into computing graph G for subtree extraction, while solving the respective discrete optimization problem can be performed in a matter of seconds. At the same time we have slightly improved the accuracy in terms of volume overlap as can be seen in Fig. 8c,d, where the same ten datasets used in (Payer et al., 2015) were evaluated again for this work. The reason for this slight improvement are minor code optimizations when extracting the vessel tree candidate paths, together with a new set of optimal parameters for subtree extraction determined by the grid search for the new subtree extraction MRF solver.

Our new, larger dataset of 25 subjects proved to be a more challenging dataset as compared to the one used

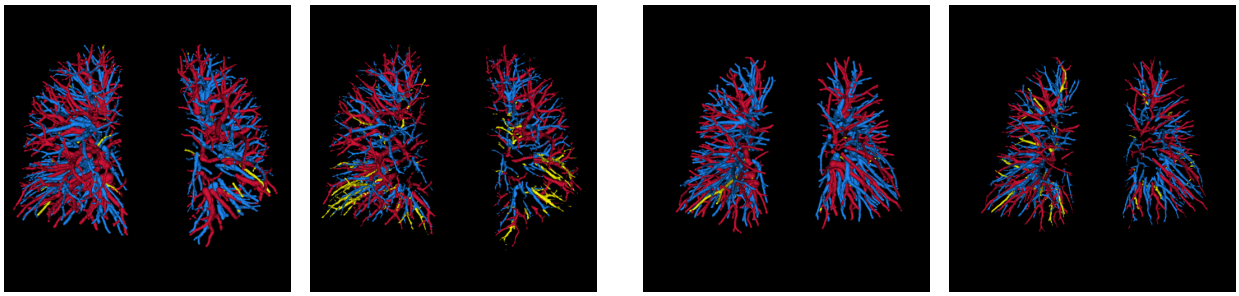

(a) Proposed # 4 (98.5%) (b) Park # 4 (90.5%) (c) Proposed # 5 (98.6%) (d) Park # 5 (93.9%)

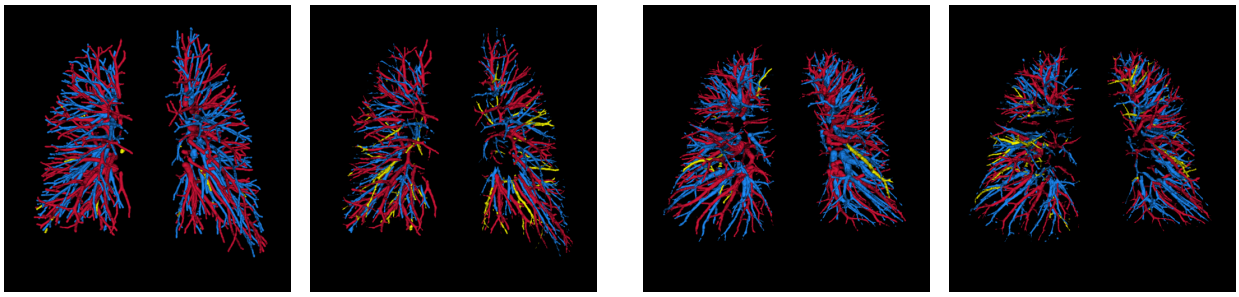

(e) Proposed # 7 (98.9%) (f) Park # 7 (91.2%) (g) Proposed # 12 (98.3%) (h) Park # 12 (94.3%)

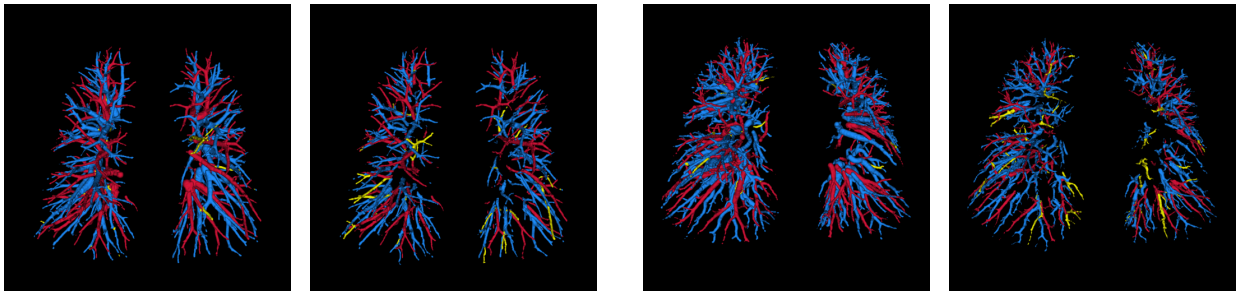

(i) Proposed # 20 (98.0%) (j) Park # 20 (94.8%) (k) Proposed # 21 (98.8%) (l) Park # 21 (93.2%)

