

Automatic, Robust, and Globally Optimal Segmentation of Tubular Structures

Simon Pezold¹(✉), Antal Horváth¹, Ketut Fundana¹, Charidimos Tsagkas²,
Michaela Andělová^{2,3}, Katrin Weier², Michael Amann²,
and Philippe C. Cattin¹

¹ Department of Biomedical Engineering, University of Basel, Basel, Switzerland
`simon.pezold@unibas.ch`

² Department of Neurology, University Hospital Basel, Basel, Switzerland

³ Department of Internal Medicine, University Hospital Motol,
Prague, Czech Republic

Abstract. We present an automatic three-dimensional segmentation approach based on continuous max flow that targets tubular structures in medical images. Our method uses second-order derivative information provided by Frangi et al.'s vesselness feature and exploits it twofold: First, the vesselness response itself is used for localizing the tubular structure of interest. Second, the eigenvectors of the Hessian eigendecomposition guide our anisotropic total variation-regularized segmentation. In a simulation experiment, we demonstrate the superiority of anisotropic as compared to isotropic total variation-regularized segmentation in the presence of noise. In an experiment with magnetic resonance images of the human cervical spinal cord, we compare our automated segmentations to those of two human observers. Finally, a comparison with a dedicated state-of-the-art spinal cord segmentation framework shows that we achieve comparable to superior segmentation quality.

Keywords: Convex optimization · Anisotropic total variation · Vesselness

1 Introduction

Segmenting tubular structures is an important task in medical image analysis; for example, for assessing vascular diseases or tracking the progress of neurological disorders that manifest in spinal cord atrophy. Especially when used in large-scale clinical trials, largely automated segmentation is desirable to reduce the workload on clinical staff. Such automated segmentation approaches, in turn, should be robust with respect to specific choices of parameterization.

In this paper, we propose a segmentation method that fulfills both criteria: it is completely automated, and it creates segmentations of similar quality over a wide range of parameter choices. Our method adapts Yuan et al.'s continuous max flow approach [10] and combines it with an anisotropic total variation (ATV)




Fig. 1. *Left to right:* T1 image I of the cervical spinal cord; closeups of the vessel directions v_1 before GVF and \hat{v}_1 after GVF (vectors scaled by the segmentation u^* for visualization); distance map D of the vessel ridge \mathcal{R} (normalized for visualization); source capacities C_s ; sink capacities C_t ; nonterminal capacities C ; segmentation u^* .

regularization term. ATV keeps changes of the segmentation’s boundary small along the course of the tubular structure. We use Frangi et al.’s well-established vesselness feature [3] as our measure of tubularity, which we exploit twofold: both for finding the location and the orientation of the structures of interest.

The directional information of vesselness, which is usually neglected, has previously been used: Manniesing et al. [6] construct an anisotropic tensor from it for accentuating vascular structures in angiography images for image enhancement. Gooya et al. [4] use this tensor in an active contour framework for blood vessel segmentation. ATV-regularized segmentation has been generally described by Olsson et al. [7] and has been used, for example, by Reinbacher et al. [8] who segment thin structures of known volume based on first-order derivatives and a volume constraint. A review of vessel segmentation is given by Lesage et al. [5]. De Leener et al. [2] review the more specific topic of spinal cord segmentation.

Our contributions lie in incorporating Hessian-based vesselness into ATV-regularized segmentation and in integrating ATV into the continuous max flow framework [10]. To the best of our knowledge, both has not been tried, so far.

2 Methods

In Sect. 2.1, we motivate our choice of ATV regularization. In Sect. 2.2, we state the ATV-regularized segmentation problem in the continuous max flow framework and propose an algorithm for solving it. In Sect. 2.3, we describe how we incorporate the vesselness feature. In Sect. 2.4, we present our choice of flow capacities. For a good general introduction to continuous max flow, see [10].

2.1 Isotropic and Anisotropic Total Variation

Segmentation on the d -dimensional image domain $\Omega \subset \mathbb{R}^d$ can be formulated as the problem of finding a binary labeling $u : \Omega \rightarrow \{0, 1\}$ for the given image $I : \Omega \rightarrow \mathcal{I}$ (e.g. with $\mathcal{I} = [0, 1]$ for a normalized single-channel image). In practice,

the problem is often relaxed such that $u : \Omega \rightarrow [0, 1]$, and the final labeling is determined by applying a threshold to the result of the relaxed problem [7, 10].

