Learning Anisotropic Metrics for Geodesic Distances via the Heat Equation for Image Segmentation

Nicolas Makaroff[1], Théo Bertrand[1], and Laurent D. Cohen[1]

CEREMADE, UMR CNRS 7534, University Paris Dauphine, PSL Research University, 75775 Paris, France {bertrand, makaroff, cohen}@ceremade.dauphine.fr

Abstract. In this paper, we investigate the computation of anisotropic metrics for geodesic distances using the heat equation and its application to image segmentation, particularly for tubular structures. Building upon the work of Bertrand et al. [2], we extend this approach to anisotropic media by incorporating spatially varying and direction-dependent diffusion tensors derived from the image's structure tensor field. Our contributions are twofold: first, we formulate anisotropic geodesic distance computation using the heat equation and integrate it within deep learning models for image segmentation; second, we propose two methods that learn the anisotropic metric directly from image data — one requiring explicit seed point selection and another eliminating the need for seed points by predicting a probability map that serves as the initial condition for heat diffusion. Experiments on synthetic and medical image datasets demonstrate the effectiveness of our methods in accurately segmenting vascular tree structures by leveraging the anisotropic properties inherent in the images without relying on manual seed point selection.

Keywords: Image Segmentation · Deep Learning · Geodesic Distance · Fast Marching · Heat Equation · Medical Image Analysis.

1 Introduction

Segmenting regions in images based on geodesic distances is a fundamental problem in computer vision and image analysis, with applications ranging from medical imaging to robotics and geophysics. Traditionally, geodesic distances and curves are computed by explicitly constructing a Riemannian metric g from data, leveraging prior knowledge of the task. This metric is then used to compute geodesic distances, which are integral to understanding the underlying geometry of the data. However, this approach introduces bias through manual parameter tuning and subjective decisions in metric selection, potentially limiting the adaptability and generalisation of the method to diverse datasets.

To address these limitations, we propose a novel framework that learns the metric tensor directly from data using deep learning techniques. Specifically, we design convolutional neural networks (CNNs) that predict the anisotropic metric and potential fields required for geodesic distance computation. By optimising the network parameters through supervised learning on training data, we eliminate the need for manual metric selection and parameter tuning. This approach allows the network to learn the intrinsic geometry of the data adaptively, enhancing segmentation accuracy and robustness.

Traditional methods, like the Fast Marching Method (FMM) introduced by Sethian [8], compute geodesic distances by solving the Eikonal equation. Several prior works have leveraged geodesic distances for segmentation tasks. Malladi and Sethian [6] used minimal path distances and fast marching methods for 3D brain image segmentation. Cohen and Deschamps [3], segmented 3D vascular trees by propagating fronts based on minimal path distances, while Benmansour [1] introduced anisotropic metrics determined dynamically during fast marching computations to segment vascular structures. These methods rely on predefined metrics and do not directly integrate metric learning from data.

Recent efforts have explored learning metrics from data. Scarvelis et al. [7] and Heitz et al. [5] investigated fitting metric tensors to spatio-temporal data to capture velocity fields and underlying geometry. However, these approaches are not tailored for end-to-end segmentation tasks and often involve complex computations unsuitable for integration within deep learning frameworks.

Our work builds upon the work of Bertrand et al. [2]. In order to extend that approach to anisotropic metric, we use the heat method introduced by Crane et al. [4], and extended to anisotropic heat equation by Yang and Cohen [11], both based on Varadhan's asymptotic formula [10], which relate the behaviour of the heat equation to geodesic distances on a Riemannian manifold. By incorporating anisotropic diffusion tensors derived from the image's structure tensor field, we extend these methods to handle anisotropic media, capturing directional dependencies essential for accurately modelling complex pathways within tubular structures.