Figure 10: Example overlap images for selected datasets. The left images show the overlap with the manual reference for the proposed method, the right images show the overlap with the manual reference for the re-implementation of (Park et al., 2013). Quantitative overlap values in brackets. Red: veins, blue: arteries, yellow: disagreement between segmentations.

in (Payer et al., 2015) due to the 15 added patients where some showed more severe lung disease or had a low signal-to-noise ratio (SNR). The most problematic datasets were Patient # 10 and # 15, which are illustrated in Fig. 11. The CT image of Patient # 10 shows a low SNR due to a high body mass. This leads to very few detected vessels that the A/V labeling part is not able to handle correctly. In Patient # 15, who has an extraordinarily large heart, almost the whole left lung is mislabeled. We observed that in this patient, the airways show low contrast, especially in the left lung, which leads to a failure of the arterialness calculation and the A/V labeling. Here the interactive method benefits from the direct correction of the mislabeled parts. Interestingly, in Patient # 15 the right lung is labeled almost perfectly, although a congenital heart defect (draining of veins into superior vena cava) is present (see arrow in Fig. 11e), which poses a problem for automatic algorithms that assume standard cardiovascular anatomy. Additionally, we observed that mislabeled subtrees often are neighbors in the generalized Voronoi diagram, due to the maximization of their contact surfaces. This may lead to switched labels of all subtrees within a lung lobe, as can be seen in Patient # 19 in Fig. 11c,f. Therefore, restricting the uniform distribution of arteries and veins to lung lobes instead of whole lungs could further improve performance and robustness. However, for the majority of datasets we get an excellent overlap with the manual reference segmentations, as can be seen in Table 2 and in the examples in Fig. 7 and Fig. 10.