A common regularization term in segmentation is the *total variation* TV, which minimizes the surface area of the segmented region, penalizing jumps between segmentation foreground ($u = 1$) and background ($u = 0$) and thus allowing for smooth segmentations even if I is noisy (here, $|\cdot|$ denotes the l_2 norm):

$$\text{TV}[u] = \int_{\Omega} |\nabla u| \, dx. \tag{1}$$

While TV is a good regularizer for many applications, it seems not optimal in the context of tubular structure segmentation. This is because TV is isotropic; that is, changes of u are penalized regardless of orientation. If we want to segment a tube, however, we would like to employ the prior knowledge that its shape ideally does not change along its course; thus we would like to penalize changes along the tube’s direction more strongly than changes perpendicular to it. In other words, we would prefer an anisotropic regularization term.

In the proposed method, we thus use *anisotropic total variation* ATV [7, 8]:

$$\text{ATV}[u; A] = \int_{\Omega} (\nabla u^T A \nabla u)^{1/2} \, dx = \int_{\Omega} |S^T \nabla u| \, dx \quad \text{with} \quad A = SS^T, \tag{2}$$

where $A : \Omega \rightarrow \mathbb{R}^{d \times d}$ is strongly positive definite in the sense of Olsson et al. [7] and S is a decomposition of A . For our particular choice of A and S , see Sect. 2.3.

If we assume, as a simple three-dimensional example, that $A = \text{diag}(1, a, a)$ with $0 < a < 1$, we see that changes along the x_1 axis will be more strongly penalized than changes along x_2 and x_3 . This would in fact be a meaningful choice if the tubular structure of interest was oriented along x_1 . From the example we can also see that ATV is a generalization of TV, as ATV becomes TV for $a = 1$.

2.2 ATV in Continuous Max Flow

The dual formulation of the max flow problem as stated in [10], with TV replaced by ATV regularization as we propose, is given by the min cut problem

$$\min_{u \in [0,1]} \int_{\Omega} (1 - u) C_s + u C_t + |S^T \nabla u| C \, dx, \tag{3}$$

with the source capacities C_s , sink capacities C_t , and nonterminal capacities C ($C : \Omega \rightarrow \mathbb{R}_{\geq 0}$). Using integration by parts and the geometric definition of the scalar product, we can show for the nonterminal flow $p : \Omega \rightarrow \mathbb{R}^d$ that

$$\int_{\Omega} |S^T \nabla u| C \, dx = \max_{|p| \leq C} \int_{\Omega} u \, \text{div}(Sp) \, dx. \tag{4}$$

Algorithm 1. Augmented Lagrangian-based max flow algorithm with ATV.

- Set bound $\hat{\epsilon}$, steps γ, c ; calculate C, C_s, C_t, S ; arbitrarily initialize p^0, p_s^0, p_t^0, u^0 .
- Starting from $n = 0$, iterate until $\frac{1}{|\Omega|} \int_{\Omega} |\epsilon^{n+1}(x)| dx < \hat{\epsilon}$:

$$\begin{cases} \tilde{p}^{n+1} = p^n + \gamma S^T \nabla \left(\operatorname{div}(Sp^n) - p_s^n + p_t^n - \frac{u^n}{c} \right) \\ p^{n+1} = \frac{\tilde{p}^{n+1}}{|\tilde{p}^{n+1}|} \min \{ |\tilde{p}^{n+1}|, C \} & \text{if } \tilde{p}^{n+1} \neq 0 \quad \text{else } 0 \\ p_s^{n+1} = \min \left\{ \left(\frac{1-u^n}{c} + \operatorname{div}(Sp^{n+1}) + p_t^n \right), C_s \right\} \\ p_t^{n+1} = \min \left\{ \left(\frac{u^n}{c} - \operatorname{div}(Sp^{n+1}) + p_s^{n+1} \right), C_t \right\} \\ \epsilon^{n+1} = c \left(\operatorname{div}(Sp^{n+1}) - p_s^{n+1} + p_t^{n+1} \right) \\ u^{n+1} = u^n - \epsilon^{n+1}. \end{cases}$$

