Linköping University Post Print

Blood vessel segmentation using multi-scale quadrature filtering

Gunnar Läthén, Jimmy Jonasson and Magnus Borga

N.B.: When citing this work, cite the original article.

Original Publication:

Gunnar Läthén, Jimmy Jonasson and Magnus Borga, Blood vessel segmentation using multi-scale quadrature filtering, 2010, Pattern Recognition Letters, (31), 8, 762-767.

<http://dx.doi.org/10.1016/j.patrec.2009.09.020>

Copyright: Elsevier Science B.V., Amsterdam.

<http://www.elsevier.com/>

Postprint available at: Linköping University Electronic Press

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-21046>

Blood vessel segmentation using multi-scale quadrature filtering

Gunnar Lthn^{a,c,1}, Jimmy Jonasson^b, Magnus Borga^{b,c}

^a*Department of Science and Technology, Linkping University, Campus Norrkping, SE-601 74 Norrkping, Sweden*

^b*Department of Biomedical Engineering, Linkping University, SE-581 85 Linkping, Sweden*

^c*Center for Medical Image Science and Visualization, Linkping University, SE-581 85 Linkping, Sweden*

Abstract

The segmentation of blood vessels is a common problem in medical imaging and various applications are found in diagnostics, surgical planning, training and more. Among many different techniques, the use of multiple scales and line detectors is a popular approach. However, the typical line filters used are sensitive to intensity variations and do not target the detection of vessel walls explicitly. In this article, we combine both line and edge detection using quadrature filters across multiple scales. The filter result gives well defined vessels as linear structures, while distinct edges facilitate a robust segmentation. We apply the filter output to energy optimization techniques for segmentation and show promising results in 2D and 3D to illustrate the behavior of our method. The conference version of this article received the best paper award in the bioinformatics and biomedical applications track at ICPR 2008.

Key words: image segmentation, blood vessels, medical imaging, multi-scale, quadrature filter, level set method

1. Introduction

The segmentation of blood vessels has been the focus of much research. The results have numerous applications in medical diagnostics, surgical planning, simulation and training. It is well accepted that no general purpose segmentation method is suitable for all applications and imaging modalities. Thus, researchers have approached this problem from many different angles, applying very different strategies. However, the basic problems of varying vessel width, bifurcations and changing contrast are common challenges in all fields. Overviews of the large body of research are presented in (Suri et al., 2002; Kirbas

Email address: `gunnar.lathen@itn.liu.se` (Gunnar Lthn)

¹Corresponding author (telephone: +46 11 36 32 53, fax: +46 11 36 32 70)

and Quek, 2004), where the authors provide categorizations and comparisons between methods. Roughly, the approaches can be classified as either skeleton based (extracting vessel centerlines) or non-skeleton based (extracting vessel walls). Among these, researchers have applied various filtering techniques, deformable models, region growing and mathematical morphology, to mention a few examples.

A popular strategy across many disciplines is to apply multi-scale approaches for dealing with the problem of varying vessel width. In this work, we also use this simple, but effective approach combined with filtering techniques. Previous work in the same direction has largely been dominated by line or Hessian based filters, see e.g. (Sato et al., 1997). The output of such filters are highly sensitive to noise and varying contrast, which are typical features in medical imaging. In contrast, (Wong and Chung, 2005) used quadrature filters for local structure estimation in a Bayesian segmentation framework. However, they did not incorporate multi-scale analysis, so problems with varying vessel width are expected. In our work, we extend this basic idea of quadrature filtering to multiple scales to handle the problem of varying vessel width. A related approach in (Li et al., 2006) proposed the use of multi-scale Gabor filters for line detection. Our work additionally gives edge detection for improving the segmentation of the vessel walls. Furthermore, instead of using multiple scales purely for noise reduction as in (Li et al., 2006), we integrate all scales to extract the maximum amount of information possible. This idea was first presented in (Lathen et al., 2008). Here we further study the properties of the method and provide more theory and experimental testing.

