

Confluent Vessel Trees with Accurate Bifurcations

Zhongwen Zhang¹ Dmitrii Marin^{1,2} Maria Drangova³ Yuri Boykov^{1,2}

¹University of Waterloo, Canada ²Vector Research Institute, Canada ³Robarts Research, Canada

Abstract

We are interested in unsupervised reconstruction of complex near-capillary vasculature with thousands of bifurcations where supervision and learning are infeasible. Unsupervised methods can use many structural constraints, e.g. topology, geometry, physics. Common techniques use variants of MST on geodesic tubular graphs minimizing symmetric pairwise costs, i.e. distances. We show limitations of such standard undirected tubular graphs producing typical errors at bifurcations where flow “directedness” is critical. We introduce a new general concept of confluence for continuous oriented curves forming vessel trees and show how to enforce it on discrete tubular graphs. While confluence is a high-order property, we present an efficient practical algorithm for reconstructing confluent vessel trees using minimum arborescence on a directed graph enforcing confluence via simple flow-extrapolating arc construction. Empirical tests on large near-capillary sub-voxel vasculature volumes demonstrate significantly improved reconstruction accuracy at bifurcations. Our code has also been made publicly available ¹.

1. Introduction

This paper is focused on unsupervised vessel tree estimation in large volumes containing numerous near-capillary vessels and thousands of bifurcations, see Figs. 1, 10. Around 80% of the vessels in such data have sub-voxel diameter resulting in *partial volume* effects such as contrast loss and gaps. Besides the topological accuracy of trees reconstructed from such challenging imagery, we are particularly interested in the accurate estimation of bifurcations due to their importance in biomedical and pharmaceutical research.

1.1. Unsupervised vasculature estimation methods

Unsupervised vessel tree estimation methods for complex high-resolution volumetric vasculature data

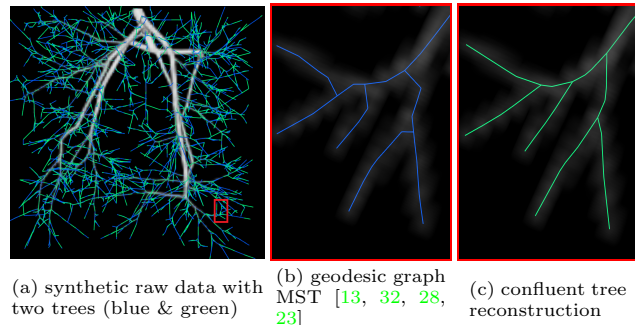


Figure 1: *Synthetic example*: (a) raw 3D data with blue & green reconstructed trees, see also zoom-ins (b,c). The blue tree (a,b) is an MST on a geodesic tubular graph. The green tree (a,c), is a *minimum arborescence* on a directed *confluent tubular graph*, see Sec. 3.

combine low-level vessel filtering and algorithms for computing global tree structures based on constraints from anatomy, geometry, physics, *etc.* Below we review the most relevant standard methodologies.

Low-level vessel estimation: Anisotropy of tubular structures is exploited by standard vessel filtering techniques, e.g. Frangi *et al.* [9]. Combined with non-maximum suppression, local tubularity filters provide estimates for vessel centerline points and tangents, see Fig. 2(a). Technically, elongated structures can be detected using intensity Hessian spectrum [9], *optimally oriented flux* models [18, 28], steerable filters [10], path operators [22] or other anisotropic models. Dense local vessel detections can be denoised using curvature regularization [24, 20]. Prior knowledge about divergence or convergence of the vessel tree (arteries vs veins) can also be exploited to estimate an *oriented* flow pattern [34], see Fig. 2(b).

Thinning: One standard approach to vessel topology estimation is via *medial axis* [26]. This assumes known vessel segmentation (volumetric mask) [21], which can be computed only for relatively thick vessels. Well-formulated segmentation of thin structures requires Gaussian- or min-curvature surface regularization that has no known practical algorithms. Segmen-

¹<https://vision.cs.uwaterloo.ca/code>.




Figure 2: *Low-level vessel estimation*: True centerline is black. Blue voxels in (a) are local maxima of some tubularity measure [9, 18, 28, 10] in the direction orthogonal to the estimated centerline tangents (red). Regularization [24, 20] can estimate subpixel centerline points (b) and oriented tangents [34] (red flow field).




Figure 3: *Global vessel tree reconstruction*: (a) geodesic tubular graph is based on low-level estimates in Fig. 2. Graph edges represent distances, geodesics, or other symmetric (undirected) properties. MST reconstruction quality (b) depends on the graph construction (nodes, neighborhoods, edge weights).

tation is particularly unrealistic for sub-voxel vessels.

Geodesics and shortest paths: Geodesics [5, 2] and shortest paths [7] are often used for *AB*-interactive reconstruction of vessels between two specified points. A vessel is represented by the shortest path with respect to some anisotropic continuous (Riemannian) or discrete (graph) metric based on a local tubularity measure. Interestingly, the minimum path in an “elevated” search space combining spatial locations and radii can simultaneously estimate the vessel’s centerline and diameter, implicitly representing vessel segmentation [19, 1]. Unsupervised methods widely use geodesics as their building blocks.

Spanning trees: The standard graph concept of a *minimum spanning tree* (MST) is well suited for unsupervised reconstruction of large trees with unknown complex topology [13, 32, 28, 23]. MST is closely related to the *shortest paths* and *geodesics* since its optimality is defined with respect to its length. Like shortest paths, globally optimal MST can be computed very efficiently. In contrast to the shortest paths, MST can reconstruct arbitrarily complex trees without user interaction.