Together with the respective equalities for C_s, C_t and the source and sink flows $p_s, p_t : \Omega \rightarrow \mathbb{R}$ (see Eqs. (18) and (19) in [10]), we derive the primal-dual formulation as $\max_{p_s, p_t, p} \min_{u \in [0,1]} E[p_s, p_t, p, u]$ with

$$E = \int_{\Omega} (1-u) p_s + u p_t + u \operatorname{div}(Sp) dx = \int_{\Omega} p_s + u (\operatorname{div}(Sp) - p_s + p_t) dx, \quad (5)$$

subject to the flow capacity constraints $p_s \leq C_s, p_t \leq C_t$, and $|p| \leq C$.

Making use of the anisotropic coarea formula in [7], it can be shown that any u^ℓ for a threshold $\ell \in (0, 1)$, given by

$$u^\ell(x) = \begin{cases} 1, & u^*(x) > \ell \\ 0, & u^*(x) \leq \ell \end{cases} \quad \text{with } u^* = \arg \min_{u \in [0,1]} E, \quad (6)$$

is a globally optimal solution for the binary problem corresponding to Eq. (3).

Following [10], we add an augmented Lagrangian term to Eq. (5), gaining

$$\max_{p_s, p_t, p} \min_{u \in [0,1]} \int_{\Omega} p_s + u (\operatorname{div}(Sp) - p_s + p_t) - \frac{c}{2} (\operatorname{div}(Sp) - p_s + p_t)^2 dx \quad (7)$$

as the final problem, which we propose to solve with Algorithm 1.

2.3 ATV Regularization with Vesselness

Frangi et al. [3] examine the Hessian matrices; that is, the second-order derivatives, in the scale space of the volumetric image I to calculate what they call vesselness. The idea is to determine from the ratios of the Hessians' eigenvalues how closely the local structure in I resembles a tube.

In particular, let \mathcal{S} be a predefined set of scales that roughly match the expected tube radii. For each scale $s \in \mathcal{S}$, let $H_s(x)$ denote its Hessian approximation in x , calculated by convolving I with Gaussian derivatives of standard deviation s . Let $\lambda_{i,s}(x)$ ($i=1, 2, 3$) denote the sorted eigenvalues

($|\lambda_{1,s}| \leq |\lambda_{2,s}| \leq |\lambda_{3,s}|$) and $v_{i,s}(x)$ corresponding eigenvectors of H_s , such that

$$H_s = V_s \Lambda_s V_s^T \quad \text{with} \quad V_s = [v_{1,s}|v_{2,s}|v_{3,s}], \quad \Lambda_s = \text{diag}(\lambda_{1,s}, \lambda_{2,s}, \lambda_{3,s}). \quad (8)$$

Note that $V_s^T = V_s^{-1}$, as H_s is symmetric. Assuming bright tubular structures on dark background, the vesselness response is $\nu(x) = \max_{s \in \mathcal{S}} \nu_s(x)$, where

$$\nu_s = \begin{cases} 0, & \lambda_{2,s} \geq 0 \text{ or } \lambda_{3,s} \geq 0 \\ \left(1 - \exp\left(\frac{-1}{2w_1^2} \frac{\lambda_{2,s}^2}{\lambda_{3,s}^2}\right)\right) \exp\left(\frac{-1}{2w_2^2} \frac{\lambda_{1,s}^2}{\lambda_{2,s}\lambda_{3,s}}\right) \left(1 - \exp\left(\frac{-\sum_i \lambda_{i,s}^2}{2w_3^2}\right)\right), & \text{else,} \end{cases} \quad (9)$$

with the weighting factors $w_i \in \mathbb{R}_{>0}$. The eigenvectors for ν are $V = [v_1|v_2|v_3]$, with $v_i = v_{i,s^*}$ and $s^* = \arg \max_{s \in \mathcal{S}} \nu_s$. In the original description of [3], no use of V is made. In our approach, we use the eigenvectors to steer the ATV regularizer. We observe that in points where ν is high, v_1 points along the local vessel orientation [6]. Recall that we want to regularize strongly along the direction of the vessel. Unfortunately, we cannot use v_1 directly for this purpose, as it reliably gives the vessel’s direction in the vessel center only, where ν is the highest. Therefore, we use the concept of gradient vector flow (GVF) [9] to first propagate the directions from places where ν is high to regions where ν is low, creating a smoothly varying vector field. The necessary steps are as follows.