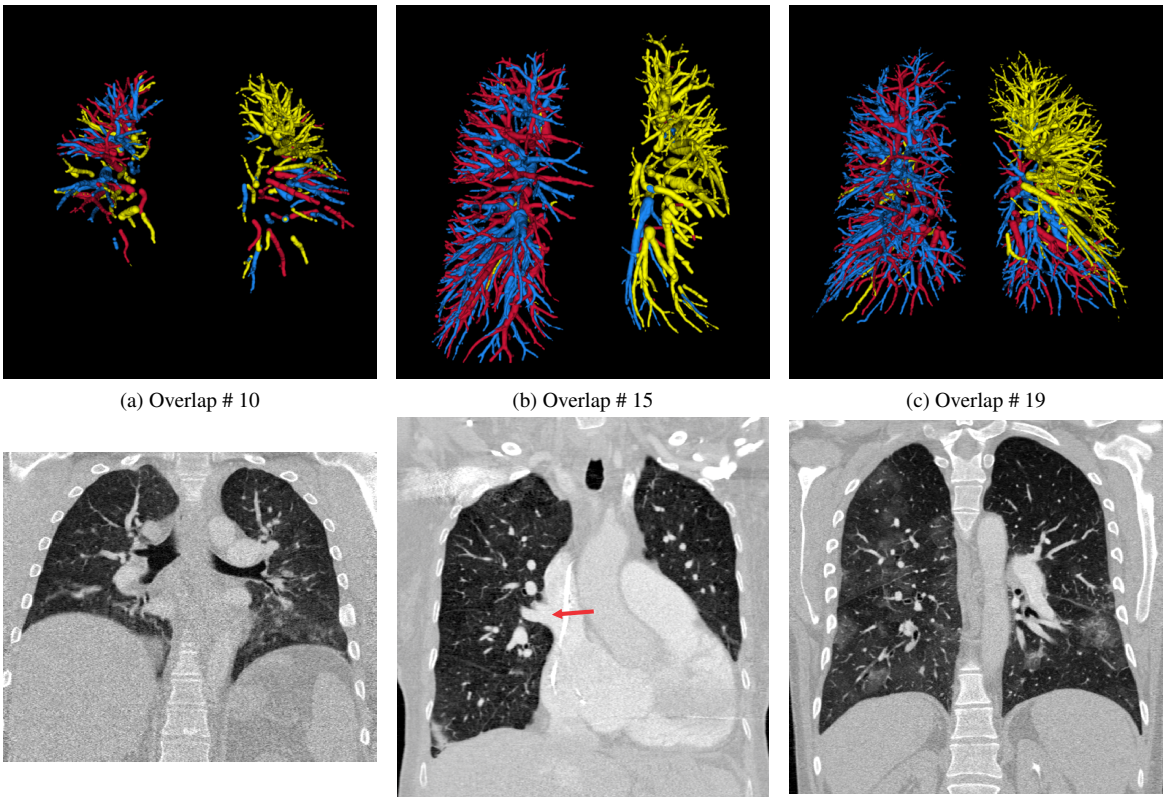

(a) Overlap # 10 (b) Overlap # 15 (c) Overlap # 19

(d) Slice # 10 (e) Slice # 15 (f) Slice # 19

Figure 11: Example overlap visualizations and CT slices, where the A/V separation had problems. In Patient # 10 only a small number of vessels are extracted, which leads to an instable A/V separation (agreement: 56.5%). In Patient # 15 the whole left side of the lung is switched, due to a hardly visible airway (agreement: 69.1%). In Patient # 19 one lobe is fully switched, due to the too strong enforcement of the uniform distribution of arteries and veins (agreement: 76.0%). Red: veins, blue: arteries, yellow: disagreement with manual reference.

We also showed on one dataset that we do not require contrast-enhancement to compute the A/V separation. Our result on the native scan of Patient # 7, despite having a slightly lower overlap compared to the contrast-enhanced scan, was well above 90% volume overlap with the manual reference. This indicates the potential of our method to be applied to native CT scans.

On our dataset of 25 patients we noticed that the choice of parameters is critical for optimal performance. We have come up with a single set of parameters that was used for all 25 datasets during evaluation. However, we have also noticed that for each dataset individually, a set of parameters exists which achieves an overlap with the manual reference above 90%. Unfortunately, these parameter sets differ between patients, where the parameters for the arterialness measure have the highest influence on the separation results. Making this measure more stable, *e.g.* by using a proper airway segmentation, could improve the robustness. We will focus future work on how to make our method more robust to overcome this parameter sensitivity. Still, our overall results are very promising even with the currently determined parameter set and due to the large number of images and their variability, we expect that this parameter set will generalize very well to new data.

We observed that our proposed method extracts very few non-vascular structures or merged subtrees, as can be seen in Table 2. Figure 12 shows two datasets with the highest amounts of non-vascular structures (Patient # 18) and merged subtrees (Patient # 11), respectively. Despite these large errors, the A/V labeling of the correctly identified vessel subtrees still works well in those two datasets, with a voxel-based overlap of over 96%.

Due to the lack of a standardized dataset and openly available A/V separation algorithms at the time of writing of this paper, our algorithm was only compared to the most recently published work in (Park et al., 2013). During the preparation of this manuscript, two very new approaches for automatic A/V separation have appeared as early access versions (Charbonnier et al., 2016; Kita-