The quality of MST vessel tree reconstruction depends on the underlying graph construction, see Figs. 3 and 7(a). Graphs designed for reconstructing thin tubular structures as their spanning tree (or sub-tree) are often called *tubular graphs*. Typically, the nodes are “anchor” points generated by low-level vessel estimators, *e.g.* see Fig. 2. Such anchors represent sparse [29] or semi-dense [20] samples from the estimated tree structure that may be corrupted by noise and outliers. Pairwise edges on a tubular graph typically represent distances or geodesics between the nodes, as in *AB*-

interactive methods discussed earlier. Such graphs are called *geodesic tubular graphs*, see Fig. 4.

There are numerous variants of tubular graph constructions designed to represent various thin structures as MST [13, 32, 28, 23] or shortest path trees [25]. There are also interesting and useful extensions of MST addressing tubular graph outliers, *e.g.* *k*-MST [30] and integer programming technique in [29]. Such approaches are more powerful as they seek minimum subtrees that can automatically exclude outliers. However, the corresponding optimization problems are NP-hard and require approximations. Such methods are expensive compared to the low-order polynomial complexity of MST. They are not practical for dense reconstruction problems in high-resolution vasculature volumes.

1.2. Motivation and contributions

We are interested in unsupervised reconstruction of large complex trees from vasculature volumes resolving near-capillary details. Common geodesic approaches can not represent asymmetric smoothness at bifurcations, which have forms sensitive to flow orientation. Hence, standard methods produce vessel tree reconstructions with significant bifurcation artifacts, see Figs. 1(b), 3(b) and 7(a). We define a general geometric property for oriented vessels, *confluence*, which is missing in prior art, and propose a practical graph-based reconstruction method enforcing it. The reconstructed confluent vessel trees have significantly better bifurcation accuracy. Our contributions are detailed below.

- We introduce *confluence* as a geometric property for overlapping oriented smooth curves in \mathbb{R}^3 , *e.g.* representing blood-flow trajectories². It is like “co-

²Confluence is known in other contexts, *e.g.* rail tracks [15].

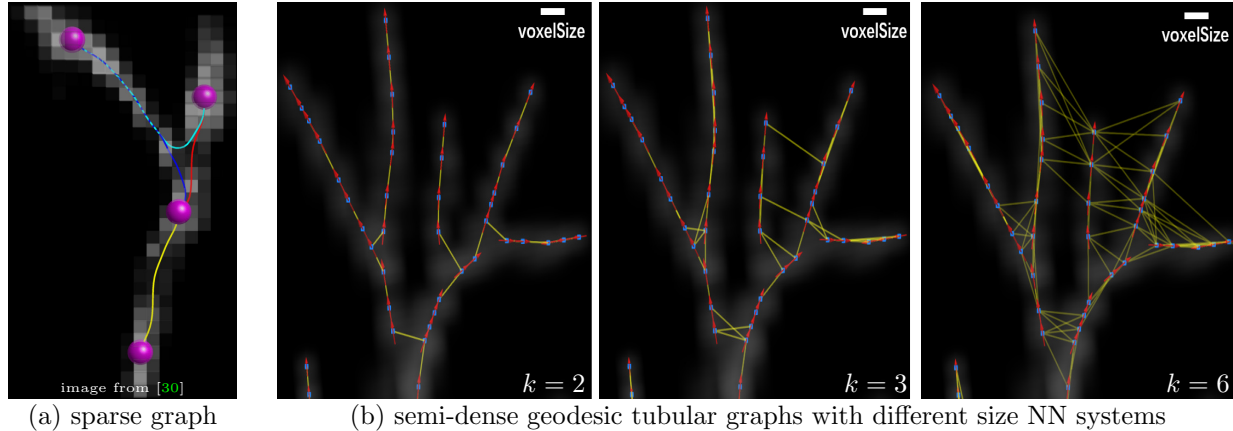


Figure 4: Examples of standard *geodesic tubular graphs* for vascular image data: A graph from [30] in (a) uses (purple) nodes connected by undirected edges corresponding to shortest paths (geodesics) w.r.t. tubularity-based Riemannian metric. Alternatively, (b) shows graphs where (blue) nodes correspond to densely-sampled voxels along the vessels [34]. Here the edges correspond to nearest neighbors (KNN) weighted by length of some spline interpolation. In both cases (a) and (b), near- or sub-voxel vessels have sparsely sampled bifurcations.

differentiability” or “co-continuity”. We define confluent vessel trees formed by overlapping oriented curves.

- We extend confluence to discrete paths and trees on directed tubular graphs where directed arcs/edges represent continuous oriented arcs/curves in \mathbb{R}^3 . We propose a simple *flow-extrapolating circular arc* construction that guarantees ε -confluence, which approximates confluence. Our confluence constraint implies *directed* tubular graph with asymmetric edge weights, which is in contrast to standard *undirected* geodesic tubular graphs [16, 13, 32, 28, 23, 30, 20, 29, 34].

- We present an efficient practical algorithm for reconstructing *confluent vessel trees*. It uses *minimum arborescence* [6, 27] on our directed *confluent tubular graph* construction.

- Our experiments on synthetic and real data confirm that confluent tree reconstruction significantly improves bifurcation accuracy. We demonstrate qualitative and quantitative improvements via standard and new accuracy measures³ evaluating tree structure, bifurcation localization, and bifurcation angles.

Our concept of confluent trees is general and our specific algorithm can be modified or extended in many ways, some of which are discussed in Sec. 3. To explicitly address outliers, minimum arborescence on our confluent tubular graph can be replaced by optimal sub-tree algorithms [30, 29]⁴ or explicit outlier detection [20], but these approximation algorithms address NP-hard problems and maybe too expensive for

³For our dataset and implementation of evaluation metrics discussed in this paper see <https://vision.cs.uwaterloo.ca/data>.