The article begins by introducing the idea of quadrature filters in Section 2. We explain how different filter directions and scales are combined to produce a multi-scale filter response which captures lines and edges with high precision. Then, Section 3 proceeds by presenting how this result can be used for segmenting the vessel walls. We stress that the result can be used for various different segmentation techniques. The most flexible tool is the level set framework, but we will also give examples using simple thresholding and traditional snakes. In Section 4 and Section 5 we present experiments and results illustrating the behavior of our method. Finally, we conclude and present ideas for future work in Section 6.

2. Quadrature filters

Quadrature filters have a long history in signal processing and have been successfully applied for local structure estimation (Granlund and Knutsson, 1995). Formally, a quadrature filter is defined in the Fourier domain by

$$F_k(\mathbf{u}) = 0, \quad \mathbf{u} \cdot \mathbf{n}_k \leq 0 \quad (1)$$

where $\mathbf{u}$ is frequency and $\mathbf{n}_k$ the filter direction. In other words, the filter is zero over one half of the Fourier domain. In the spatial domain, the filter can be viewed as a complex filter pair, where the real and imaginary parts are line and

Figure 1: Quadrature filter components.

edge filters respectively. Typical one-dimensional filter kernels are illustrated in Figure 1(a) and Figure 1(b). When the filter encounters a structure matching the line filter, the response will be dominantly real, while any edge matches will yield an imaginary response. Thus, the type of structure (line/edge) can be determined by the argument θ of the complex response, as indicated in Figure 1(c). This argument is typically referred to as the *local phase*. In addition, the strength of the structure is given by the magnitude of the response. The separation of structure determination in θ and confidence measure in the magnitude yields a line/edge detector invariant to image intensity, giving increased robustness compared to simpler line filters or gradient based approaches.

2.1. Combining filter directions

For quadrature filtering of 2D and 3D images, the filter kernel is applied in at least 3 and 6 uniformly distributed directions respectively. This produces several filter responses, in which the local phase is a characteristic along the different directions. When combining the filter responses to produce an orientation invariant phase map, the different directions introduce an ambiguity. A filter oriented along the image gradient, gives a phase tending to 90° . On the contrary, a filter oriented against the image gradient, gives a phase tending to -90° . Thus, the local phase of the two filter directions will produce cancellation effects, which will arise as discontinuities in the combined orientation invariant phase map. To counteract this effect, we compute the orientation of the dominant structure at each point, and flip the phase along the real axis for a filter with direction opposing the orientation. In practice, this will yield a phase map with strictly non-negative imaginary components. As a result, we can produce an orientation invariant phase map by summing the filter responses for all directions.

2.2. Multi-scale integration

Our method uses the common approach of multi-scale filtering to handle vessels of varying width. In order to use the result for segmentation, we want to combine all scales in a *global phase map*. We achieve this by a weighted summation, favoring scales of high strength. Formally, the global phase map q

Figure 2: Filtering of a test image on multiple scales.

is computed as:

$$q = \frac{\sum_{i=1}^N |q_i|^\beta q_i}{\sum_{i=1}^N |q_i|^\beta} \quad (2)$$

where N is the number of scales, q_i is the phase map for each scale and β is a weight parameter. As an example, we created a synthetic test image in Figure 2(a) which illustrates a vessel of varying width and orientation. This example contains normally distributed noise with a signal-to-noise ratio (SNR) of 10dB. The filtering on multiple scales and the integrated result are shown in Figure 2, where we used the weight parameter $\beta = 3$. If not stated otherwise, we use four filters of directions $\{0^\circ, 45^\circ, 90^\circ, 135^\circ\}$, with kernel size of 15×15 , bandwidth of 4 octaves and center frequency of $5\pi/7$ for all experiments in 2D. The color of the phase map represents the local phase, illustrated in Figure 1(c). Assuming bright intensity objects and dark background, line structures are indicated by green colors while edges are blue. Note that the result in Figure 2(f) shows a clear line structure and distinct edges, even for an object with large width variation. The choice of β has implications for the segmentation results, which will be further elaborated in Section 3.1.

3. Segmentation

For segmenting the vessels, we use the fact that a local phase of 90° indicates edge structures. Thus, the key idea is to only consider the *real* part of the filter response. This will yield positive and negative values for the inside and outside of line structures, while edges are found at the zero-crossing. Thus, the simplest method for segmenting the vessels is to extract the zero level set from the real part of the phase map. We will use this simple threshold approach to discuss an appropriate choice of weight parameter β for the multi-scale integration.

