

DESIGN AND STUDY OF FLUX-BASED FEATURES FOR 3D VASCULAR TRACKING

David Lesage^{*†}, Elsa D. Angelini[†], Isabelle Bloch[†] and Gareth Funka-Lea^{*}

^{*}Siemens Corporate Research, Imaging and Visualization dept., Princeton NJ, USA.

[†]Institut Télécom, Télécom ParisTech, CNRS LTCI, Paris, France.

ABSTRACT

In this paper, we present and study two local features for the tracking of vascular structures on 3D angiograms. The first one, FLUX , measures the inward gradient flux through circular cross-sections. The second one, MFLUX , introduces a non-linear penalization of asymmetric flux contributions to reduce false positive responses.

Through a series of experiments on synthetic and real cardiac CT data, we discuss the properties of these features with respect to their parameters. We compare them to a selection of published vessel-dedicated features. We show that MFLUX induces a particularly discriminative response landscape, which is a desirable property for tracking purposes on such large search spaces.

A key characteristic of the proposed features is their simplicity of implementation and their high computational efficiency, enabling their practical use for advanced tracking strategies.

Index Terms— Vascular features, gradient flux, vessel tracking

1. INTRODUCTION

3D angiographic acquisitions raise numerous challenges for vascular segmentation algorithms. In complex applications, vessels of interest are thin, branching structures of varying radius and curvature, embedded in large volumes and surrounded by non-vascular structures of similar luminance. Intricate models and extraction schemes are often required to ensure sufficient robustness, to the detriment of computational cost. *Tracking* strategies are appealing for the efficient exploration of such large search spaces. Recent works proposed *multi-hypotheses* schemes to alleviate limitations of the local optimization process [1, 2, 3, 4]. The combination of all these factors exacerbates the need for selective and computationally efficient features to evaluate shape and appearance hypotheses on the image.

In this paper, we propose and study two local vascular features. The first one, FLUX , is formulated as the evaluation of the inward gradient flux through a 3D circular cross-section. The second one, MFLUX , introduces a non-linear modification to reduce particular shortcomings of FLUX . These features are designed for the local evaluation of generalized cylinder models [5] with variable radius r and orientation d (left of Fig. 1). Cross-sections are modeled as circular, a reasonable approximation for small scale vessels. The local centerline position p completes the description of the 7D parameter space. Tracking schemes typically evolve such a model *sequentially*, adapting the parameters at discrete steps and evaluating them on the image through features such as FLUX and MFLUX .

For both features, we highlight the links with existing works, discuss their computational efficiency and study their properties with respect to their parameters, in comparison to other selected features. Finally, we illustrate their behavior on 3D coronary CT data.




Fig. 1. Geometric and appearance models. Left: generalized cylinder model and feature parameterization. Right: views of a synthetic image generated from a bar-convolved model of radius 4 voxels and standard deviation $\sigma_m = 1$ voxel.

2. FLUX-BASED FEATURES

2.1. Definitions

Flux-based segmentation methods exploit the orientation of the gradient vectors by computing the gradient flux through the surface of the extracted object. As demonstrated in [6] with a level-set evolution, such approaches are well adapted to the extraction of thin, low-contrast vessels. The *inward*¹ flux is maximized when the surface is aligned with the gradient vector field. In [6], flux measures are approximated through local spheres. An efficient implementation for such spherical flux can be found in [7]. For our parametric model (Fig. 1, left), we can alleviate that approximation by considering circular cross-sections as local surface patches. For slowly narrowing or widening vessels, radial directions give a reasonable approximation of the local normals to the surface. After equi-angular discretization of the cross-section into N points x_i , we obtain our first feature measuring the local, cross-sectional flux contribution:

$$\text{Flux}(p, r, d) = \frac{1}{N} \sum_{i=1}^N \langle \nabla I(x_i), u_i \rangle$$

with $\nabla I(x_i)$ the gradient vector at point x_i and $u_i = \frac{p-x_i}{|p-x_i|}$ the inward radial direction (Fig. 1). As a linear feature, $\text{Flux}(p, r, d)$ is prone to step-edge responses [8]. This effect is especially visible at the surface of large scale hyper-intense structures such as heart chambers in cardiac CT data (Fig. 5). It is characterized by *asymmetric* contributions to the flux along the cross-section. Following the seminal idea of non-linear combination from [8] in 2D, we propose to pair diametrically opposed points (x_i, x_i^π) (Fig. 1) and retain the *minimal* flux contribution per pair. We obtain our second feature:

$$\text{MFlux}(p, r, d) = \frac{2}{N} \sum_{i=1}^{\frac{N}{2}} \min(\langle \nabla I(x_i), u_i \rangle, \langle \nabla I(x_i^\pi), u_i^\pi \rangle)$$

with $x_i^\pi = x_{\frac{N}{2}+i}$ for an even number N of cross-sectional points. In [8], the geometric mean is proposed as a more tolerant alternative to the \min operator. The 3D extension of [8] tests only two

¹We assume hyper-intense vessels.

cross-sectional pairs and again combines responses with a \min . In contrast, we test more points (4 pairs, for $N = 8$ points) and find *averaging* pair responses to be a more adequate, less strict tradeoff between robustness and suppression of step-edge responses.

2.2. Gradient Field Regularization and Related Work

It is preferable to pre-regularize the gradient field for increased robustness. We integrate regularization by isotropic Gaussian smoothing in the numerical scheme for gradient vector computation (Sec. 2.3). We use a standard deviation σ_s of the order of the image resolution, so as to preserve the smallest vessels.

By design, `Flux` and `MFlux` are closely related to other cross-section-based features sharing the seminal ideas from [8]. In [9, 10], Gaussian-smoothed gradient vectors are also projected on radial directions. The *core* feature from [11] computes Gaussian derivatives *oriented* along radial rays, which is theoretically equivalent to the projection method from a continuous point of view. In fact, the main difference between `Flux` and the features from [9, 11] lies in the use of regularization. From the flux point of view, regularization is considered as preprocessing of the gradient vector field and is not necessarily performed by Gaussian filtering. In [8, 9, 10, 11], the Gaussian standard deviation varies with the targeted vessel radius, following principles of Gaussian linear scale-space theory. Regularization induces a tradeoff between detection robustness and discriminative power [11]. Stronger regularization may be desirable for acquisitions with high noise levels or dealing with large-scale vessels, requiring some tolerance to deviations from circular cross-sections. Increased robustness comes however at the detriment of accuracy, as discussed in [11, 12]. Using fixed, limited regularization induces a more discriminative behavior over parameters (see Sec. 3), which is particularly desirable for tracking tasks. It also leads to important computational benefits, as discussed thereafter.

From their oriented flux formulation, our features are also related to the *optimally oriented flux* feature from [13], which can be seen as a flux-based anisotropy descriptor. Differences lie mainly in applicative scopes. `Flux` and `MFlux` are designed to explicitly test given parameters (position, scale, orientation) and explore the data sparsely and efficiently, in tracking setups. On the other hand, the frequency domain implementation of [13] is well adapted to the filtering of whole volumes, using a limited set of pre-selected scales.

2.3. Implementation Details and Computational Performance

The straightforward, naive implementation can be optimized in a number of ways. The 3D gradient field can be jointly computed and regularized using an efficient, recursive implementation of Gaussian derivatives [14]. For a discrete set of orientations and radiuses, offsets to tested cross-sectional points can be pre-computed. We use $N = 8$ points as a tradeoff between speed and robustness. Core computations then reduce to dot products. With our C++ implementation, the evaluation of a cross-section takes on average less than 350 clock cycles, for both `Flux` and `MFlux`, on a 2.16GHz Core Duo processor. This corresponds to more than 6 million cross-sections per second. In fact, the main bottleneck on large datasets is memory access. This cost can be alleviated by a cache-friendly organization of the gradient maps.

3. EXPERIMENTS

We now present experiments illustrating the properties of our features with respect to the parameters. We compare `Flux` and `MFlux`

to the Hessian-based *Vesselness* measure from [15] (with $\alpha = \beta = 0.5$), the *Ribbon* measure from [1] and the *Core* tubular feature from [11]. The *Ribbon* feature measures the local contrast as the difference of mean intensity between a 2D cross-section and an outer annulus. We restrict ourselves to circular cross-sectional areas, where [1] also considered ellipsoids. For the *Core* feature, the standard deviation of the Gaussian used for derivations is set proportionally to the radius as $\sigma = 0.25r$, with the restriction that $\sigma \geq 1$ voxel, as in [11]. For *Ribbon*, *Core*, `Flux` and `MFlux`, negative responses are set to zero². *Vesselness* is positive by design. All responses are normalized between 0 and 1 to ease comparisons.