In contrast, our approach integrates metric learning and geodesic distance computation within a deep learning framework, enabling end-to-end training and inference. We propose two distinct methods: 1. Seed Point-Based Anisotropic Geodesic Distance Computation: Extending prior isotropic models, we predict both the anisotropic metric tensor and a Gaussian potential that initiates heat propagation from predefined seed points. This method leverages the strengths of deep learning for feature extraction and analytical computations of geodesic distances, enhancing precision in heterogeneous media. 2. Seedless Anisotropic Geodesic Distance Computation: To address challenges associated with seed point selection in non-convex regions, we introduce a seedless approach. Our neural network comprises two branches: one predicts the anisotropic diffusion tensor, and the other predicts a probability map serving as an initial condition for the heat equation. This design eliminates the need for explicit seed points, using an attention mechanism to guide heat propagation and refine segmentation.

Our contributions are as follows:

- We formulate anisotropic geodesic distance computation using the heat equation, integrating it within deep learning models for image segmentation.

- We propose two methods that learn the anisotropic metric directly from data, with and without the need for explicit seed point selection.
- We demonstrate the effectiveness of our methods on synthetic and real-world datasets, particularly in segmenting tubular structures in medical images.

The remainder of the paper is organised as follows. Section 2 provides background on geodesic distance computation, anisotropic heat diffusion, and the structure tensor field. Section 3 details our two proposed methodologies, distinguishing between the seed point-based and seedless approaches. Section 4 presents experimental results for both methods, including data descriptions, training procedures, and performance evaluations.

2 Background

2.1 Fast Marching Energy CNN for Image Segmentation

Bertrand et al. introduced the Fast Marching Energy CNN (FMECNN), integrating geodesic distance computations with convolutional neural networks for image segmentation, particularly in medical imaging. This framework learns isotropic metrics via a CNN to impose geometric and topological constraints on the segmentation process. By predicting both the potential function and the barycenter (seed point) of target regions, it generates segmentation masks corresponding to geodesic balls, ensuring path-connected and coherent outputs.

FMECNN extends traditional segmentation methods by incorporating a fast marching module within the network architecture, enabling end-to-end learning of geometrically meaningful segmentations. This addresses the limitations of conventional CNN-based methods, which often lack explicit structural priors and may produce irregular or disconnected segmentations. Applied to brain tumour segmentation using MRI data, FMECNN demonstrated competitive performance compared to architectures like U-Net and ResNet-U-Net, while enforcing geometric consistency.

However, FMECNN faces challenges that limit its applicability:

- Seed Point Prediction: The method relies on accurately predicting seed points from which geodesic distances are computed. Identifying optimal seed points in complex or non-convex regions is difficult, and inaccuracies can degrade segmentation performance.
- Isotropic Metric Assumption: Assuming isotropic diffusion restricts effectiveness in images with anisotropic properties, such as tubular structures (e.g., blood vessels), where diffusion varies with direction.
- Computational Overhead: Incorporating the fast marching algorithm introduces computational overhead, especially when dealing with multiple seed points or high-resolution images, hindering scalability and efficiency.

These challenges highlight the need for methods that handle anisotropic media, eliminate dependency on explicit seed point selection, and maintain computational efficiency.

Figure 1 illustrates FMECNN applied to vascular tree segmentation using predefined seed points. The ground truth shows an intricate network of vessels, including small branches. While the model successfully segments larger vessels, it misses many small vessels and fine branches. The learned potential captures primary vascular pathways but struggles with detailed branching patterns. This indicates that FMECNN, in its current form, is not well-suited for tasks requiring the segmentation of small and intricate vascular structures.

Relying on predefined seed points and isotropic metrics poses significant challenges in complex, anisotropic vascular networks. The isotropic assumption fails to account for the directional diffusion inherent in blood vessels, leading to inadequate segmentation of smaller vessels. Manually selecting seed points is impractical and may not cover all regions of interest. These limitations underscore the need for advanced methods that can learn anisotropic metrics directly from data and eliminate the need for explicit seed point selection. Our proposed methods aim to address these challenges by incorporating anisotropic diffusion tensors and seedless segmentation approaches.