3.1. Optimal weight parameter for multi-scale integration

Since we know the correct location of the edges for our synthetic test image in Figure 2(a), we can quantify the error produced by different choices of β in Eq. (2). Given a threshold segmentation of the real part of the phase map, we compute the error metric by counting the number of incorrectly labelled pixels within a four-pixel wide band around the true edge. While varying the noise

Figure 3: Plot of error produced by different choices of β in (a). Figures (b)-(d) display the input image with SNR of 5.2dB, 2.2dB and 0.5dB (from left to right).

Figure 4: Illustration of the weight parameter β for the multi-scale integration. The examples show the real part of the filter response for different β accompanied by a threshold segmentation to display the zero-crossing. The zoomed images display the true edge for reference.

level in the image, this produces the error curves in Figure 3. We see that for low levels of noise we can identify optimal choices of $0 < \beta < 5$, while $\beta = 0$ is suitable for higher noise levels. Also note that since the true edge is smooth, while the threshold segmentation is binary, the results suffer from quantization issues. However, this problem is consistent for all β , and does not affect the shape of the error curves.

To discuss the effect of different β , we show a set of results at noise level 10dB in Figure 4. First note that the choice of $\beta = 0$ gives an unweighted summation over all scales in Eq. (2). Thus, the real part of the sum (used for segmentation) can be expressed as the sum of the real parts of each scale. This means that we only consider the real part of the quadrature filter, which is a line filter. In other words, the choice of $\beta = 0$ gives the result of filtering with a line filter, or Laplacian, over multiple scales. By intuition, this should give good line detection performance, but fail in capturing edges. For low noise levels, as can be seen in Figure 4(a) and Figure 4(d), the hypothesis is true for thin lines, although wider lines are more precise. However, stepping away from this simple model to $\beta > 0$ gives visible improvements as shown in Figure 4(e). On the other hand, the extreme value of $\beta = 20$ acts as a “max” operation, only selecting the scale of highest magnitude for each pixel. This explains the poor result for the vertical edge in Figure 4(f), where the edge filter of finer scales is discarded in favor of a coarse scale.

Considering increasing levels of noise, Figure 3 shows that the simpler model of $\beta = 0$ is the best choice. This can be explained by the increased magnitude of noise detected in fine scales. When $\beta > 0$, the larger noise magnitude can dominate and suppress the response from coarser scales, which contain more accurate structural information. For real images, this problem can be handled by pre-filtering to remove noise, or tuning of filter parameters, to mention a few possible approaches. For this experiment however, the aim was to study the relation between the segmented result and β with as few parameters as possible. Note that we kept several parameters fixed, such as filter bandwidth, center frequency, the number and orientation of filter kernels, and possible rotations of the input image. For real applications, the filter parameters are usually optimized to suit the particular problem at hand. Thus, assuming parameters optimized for the noise level, it is expected that choices in the range $0 < \beta < 5$ performs well.

3.2. Segmentation using energy optimization

As can be seen in Figure 4, the segmented results when using simple thresholding is very susceptible to noise. The standard solution to this problem is to incorporate means of regularization. This can be elegantly solved using energy optimization techniques. For 2D images, we use an energy well studied in the field of variational methods for image segmentation:

$$E(C(s)) = \iint_{\Omega_C} f(x, y) dx dy - \alpha \oint_C ds \quad (3)$$

Figure 5: Snakes segmentation on the test image.

where C is a one-dimensional curve representing the segmented boundary and Ω_C is the interior of this curve. When maximizing this energy, the goal is to find a curve C which encloses all positive values of the function f (first integral), while at the same time the length of the curve is minimized (second integral). For our application, the target function f is the real part of the phase map, so we set $f(x, y) = \text{Re}(q)$. The α parameter controls the amount of regularization and should be set based on the noise level in the image.

