

Vessel Scale-Selection using MRF Optimization

Hengameh Mirzaalian^{1,2} and Ghassan Hamarneh¹

¹Medical Image Analysis Lab, Simon Fraser University, BC, Canada.

²Photomedicine Institute, Department of Dermatology and Skin Science, University of British Columbia, and Vancouver Coastal Health Research Institute, BC, Canada.

Abstract

Many feature detection algorithms rely on the choice of scale. In this paper, we complement standard scale-selection algorithms with spatial regularization. To this end, we formulate scale-selection as a graph labeling problem and employ Markov random field multi-label optimization. We focus on detecting the scales of vascular structures in medical images. We compare the detected vessel scales using our method to those obtained using the selection approach of the well-known vesselness filter (Frangi et al 1998). We propose and discuss two different approaches for evaluating the goodness of scale-selection. Our results on 40 images from the Digital Retinal Images for Vessel Extraction (DRIVE) database show an average reduction in these error measurements by more than 15%.

1. Introduction

Many feature detection algorithms rely on the choice of scale. The importance of the scale-selection originates from the basic observation that real-world objects are composed of different structures at different scales. For example, the concept of a *tree* is appropriate at the scale of meters, while concepts such as leaves and molecules are more appropriate at finer scales.

For a computer vision system analyzing an image without a priori knowledge of the scales of interesting structures, it is reasonable to consider image descriptions at multiple scales, in order to capture the unknown scale variations that may occur. Taken to the limit, a scale-space representation considers an image at all possible scales.

The most common scale-space representation of an image $I(x, y)$ is achieved by convolving the image with a Gaussian kernel g_s of a given width or scale s :

$$g_s(x, y) = \frac{1}{2\pi s} \exp(-(x^2 + y^2)/2s) \quad (1)$$

The scale-space representation of the image is then given by:

$$L(x, y; s) = g_s(x, y) * I(x, y) \quad (2)$$

The application of local derivative operators to the image at any scale in scale-space is given by:

$$L_{x^m y^n}(x, y; s) = \partial_{x^m y^n} L(x, y; s) \quad (3)$$

which is equivalent to convolving the original image $I(x, y)$ with Gaussian derivative operators, i.e.

$$L_{x^m y^n}(x, y; s) = \partial_{x^m y^n} g_s(x, y) * I(x, y) \quad (4)$$

These Gaussian derivative operators can in turn be combined linearly or non-linearly into a larger variety of different types of feature detectors. For example, edge, blob, corner, and ridge multi-scale feature detectors expressed through first and second order spatial derivatives in scale-space [21, 33]. Of particular interest in this work is the detection of tube-like structures through the analysis of the Hessian matrix at multiple scales (section 2.2).

Selecting the appropriate scales for the various features in an image is an important problem with many applications in computer vision and image processing [21, 10, 27, 30]. Further, a strong link between scale-space/scale-selection theory and biological vision has been demonstrated in many earlier works [8, 9].

The standard scale-selection approach typically employed is simply to choose, at every pixel in the image, the scale with the highest feature detector response at that pixel. However, since the goal is to assign a scale to each feature of interest in the image, and given that any feature will typically occupy multiple neighboring pixels, it is reasonable to assume that most neighboring pixels in the images should be assigned equal scales. In other words, we expect the scale-selection to give a regularized scale field.

Since the early work of Geman and Geman in 1993 [13], Markov Random Field (MRF) theory has proven useful for encoding neighborhood properties in images. Many computer vision tasks now rely on an underlying MRF model and graph representations encoding spatial relationships between pixels. The computer vision task is subsequently transformed into a MRF multi-label optimization problem, where the goal is to optimally assign one of many labels to

each of the vertices of the graph, which in turn provides the solution for image segmentation [35, 4, 6, 3], image registration [23, 38], stereo matching [32] or object recognition problems [22].

In this paper, we adopt the MRF multi-label optimization approach to address the spatial regularization requirement in scale-selection problems. Our work benefits from the large body of research into MRF optimization e.g. [5, 19, 6, 18, 15] (see [37] for a recent survey).

Given the importance of vasculature in understanding, diagnosing, and treating diseases [7, 26], the particular focus of this work is the analysis of medical images of vasculature through the automatic detection of vessel scales [12, 20]. Multi-scale approaches to vessel enhancement have been proposed in the past, including cores [2], steerable filters [17], and assessment of local orientation via eigenvalue analysis of the Hessian matrix [24, 33]. A well-known multiscale vessel enhancement filtering method has been proposed by Frangi et al. [12] and has gained large popularity for the analysis of medical angiographic images [25, 28, 31, 39].