2.2 Isotropic Heat Diffusion

The heat equation is a partial differential equation (PDE) that describes heat propagation in a medium over time. In the isotropic case, where the thermal diffusivity is constant and the same in all directions, the heat equation in $\mathbb{R}^2$ is given by:

$$\frac{\partial u}{\partial t} = \alpha \Delta u, \tag{1}$$

where $u(x,t)$ represents the temperature at position $x \in \mathbb{R}^2$ and time $t \geq 0, \alpha > 0$ is the thermal diffusivity coefficient, and Δ is the Laplace operator.

The Laplace operator in two dimensions is:

$$\Delta u = \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2}. \tag{2}$$

In homogeneous media, the thermal diffusivity α is constant, implying uniform heat propagation in all directions.

2.3 Spatially Varying Thermal Diffusivity

In heterogeneous media, the thermal diffusivity α can vary spatially. Following Yang et al. [11], we define α as:

$$\alpha(x) = |1 - |p(x_0) - p(x)||^d + \varepsilon, \tag{3}$$

where $p(x)$ represents a potential function (e.g., image intensity), x_0 is the heat source location, $d > 0$ controls the contrast between features, and ε ensures positivity.

Varying α implies that the rate of heat diffusion varies spatially within the domain, allowing for modelling complex propagation scenarios, such as barriers or preferred diffusion paths.

2.4 Anisotropic Heat Diffusion

To model anisotropic diffusion, we introduce the diffusion tensor $D(x)$, leading to the anisotropic heat equation:

$$\frac{\partial u}{\partial t} = \nabla \cdot (D(x)\nabla u), \tag{4}$$

where $D(x)$ is a symmetric positive-definite matrix at each point x. This allows us to control the diffusion rate in different directions.

By constructing $D(x)$ using features derived from the image, we can guide the heat propagation along preferred directions, such as the principal orientations of tubular structures.

2.5 Structure Tensor Field

The structure tensor field $T(x)$ captures local image orientations and anisotropies. It is defined as:

$$T(x) = \begin{pmatrix} \left(\frac{\partial I}{\partial x}\right)^2 & -\frac{\partial I}{\partial x}\frac{\partial I}{\partial y} \\ -\frac{\partial I}{\partial y}\frac{\partial I}{\partial x} & \left(\frac{\partial I}{\partial y}\right)^2 \end{pmatrix} * G_\sigma, \tag{5}$$

where $I(x)$ is the image intensity function, G_σ is a Gaussian kernel with standard deviation σ, and $*$ denotes convolution.

The structure tensor is diagonalised to obtain its eigenvalues $\mu_1(x) \geq \mu_2(x)$ and corresponding eigenvectors $e_1(x)$ and $e_2(x)$, representing the principal orientations and magnitudes of local variations:

$$T(x) = \mu_1(x)e_1(x)e_1(x)^\top + \mu_2(x)e_2(x)e_2(x)^\top. \tag{6}$$

This decomposition allows us to construct the tensor $T(x)$ to model anisotropic diffusion aligned with the image features.

2.6 Varadhan's Asymptotic Formula

Varadhan [10] established a relationship between the behaviour of the heat kernel and the geodesic distance on a Riemannian manifold. Specifically, as $t \to 0$, the heat kernel satisfies:

$$\lim_{t\to 0} (-4t \log u(x,y,t)) = d_\phi^2(x,y), \tag{7}$$

where $u(x,y,t)$ is the heat kernel, and $d_\phi(x,y)$ is the geodesic distance between points x and y with respect to the metric defined by ϕ.

This formulation allows us to compute anisotropic geodesic distances by solving the heat equation and analysing the behaviour of the solution as $t \to 0$.

3 Methodology

This section presents two methods for computing anisotropic geodesic distances using the heat equation integrated within deep learning frameworks for image segmentation. The first method extends prior work by incorporating seed points and predicting the anisotropic metric and the initial potential field. The second method addresses the challenges of seed point selection by introducing a seedless approach that leverages an attention mechanism within the neural network.