From the calculus of variations, stationary points for the energy in Eq. (3) are found by the Euler-Lagrange equation. The standard procedure for finding such stationary points is by an iterative search, where the curve C is evolved in the gradient descent direction of the energy. Following (Kimmel, 2003), the curve evolution given Eq. (3) is:

$$\frac{\partial C}{\partial t} = -\text{Re}(q)\mathbf{n} + \alpha\kappa\mathbf{n} \quad (4)$$

where t is an artificial time parameter, $\mathbf{n}$ is the curve normal and κ is the curvature of the curve. Using a parameterized curve representation for C , this evolution can be directly implemented as a traditional snakes model (Kass et al., 1988). To exemplify this, we used an implementation by Xu and Prince available at <http://iacl.ece.jhu.edu/projects/gvf/>. However, this implementation only allows for a uniform motion in the normal direction of the curve (pressure), so a slight modification was made to allow for a non-uniform pressure field. Targeting the test image in Figure 2(a) with SNR = 10dB, $\beta = 3$, regularization parameter (snake elasticity) $\alpha = 0.1$ and the snake pressure field given by $\text{Re}(q)$, a sequence of iterations is displayed in Figure 5. Compared to the threshold segmentation in Figure 4(b) we note that the regularization has successfully eliminated the influence of noise.

Because of the parameterized curve representation, it is well known that the snakes model suffers from issues with self-intersection and reparameterization. Furthermore, these issues are even more problematic for surfaces in 3D. A more robust solution is given by the level set method, where the curve is represented implicitly as the zero level set of a signed distance function, usually referred to as the level set function. This representation generalizes to any dimension, while self-intersections are handled naturally. To deform the contour, the level

set function ϕ is evolved in time according to a given PDE. As described in (Kimmel, 2003), the curve evolution in Eq. (4) translates to the level set PDE:

$$\frac{\partial \phi}{\partial t} = -\text{Re}(q) |\nabla \phi| + \alpha \kappa |\nabla \phi| \quad (5)$$

We will use this level set model for the experiments described in Section 4.

4. Experiments

The method described in this article has been implemented in Matlab. We use a standard type of level set implementation based on (Osher and Fedkiw, 2003) and (Peng et al., 1999). Reference code and examples can be found online at <http://dmforge.itn.liu.se/pr109/>. For real images with vessels of varying intensity, we apply a normalization procedure to the phase map magnitude in order to make the problem more regular. This is defined as:

$$\hat{a}(\sigma) = 1/(1 + (\sigma/a)^2) \quad (6)$$

where $a = |q|$, q is the global phase map and σ is a data dependent threshold parameter. By applying this normalization we remove scaling issues associated with different inputs, so σ should be viewed as a replacement for more complicated parameters compensating for scaling. Currently, σ is set manually, but can be automatically computed based on analysis of the noise.

For a qualitative comparison, we implemented the “flux maximizing flow” (FMF) method (Vasilevskiy and Siddiqi, 2002), which also relies on a multi-scale approach. Similar to our method, the output gives vessel walls located on zero-crossings, so we can directly use the same energy optimization technique for segmentation.

Our first experiment is a 2D retinal image of size 458×265 from the DRIVE database (Staal et al., 2004), displayed in Figure 6(a) along with the initial curve used for the segmentation. The output of FMF and our phase map are shown in Figure 6(b) and Figure 6(c) respectively. The gray-levels are not visually comparable, but positive and negative values are indicated by bright and dark colors respectively. For the phase map we used 4 filters of size 15×15 with bandwidths of 4 octaves and center frequencies of $5\pi/7$. We used 3 scales by subsampling the image with a factor of $1/2$. The normalization parameter in Eq. (6) was set to $\sigma = 3$. For the level set evolution we used the regularization parameter $\alpha = 0.01$ for both methods.

Our second experiment is a 3D CT dataset of size $512 \times 512 \times 821$, where we target the blood vessels in the liver. To restrict our result to the targeted vessels, a manual segmentation was performed on the liver and used as a mask for the segmentation. For this experiment we used 6 filters of size $7 \times 7 \times 7$, with bandwidths of 2 octaves and center frequencies of $\pi/2$. We used 3 scales with ratios of 2 and parameters $\beta = 1$, $\sigma = 80$ and $\alpha = 0.1$. A sequence of iterations are shown in Figure 7 along with one slice to visually validate the result.

