

Tracing Retinal Blood Vessels by Matrix-Forest Theorem of Directed Graphs

Li Cheng^{1,3}, Jaydeep De^{1,2}, Xiaowei Zhang¹, Feng Lin², and Huiqi Li⁴

¹ Bioinformatics Institute, A*STAR, Singapore

² School of Computer Engineering, Nanyang Technological University, Singapore

³ School of Computing, National University of Singapore, Singapore

⁴ Beijing Institute of Technology, China

Abstract. This paper aims to trace retinal blood vessel trees in fundus images. This task is far from being trivial as the *crossover* of vessels are commonly encountered in image-based vessel networks. Meanwhile it is often crucial to separate the vessel tree structures in applications such as diabetic retinopathy analysis. In this work, a novel directed graph based approach is proposed to cast the task as label propagation over directed graphs, such that the graph is to be partitioned into disjoint sub-graphs, or equivalently, each of the vessel trees is traced and separated from the rest of the vessel network. Then the tracing problem is addressed by making novel usage of the matrix-forest theorem in algebraic graph theory. Empirical experiments on synthetic as well as publicly available fundus image datasets demonstrate the applicability of our approach.

1 Introduction

Topological and geometrical properties of retinal blood vessels in fundus images can provide valuable clinical information in diagnosing diseases. In particular, vascular anomaly in retina is one of the clinical manifestations of retinal diseases such as diabetic retinopathy, glaucoma, and hypertensive retinopathy. Take diabetic retinopathy as an example, it is a leading cause of blindness in the working-age population of most developed countries. Diabetic retinopathy is the result of progressive damage to the network of tiny blood vessels that supply blood to the retina. Proliferative diabetic retinopathy is specifically characterized by the formation of newly formed vessels in the retina [1]. This thus requires the description of blood vessel tree structure in clinical diagnosis, and as a result, calls for the tracing and separation of each vessel tree in the fundus images.

Existing efforts in retinal vessel analysis can be roughly categorized into segmentation-based and tracking-based. The segmentation-based methods (e.g. [2]) often use pixel classification to produce a binary segmentation (i.e a pixel is classified into vessel or non-vessel). The tracking-based methods (e.g. [3]) usually start with a seed, and track the intended vessel based on local intensity or texture information. Segmentation-based methods tend to produce many disconnected and isolated segments, less favourable for retaining the important




Fig. 1. Overview of our tracing pipeline where our focus is the tracing step.

topological properties of vessel networks. Meanwhile vessel tracking methods often better preserve the connectivity structure of vessel segments. Nonetheless they encounter great difficulties with the occurrence of *crossover* at the junction points. Current methods often fail to trace properly, as it is nontrivial to predict whether the vessel segments at a junction point belong to one tree or multiple trees, and for the later case, to which tree each segment belongs. In this paper, we dedicate our attention to addressing this bottleneck *crossover* issue.

We consider a different tracing approach that can take into account both local and global contextual information of the vessel network, as summarized below: After initial pixel-based segmentation and skeleton extraction, a novel *directed graph* (or digraph) representation is formed, where each segment in the skeleton map becomes a node, and a direct contact between two adjacent segments corresponds to an edge of the two corresponding nodes. The segments in the skeleton map touching the optic disk area are considered as the root nodes. The number of trees to-be-found in the the vessel network thus equals the number of root nodes. The tracing problem is now formulated as label propagation on directed graphs or digraphs: The goal is to propagate tree labels from known root nodes to the rest of the graph, such that the digraph will be split into disjoint sub-graphs, which corresponds to trees of the vessel network. This allows us to consider and make novel usage of the recent development of matrix-forest theorem [4] studied in algebraic digraph theory.

In term of major contributions, our approach offers a principled way of addressing the tracing with crossover problem. By connecting to the well-established algebraic directed graph theory [4], local and global contextual information can be both considered explicitly. We expect the digraph theoretical representation can open the door to some insightful understanding of the characteristics of crossover sections in vessel networks. Finally, our algorithm is also simple and easy to implement.