Synthetic images are generated from a bar-convolved model [9] with a Gaussian standard deviation $\sigma_m = 1$ voxel. This model, illustrated in Fig. 1, is faithful to the appearance of vessels and typical point-spread function of our application of choice, CT coronary segmentation, for which illustrations on real datasets are also presented.

3.1. Scale Parameter vs. Theoretical Radius

All the features of our comparative study have scale-related parameters, linked but not necessarily equal to the theoretical radius of the vessel. For each feature, we numerically derived the mapping between the theoretical radius (from the bar-convolved model) and the optimally responding parameter value, as done in [9]. Besides a more accurate estimation of the vessel radius, this allows fairer comparisons between the different features in our experiments.

3.2. Localization

The first experiment (Fig. 2) illustrates the sensitivity to location eccentricity. Scale and orientation parameters are kept fixed to their optimal values. As expected, all feature responses are maximum at central positions and drop to zero about or slightly outside the vessel wall. More interesting are the differences of behaviors as a function of the target vessel radius. *Vesselness* shows a consistent, scale-independent behavior, a desirable property for its seminal purpose, vessel enhancement [15]. In contrast, the decrease rate of `Flux` mainly depends on the contour blur (σ_m for the bar-convolved model) and on the scale-independent regularization of the gradient field. As a result, `Flux` appears as increasingly discriminative relative to the vessel size. `MFlux` magnifies this advantage by penalizing asymmetric flux contributions arising with eccentricity.

3.3. Radius and Orientation

The top left of Fig. 3 illustrates responses over tested radiuses for central positions and fixed orientation. In such symmetric situations, `Flux` and `MFlux` behave identically. Their responses are roughly symmetric about the correct radius. Because of their contour-based scheme and their scale-independent regularization, they drop faster than the other features for tested radiuses larger than the true value.

Responses for varying orientation are given on the top right of Fig. 3, for central positions and fixed tested radius. *Vesselness* is left out of this experiment for its lack of an orientation parameter³. `Flux` and `MFlux` seem marginally more sensitive to orientation changes than *Ribbon* or *Core*, although none shows a high discriminative power. Sensitivity to orientation could be increased by

²For data such as cardiac CTA, negative responses happen only sparsely, for particular geometric configurations.

³One could imagine a modified *Vesselness* measure to explicitly test a given orientation by orientating the Hessian computation. In its classical implementation, the Hessian basis is fixed by the volume axes and estimations of the principal vessel orientations are given by the matrix inversion.




Fig. 2. Responses over positions. From left to right: bar-convolved straight cylinder of radius 2, 4 and 8 voxels, respectively. Please note the differences in x-axis scales.




Fig. 3. Responses over radii and orientations. Bar-convolved straight cylinder of radius 4. Top left: response over radii. Top right: response over orientations (angle deviation). Bottom row: joint variations of radii and orientations.

testing several successive cross-sections instead of a single one, *i.e.*, by exploiting the local *elongation* of the vessel. However, such a local cylindrical pattern deviates from the generalized cylinder model of Fig. 1 and can cause drops of response for highly curved vessels.

Joint variations of scale and orientation are depicted in the bottom of Fig. 3. Besides further illustration of the higher selectivity of Flux and MFlux, one can note the trend of increasing scale and increasing angle deviation compensating each other.

3.4. Joint Variations of Radius and Position

Fig. 4 shows the response landscapes obtained when jointly varying the tested radius and the tested position. Of particular interest, this experiment illustrates the step-edge issue previously mentioned. It arises for small scales and eccentric positions and is particularly visible for Core and Flux. A similar effect is found for large scales and eccentric positions, more markedly for Ribbon and Core. The non-linear component introduced by MFlux dramatically reduces these issues, resulting in a much tighter response landscape.

3.5. Illustrations on Real Data

In Fig. 5, we present results on subvolumes from real cardiac CT data. For each voxel of the result images, the intensity encodes the maximum response over tested radii and orientations (just radii for Vesselness). This corresponds to a *projection* of the response landscape, equivalent to the filtering setup in which Vesselness is classically used. We emphasize that we do not advocate the use of our features for enhancement filtering. The method of [13] would be more adapted. Our purpose here is simply to illustrate the selective behavior of the features. We use 10 radii, equally spaced between 1 and 10 voxels, corresponding to the typical radius range of the coronary arteries in CT data⁴. 500 test orienta-

⁴With intra-slice resolutions of 0.3×0.3 mm, average coronary branches typically exhibit radii around 1.2mm/4 voxels.

tions are pseudo-uniformly distributed on the unit hemisphere.