⁴IP solver in [29] uses *minimum arborescence* as a subroutine.




Figure 5: Confluence for oriented curves α and β .

large semi-dense tubular graphs we study in this work. While outlier detection is relevant, this work is not focused on this problem.

2. Confluence of Oriented Curves

This section introduces geometrically-motivated concept of smoothness for objects containing multiple oriented curves, such as vessel trees. We define *confluence* as follows.

DEFINITION 1 (confluence at a point). *Two differentiable oriented curves $\alpha(t)$ and $\beta(\tau)$ are called confluent at a shared point p if for some $k > 0$*

$$\alpha'(t_p) = k \beta'(\tau_p) \quad \text{where}$$

t_p and τ_p are s.t. $\alpha(t_p) = \beta(\tau_p) = p$, see Fig. 5(a).

We will call two oriented curves *confluent* if they are confluent at all points they share, see Fig. 5(b).

Our concept of confluence is closely related to the geometric \mathcal{G}^1 -continuity [4, 8]. A curve α is called \mathcal{G}^1 -continuous if at any point on the curve the slope orientation is continuous. Incidentally, the differentiability

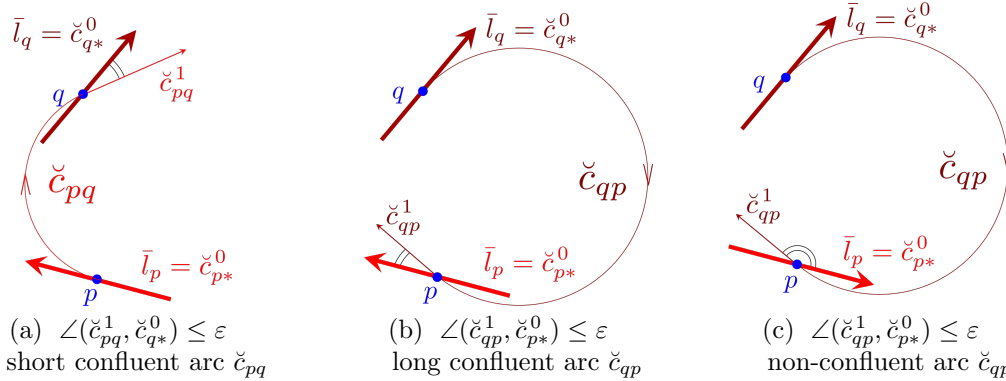


Figure 6: Examples of directed flow-extrapolating circular arcs: (a) flow \check{c}_{pq}^1 extrapolated by *confluent* arc \check{c}_{pq} from p to q is consistent with the local flow estimate $\bar{l}_q = \check{c}_{q*}^0$ at point q as the angle between two vectors is small. Flow extrapolation from q to p in (b) requires another circular arc \check{c}_{qp} that belongs to a different circle defined by tangent \bar{l}_q . Two arcs \check{c}_{pq} and \check{c}_{qp} are not even co-planar if tangents \bar{l}_p and \bar{l}_q are not. The extrapolated flow \check{c}_{qp}^1 is also consistent with the local flow estimate $\bar{l}_p = \check{c}_{p*}^0$ at p , so that the reverse arc \check{c}_{qp} in (b) is confluent as well. In (c) the local flow estimate \bar{l}_p at p is flipped and arc \check{c}_{qp} becomes non-confluent since the angle between \check{c}_{qp}^1 and $\bar{l}_p = \check{c}_{p*}^0$ is large. Note that arc \check{c}_{pq} in (c) differs from (a) but it must be non-confluent as well, see Theorem 1.

classes C^k are too restrictive as a \mathcal{G}^1 -continuous curve can easily be not C^1 due to the curve parameterization. Note that \mathcal{G}^1 -continuity is only defined for a *single* curve while our confluence extends it for a pair of curves and can be seen as “co- \mathcal{G}^1 -continuity”.

Our concept of confluence allows defining arbitrarily complex (continuous) *confluent vessel trees*. Such trees are formed by multiple oriented curves representing motion trajectories of blood particles from the common root to an arbitrary number of leaves where each pair of curves must be confluent. Fig. 5(b) shows a simple example of a tree formed by two confluent curves with one bifurcation, which can be formally defined.

3. Confluent Tubular Graphs

Our discrete approach to reconstructing *confluent vessel trees* is based on efficient algorithms for directed graphs. Our “tubular” graph nodes correspond to a finite set of detected vessel points. We use discrete representation of oriented vessels as paths along directed edges or *directed arcs* (p, q) connecting the graph nodes. Each arc (p, q) on our tubular graph represents an oriented continuous “flow-extrapolating” curve in \mathbb{R}^3 from p to q . Such curves could be obtained from physical models based on fluid dynamics. For simplicity, this paper is focused on oriented *circular arcs*, see Sec. 3.1, motivated as the lowest-order polynomial splines capable of enforcing \mathcal{G}^1 -continuity and confluence. In general, our confluent tubular graph construction can use higher-order flow-extrapolation models, *e.g.* cubic Hermite splines that are common in computer graphics and

geometric modeling of motion trajectories.

By using circular *arcs* as flow-extrapolating curves, we introduce some ambiguity with “arcs” as the standard term for graph edges. However, this should not create confusion since there is a one-to-one relation between *directed arcs* on our tubular graph and the corresponding (circular) *oriented arcs* in \mathbb{R}^3 . Note that both interpretations are oriented/directed. In all technically formal sentences, continuous or discrete interpretation of the “arc” is clear from the context. In more informal settings, both interpretations are often equally valid.