2 Our Approach

The problem of vessel tracing is to trace blood vessels by separating them into disjoint vessel trees, each starting from a unique root segment within the optic disk, as illustrated in red dots at Figure 1(b). Figure 1 describes the pipeline

of our approach that consists of *two* main steps: The *preprocessing* step mainly consists of segmentation, skeleton extraction, and digraph construction; The *tracing* step then focuses on digraph-based label propagation using Matrix-forest theorem — the main focus of this paper.

2.1 Preprocessing

The preprocessing step is comprised of the following three modules: *Segmentation*: As illustrated in Figure 1 (a)→(b), an input retinal image is segmented into a binary image, with vessel pixels being foreground and the rest as background. Note this step is skipped for synthetic retinal images as they are already binary images. *Skeleton map*: Build a skeleton map from the binary image, and remove the optic disk area as marked within red ellipse in Figure 1(b). The cusps attached to the removed optic disk are the tips of root segments, which are also presented as red dots in Figure 1(c), a zoom-in subset of (b). *Skeleton to digraph*: A segment is defined in the skeleton as the group of connected pixels that ends in either a junction or a tip. This segment corresponds to a node in the resulting digraph, as shown in Figure 1(c)→(d). Two nodes are then linked with a directed edge, when the two coinciding segments from the skeleton map contact and satisfy the ordering criteria, a modification of the well-known stream ordering method from the hydrology community (http://en.wikipedia.org/wiki/Strahler_number). More details can be found at our previous work [5].

2.2 Tracing by Matrix-Forest Theorem on Digraphs

The tracing problem becomes that of separating the vessel trees with only tree roots known, which can be equivalently formulated as a digraph-based label propagation problem with one labeled node per class (vessel tree). That is, all the source nodes are labeled in this problem, each with a unique label (tree), and the task is to make predictions on the remaining unlabeled nodes by propagating the class labels following the underlying digraph structure.

Problem Set-up Let $G = (\mathcal{V}, \mathcal{E}, W)$ denote a digraph, where $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ is the set of nodes, \mathcal{E} the set of directed edges each connecting two adjacent nodes, and $W = [w_{ij}] \in \mathbb{R}^{n \times n}$ the asymmetric non-negative matrix with $w_{ij} \geq 0$ being the weight of the directed edge from v_i to v_j . The out-degree of each node v_i is computed as $d_i^+ = \sum_{j=1}^n w_{ij}$. Denote the out-degree matrix $D = \text{Diag}(d_1^+, \dots, d_n^+)$, that is, $D = \text{Diag}(W\mathbf{1})$, with $\mathbf{1}$ an all-one column vector. Define the digraph Laplacian $L = D - W$. A row-stochastic transition probability matrix $P = [p_{ij}]$ can be constructed as $p_{ij} = \frac{w_{ij}}{d_i^+}$, or equivalently as $P = D^{-1}W$. Note undirected graphs can be regarded as special digraphs characterized algebraically by their symmetric weight matrix W , i.e. the symmetric pair w_{ij} & w_{ji} correspond to bi-directional edges with equal weights. We focus here on a transductive inference scenario where labels from the set of few labeled nodes \mathcal{V}_l are to be propagated to the rest unlabeled nodes \mathcal{V}_u , with $\mathcal{V} = \mathcal{V}_l \cup \mathcal{V}_u$. The labels are multiclass, each corresponds to a separate vessel tree. To simplified the notation we assume \mathcal{V}_l contains the first l nodes, $\mathcal{V}_l = \{v_1, \dots, v_l\}$. To accommodate label information, define a label matrix Y of size $n \times K$ (assuming there are K class

labels available), with each entry Y_{ik} containing 1 if node i belongs to \mathcal{V}_l and is labeled with class k , and 0 otherwise. Also define the length n ground-truth label vector \mathbf{y} that includes two disjoint parts \mathbf{y}_l and \mathbf{y}_u : \mathbf{y}_l is the input label vector of length l over the set of labeled nodes, with each entry y_i for the input class assignment of node $v_i \in \mathcal{V}_l$; \mathbf{y}_u is the hold-out ground-truth label for the unlabeled nodes, i.e. a vector of length $n - l$. Similarly, define the initial label vector $\hat{\mathbf{y}}$ containing also two parts, $\hat{\mathbf{y}}_l := \mathbf{y}_l$ and $\hat{\mathbf{y}}_u = \mathbf{0}$, where $\mathbf{0}$ is an all-zero vector of length $n - l$. Define the prediction vector \mathbf{y}^* with also two parts $\mathbf{y}_l^* := \mathbf{y}_l$, as well as \mathbf{y}_u^* of length $n - l$, containing the prediction results, where each y_i^* denotes the predicted class assignment for a node $v_i \in \mathcal{V}_u$.