As expected, Vesselness produces the smoothest result, as its Hessian-based, Gaussian scale-space design is well adapted to filtering applications. One can note nonetheless the well-known drops of response at branchings and the weak response for smallest branches. For cross-sectional features, drops at branchings are more limited and smaller branches are better preserved. High responses are more tightly concentrated around the vessel centerlines, confirming the experiments on synthetic data and their increased potential for tracking tasks. Responses of Ribbon and Core appear smoother and less concentrated than those of Flux and MFlux. This is due, for Ribbon, to its region-based and averaging nature, and for Core, to its use of increased regularization for larger radii. This correlates with their generally lower sensitivity over positions (Fig. 2) and over tested radii (Fig. 3). The row corresponding to multi-planar reformation views (MPR 1) clearly exhibits the step-edge issue arising for Ribbon, Core and Flux at the surface of heart chambers. For MFlux, this effect is dramatically reduced while the response at true vessel locations is nearly unaffected. MFlux still responds in some non-vascular areas, especially at elongated tips of ventricles.

Noise robustness depends mostly on the regularization of the gradient field. Experimentally, we found $\sigma_s = 1$ voxel to be satisfactory until an additive Gaussian noise level of 200 H.U., about 4 times the typical level of our data. From the tracking point of view, it is also worth noting that noise is essentially an issue for very small vessels.

4. CONCLUSION AND PERSPECTIVES

In this paper, we have presented two features for 3D vascular tracking. The first one, Flux, is formulated as the measure of inward gradient flux through a local circular cross-section. The second one, MFlux, introduces a non-linear component to reduce false positive responses in situations such as step-edges. We discussed the links and differences of these features with existing designs and illustrated their properties on both synthetic and real cardiac CT data. We notably showed that MFlux dramatically reduces false positives and compares favorably to a selection of existing works in terms of discriminative behavior over the parameter space.

The implementation of these features is straightforward and highly efficient. This key characteristic opens very interesting perspectives for the design of advanced, massively multi-hypotheses yet practical tracking strategies, which we are currently investigating.

5. REFERENCES

- [1] C. Florin, N. Paragios, and J. Williams, “Particle filters, a quasi-Monte Carlo solution for segmentation of coronaries.” in *Med. Image Comput. Assist. Interv.*, 2005, pp. 246–253.
- [2] O. Friman, M. Hindennach, and H.-O. Peitgen, “Template-based multiple hypotheses tracking of small vessels.” in *Proc. IEEE Int. Symp. Biom. Imaging*, 2008, pp. 1047–1050.




Fig. 4. Joint variations of tested radius and position. Test for a bar-convolved straight cylinder of radius 4. From left to right: Vesselness, Ribbon, Core, Flux, MFlux.