The rest of this Section is as follows. Oriented flow-extrapolating circular arcs between tubular graph nodes are introduced in Sec. 3.1 where ε -confluence constraint is defined in the context of such arcs. We also define directed arc weights to represent the confluence constraint and the local costs of sending flow along these arcs. Geometric properties of confluent circular arcs are discussed in Sec. 3.2. The algorithm estimating confluent vessel trees via *minimum arborescence* on our directed tubular graph is presented in Sec. 3.3.

3.1. Confluent flow-extrapolating arcs

Formally, tubular graph $G = \langle V, A \rangle$ is based on a set of nodes/points V embedded in \mathbb{R}^3 representing semi-densely sampled centerlines of a tubular structure. $A \subseteq V^2$ is a set of directed arcs. For our tubular graph construction, each directed arc $(p, q) \in A$ represents some continuous oriented curve in \mathbb{R}^3 modelling flow-extrapolation from point p to point q , see Fig. 6. As discussed earlier, this paper is focused on oriented cir-

cular arcs as the simplest geometric model that can represent confluent vessels, even though higher-order geometric splines or physics-motivated curves are possible. Our specific construction uses flow-extrapolating circular arcs based on a set of unit vectors $\bar{L} = \{\bar{l}_p\}_{p \in V} \subset S^2$ representing flow direction estimates at the nodes, see Fig. 2(b). Oriented circular arc \check{c}_{pq} is fit into starting point p , its flow orientation estimate \bar{l}_p , and the ending point q , see Fig. 6. Formally, curve \check{c}_{pq} corresponds to a differentiable function

$$\check{c}_{pq} : [0, 1] \rightarrow \mathbb{R}^3$$

traversing points on a circular arc in the plane spanned by p , q , and vector \bar{l}_p so that

$$\check{c}_{pq}(0) = p, \quad \check{c}_{pq}(1) = q, \quad \frac{\check{c}'_{pq}(0)}{\|\check{c}'_{pq}(0)\|} = \bar{l}_p \quad (1)$$

where derivative $\check{c}'_{pq}(s)$ gives an oriented tangent.

For shortness, we define (oriented) unit tangents at the beginning and the end points of any flow-extrapolating arc \check{c} as

$$\check{c}^0 := \frac{\check{c}'(0)}{\|\check{c}'(0)\|}, \quad \check{c}^1 := \frac{\check{c}'(1)}{\|\check{c}'(1)\|}. \quad (2)$$

The definition of arc \check{c}_{pq} in (1) implies $\check{c}_{pq}^0 \equiv \bar{l}_p$ so that tangent \check{c}_{pq}^0 is the same for any arc starting at given point p regardless of its end point q . Thus,

$$\check{c}_{p*}^0 \equiv \bar{l}_p \quad (3)$$

where the star $*$ represents an arbitrary end point. On the other hand, tangent \check{c}_{pq}^1 at the end point q depends on the arc's starting point p . That is, generally,

$$\angle(\check{c}_{pq}^1, \check{c}_{rq}^1) \neq 0 \quad \text{if } p \neq r.$$

ε -Confluence Constraint: To constrain our tubular graph so that all feasible vessel trees are confluent, it suffices to enforce confluence of the arcs at the nodes where they meet. However, our simple flow-extrapolating circular arcs (1) can not be used to enforce the exact confluence. We use some threshold ε to introduce a relaxed version of confluence in Definition 1 for an arbitrary pair of adjacent arcs $\check{c}_{pq}, \check{c}_{qr}$ connecting points p, q and r

$$\angle(\check{c}_{pq}^1, \check{c}_{qr}^0) \leq \varepsilon.$$

In general, this is a high-order (triple clique) constraint. But, property (3) of our flow-extrapolating arc construction shows that the end point of the second arc \check{c}_{qr} is irrelevant. Indeed,

$$\angle(\check{c}_{pq}^1, \check{c}_{qr}^0) = \angle(\check{c}_{pq}^1, \bar{l}_q) \equiv \angle(\check{c}_{pq}^1, \check{c}_{q*}^0)$$

implying that our specific tubular graph construction allows to express confluence as a pairwise constraint

$$\angle(\check{c}_{pq}^1, \check{c}_{q*}^0) \leq \varepsilon \quad (4)$$

for any pair of points p, q . In essence, this becomes a constraint for our flow extrapolating arcs \check{c}_{pq} that can be called confluent if $\angle(\check{c}_{pq}^1, \bar{l}_q) \leq \varepsilon$, see Fig. 6.

To enforce ε -confluence constraint, our tubular graph can simply drop all non-confluent arcs. Thus, any directed vessel tree on our graph will be confluent by construction. This paper explores the simplest approach to reconstructing confluent vessel trees as the *minimum arborescence* on our directed tubular graph. In this case, instead of dropping non-confluent arcs, one can incorporate ε -confluence constraint directly into the cost of the corresponding directed graph arcs

$$w_{pq} := \begin{cases} \text{length}(\check{c}_{pq}) & \text{if } \angle(\check{c}_{pq}^1, \check{c}_{q*}^0) \leq \varepsilon \\ \infty & \text{otherwise.} \end{cases} \quad (5)$$

The reverse edge on our tubular graph has different weight $w_{qp} \neq w_{pq}$ because it corresponds to a different flow extrapolating arc \check{c}_{qp} that has a different length, see Fig. 6(b). As an extension, our approach also allows “elastic” arc weights by adding integral of arc's curvature to its length in (5). It is also possible to impose soft penalties for the discrepancy between the extrapolated flow \check{c}_{pq}^1 and flow estimate $\check{c}_{q*}^0 \equiv \bar{l}_q$ in (5) based on physical, physiological, or other principles.