Let \mathcal{R} be the set of vesselness ridge points; that is, the local maxima of ν , down to a noise threshold. As both $-v_1$ and v_1 are valid eigenvectors, we have to make sure that the vectors of neighboring points approximately point in the same rather than the opposite direction, so that they don’t cancel each other out when diffusing them via GVF. Thus, we fix their signs beforehand, gaining \bar{v}_1 : We calculate the minimum spanning tree over the ridge points \mathcal{R} , select a root point, keep its sign, and traverse the tree. For each child point, we choose \bar{v}_1 as either $-v_1$ or v_1 , depending on which one maximizes the dot product (i.e. minimizes the angle) with its parent’s \bar{v}_1 . After traversal, the signs of the v_1 for all remaining domain points $x \in \Omega \setminus \mathcal{R}$ are fixed w.r.t. their closest point in \mathcal{R} , following the same rule. We scale all \bar{v}_1 with ν , apply GVF, and scale the resulting vectors back to unit length, gaining \tilde{v}_1 . A comparison of the vector field before and after sign adjustment and GVF is shown in Fig. 1.

Finally, we recomplete \tilde{v}_1 to an orthonormal basis $\tilde{V} = [\tilde{v}_1|\tilde{v}_2|\tilde{v}_3]$. The particular choice of \tilde{v}_2, \tilde{v}_3 does not matter, as we will treat all directions perpendicular to \tilde{v}_1 the same when regularizing. From \tilde{V} , we construct A, S for Eq. (2) as

$$A = \tilde{V} \tilde{A} \tilde{V}^T \quad \text{and} \quad S = \tilde{V} \tilde{A}^{1/2} \quad \text{with} \quad \tilde{A} = \text{diag}(1, a, a) \quad \text{and} \quad 0 < a \leq 1. \quad (10)$$

Notice the similarity to A in the example at the end of Sect. 2.1: The idea of regularizing one direction stronger than the others remains the same; however, as we now scale with $\tilde{A} = \text{diag}(1, a, a)$ in the new basis \tilde{V} , we target the actual local vessel direction \tilde{v}_1 rather than a fixed axis.

2.4 Flow Capacities

For the source and sink capacities C_s, C_t of Eqs. (3) and (5), we use a combination of the normalized image intensities and the distances to the vessel ridge points \mathcal{R} . Intuitively, using the intensities enables the distinction of foreground (i.e. the vessel) and background (i.e. everything else), while the distances w.r.t. \mathcal{R} isolate the vessel surrounding. This helps avoiding oversegmentations in case other structures have intensities similar to those of the vessel of interest. More formally, let $D : \Omega \rightarrow \mathcal{D} = \mathbb{R}_{\geq 0}$ be a Euclidean distance map of \mathcal{R} , and let $I : \Omega \rightarrow \mathcal{I} = [0, 1]$ be the normalized image. Let $p_b^D, p_f^D, p_b^I, p_f^I$ be predefined estimates of the background and foreground probability densities for \mathcal{D} and \mathcal{I} , with $p_b^D : \mathcal{D} \rightarrow \mathbb{R}_{\geq 0}$ and $p_f^I : \mathcal{I} \rightarrow \mathbb{R}_{\geq 0}$. We calculate C_s, C_t as

$$C_s(x) = \frac{1}{q} \max\{r(x), 0\}, \quad \text{with } r(x) = \ln \left(\frac{p_f^D(D(x)) \cdot p_f^I(I(x)) + \varepsilon}{p_b^D(D(x)) \cdot p_b^I(I(x)) + \varepsilon} \right), \quad (11)$$

$$C_t(x) = \frac{1}{q} \max\{-r(x), 0\}, \quad q = \ln \left(\frac{\max\{\hat{p}_b^D, \hat{p}_f^I, \hat{p}_b^D \cdot \hat{p}_f^I\} + \varepsilon}{\varepsilon} \right), \quad (12)$$

where $\hat{p} = \max p$. The small positive constant ε avoids zero logarithms and zero divisions in r . Normalization with q ensures that $C_s, C_t \in [0, 1]$, which eases their balancing with the nonterminal capacities C of Eqs. (3) and (5).