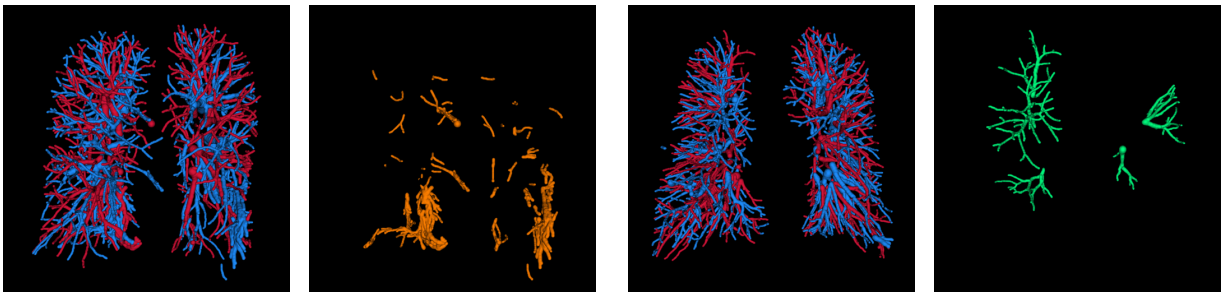

(a) Patient # 18 segmentation (b) Patient # 18 non-vascular (c) Patient # 11 segmentation (d) Patient # 11 merged subtrees

Figure 12: Example segmentations, where either many non-vascular structures are extracted or many subtrees are merged. (a) and (c) show the automatic segmentation results, while (b) and (d) show only the non-vascular structures (10.8% of total volume) or the merged subtrees (8.3% of total volume), respectively. Red: veins, blue: arteries, orange: non-vascular structures, green: merged subtrees.

mura et al., 2016). A comparison to these approaches would be an interesting task for future work, especially since (Charbonnier et al., 2016) propose a publicly available challenge data set for evaluating artery/vein separation algorithms.

Another interesting idea for future work is to combine subtree extraction and labeling into a single integer program, as was done in (Robben et al., 2014) for segmentation and labeling of brain vascular structures. However, this will lead to a very complex problem to solve and may prove computationally intractable.

7. Conclusions

In this work a novel automatic A/V separation algorithm based on the solution of two subproblems, vessel subtree extraction and A/V labeling of subtrees was presented. Our main contributions were the use of bottom-up and top-down knowledge on the geometry and anatomy of vascular trees to constrain this labeling problem formulated as integer programs, and the efficient solution of the more costly subtree extraction problem using an MRF solver. We improved on the A/V separation accuracy compared to our previous work in (Payer et al., 2015) on a larger, more challenging dataset. Except for some cases of low SNR and severe lung disease, we also outperform a related semi-automatic method from (Park et al., 2013) without requiring any user interaction. A drawback of our method is parameter sensitivity, mostly due to the computation of the arterialness measure, which needs to be addressed more extensively in future work. We conclude, that our novel method provides an opportunity to become an integral part of CAD of a number of lung diseases, where the separation of arteries and veins is of importance.

Acknowledgements

We acknowlede the reviewer comments that helped in improving our manuscript. We thank Peter Kullnig for radiological support, Vasile Fóris and Gabor Kovacs for medical support, and Martin Wurm and Dietmar Scherjau for help in some manual artery-vein separations. A. Shekhovtsov was supported by the European Research Council under the Horizon 2020 program, ERC starting grant agreement 640156. C. Payer was partly supported by the project FAME of the Austrian Science Fund (FWF): P28078-N33.

Appendix A. MRF ILP Formulation

Proposition 1. *The Standard ILP formulation for MRF model* (5)-(6) *can be written as:*

$$\min \sum_{(i,j),(j,k)\in E'} w_{ijk} t_{ijjk} \tag{A.1a}$$

$$t_{ij} \in \{0,1\},\ t_{ijj'k} \in \{0,1\}\ \ \forall (i,j),(j,k),(j',k) \in E'; \tag{A.1b}$$

$$\sum_{i\,|\,(i,j)\in E'} t_{ij} = 1\ \ \forall j \in V; \tag{A.1c}$$

$$t_{ij} + t_{ji} \leq 1\ \ \forall (i,j) \in E; \tag{A.1d}$$

$$\sum_{j'\,|\,(j',k)\in E'} t_{ijj'k} = t_{ij}\ \ \forall (i,j),(j,k) \in E'; \tag{A.1e}$$

$$\sum_{i\,|\,(i,j)\in E'} t_{ijj'k} = t_{j'k}\ \ \forall (j',k),(j,k) \in E'. \tag{A.1f}$$

Variables t_{ij} and $t_{j'k}$ correspond to the indicators of relations $x_j = i$ and $x_k = j'$, respectively. Variable $t_{ijj'k}$ represents the lifting of the product $t_{ij}t_{j'k}$. This lifting is of one order higher than would be sufficient to linearize objective (2). However, it is useful for the tightness of the subsequent relaxation. Since MRF model (5) is

equivalent to IQP model (2)-(4), formulation (A.1a) is also equivalent to IQP and variables t_{ij} have in fact the same meaning as in IQP. Consequently, LP relaxation of (A.1a) is also an LP relaxation for IQP. Experimentally we found it to be tighter then the relaxation constructed by Gurobi optimizer, while the use of a specialized solver allows to handle it efficiently.