Note that higher-order (non-circular) extrapolation arcs \check{c}_{pq} can be constructed to fit the flow orientation estimates at both ends exactly, implying an exactly confluent graph. However, some non-trivial physiological constraints have to be imposed on the smoothness/curvature of such (non-circular) confluent arcs which should result in very long curves in cases like Fig. 6(c). Thus, the confluence constraint will manifest itself similarly to the second line in (5).

3.2. Confluence and co-circularity

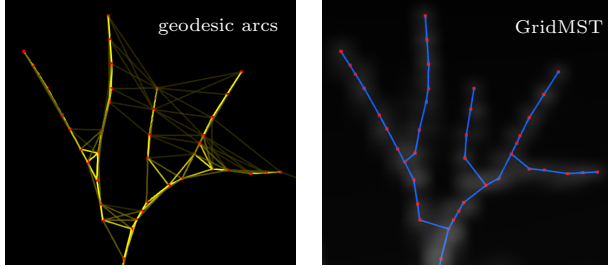
Specifically for oriented circular arcs, confluence implies several interesting properties and can be juxtaposed with the standard concept of *co-circularity* [24]. Assume some circular flow-extrapolating arc \check{c}_{pq} and its reverse \check{c}_{qp} defined by two oriented tangents \bar{l}_p, \bar{l}_q , see Sec. 3.1 and Fig. 6(a,b).

PROPERTY 1. *The angle between \check{c}_{pq}^1 and \bar{l}_q is equal to the angle between \check{c}_{qp}^1 and \bar{l}_p . That is,*

$$\angle(\check{c}_{pq}^1, \check{c}_{qp}^0) \equiv \angle(\check{c}_{qp}^1, \check{c}_{pq}^0). \quad (6)$$


While not immediately obvious, particularly in 3D, this property is not difficult to prove, see supplementary materials or Appendix A in [35]. Identity (6) implies the following.

vessel tree reconstruction using
undirected Geodesic Tubular Graph (standard)



(a) MST (blue) for *Geodesic Tubular Graph* (yellow) [16, 30, 34]

vessel tree reconstructions using
directed Confluent Tubular Graph (our)



(b) Min Arborescence (green) for *Confluent Tubular Graph* (yellow)

Figure 7: Typical tree reconstruction examples for standard geodesic (a) and our *confluent* (b) tubular graphs. Arc weights are represented via thickness (yellow). The data is a (representative) crop with near-capillary vessels at a periphery of large volumes, *e.g.* Fig. 1, 10. Sub-voxel vessels have bifurcations sparsely sampled by tubular graph nodes, as in Fig. 4. MST (blue) on geodesic graph “short-cuts” most bifurcations. *Minimum arborescence* (green) on a directed tubular graph with *confluent* arcs (b), see Sec. 3.1, reconstructs flow-consistent bifurcations.

Theorem 1. *For circular flow extrapolating arcs, \check{c}_{pq} is confluent iff the reverse arc \check{c}_{qp} is confluent.*

This theorem shows that confluence of \check{c}_{pq} and \check{c}_{qp} in Fig. 6(a,b) is not a coincidence. However, in general, such symmetry does not hold for non-circular flow-extrapolating arcs (higher order polynomial curves, etc). Also, Theorem 1 does not imply “undirectedness” of our confluent tubular graph construction using simple circular arcs. As follows from (5), $w_{pq} \neq w_{qp}$ since the reverse arcs \check{c}_{pq} and \check{c}_{qp} have different lengths regardless of confluence, see Fig. 6(a,b).

Interestingly, our confluence constraint in case of *circular* oriented arcs can be related to an “oriented” generalization of *co-circularity* that was originally defined in [24] for 2D curves. In \mathbb{R}^n co-circularity constraint can be defined for two *unoriented* tangent lines l_p and l_q at points p and q in a way similar to our definition of confluence for \check{c}_{pq} that is based on *oriented* tangents \bar{l}_p and \bar{l}_q . Assume *unoriented* circle c_{pq} uniquely defined in \mathbb{R}^n by a pair of points p, q and tangent l_p at the first point. If we use c^x to denote an unoriented unit tangent of circle c at any given point x , then circle c_{pq} is uniquely defined by three conditions

$$c_{pq} \quad : \quad p \in c_{pq}, \quad q \in c_{pq}, \quad c_{pq}^p = l_p. \quad (7)$$

In general, circle c_{qp} is different as it is defined by tangent l_q at point q , that is $c_{qp}^q = l_q$. Then, co-circularity constraint for l_p and l_q can be defined as

$$\angle(c_{pq}^q, c_{qp}^q) \equiv \angle(c_{qp}^p, c_{pq}^p) \leq \epsilon \quad (8)$$

where $\angle(\cdot, \cdot)$ is the angle between two lines in contrast to the angle between vectors in the similar identity (6).

The difference between confluence for \bar{l}_p, \bar{l}_q and co-circularity for l_p, l_q can be illustrated by the examples in Fig. 6. Note that unoriented versions of \bar{l}_p, \bar{l}_q are identical in all three examples (a,b,c) as they do not depend of the flip of orientation in (c). Thus, they are equally co-circular in (a,b,c). At the same time, oriented tangents are confluent in (a,b) while flipping orientation for \bar{l}_p results in non-confluence in (c). The properties discussed above imply that confluence can be seen as oriented generalization of co-circularity [24].