The proposed label propagation algorithm (shown in Algorithm 1 and referred to as MFTD) is derived based on matrix-forest theorem [4] of algebraic digraph theory [6], as follows. Let w_{max} denote the entry in W containing the strongest signal, i.e. $w_{max} = \max_{i,j} |w_{ij}|$. The forest matrix is defined as

$$S_1 := (I + \alpha L)^{-1}, \quad (1)$$

a normalized forest matrix where each (i, j) -th entry denotes the number of trees rooted at node i that also include the j -th node, as in Theorem 4 of [4]. It can be viewed as a generalization of the celebrated matrix-tree theorem (e.g. [7]) for undirected graphs to digraphs. Further, let $\tilde{L} := \lim_{\alpha \rightarrow \infty} (I + \alpha L)^{-1}$, which is a matrix of normalized spanning forests. Both S_1 and \tilde{L} has a number of interesting properties [8]: Each entry of both matrices is non-negative, and both matrices are row-stochastic; \tilde{L} resides in the null space of digraph Laplacian L , as $L\tilde{L} = \tilde{L}L = 0$; $\text{rank}(L) = n - \text{rank}(\tilde{L})$; $L + \beta\tilde{L}$ is non-singular for any $\beta > 0$, and is the ‘‘complementary perturbation of L ’’ [9]. Indeed, this brings forward the second forest matrix,

$$S_2 := (L + \beta\tilde{L})^{-1}, \quad (2)$$

which is also termed the matrix of dense forest in [4]. As presented in what follows, varying the preprocessing schemes of normalizing W , we have two variants: MFTD_a starts with a preprocessing effort to normalize W , $W \leftarrow \frac{W}{w_{max}}$; MFTD_b considers a different normalization of W as $W \leftarrow D^{-1}W$ instead, i.e. $P = W$.

Proposition 1. *Under normalization scheme of $W \leftarrow D^{-1}W$, the forest matrix becomes $S_1 = (1 - \tau)(I - \tau P)^{-1}$, with $\tau = \frac{\alpha}{1 + \alpha}$.*

When applied on the second forest matrix S_2 using (2), this clearly leads to two additional variants that are denoted as MFTD_c & MFTD_d, respectively.

Algorithm 1 Label Propagation by Matrix-Forest Theorem of Digraphs (MFTD)

Input: A digraph $G = (\mathcal{V}, \mathcal{E}, W)$, label information Y , \mathbf{y}_l , and $\alpha \in (0, 1)$.

Output: \mathbf{y}_u^*

Compute the out-degree matrix D .

Compute the affinity matrix by $A = SY$ and (1) (or (2)).

Predict \mathbf{y}_u^* : Compute the i -th entry by (3), for any unlabeled node $v_i \in \mathcal{V}_u$.

One can interpret the (i, j) -th entry S_{ij} of the forest matrix S (being either $S := S_1$ or $S := S_2$) as quantifying the accessibility of a particle from a node v_i to visit node

v_j along the digraph structure. This provides a notion of affinity from state i to j . The intuition is, if a state j is close to the initial state i in terms of graph structure, it will be visited by the particle more often than if it is far away from initial state, i.e., we visit our close relatives more often than our distant ones. Now define the affinity matrix as $A = SY$, a matrix of size $n \times K$, with each entry a_{ik} being associated with an affinity score of state i belonging to class k . To infer \mathbf{y}_u^* of the unlabeled states \mathcal{V}_u , our algorithm predicts each entry's class assignment by identifying a label with the largest affinity score,

$$\mathbf{y}_i^* = \arg \max_k a_{ik}, \quad \forall v_i \in \mathcal{V}_u. \quad (3)$$