Appendix B. Handling Cycles

Enforcing that the solution, given by the selected edges $X \subseteq E$, is a tree or a forest, *i.e.* is free of cycles, makes the problem significantly more difficult. Assume $S \subseteq X$ is a cycle, then one can add the constraint

$$\sum_{(i,j)\in S} t_{ij} \leq |S| - 1, \tag{B.1}$$

which enforces that not all edges may be selected simultaneously in cycle S, known as *subtour elimination* in the context of the traveling salesman problem (TSP)[6] or *packing* constraint. A cutting plane approach can then be used to find a cycle-free solution (adding such constraints to the problem as long as the solution contains cycles).

Another possibility proposed by (Türetken et al., 2013) is to apply flow constraints similarly to (Duhamel et al., 2008), essentially enforcing that there is a path from the root node (in our case the virtual root node) to every other node. The existence of a path is formulated using flows. In total, there are $|E||V|$ additional flow variables, $|E|$ for every node. Unlike packing constraints, flow constraints are only polynomially many. However, for large problems, introducing all of them from the start would lead to a prohibitively large problem to solve.

In the MRF formulation, we could enforce subtour elimination constraint (B.1) by adding a cycle subproblem in the dual decomposition manner, *c.f.* (Wang and Koller, 2013) with one forbidden configuration. This however leaves open a problem of rounding a relaxed solution to a cycle-free integer one, and thus does not fully resolve the issue.

Instead, we propose the following heuristic to obtain a good approximate solution. It can be applied with both IQP and MRF formulations. When a solution is found to contain a cycle, we detect the edge in the cycle with the worst radius ratio (indicating an edge going in the wrong direction), *i.e.* maximize the ratio of $r^{start}(p_{ij})$ and $r^{end}(p_{ij})$, the radii at the start and end of the tubular path p_{ij}, respectively, as defined in Sec. 3.1.3. The cycle is cut by excluding this edge from the graph, after that the problem can be resolved. This is obviously a greedy procedure, however, since we do not expect cycles to occur often, it is reasonable for our problem. In fact, it is intuitively similar to what a human operator would do if he/she is presented a cycle for correction.

[6] TSP is known to be one of the most difficult optimization problems. In particular, there is no polynomial time approximation algorithm with a sub-exponential approximation factor unless P = NP.

References

Agam, G., Armato III, S.G., Wu, C., 2005. Vessel Tree Reconstruction in Thoracic CT Scans With Application to Nodule Detection. IEEE Transactions on Medical Imaging 24, 486–499.

Aurenhammer, F., 1991. Voronoi Diagrams – A Survey of a Fundamental Data Structure. ACM Computing Surveys 23, 345–405.

van Bemmel, C.M., Spreeuwers, L.J., Viergever, M.A., Niessen, W.J., 2003. Level-set-based artery-vein separation in blood pool agent CE-MR angiograms. IEEE Transactions on Medical Imaging 22, 1224–1234.

Benmansour, F., Türetken, E., Fua, P., 2013. Tubular Geodesics using Oriented Flux: An ITK Implementation. The Insight Journal - 2013 January-December , 1–8.

Bredies, K., Kunisch, K., Pock, T., 2010. Total Generalized Variation. SIAM Journal of Imaging Sciences 3, 492–526.

Bülow, T., Wiemker, R., Blaffert, T., Lorenz, C., Renisch, S., 2005. Automatic extraction of the pulmonary artery tree from multi-slice CT data, in: Proc. SPIE 5746, Med. Imaging Physiol. Funct. Struct. from Med. Images, pp. 730–740.

Charbonnier, J.P., Brink, M., Ciompi, F., Scholten, E., Schaefer-Prokop, C., van Rikxoort, E., 2016. Automatic Pulmonary Artery-Vein Separation and Classification in Computed Tomography Using Tree Partitioning and Peripheral Vessel Matching. IEEE Transactions on Medical Imaging 35, 882–892.