3.1 First Method: Seed Point-Based Anisotropic Geodesic Distance Computation

Overview Our first method builds upon the framework introduced in the isotropic case, merging deep learning with geodesic distance computation for segmentation tasks. We extend this approach to handle anisotropic geodesic distances by incorporating the heat method based on Varadhan's formulation [10, 11]. The key idea is to predict the anisotropic metric tensor and a Gaussian potential that initiates heat propagation from predefined seed points.

Deep Learning Model We employ a convolutional neural network (CNN) based on the U-Net architecture, renowned for its effectiveness in medical image segmentation. The network is designed to predict:

- Anisotropic Metric $\phi(x)$: Ensuring positivity and non-zero values by defining $\phi(x) = f_\theta(u)^2 + \varepsilon$, where $f_\theta(u)$ is the network output for input image u, θ represents the network parameters, and ε is a small constant to maintain numerical stability.
- Gaussian Potential Field: Serving as the source for heat propagation, the network predicts a Gaussian potential field that determines the initial conditions for the diffusion process.

The model comprises convolutional layers with decreasing filter sizes, followed by batch normalisation and ReLU activation functions. The final layer uses a softmax activation to produce a probability distribution, ensuring that the output values sum to one and can be interpreted as a potential field for initiating heat propagation.

Anisotropic Geodesic Distance Computation Using the predicted anisotropic metric $\phi(x)$ and the Gaussian potential field, we construct a structure tensor $D(x)$ that models the directional heat flow properties of the medium. The structure tensor is crucial for accurately simulating heat diffusion in media where thermal properties, such as tubular structures, vary with direction. We use the formulation of (6) to construct the structure tensor. We solve the anisotropic heat equation:

$$\frac{\partial u}{\partial t} = \nabla \cdot (D(x)\nabla u), \tag{8}$$

With the initial condition provided by the Gaussian potential field. As $t \to 0$, Varadhan's formulation allows us (see [11]) to compute the geodesic distance map:

$$d_\phi(x_0, x) = \lim_{t\to 0} \sqrt{-4t \log u(x,t)}, \tag{9}$$

where x_0 are the seed points from which heat propagation begins.

Segmentation with Multiple Seed Points We accommodate multiple seed points $S = \{x_0^i\}_{1 \leq i \leq q}$, q the number of seed points, for complex tubular networks with branching paths or intersections. The geodesic distance to the closest seed point is computed as:

$$d_\phi(S, x) = \min_{x_0^i \in S} d_\phi(x_0^i, x), \forall i \in \{1, \cdots, q\}. \tag{10}$$

This computation efficiently extends the model's applicability to tubular networks of varied complexity without imposing additional computational burdens on the heat method.

Mask Generation To generate the segmentation mask, we use a differentiable approximation of the characteristic function of a geodesic ball:

$$\chi^\delta(d_\phi(S, x)) = 1 - \frac{1}{1 + \exp\left(-\frac{d_\phi(S,x)-1}{\delta}\right)}, \tag{11}$$

where δ is a small positive parameter related to the image resolution, controlling the smoothness of the transition. This approach allows for gradient-based optimisation during training, facilitating the integration of the geodesic distance computation within the neural network framework.

Advantages and Limitations This method effectively leverages seed points to guide the segmentation process, capturing the complex pathways within tubular structures by accurately modelling the anisotropic properties of the medium. However, selecting appropriate seed points can be challenging in practice, especially in non-convex regions, and may introduce computational overhead when dealing with multiple seeds.

3.2 Second Method: Seedless Anisotropic Geodesic Distance Computation

Motivation While the first method relies on predefined seed points, this requirement poses challenges in non-convex regions like vascular networks, where identifying a unique optimal seed point is difficult. Additionally, computing geodesic distances from multiple seed points can be computationally expensive. To address these issues, we propose a second method that avoids explicit seed point selection by modifying the initialisation of the heat flow.