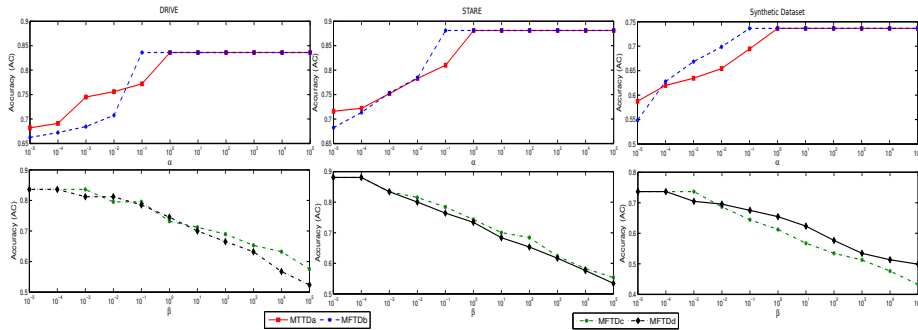


Fig. 2. Performance (AC) as a function of varying α or β of MFTD_{a-d} .

PAC-Bayesian Label-propagation Bound We report an investigation of the generalization bound of our approach on unseen data, which is an adaptation of [10]. We start by reformulating our algorithm (i.e. both (1) and (2)) as an equivalent representation $\mathbf{h} = S\hat{\mathbf{y}}$, where $\hat{\mathbf{y}}$ is the initial label vector with partial labels $\hat{y}_i \in \{\pm 1\}$ for $v_i \in \mathcal{V}_l$, and $\hat{y}_i = 0$ otherwise. The obtained \mathbf{h} is the “soft” label vector with h_i being the “soft” label for node v_i , which will be assigned with class label $\text{sign}(h_i)$ when making predictions. The hypothesis space is defined as $\mathcal{H} := \left\{ \mathbf{h} \mid \mathbf{h} = S\hat{\mathbf{y}}, \|\hat{\mathbf{y}}\|_2 \leq \sqrt{l} \right\}$. For any label vector \mathbf{h} , define the *test error* as $\mathcal{L}_{l,n}(\mathbf{h}) := \frac{1}{n-l} \sum_{i=l+1}^n \ell(\mathbf{h}_i, y_i)$ w.r.t. its 0/1 loss function ℓ satisfying $\ell(\mathbf{h}_i, y_i) = 1$ if $h_i \neq y_i$ and 0 otherwise, and let the *empirical error* of \mathbf{h} be $\hat{\mathcal{L}}_{l,n}(\mathbf{h}) := \frac{1}{l} \sum_{i=1}^l \ell(\mathbf{h}_i, y_i)$.

Theorem 1. For any $\delta \in (0, 1)$, with probability at least $1 - \delta$ over random draws of \mathcal{V}_l from \mathcal{V} , the following bound holds for any $\mathbf{h} \in \mathcal{H}$

$$\mathcal{L}_{l,n}(\mathbf{h}) \leq \hat{\mathcal{L}}_{l,n}(\mathbf{h}) + \sqrt{\left(\frac{2\hat{\mathcal{L}}_{l,n}(\mathbf{h})n}{n-l} \right) \frac{\ln \frac{1}{\delta} + 7ln(n+1)}{l-1}} + \frac{2 \left(\ln \frac{1}{\delta} + 7ln(n+1) \right)}{l-1}.$$

3 Experiments

Our approach is evaluated in our in-house synthetic dataset ⁵, as well as two testbeds, DRIVE [11] and STARE [12]. The synthetic dataset contains 17,000 synthesized retinal images with varying densities of blood vessels. Meanwhile, DRIVE dataset contains 40 retinal fundus images, and STARE has 20 fundus images. Exemplar images of the three datasets are illustrated in Figure 3.

		Synthetic Dataset					
		CDRN	WVRN	CTK _d	CTK _u	SGL	MFTD
AC		0.63	0.65	0.63	0.71	0.71	0.74
DS		0.62	0.62	0.61	0.68	0.64	0.71
		DRIVE [11]					
		CDRN	WVRN	CTK _d	CTK _u	SGL	MFTD
AC		0.69	0.67	0.73	0.79	0.76	0.83
DS		0.68	0.64	0.72	0.74	0.75	0.77
		STARE [12]					
		CDRN	WVRN	CTK _d	CTK _u	SGL	MFTD
AC		0.71	0.73	0.75	0.83	0.79	0.88
DS		0.68	0.69	0.74	0.78	0.76	0.82

Table 1. Comparison with leading label propagation methods. See text for details.