Duhamel, C., Gouveia, L., Moura, P., de Souza, M.C., 2008. Models and heuristics for a minimum arborescence problem. Networks 51, 34–47.

Estépar, R.S.J., Kinney, G.L., Black-Shinn, J.L., Bowler, R.P., Kindlmann, G.L., Ross, J.C., Kikinis, R., Han, M.K., Come, C.E., Diaz, A.A., Cho, M.H., Hersh, C.P., Schroeder, J.D., Reilly, J.J., Lynch, D.A., Crapo, J.D., Wells, J.M., Dransfield, M.T., Hokanson, J.E., Washko, G.R., 2013. Computed tomographic measures of pulmonary vascular morphology in smokers and their clinical implications. Am. J. Respir. Crit. Care Med. 188, 231–239.

Frangi, A.F., Niessen, W.J., Hoogeveen, R.M., van Walsum, T., Viergever, M.A., 1999. Model-based quantitation of 3-D magnetic resonance angiographic images. IEEE Transactions on Medical Imaging 18, 946–956.

Gao, Z., Grout, R.W., Holtze, C., Hoffman, E.A., Saha, P.K., 2012. A new paradigm of interactive artery/vein separation in noncontrast pulmonary CT imaging using multiscale topomorphologic opening. IEEE Transactions on Biomedical Engineering 59, 3016–3027.

Helmberger, M., Pienn, M., Urschler, M., Kullnig, P., Stollberger, R., Kovacs, G., Olschewski, A., Olschewski, H., Bálint, Z., 2014. Quantification of tortuosity and fractal dimension of the lung vessels in pulmonary hypertension patients. PLoS One 9, e87515.

Hoeper, M.M., Bogaard, H.J., Condliffe, R., Frantz, R., Khanna, D., Kurzyna, M., Langleben, M., Manes, A., Satoh, T., Torres, F., Wilkins, M.R., Badesch, D.B., 2013. Definitions and diagnosis

of pulmonary hypertension. Journal of the American College of Cardiology 62(25 SUPPL), D42–D50.

Jin, D., Iyer, K.S., Hoffman, E.A., Saha, P.K., 2014. Automated Assessment of Pulmonary Arterial Morphology in Multi-row Detector CT Imaging Using Correspondence with Anatomic Airway Branches, in: Proceedings 10th International Symposium Advances in Visual Computing, ISVC 2014, Springer LNCS 8887. pp. 521–530.

Joshi, V.S., Reinhardt, J.M., Garvin, M.K., Abramoff, M.D., 2014. Automated Method for Identification and Artery-Venous Classification of Vessel Trees in Retinal Vessel Networks. PLoS One 9, e88061.

Karp, R.M., 1972. Reducibility among combinatorial problems, in: Proceedings of a Symposium on the Complexity of Computer Computations, The IBM Research Symposia Series, pp. 85–103.

Kasahara, Y., Tuder, R.M., Cool, C.D., Lynch, D.A., Flores, S.C., Voelkel, N.F., 2001. Endothelial cell death and decreased expression of vascular endothelial growth factor and vascular endothelial growth factor receptor 2 in emphysema. Am. J. Respir. Crit. Care Med. 163, 737–744.

Kirbas, C., Quek, F., 2004. A review of vessel extraction techniques and algorithms. ACM Computing Surveys 36, 81–121.

Kitamura, Y., Li, Y., Ito, W., Ishikawa, H., 2013. Adaptive higher-order submodular potentials for pulmonary artery-vein segmentation, in: Proc. Fifth International Workshop on Pulmonary Image Analysis, pp. 53–61.

Kitamura, Y., Li, Y., Ito, W., Ishikawa, H., 2016. Data-Dependent Higher-Order Clique Selection for Artery–Vein Segmentation by Energy Minimization. International Journal of Computer Vision 117, 142–158.

Kolmogorov, V., 2006. Convergent tree-reweighted message passing for energy minimization. IEEE Transactions on Pattern Analysis and Machine Intelligence 28, 1568–1583.