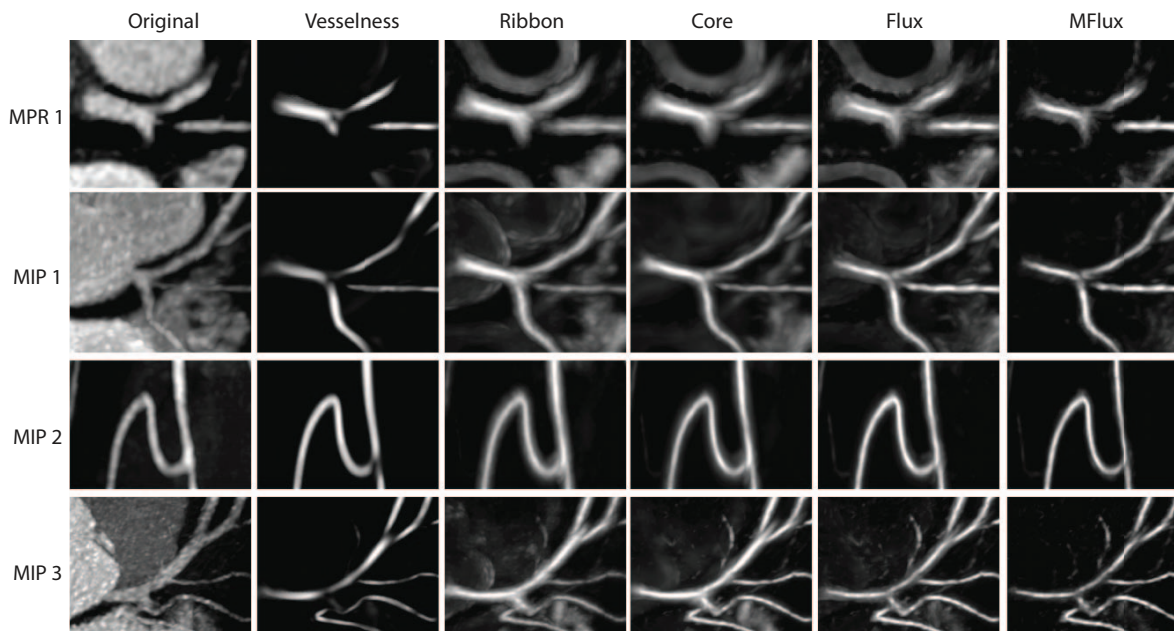


Fig. 5. Illustration of the responses on cardiac CT data. Results for 3 volumes of interest ($128 \times 128 \times 128$ voxels) from 3 different datasets. For dataset 1, multi-planar reformation views (MPR 1) are obtained through the center of the volume of interest. For the three datasets, maximum intensity projection (MIP) views are shown.

- [3] D. Lesage, E.D. Angelini, I. Bloch, and G. Funka-Lea, "Medial-based bayesian tracking for vascular segmentation: Application to coronary arteries in 3d CT angiography," in *Proc. IEEE Int. Symp. Biom. Imaging*, 2008, pp. 268–271.
- [4] M. Schaap, R. Manniesing, I. Smal, T. van Walsum, A. van der Lugt, and W. J. Niessen, "Bayesian tracking of tubular structures and its application to carotid arteries in CTA," in *Med. Image Comput. Assist. Interv.*, 2007, pp. 562–570.
- [5] T. O'Donnell, T. Boult, X. Fang, and A. Gupta, "The extruded generalized cylinder: A deformable model for object recovery," in *Proc. IEEE Conf. Comput. Vision and Pattern Recognit.*, 1994, pp. 174–181.
- [6] A. Vasilevskiy and K. Siddiqi, "Flux maximizing geometric flows," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 12, pp. 1565–1578, 2002.
- [7] MWK Law and ACS Chung, "Efficient implementation for spherical flux computation and its application to vascular segmentation," *IEEE Trans. Image Process.*, vol. 18, no. 3, pp. 596–612, 2009.
- [8] T. Koller, G. Gerig, G. Székely, and D. Dettwiler, "Multiscale detection of curvilinear structures in 2D and 3D image data," in *Proc. IEEE Int. Conf. Comput. Vision*, 1995, pp. 864–869.
- [9] K. Krissian, G. Malandain, N. Ayache, R. Vaillant, and Y. Troussset, "Model-based detection of tubular structures in 3D images," *Comput. Vision and Image Underst.*, vol. 80, no. 2, pp. 130–171, 2000.
- [10] O. Wink, W. J. Niessen, and M. A. Viergever, "Fast delineation and visualization of vessels in 3D angiographic images," *IEEE Trans. Med. Imaging*, vol. 19, no. 4, pp. 337–346, 2000.
- [11] Y. Fridman, *Extracting Branching Object Geometry via Cores*, Ph.D. thesis, University of North Carolina, 2004.
- [12] M. A. Gulsun and H. Tek, "Robust vessel tree modeling," in *Med. Image Comput. Assist. Interv.*, 2008.
- [13] M.W.K. Law and A.C.S. Chung, "Three dimensional curvilinear structure detection using optimally oriented flux," in *Eur. Conf. Comput. Vision*, 2008, pp. 368–382.
- [14] R. Deriche, "Recursively implementing the Gaussian and its derivatives," in *Proc. IEEE Int. Conf. Image Process.*, 1992, pp. 263–267.
- [15] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever, "Multiscale vessel enhancement filtering," in *Med. Image Comput. Assist. Interv.*, 1998, vol. 1496, pp. 130–137.