Our approach is compared with the following label propagation methods: Class Distribution Relational Neighbor classifier (CDRN) [13], Weighted Vote Relational Neighbor classifier (WVRN) [13], Digraph variant of the Commute Time Kernel classifier (CTK_d), and the original Commute Time Kernel classifier for undirected graphs (CTK_u) [14], and Symmetrized Graph Laplacian (SGL) [15]. To summarize, CTK_u is an undirected graph-based method, CTK_d, SGL, and the proposed MFTD are digraph-based methods, while the rest methods are not graph-theoretical. To ensure fair evaluations, the internal parameters of the comparison methods are either set to as is from the authors’ original source code, or as suggested in the papers. In terms of evaluation metric, the micro-averaged accuracy (AC) is utilized, which is the sum of all true positive counts divided by the total number of instances. Besides, the DIADEM score (DS) [16] is also employed, being a dedicated measure that has been widely used by the biological tracing community.

Effect of Varying α (or β) of Our Approach Our first experiment is to evaluate the effect of varying the value of our algorithmic parameter, namely α in MFTD_{a-b}, or β in MFTD_{c-d}. This is performed on all three datasets. As presented in Figure 2, the performance (AC) is displayed as a function of varying parameter value (α in row 1 & β in row 2) of our proposed algorithms MFTD_{a-d}, with x-axis being in log-scale. Surprisingly, all four variants of our MFTD framework produces exactly the same results when $\alpha \geq 1$ and $\beta \leq 10^{-4}$, which is also verified under the AC criterion as in the figure. This is very interesting as despite their differences in algebraic forms and graph-theoretical interpretations, effectively these variants are *equivalent* characterized by

⁵ Downloadable at <http://web.bii.a-star.edu.sg/~jaydeepd/tracing.htm>.

their ability of tracing retinal blood vessels. To avoid redundancies, we will collectively refer to the performance of all these four variants as MFTD, and fix $\alpha = 1$ and $\beta = 10^{-4}$ during the rest experiments.

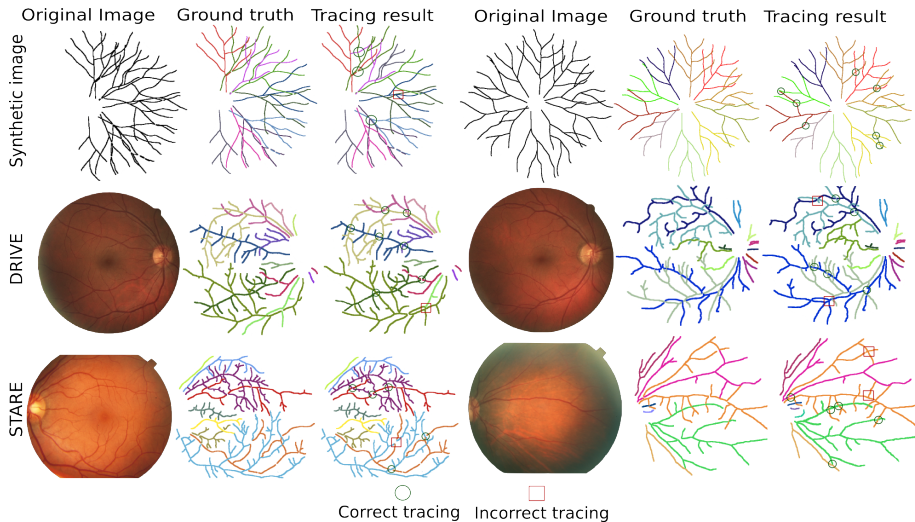


Fig. 3. Each row presents two exemplar retinal tracing results on Synthetic dataset, DRIVE, and STARE, respectively. Segments with the same color form a distinct vessel tree. Thus the number of colors equal to the number of classes (vessel trees). Selected correct (wrong) tracing segments are shown in green circles (red squares).