Figure 6: Experiment on 2D retinal image, comparing our approach with the “flux maximizing flow”.

5. Results

Our first experiment included a comparison with the FMF method in Figure 6. Although similar results, we note that our approach exhibits less leakage for small vessels. It can be seen in Figure 6(b) and Figure 6(c) that the phase map contains a stronger “negative force” around the vessels, which prohibits the curve from entering that region. In addition, the computational time for generating the FMF output and the phase map differs by several orders of magnitude in favor of the phase map (31 minutes vs. 1.9 seconds).

Our second experiment applies our method to 3D data. This example shows the ability to grow a large portion of a vessel tree from a small initial seed. Studying the segmentation in slices (one example is shown in Figure 7(e)), the result is visually plausible. Future work includes validating the results more systematically. Note that the data is captured in the arterial phase of the liver, so we are not expected to see the complete vessel tree of the liver in the image.

6. Conclusions and future work

We have presented an approach based on multi-scale quadrature filtering for detecting both clear vessels and distinct vessel walls. Combined with energy optimization techniques for segmentation, we show promising results in 2D and 3D for typical medical images. A comparison with “flux maximizing flow” shows that our method tends to produce more robust segmentations in terms of boundary leakage. Furthermore, our filtering method is easily implemented and the computational complexity is low. Since the results are proof-of-concept, future work includes a systematic validation of the segmentations and quantitative comparisons with other recent methods. In addition, we will study other types of regularization methods for segmentation and apply a wider range of filter parameters to different images.

Figure 7: Experiment on 3D CT dataset targeting the blood vessels in the liver.

References

- Granlund, G. H., Knutsson, H., 1995. Signal Processing for Computer Vision. Kluwer Academic Publishers, Netherlands.
- Kass, M., Witkin, A., Terzopoulos, D., January 1988. Snakes: Active contour models. International Journal of Computer Vision 1 (4), 321–331.
- Kimmel, R., 2003. Geometric Level Set Methods in Imaging, Vision and Graphics. Springer Verlag, Ch. Fast edge integration.
- Kirbas, C., Quek, F., 2004. A review of vessel extraction techniques and algorithms. ACM Comput. Surv. 36 (2), 81–121.
- Lathen, G., Jonasson, J., Borga, M., December 2008. Phase based level set segmentation of blood vessels. In: Pattern Recognition, 2008. ICPR 2008. 19th International Conference on. IAPR, Tampa, FL, USA.
- Li, Q., You, J., Zhang, L., Zhang, D., Bhattacharya, P., 2006. A new approach to automated retinal vessel segmentation using multiscale analysis. In: ICPR. Vol. 4. pp. 77–80.
- Osher, S., Fedkiw, R., 2003. Level Set and Dynamic Implicit Surfaces. Springer-Verlag New York Inc.

- Peng, D., Merriman, B., Osher, S., Zhao, H.-K., Kang, M., 1999. A pde-based fast local level set method. *Journal of Computational Physics* 155 (2), 410–438.
- Sato, Y., Nakajima, S., Atsumi, H., Koller, T., Gerig, G., Yoshida, S., Kikinis, R., 1997. 3d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In: *CVRMed-MRCAS '97: Proceedings of the First Joint Conference on Computer Vision, Virtual Reality and Robotics in Medicine and Medial Robotics and Computer-Assisted Surgery*. Springer-Verlag, London, UK, pp. 213–222.
- Staal, J., Abramoff, M., Niemeijer, M., Viergever, M., van Ginneken, B., 2004. Ridge based vessel segmentation in color images of the retina. *IEEE Transactions on Medical Imaging* 23 (4), 501–509.
- Suri, J. S., Liu, K., Reden, L., Laxminarayan, S., 2002. A review on mr vascular image processing: skeleton versus nonskeleton approaches: part ii. *IEEE Trans Inf Technol Biomed* 6 (4), 338–350.
- Vasilevskiy, A., Siddiqi, K., 2002. Flux maximizing geometric flows. *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (12), 1565–1578.
- Wong, W., Chung, A., Oct. 2005. Bayesian image segmentation using local iso-intensity structural orientation. *Image Processing, IEEE Transactions on* 14 (10), 1512–1523.