PROPERTY 2. *Confluence of oriented circular arcs \check{c}_{pq} or \check{c}_{qp} , which are defined by oriented tangents \bar{l}_p, \bar{l}_q , implies co-circularity of the corresponding unoriented tangents l_p, l_q , but not the other way around.*

Note that co-circularity constraint for *unoriented* circular arcs along a path on a tubular graph can enforce \mathcal{G}^1 -smoothness within a single vessel branch. But, unoriented co-circularity enforces smoothness indiscriminately in all directions from a bifurcation point without resolving conflicts between multiple branches. This leads to artifacts observed on geodesic tubular graphs, see Fig. 7(a). In contrast, the confluence constraint discriminates orientations of branches when enforcing smoothness at bifurcations, see Fig. 7(b).

3.3. Confluent tree reconstruction algorithm

Our Confluent Tree Reconstruction Algorithm 1 is discussed below. It inputs raw volumetric data with a marked root of the tree. The algorithm has four steps. First, it runs a subroutine that estimates a set of points on the tree centerline V and oriented flow pattern at these points \bar{L} . We use a standard vector field esti-

Algorithm 1 Confluent Tree Reconstruction

Require: Raw volumetric data and root location

- 1: Estimate a set of centerline points V and directed flow estimates $\bar{L} = \{\bar{l}_p | p \in V\}$.
 - 2: Build a set of oriented arcs $A \subseteq V \times V$, Sec. 3.1
 - 3: Build a *confluent tubular graph* G by computing weights w_{pq} for $(p, q) \in A$ using (5).
 - 4: Return the *Minimum Arborescence* of G .
-

mation method [34] based on non-negative divergence constraint and regularization over the voxel-grid neighborhood. Since 80% of our large vasculature volumes are near-capillary vessels, the weak sub-voxel signal often results in missing data points and grid-based regularization fails to produce consistent flow orientations, particularly at the tree periphery. We modified [34] by (anisotropically) enlarging their regularization neighborhood, see supplementary materials or Appendix B in [35], improving the quality of flow estimates \bar{L} that helps to reconstruct confluent vessel trees.

Second, we build a set of oriented arcs between the points in V that correspond to directed edges on our tubular graph G . Our confluence constraint works well even with a complete graph $A = V \times V$. But, for efficiency, we restrict the neighborhood to K nearest neighbors (KNN). The running time is $\mathcal{O}(K|V| \log |V|)$ with k -d trees.

The last two steps compute a directed weight w_{pq} for all arcs $(p, q) \in A$ as described in Sec. 3.1, and invoke a standard minimum arborescence algorithm that has complexity $\mathcal{O}(|A| + |V| \log |V|)$ [11]. In practice, the overall running time of our method for vessel tree reconstruction is dominated by the centerline localization and flow pattern estimation in the first step.

4. Experimental results

We use two baselines which we call NMS-MST and GridMST. NMS-MST uses the Frangi method [9] along with Non-Maximum Suppression (NMS) to obtain the centerline points and local unoriented tangent estimates. On top of these, NMS-MST uses the KNN ($K = 500$) graph of the centerline points to build the MST. Note that the KNN graph is symmetric such that a pair of nodes have an arc as long as one is a neighbor of the other. Here, the undirected edge weight is computed by the sum of two shorter arc lengths. GridMST uses [34] for estimating the centerline points and flow direction. Then, it also uses the KNN graph of the centerline points to build the MST.

GridArb uses the set of centerline points and flow estimates produced by [34], but it uses the confluent tubular graph to build the minimum arborescence (dis-

cussed in Sec. 3.3). MARb exploits a modified version of [34] (see supplementary materials or Appendix B in [35]) and also uses the confluent tubular graph to build the minimum arborescence. We set $\varepsilon = \frac{\pi}{2}$ in (5) for all our experiments.

Method	Flow estimates	Graph Weights	Tree Extraction
NMSMST	Frangi <i>et al.</i> [9]	standard	MST
GridMST	Zhang <i>et al.</i> [34]	geodesic	
GridArb	Zhang <i>et al.</i> [34]	our	minimum
MARb	Modified [34], see suppl materials	confluent (5)	arborescence

4.1. Validation Measures

Many validation measures rely on matching between the ground truth and the predicted tree. Matching algorithms could be separated into several groups. First, match the nodes of the trees independently based on a distance measure, *e.g.* [34, 28, 31], partial (local) subtree matching [12], or global tree matching approaches [17, 33, 3]. We base our evaluation approach on the first group of methods due to their efficiency and the size of our problem.

Centerline reconstruction quality. Our reconstructed tree is ideally the centerline of the vasculature. We compute the recall and fall-out statistics of the centerline points to evaluate the reconstruction quality. To obtain the *centerline receiver operating characteristic (ROC) curve*, we generate a sequence of *recall/fall-out* points by varying the *detection threshold* parameter for the low-level vessel filter of Frangi *et al.* [9].

Similarly to [34], a specific point on the ground truth centerline is considered detected correctly (recall) iff it is located within $\max(r, \zeta)$ distance of a reconstructed tree where r is the radius of the corresponding ground truth vessel segment and $\zeta = \frac{\sqrt{2}}{2}$ voxel-size. A point on the reconstructed tree that is farther away than this distance is considered incorrectly detected (fall-out). Before computing the ROC curve we re-sample uniformly both the ground truth and reconstructed trees.

Bifurcation reconstruction quality. We introduce two separate metrics. First, we compute the ROC curve for only bifurcation points to assess the quality of detection. Second, we measure the median angular error at the reconstructed bifurcations to assess the accuracy, where we match *all* ground truth bifurcations to closest branching points on the detected tree regardless of their proximity and use the median rather than the average for greater stability. The difference between our angular error measure and [34] is discussed in supplementary materials and Appendix C in [35].