Krissian, K., Malandain, G., Ayache, N., Vaillant, R., Trousset, Y., 2000. Model Based Detection of Tubular Structures in 3D Images. Computer Vision and Image Understanding 80, 130–171.

Law, M.W.K., Chung, A.C.S., 2008. Three dimensional curvilinear structure detection using optimally oriented flux, in: European Conference on Computer Vision (ECCV), pp. 368–382.

Lei, T., Udupa, J.K., Saha, P.K., Odhner, D., 2001. Artery-Vein Separation via MRA - An Image Processing Approach. IEEE Transactions on Medical Imaging 20, 689–703.

Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G., 2009. A review of 3D vessel lumen segmentation techniques: Models, features and extraction schemes. Medical Image Analysis 13, 819–845.

Linguraru, M.G., Pura, J.A., Van Uitert, R.L., Mukherjee, N., Summers, R.M., Minniti, C., Gladwin, M.T., Kato, G., Machado, R.F., Wood, B.J., 2010. Segmentation and quantification of pulmonary artery for noninvasive CT assessment of sickle cell secondary pulmonary hypertension. Medical Physics 37, 1522–1532.

Masutani, Y., MacMahon, H., Doi, K., 2002. Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis. IEEE Transactions on Medical Imaging 21, 1517–1523.

Matsuoka, S., Washko, G.R., Dransfield, M.T., Yamashiro, T., San Jose Estepar, R., Diaz, A., Silverman, E.K., Patz, S., Hatabu, H., 2010a. Quantitative CT measurement of cross-sectional area of small pulmonary vessel in COPD: correlations with emphysema and airflow limitation. Academic Radiology 17, 93–99.

Matsuoka, S., Washko, G.R., Yamashiro, T., San Jose Estepar, R., Diaz, A., Silverman, E.K., Hoffman, E., Fessler, H.E., Criner, G.J., Marchetti, N., Scharf, S.M., Martinez, F.J., Reilly, J.J., Hatabu, H., 2010b. Pulmonary Hypertension and Computed Tomography Measurement of Small Pulmonary Vessels in Severe Emphysema. Am. J. Respir. Crit. Care Med. 181, 218–225.

McNeil, K., Dunning, J., 2007. Chronic thromboembolic pulmonary hypertension (CTEPH). Heart 93, 1152–1158.

Murphy, K., van Ginneken, B., Schilham, A.M.R., de Hoop, B.J., Gietema, H.A., Prokop, M., 2009. A large-scale evaluation of automatic pulmonary nodule detection in chest CT using local image features and k-nearest-neighbour classification. Medical Image Analysis 13, 757–770.

Park, S., Bajaj, C., Gladish, G., 2006. Artery-vein separation of human vasculature from 3D thoracic CT angio scans, in: Proc Int Symp CompIMAGE 2006, pp. 331–336.

Park, S., Lee, S.M., Kim, N., Seo, J.B., Shin, H., 2013. Automatic reconstruction of the arterial and venous trees on volumetric chest CT. Medical Physics 40, 071906.

Payer, C., Pienn, M., Balint, Z., Olschewski, A., Olschewski, H., Urschler, M., 2015. Automatic artery-vein separation from thoracic CT images using integer programming. Medical Image Computing and Computer Assisted Intervention (MICCAI) Springer LNCS 9350, 36–43.

Rempfler, M., Schneider, M., Ielacqua, G.M., Xiao, X., Stock, S.R., Klohs, J., Szekely, G., Andres, B., Menze, B.H., 2015. Reconstructing cerebrovascular networks under local physiological constraints by integer programming. Medical Image Analysis 25, 86–94.

van Rikxoort, E.M., van Ginneken, B., 2013. Automated segmentation of pulmonary structures in thoracic computed tomography scans: a review. Physics in Medicine and Biology 58, R187–R220.

Robben, D., Türetken, E., Sunaert, S., Thijs, V., Wilms, G., Fua, P., Maes, F., Suetens, P., 2014. Simultaneous Segmentation and Anatomical Labeling of the Cerebral Vasculature. Medical Image Computing and Computer Assisted Intervention (MICCAI) Springer LNCS 8673, 307–314.