Comparison with the State-of-the-art Label Propagation Methods The performance of our approach is evaluated on three scenarios and is compared with nine label propagation methods as reported in Table 1. Overall our approach consistently outperforms the other methods by a large margin. The second best performer is usually CTK_u , which are followed by SGL and CTK_d . To facilitate visual inspection, Figure 3 presents exemplar images and comparison results. It suggests that empirically our approach delivers visually plausible tracing results when comparing to the ground-truths side-by-side, and the errors occur at those challenging spots that are often also difficult for human observers.

Comparison with State-of-the-art Tracing Systems We also compare with GOGP of [3] on DRIVE. DIADEM scores (DS) of the proposed MFTD on image no. 19 of DRIVE is 0.81, which is significantly higher than that of 0.71 obtained on the same image by [3]. We note in the passing that MFTD slightly outperforms the 0.765 of our earlier work [5], meanwhile we expect further gain would be achieved with the adoption of a better skeleton to digraph conversion.

4 Conclusion and Outlook

A novel approach is proposed for tracing vessels in fundus images. The tracing problem is solved by utilizing matrix-forest theorem of digraphs. Empirical evaluation demonstrates the superior performance of our approach. For future direction we plan to work with broader applications such as neurite tracing.

Acknowledgements

This research was partially supported by A*STAR JCO grants, as well as NSFC (No. 81271650) and NCET-10-0041.

References

1. Viswanath, K., McGavin, D.: Diabetic retinopathy: Clinical findings and management. *Community Eye Health* **16**(46) (2003) 21–24
2. Becker, C., Rigamonti, R., Lepetit, V., Fua, P.: Supervised feature learning for curvilinear structure segmentation. In: MICCAI. (2013)
3. Turetken, E., Gonzalez, G., Blum, C., Fua, P.: Automated reconstruction of dendritic and axonal trees by global optimization with geometric priors. *Neuroinformatics* **9** (2011) 279–302
4. Agaev, R.P., Chebotarev, P.Y.: Spanning forests of a digraph and their applications. *Automation and Remote Control* **62**(3) (2001) 443–466
5. De, J., Li, H., Cheng, L.: Tracing retinal vessel trees by transductive inference. *BMC Bioinformatics* **15**(20) (2014) 1–20
6. Harary, F., Norman, R., Cartwright, D.: *Structural models : an introduction to the theory of directed graphs*. Wiley (1965)
7. Brualdi, R., Ryser, H.: *Combinatorial Matrix Theory*. Cambridge Uni. Press (1991)
8. Chebotarev, P.Y., Agaev, R.P.: Forest matrices around the laplacian matrix. *Linear Algebra and its Applications* **356**(1-3) (2002) 253–247
9. Meyer, C., Stadelmaier, M.: Singular m-matrices and inverse positivity. *Linear Algebra and its Applications* **22** (1978) 139–156
10. Derbeko, P., El-Yaniv, R., Meir, R.: Explicit learning curves for transduction and application to clustering and compression algorithms. *J. Artif. Intell. Res.* **22** (2004) 117–142
11. Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., van Ginneken, B.: Ridge based vessel segmentation in color images of the retina. *IEEE Trans. Med. Imag.* **23**(4) (2004) 501–509
12. Hoover, A., Kouznetsova, V., Goldbaum, M.: Locating blood vessels in retinal images by piecewise threshold probing of a matched filter response. *IEEE Trans Med Imag* **19**(3) (2000) 203–210
13. Macskassy, S., Provost, F.: Classification in networked data: A toolkit and a univariate case study. *JMLR* **8** (2007) 935–983
14. Fousf, F., Francoise, K., Yen, L., Pirotte, A., Saerens, M.: An experimental investigation of kernels on graphs for collaborative recommendation and semisupervised classification. *Neural Network* **31** (2012) 53–72
15. Zhou, D., Huang, J., Schölkopf, B.: Learning from labeled and unlabeled data on a directed graph. In: ICML. (2005)
16. Gillette, T.A., Brown, K.M., Ascoli, G.A.: The diadem metric: comparing multiple reconstructions of the same neuron. *Neuroinformatics* **9**(2-3) (2011) 233–245