Deep Learning Model Our alternative approach employs a neural network composed of two branches, as illustrated in Figure 1:

1. Metric Prediction Branch: Predicts the anisotropic diffusion tensor $D(x)$, which defines the heat flow and is related to the inverse of the classic metric tensor.
2. Probability Map Prediction Branch: Predicts a 2D probability map $\mu(x)$, indicating likely vascular landmarks or regions of interest that guide the heat flow and refine the segmentation.

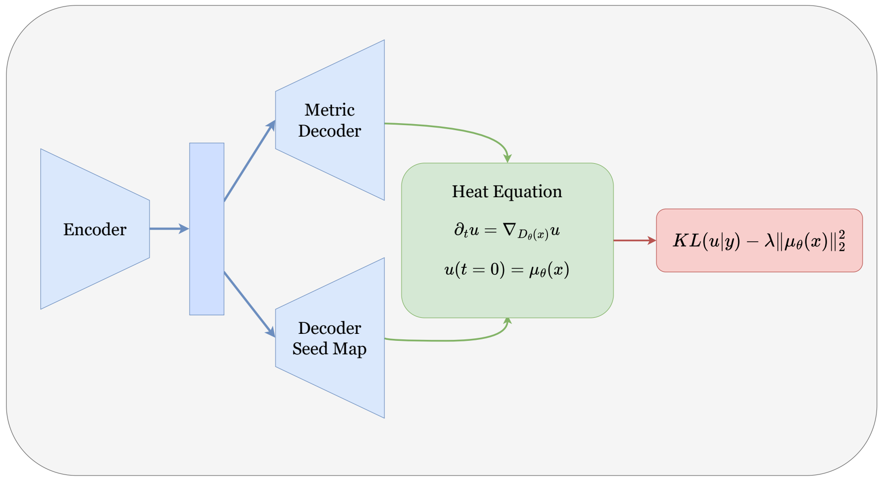

Fig. 1: Architecture of the alternative approach for anisotropic tubular structure segmentation. The network comprises two branches: predicting the diffusion tensor $D(x)$ and predicting the probability map $\mu(x)$.

This design allows the network to learn the medium's anisotropic properties and the areas from which heat should propagate without specifying exact seed points.

Anisotropic Heat Flow without Seed Points We initialise the heat equation with the predicted probability map $\mu(x)$ instead of a Dirac delta function at seed points:

$$\frac{\partial u}{\partial t} = \nabla \cdot (D(x)\nabla u), \quad u(x, 0) = \mu(x). \tag{12}$$

This formulation enables the heat to propagate from high-probability regions, effectively serving as an attention mechanism focusing on areas of interest within the image. The heat flow $u(x, t)$ captures the influence of the anisotropic diffusion and the initial probability distribution $\mu(x)$.

Loss Function and Training Given input images x and their corresponding normalised ground truth masks y, we define the loss function as:

$$\mathcal{L} = \mathcal{L}_{\text{seg}}(\Phi_t^{D_{\theta(x)}}(\mu_\theta(x)), y) - \lambda\|\mu_\theta(x)\|_2^2, \tag{13}$$

where:

- $\mathcal{L}_{\text{seg}}$ is the segmentation loss (e.g., Dice loss or cross-entropy loss) between the heat flow output $\Phi^{D\theta(x)}t(\mu\theta(x))$ and the ground truth mask y.
- $\Phi_t^D(\mu)$ represents the solution to the heat equation (heat flow) applied to μ until time t, using the diffusion tensor D.

– λ is a regularisation parameter promoting sparsity in $\mu(x)$, encouraging the network to focus on the most relevant regions.

The sparsity-promoting penalty $\|\mu_\theta(x)\|_2^2$ prevents the probability map from spreading too diffusely across the image, which helps in refining the segmentation by concentrating the heat propagation in areas of interest.

Advantages and Limitations This seedless approach simplifies the segmentation pipeline by eliminating the need for explicit seed point selection, reducing computational complexity. Using a probability map as an initial condition allows the network to learn where to focus the heat propagation, effectively functioning as an attention mechanism. However, this method relies on the network's ability to predict accurate probability maps, and the interpretation of the probability map may differ from explicit seed point methods.