Using C , we try to move the segmentation boundary to image edges by making C small where the intensity gradient magnitude is high and vice versa:

$$C(x) = w \exp\left(-1/\varsigma^2 |\nabla I(x)|^2\right) \quad \text{with } w, \varsigma \in \mathbb{R}_{>0}, \quad (13)$$

where ς controls C 's sensitivity regarding the size of $|\nabla I|$ and w balances C and C_s, C_t . For an example of the capacities, see Fig. 1.

3 Experiments and Results

Implementation: Most of the method's steps can be described as *embarrassingly parallel*, which means they can be calculated independently for different voxels. This is true, for example, for the vesselness ν , GVF, the capacities C_s, C_t, C , and large parts of Algorithm 1. For this reason, we ported code to the GPU wherever possible. To reduce memory consumption, which is still a limiting factor for GPU programming, we represent the pointwise basis matrices \hat{V} in Eq. (10) as equivalent unit quaternions. Derivatives are approximated using forward (∇ , Algorithm 1), backward (div, Algorithm 1), and central (∇ , Eq. 13) differences. *Parameterization:* For all experiments, the following parameters were chosen. Algorithm 1: $\hat{\varepsilon} = 10^{-6}$, $\gamma = 0.11\Delta x^2$, $c = 0.2\Delta x^2$ (Δx : minimum voxel edge length in mm); GVF: 316 iterations with a regularization parameter of $\mu = 3.16$ and step size determined as defined and described in [9]; Eq. (6): $\ell = 0.5$; Eq. (9): $w_1 = w_2 = 0.5$ as suggested in [3], w_3 determined following [3]; Eq. (12): $\varepsilon = 10^{-9}$.




Fig. 2. Phantom experiment. *Left to right:* Mean intensity projection at $\sigma = 1.5$; segmentation with ATV; segmentation with TV; Dice coefficients w.r.t. ground truth (GT) for all noise levels σ .




Fig. 3. Spinal cord experiment. *Left to right:* Mean surface distances for ATV/TV/*PropSeg* [1] vs. E_1 and E_2 ; corresponding Hausdorff distances; corresponding Dice coefficients; T1 Dice coefficients for ATV vs. E_1 and E_2 with varying parameter values (central bar: median, circle: mean, box limits: 25/75th percentile, whiskers: extrema).

Helical Phantom. In this simulation experiment, we rendered images of a synthetic helical phantom with values in $[0, 1]$ to which we added Gaussian noise of standard deviation σ (Fig. 2). The phantom’s tube radius varied between 3 mm and 6 mm, so we set $\mathcal{S} = [2.5 \text{ mm}, 7.2 \text{ mm}]$ (16 scales) in Eq. (9). We segmented the images with TV and ATV regularization, setting $a = 1$ and $a = 0.03$ in Eq. (10), respectively. We modeled p^I as normal distributions, using the true background, foreground, and noise level values. For the sake of simplicity, we set $p^D = 1$ always. For w and ζ in Eq. (13), we made grid searches for each noise level, using the Dice coefficients w.r.t. the ground truth as optimization criterion.

Figure 2 shows the Dice coefficients for the best w, ζ combinations. The advantage of ATV becomes apparent as soon as the noise level increases.