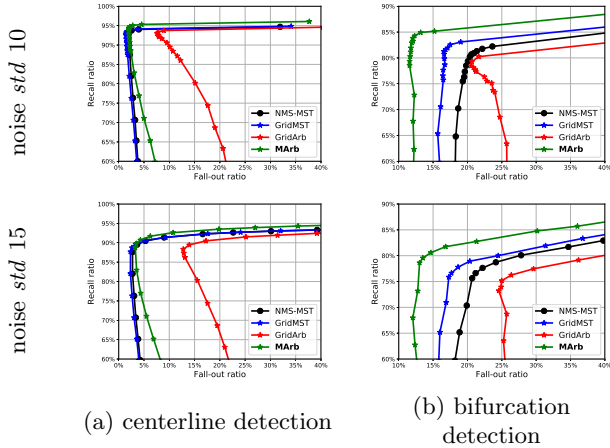


Figure 8: Quantitative comparison. Our methods are denoted by MARb and GridArb. GridMST is the best result from [34].




Figure 9: Median branching angular error for our methods (GridArb and MARb from Sec. 3) outperform all competitors (NMS-MST and GridMST [34]).

4.2. Synthetic Data with Ground Truth

One of the major challenges in large-scale vessel tree reconstruction is the lack of ground truth. That complicates many interactive and supervised learning methods and makes evaluation hard. Zhang *et al.* [34] generated and published a dataset with ground truth using [14]. We used our newly generated 15 volumes (see supplementary materials or Appendix D in [35]) $100 \times 100 \times 100$ with intensities between 0 and 512. Our new dataset has a larger variance of bifurcation angles. The voxel size is 0.046 mm. We add Gaussian noise with *std* 10 and 15.

Fig. 8 compares the results of our methods with two competitors. One is the method of [34], another baseline is simple MST computed over non-maximum suppression of vessel filter output. All methods use essentially the same detection mechanism, *i.e.* Frangi *et al.* filter, so the centerline extraction quality does not differ much. On the other hand, our method significantly outperforms in the quality of bifurcation detection, see Fig. 8 (b). This result is complemented by superior




Figure 10: Real data micro-CT volume and vessel centerline reconstruction (green) obtained by our method and (blue) by [34]. To reduce clutter we show the result of Frangi *et al.* filter [9] instead of raw input data. The real data zoom-in is at the top left.

angular errors in Fig. 9. We attribute this to the sub-voxel accuracy and better reconstruction of bifurcation. A typical example is shown in Fig. 1(b,c).

The GridArb performs competitively in terms of angular errors but gives the worst results in terms of centerline quality. This is due to the artifact caused by some inconsistent flow estimates near the tree periphery (see supplementary materials or Appendix B in [35] for concrete examples). In Sec. 3.3 we argue that enlarging the regularization neighborhood helps improve the estimation of the flow orientation.

4.3. High Resolution Microscopy CT

We use challenging microscopy computer tomography volume (micro-CT) of size $585 \times 525 \times 892$ voxels to qualitatively demonstrate the advantages of our approach. The data is a high-resolution image of mouse heart obtained *ex vivo* with the use of contrast. The resolution allows detecting nearly capillary level vessels, which are partially resolved (partial volume). Fig. 10 shows the whole volume and reconstruction.