3.3 Comparison of Methods

The two methods present distinct strategies for integrating anisotropic geodesic distance computation into Deep Learning models for image segmentation. The first method, a seed point-based approach, directly incorporates seed points, granting explicit control over the initiation of heat propagation. It efficiently handles multiple seed points for complex structures and leverages the predicted anisotropic metric and initial potential to capture the geometry of tubular formations. However, this method requires selecting or predicting appropriate seed points, which can be challenging in practice and may introduce computational overhead when numerous seed points are involved.

In contrast, the second method employs a seedless approach using an attention mechanism. It eliminates the need for explicit seed point selection, thereby simplifying the segmentation process. Using a probability map to guide heat propagation effectively focuses on regions of interest and reduces the computational complexity associated with handling multiple seed points. Nonetheless, this method relies on the network's ability to accurately predict the probability map $\mu(x)$. The initial probability distribution may not precisely represent discrete seed locations, which can potentially affect segmentation accuracy in some cases.

Both methods effectively integrate anisotropic diffusion within a deep learning framework to enhance the segmentation of complex structures like vascular networks. The choice between these approaches depends on the application's specific requirements and the data's nature. The first method offers explicit control over heat propagation when seed point selection is feasible and advantageous. Conversely, the second method provides a flexible and computationally efficient alternative in scenarios where seed point selection is impractical.

4 Experiments

4.1 Data

We conducted experiments on synthetic datasets and real-world medical images.

– Synthetic Tree Structures: Designed to test the model's capability in detecting arboreal structures. The dataset includes 20 images for training and 20 for testing, with key points represented by Gaussian heatmaps.
– DRIVE Dataset [9]: This dataset comprises 40 colour retinal images with corresponding manual vessel annotations provided by experts. We used only the green channel of the images, as it offers the highest contrast for blood vessels.

The model was trained using the Adam optimiser with a learning rate 0.01. Data augmentation techniques were applied, including horizontal and vertical flipping, random rotations, Geometric transformations (shifting and scaling).

These augmentations expose the model to various structural orientations and scales, improving generalisation.

4.2 Results

Our methods effectively segmented tubular structures, accurately capturing complex geometries and anisotropic properties.

Isotropic vs. Anisotropic Heat Propagation Figure 2 illustrates the difference between isotropic and anisotropic heat propagation across a fingerprint image. The anisotropic diffusion tensors guide the heat along the fingerprint ridges, adapting to the local geometry.

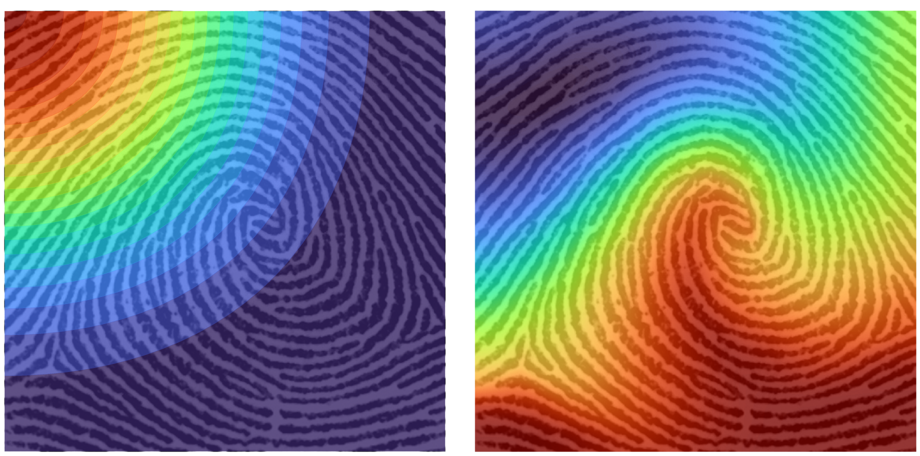

(a) Isotropic heat propagation. (b) Anisotropic heat propagation.