Rudyanto, R.D., Kerkstra, S., van Rikxoort, E.M., Fetita, C., Brillet, P.Y., Lefevre, C., Xue, W., Zhu, X., Liang, J., Öksüz, I., Ünay, D., Kadipasaoglu, K., Estépar, R.S.J., Ross, J.C., Washko, G.R., Prieto, J.C., Hoyos, M.H., Orkisz, M., Meine, H., Hüllebrand, M., Stöcker, C., Mir, F.L., Naranjo, V., Villanueva, E., Staring, M., Xiao, C., Stoel, B.C., Fabijanska, A., Smistad, E., Elster, A.C., Lindseth, F., Foruzan, A.H., Kiros, R., Popuri, K., Cobzas, D., Jimenez-Carretero, D., Santos, A., Ledesma-Carbayo, M.J., Helmberger, M., Urschler, M., Pienn, M., Bosboom, D.G.H., Campo, A., Prokop, M., de Jong, P.A., Ortiz-de Solorzano, C., Muñoz-Barrutia, A., van Ginneken, B., 2014. Comparing algorithms for automated vessel segmentation in computed tomography scans of the lung: the VESSEL12 study. Medical Image Analysis 18, 1217–32.

Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J., 1999. Nonrigid Registration Using Free-Form Deformations: Application to Breast MR Images. IEEE Transactions on Medical Imaging 18, 712–721.

Saha, P.K., Gao, Z., Alford, S.K., Sonka, M., Hoffman, E.A., 2010. Topomorphologic separation of fused isointensity objects via multiscale opening: Separating arteries and veins in 3-D pulmonary CT. IEEE Transactions on Medical Imaging 29, 840–851.

Santos, S., Peinado, V.I., Ramirez, J., Morales-Blanhir, J., Bastos, R., Roca, J., Rodriguez-Roisin, R., Barbera, J.A., 2003. Enhanced expression of vascular endothelial growth factor in pulmonary arteries of smokers and patients with moderate chronic obstructive pulmonary disease. Am. J. Respir. Crit. Care Med. 167, 1250–1256.

Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., Kikinis, R., 1998. Three-dimensional multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. Medical Image Analysis 2, 143–168.

Shekhovtsov, A., Swoboda, P., Savchynskyy, B., 2015. Maximum Persistency via Iterative Relaxed Inference with Graphical Mod-

els, in: Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 521–529.

Sluimer, I., Schilham, A., Prokop, M., van Ginneken, B., 2006. Computer Analysis of Computed Tomography Scans of the Lung: A Survey. IEEE Transactions on Medical Imaging 25, 385–405.

Stoecker, C., Welter, S., Moltz, J.H., Lassen, B., Kuhnigk, J.M., Krass, S., Peitgen, H.O., 2013. Determination of lung segments in computed tomography images using the Euclidean distance to the pulmonary artery. Medical Physics 40, 091912.

Sugiura, T., Tanabe, N., Matsuura, Y., Shigeta, A., Kawata, N., Juio, T., Yanagawa, N., Sakao, S., Kasahara, Y., Tatsumi, K., 2013. Role of 320-Slice CT Imaging in the Diagnostic Workup of Patients With Chronic Thromboembolic Pulmonary Hypertension. Chest 143, 1070–1077.

Tawhai, M.H., Clark, A.R., Burrowes, K.S., 2011. Computational models of the pulmonary circulation: Insights and the move towards clinically directed studies. Pulmonary Circulation 1, 224–238.

Tizon, X., Smedby, O., 2002. Segmentation with gray-scale connectedness can separate arteries and veins in MRA. Journal of Magnetic Resonance Imaging 15, 438–445.

Türetken, E., Benmansour, F., Andres, B., Pfister, H., Fua, P., 2013. Reconstructing loopy curvilinear structures using integer programming, in: Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1822–1829.

Wang, H., Koller, D., 2013. A fast and exact energy minimization algorithm for cycle MRFs., in: Proc. Int. Conference on Machine Learning (ICML), pp. 190–198.

Zhou, C., Chan, H.P., Sahiner, B., Hadjiiski, L.M., Chughtai, A., Patel, S., Wei, J., Ge, J., Cascade, P.N., Kazerooni, E.A., 2007. Automatic multiscale enhancement and segmentation of pulmonary vessels in CT pulmonary angiography images for CAD applications. Medical Physics 34, 4567–4577.