References

- [1] Fethallah Benmansour and Laurent D Cohen. Tubular structure segmentation based on minimal path method and anisotropic enhancement. *IJCV*, 92(2):192–210, 2011. [2](#)
- [2] Da Chen, Jean-Marie Mirebeau, and Laurent D Cohen. Global minimum for a finlser elastica minimal path approach. *International Journal of Computer Vision*, 122(3):458–483, 2017. [2](#)
- [3] Egor Chesakov. Vascular tree structure: Fast curvature regularization and validation. *Electronic Thesis and Dissertation Repository. The University of Western Ontario*, (3396), 2015. Master of Science thesis. [7](#)
- [4] Anthony D DeRose. Geometric continuity: a parametrization independent measure of continuity for computer aided geometric design. Technical report, CA Univ Berkeley Dept of Electrical Engineering and Computer Sciences, 1985. [3](#)
- [5] Thomas Deschamps and Laurent D. Cohen. Fast extraction of minimal paths in 3d images and applications to virtual endoscopy. *Medical Image Analysis*, 5(4):281 – 299, 2001. [2](#)
- [6] Jack Edmonds. Optimum branchings. *J. Res. Nat. Bur. Standards*, 71B(4), October- December 1967. [3](#)
- [7] M. A. T. Figueiredo and J. M. N. Leitao. A non-smoothing approach to the estimation of vessel contours in angiograms. *IEEE Transactions on Medical Imaging*, 14(1):162–172, March 1995. [2](#)
- [8] AH Fowler and CW Wilson. Cubic spline: A curve fitting routine. Technical report, Union Carbide Corp., Oak Ridge, Tenn. Y-12 Plant, 1966. [3](#)
- [9] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. Multiscale vessel enhancement filtering. In *MICCAI'98*, pages 130–137. Springer, 1998. [1](#), [2](#), [7](#), [8](#)
- [10] William T. Freeman and Edward H Adelson. The design and use of steerable filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (9):891–906, 1991. [1](#), [2](#)
- [11] Harold N Gabow, Zvi Galil, Thomas Spencer, and Robert E Tarjan. Efficient algorithms for finding minimum spanning trees in undirected and directed graphs. *Combinatorica*, 6(2):109–122, 1986. [7](#)
- [12] Todd A Gillette, Kerry M Brown, and Giorgio A Ascoli. The diadem metric: comparing multiple reconstructions of the same neuron. *Neuroinformatics*, 9(2-3):233, 2011. [7](#)
- [13] Germán González, François Fleuret, and Pascal Fua. Automated delineation of dendritic networks in noisy image stacks. In *European Conference on Computer Vision*, pages 214–227. Springer, 2008. [1](#), [2](#), [3](#)
- [14] Ghassan Hamarneh and Preet Jassi. Vascusynth: simulating vascular trees for generating volumetric image data with ground-truth segmentation and tree analysis. *Computerized medical imaging and graphics*, 34(8):605–616, 2010. [8](#)
- [15] Peter Hui, Michael J Pelsmajer, Marcus Schaefer, and Daniel Stefankovic. Train tracks and confluent drawings. *Algorithmica*, 47(4):465–479, 2007. [2](#)
- [16] Julien Jomier, Vincent LeDigarcher, and Stephen R Aylward. Automatic vascular tree formation using the mahalanobis distance. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 806–812. Springer, 2005. [3](#), [6](#)
- [17] Philip N. Klein. Computing the edit-distance between unrooted ordered trees. pages 91–102, 1998. [7](#)
- [18] Max W K Law and Albert C S Chung. Three dimensional curvilinear structure detection using optimally oriented flux. In *European conference on computer vision*, pages 368–382. Springer, 2008. [1](#), [2](#)
- [19] Hua Li and Anthony Yezzi. Vessels as 4-d curves: Global minimal 4-d paths to extract 3-d tubular surfaces and centerlines. *IEEE transactions on medical imaging*, 26(9):1213–1223, 2007. [2](#)
- [20] Dmitrii Marin, Yuchen Zhong, Maria Drangova, and Yuri Boykov. Thin structure estimation with curvature regularization. In *International Conference on Computer Vision (ICCV)*, 2015. [1](#), [2](#), [3](#)
- [21] Odyssee Merveille, Benoît Naegel, Hugues Talbot, and Nicolas Passat. nD variational restoration of curvilinear structures with prior-based directional regularization. *IEEE Transactions on Image Processing*, 28(8):3848–3859, 2019. [1](#)
- [22] Odyssee Merveille, Hugues Talbot, Laurent Najman, and Nicolas Passat. Curvilinear structure analysis by ranking the orientation responses of path operators. *IEEE transactions on pattern analysis and machine intelligence*, 40(2):304–317, 2017. [1](#)
- [23] Stefano Moriconi, Maria A Zuluaga, H Rolf Jäger, Parashkev Nachev, Sébastien Ourselin, and M Jorge Cardoso. Inference of cerebrovascular topology with geodesic minimum spanning trees.

- IEEE transactions on medical imaging*, 38(1):225–239, 2018. 1, 2, 3
- [24] Pierre Parent and Steven W Zucker. Trace inference, curvature consistency, and curve detection. *PAMI*, 11:823–839, 1989. 1, 2, 5, 6
- [25] Hanchuan Peng, Fuhui Long, and Gene Myers. Automatic 3d neuron tracing using all-path pruning. *Bioinformatics*, 27(13):i239–i247, 2011. 2
- [26] Kaleem Siddiqi and Stephen Pizer. *Medial representations: mathematics, algorithms and applications*, volume 37. Springer Science & Business Media, 2008. 1
- [27] R. E. Tarjan. Finding optimum branchings. *Networks*, 7(1):25–35, 1977. 3
- [28] Engin Turetken, Carlos Becker, Przemyslaw Glowacki, Fethallah Benmansour, and Pascal Fua. Detecting irregular curvilinear structures in gray scale and color imagery using multi-directional oriented flux. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1553–1560, 2013. 1, 2, 3, 7
- [29] Engin Turetken, Fethallah Benmansour, Bjoern Andres, Przemyslaw Glowacki, Hanspeter Pfister, and Pascal Fua. Reconstructing curvilinear networks using path classifiers and integer programming. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(12):2515–2530, December 2016. 2, 3
- [30] Engin Turetken, German Gonzalez, Christian Blum, and Pascal Fua. Automated reconstruction of dendritic and axonal trees by global optimization with geometric priors. *Neuroinformatics*, 9(2-3):279–302, 2011. 2, 3, 6
- [31] J. A. Tyrrell, E. di Tomaso, D. Fuja, R. Tong, K. Kozak, R. K. Jain, and B. Roysam. Robust 3-d modeling of vasculature imagery using superellipsoids. *IEEE Transactions on Medical Imaging*, 26(2):223–237, Feb 2007. 7
- [32] Jun Xie, Ting Zhao, Tzumin Lee, Eugene Myers, and Hanchuan Peng. Automatic neuron tracing in volumetric microscopy images with anisotropic path searching. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 472–479. Springer, 2010. 1, 2, 3
- [33] Kaizhong Zhang and Dennis Shasha. Simple fast algorithms for the editing distance between trees and related problems. *SIAM journal on computing*, 18(6):1245–1262, 1989. 7
- [34] Zhongwen Zhang, Dmitrii Marin, Egor Chesakov, Marc Moreno Maza, Maria Drangova, and Yuri Boykov. Divergence prior and vessel-tree reconstruction. In *IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, Long Beach, California, June 2019. 1, 2, 3, 6, 7, 8
- [35] Zhongwen Zhang, Dmitrii Marin, Maria Drangova, and Yuri Boykov. Confluent vessel trees with accurate bifurcations. In *arXiv:2103.14268*, 2021. 5, 7, 8