Fig. 2: Comparison of isotropic and anisotropic heat propagation across a fingerprint image.

Segmentation Results Figure 3 demonstrates the results of anisotropic tree structure analysis. In the first image (a), the output of the anisotropic model on the tree structure is shown, illustrating the heat propagation pattern. In (b), the potential field generated from the anisotropic diffusion is visualised, highlighting how the heat flows through the tree branches. Image (c) shows the modulation of the geodesic distance with the coefficient α, the thermal diffusion, affecting the diffusion process based on local variations. Finally, in (d), the predicted segmentation or outline of the tree structure is displayed, highlighting the model's ability to capture the geometric features of the tree.

(a) Output of the proposed model. (b) Predicted potential field. (c) Modulated geodesic distance. (d) Predicted segmentation.

Fig. 3: Anisotropic segmentation results for a tree structure.

4.3 Second Model Results

In Figure 4, the barycenter predicted by the Seed decoder branch of our model does not coincide with the barycenter maps produced by the original FMECNN model. Instead, it predicts segmentations of the vascular network. This divergence occurs because the second branch of our network is not explicitly trained to predict seed points, and the sparsity-promoting penalty in the loss function does not collapse the barycenter to a set of discrete points, even for larger values of λ. This behaviour can be interpreted as an attention mechanism, where the solution to the heat equation, computed from the two network outputs, refines the segmentation map.