Spinal Cord. In this experiment with real data, two clinical experts (E_1, E_2) manually segmented 10 MR scans (5 T1, 5 T2) of the healthy human cervical spinal cord over the C1–C3 region. For each image, we then used the remaining four images of the same sequence (T1/T2) to estimate p^I and to find an optimal parameter combination for a, w, ζ . The distributions p^I were estimated from the manual labelings of the remaining four images, modeling p_b as mixtures of four Gaussians and p_f as normal distributions. We set $\mathcal{S} = [2 \text{ mm}, 4 \text{ mm}]$ (16 scales) in Eq. (9). The parameters a, w, ζ were optimized via grid search, using the mean

Dice coefficients of the remaining four images w.r.t. their manual segmentations as optimization criterion. These distributions and the determined optimum parameterization were then used to segment the left-out image for evaluation of the method. For comparison with the state of the art, we also segmented all images with *PropSeg* [1].

Figure 3 shows the averaged mean surface distances, Hausdorff distances, and Dice coefficients w.r.t. their manual segmentations. Especially the T1 images profit from ATV, as they have both a lower resolution (T1: $1 \times 1 \times 1 \text{ mm}^3$, T2: $0.75 \times 0.38 \times 0.38 \text{ mm}^3$) and a lower contrast-to-noise ratio (about one eighth) than the T2 images. On the right, Fig. 3 shows the Dice coefficients for applying a wide range of parameter values to the T1 images, demonstrating the robustness of our method w.r.t. parameterization. For each varied parameter, the others were kept constant ($a = 0.03$, $w = 0.32$, $\varsigma = 1.00$).

4 Discussion and Conclusion

We presented a fully automated method for the segmentation of tubular structures. For a single image of 256^3 voxels, the complete process of calculating the vesselness, GVF, capacities, and segmentation takes about 1 to 1.5 min (GPU: Nvidia GeForce GTX 770). Although image segmentation in general and tubular structure segmentation in particular have often been addressed, the results of comparison with a state-of-the-art approach lead us to believe that our method may be of value to the scientific community. Future experiments will have to show in more detail how strong is the dependence of the segmentation quality on the outcome of the Frangi vesselness response and what is the influence of a particular GVF parameterization in this context. Furthermore, the use of alternative vesselness indicators will have to be considered. We provide our reference implementation at <https://github.com/spezold/miccai2016>.

References

1. De Leener, B., Kadoury, S., Cohen-Adad, J.: Robust, accurate and fast automatic segmentation of the spinal cord. *NeuroImage* **98**, 528–536 (2014)
2. De Leener, B., Taso, M., Cohen-Adad, J., Callot, V.: Segmentation of the human spinal cord. *Magn. Reson. Mater. Phys., Biol. Med.* **29**(2), 125–153 (2016)
3. Frangi, A.F., Niessen, W.J., Vincken, K.L., Viergever, M.A.: Multiscale vessel enhancement filtering. In: Wells, W.M., Colchester, A.C.F., Delp, S.L. (eds.) *MIC-CAI 1998*. LNCS, vol. 1496, pp. 130–137. Springer, Heidelberg (1998)
4. Gooya, A., Liao, H., Sakuma, I.: Generalization of geometrical flux maximizing flow on Riemannian manifolds for improved volumetric blood vessel segmentation. *Comput. Med. Imaging Graph.* **36**(6), 474–483 (2012)
5. Lesage, D., Angelini, E.D., Bloch, I., Funka-Lea, G.: A review of 3D vessel lumen segmentation techniques: models, features and extraction schemes. *Med. Image Anal.* **13**(6), 819–845 (2009)
6. Manniesing, R., Viergever, M.A., Niessen, W.J.: Vessel enhancing diffusion: a scale space representation of vessel structures. *Med. Image Anal.* **10**(6), 815–825 (2006)

7. Olsson, C., Byrod, M., Overgaard, N., Kahl, F.: Extending continuous cuts: anisotropic metrics and expansion moves. In: 2009 IEEE 12th International Conference on Computer Vision, pp. 405–412 (2009)
8. Reinbacher, C., Pock, T., Bauer, C., Bischof, H.: Variational segmentation of elongated volumetric structures. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3177–3184 (2010)
9. Xu, C., Prince, J.L.: Snakes, shapes, and gradient vector flow. *IEEE Trans. Image Process.* **7**(3), 359–369 (1998)
10. Yuan, J., Bae, E., Tai, X.C.: A study on continuous max-flow and min-cut approaches. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2217–2224 (2010)