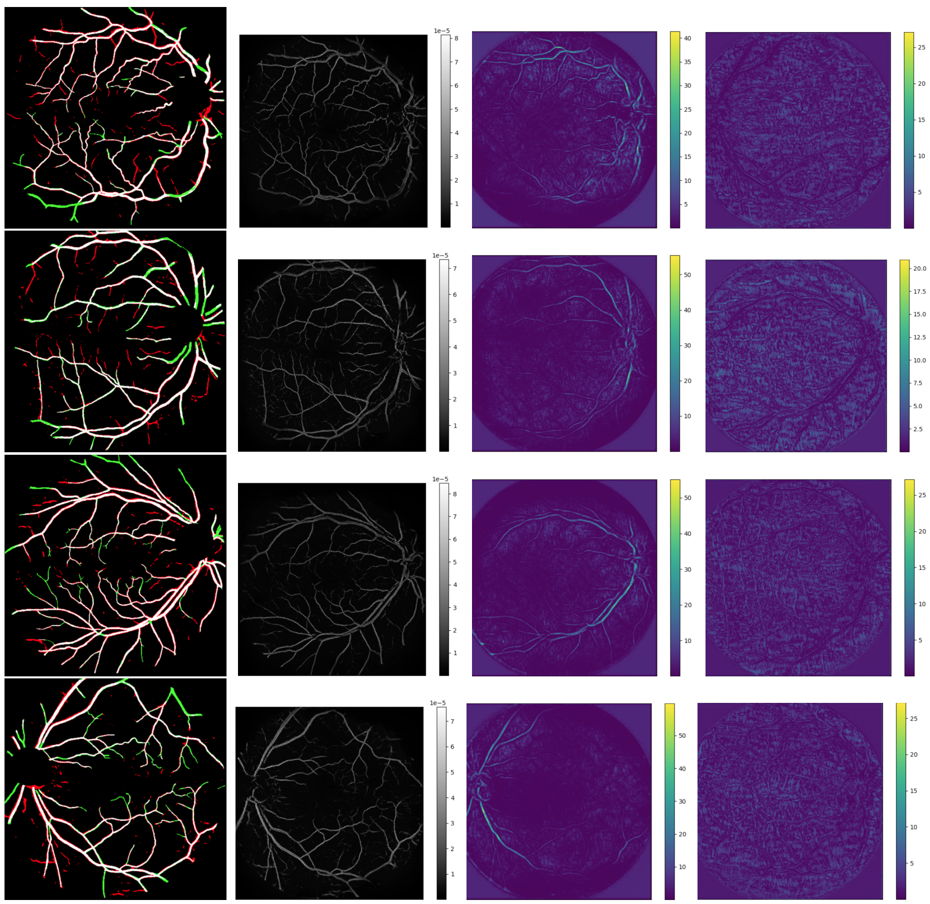

Fig. 4: Output of our method on a sample from the IOSTAR dataset. Left: Comparison of proposed segmentation versus Ground Truth. Center Left: Barycenter map output by the network. Center Right: Sum of the metric elements in both directions. Right : (log of) Anisotropy factor.

5 Conclusion

We introduced two methods for integrating anisotropic geodesic distance computation into deep learning models for image segmentation of complex structures like vascular networks. The first method employs seed points to control heat propagation, capturing intricate tubular geometries but requiring accurate seed placement. The second method is seedless, using a probability map to guide heat propagation, simplifying the segmentation process but relying on the network's ability to predict this map accurately. Both methods effectively segment complex, anisotropic structures without manual seed selection. Our experiments demonstrate that learning anisotropic metrics directly from data and integrating geodesic computation into deep learning models advance image segmentation techniques for complex, directionally dependent structures.

Acknowledgments. This work is in part supported by the French government under management of Agence Nationale de la Recherche as part of the "Investissements d'avenir" program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute).

Disclosure of Interests. The authors have no competing interests to declare that they are relevant to the content of this article.

References

1. Benmansour, F., Cohen, L.D.: Tubular Structure Segmentation Based on Minimal Path Method and Anisotropic Enhancement. International Journal of Computer Vision **92**(2), 192–210 (Apr 2011). https://doi.org/10.1007/s11263-010-0331-0, http://link.springer.com/10.1007/s11263-010-0331-0
2. Bertrand, T., Makaroff, N., Cohen, L.D.: Fast Marching Energy CNN. In: International Conference on Scale Space and Variational Methods in Computer Vision (SSVM'23. pp. 276–287. Springer (2023)
3. Cohen, L.D., Deschamps, T.: Segmentation of 3D Tubular Objects With Adaptive Front Propagation and Minimal Tree Extraction for 3D Medical Imaging. Computer Methods in Biomechanics and Biomedical Engineering **10**(4), 289–305 (2007)
4. Crane, K., Weischedel, C., Wardetzky, M.: Geodesics in Heat: A New Approach to Computing Distance Based on Heat Flow. ACM Transactions on Graphics (TOG) **32**(5), 1–11 (2013)
5. Heitz, M., Bonneel, N., Coeurjolly, D., Cuturi, M., Peyré, G.: Ground Metric Learning on Graphs. Journal of Mathematical Imaging and Vision **63**, 89–107 (2021)
6. Malladi, R., Sethian, J.A.: A Real-Time Algorithm for Medical Shape Recovery. In: Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271). pp. 304–310. IEEE (1998)
7. Scarvelis, C., Solomon, J.: Riemannian Metric Learning via Optimal Transport. arXiv Preprint arXiv:2205.09244 (2022)
8. Sethian, J.A.: A Fast Marching Level Set Method for Monotonically Advancing Fronts. Proceedings of the National Academy of Sciences **93**(4), 1591–1595 (1996)
9. Staal, J., Abràmoff, M.D., Niemeijer, M., Viergever, M.A., Van Ginneken, B.: Ridge-based vessel segmentation in color images of the retina. IEEE transactions on medical imaging **23**(4), 501–509 (2004)
10. Varadhan, S.R.S.: On the Behavior of the Fundamental Solution of the Heat Equation With Variable Coefficients. Communications on Pure and Applied Mathematics **20**(2), 431–455 (1967)
11. Yang, F., Cohen, L.D.: Geodesic Distance and Curves Through Isotropic and Anisotropic Heat Equations on Images and Surfaces. Journal of Mathematical Imaging and Vision **55**, 210–228 (2016)