To the best of our knowledge, the proposed method is the first to incorporate an MRF-based spatial regularization in scale-selection. We evaluate our MRF-based vessel scale-selection method on synthetic data as well as 40 2D-retinal images from the Digital Retinal Images for Vessel Extraction (DRIVE) database [36]. The experimental results show that we obtain improved scale estimates compared to the maximum likelihood scale estimates at each pixel as performed in [12].

2. Method

We solve the scale-selection problem by MRF multi-label optimization, where the goal is to assign to each pixel in the image a ‘scale label’ out of several possible scale labels, while taking into account label-interaction between neighbouring pixels. We begin by summarizing the MRF labeling formulation and the multi-scale vesselness filter.

2.1. MRF Multi-Label Optimization

We represent the pixel connectivity in image I via a graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, in which vertices $v \in \mathcal{V}$ correspond to the pixels of I and edges $e_{v_p v_q} = e_{pq} \in \mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ connect two neighbouring pixels (p, q) and are assigned nonnegative weight w_{pq} . In MRF-optimization, we seek the labeling f_p of each vertex v_p that minimizes an energy of the form:

$$E(f) = (1 - \eta) \sum_{p \in \mathcal{V}} \varphi_p(f_p) + \eta \sum_{(p,q) \in \mathcal{E}} \varphi_{pq}(f_p, f_q) \quad (5)$$

where φ_p is the unary data (or image fidelity) term, which measures the likelihood of labeling a pixel with a specific label, disregarding the labels of any of the neighbors, φ_{pq}

is the binary spatial regularization term penalizing different label configurations of neighboring vertices, and η controls the trade-off between the two terms.

There are different possible choices for the label interaction term φ_{pq} in (5). φ_{pq} can be seen as a metric on the space of labels $\varphi_{pq} = \varphi_{pq}(f_p, f_q)$ or may be chosen to depend on the underlying data $\varphi_{pq} = \varphi_{pq}(I(p), I(q))$, or both $\varphi_{pq} = \varphi_{pq}(I(p), I(q), f_p, f_q)$. Often it is preferred that neighboring pixels are assigned similar or equal labels. Various label interaction penalties have been proposed, including linear ($|f_p - f_q|$), quadratic ($(f_p - f_q)^2$), truncated ($\min\{T, |f_p - f_q|\}$ or $\min\{T, (f_p - f_q)^2\}$) [14, 34]), or data-dependent, e.g. Gaussian ($\exp(-(I(p) - I(q))^2)$). Setting the regularization weight η in (5) can be found empirically using training data.

2.2. Multi-Scale Vesselness

According to Frangi et al. [12], the vesselness measure at pixel $\mathbf{x} = (x, y)$ is obtained by:

$$\nu(\mathbf{x}, s) = \begin{cases} 0 & \lambda_2(\mathbf{x}, s) > 0 \\ e^{-\frac{R_B^2(\mathbf{x}, s)}{2\beta^2}} \left(1 - \exp\left(-\frac{S^2(\mathbf{x}, s)}{2c^2}\right)\right) & o/w \end{cases} \quad (6)$$

with:

$$R_B(\mathbf{x}, s) = \frac{\lambda_1(\mathbf{x}, s)}{\lambda_2(\mathbf{x}, s)} \quad (7)$$

$$S(\mathbf{x}, s) = \sqrt{\sum_{i=2} \lambda_i^2(\mathbf{x}, s)} \quad (8)$$

where $\lambda_i(\mathbf{x}, s)$, $i = 1, 2$ ($|\lambda_1| \leq |\lambda_2|$) are the eigenvalues of the Hessian matrix of image I computed at scale s , i.e. $H(L(\mathbf{x}; s))$. R_B and S are measures of blobness and second order structureness, respectively.

Given K different possible scales, i.e. $s \in \{s^1, \dots, s^K\}$, equation (6) can be evaluated at each of the K scales. Frangi et al. estimate the vesselness ν at every pixel \mathbf{x} using:

$$\nu(\mathbf{x}) = \max_{s \in \{s^1, \dots, s^K\}} \nu(\mathbf{x}, s) \quad (9)$$

which implies that the selected scale \hat{s}_{MVR} , giving the maximum vesselness response (MVR) at \mathbf{x} , is given by:

$$\hat{s}_{MVR}(\mathbf{x}) = \arg \max_{s \in \{s^1, \dots, s^K\}} \nu(\mathbf{x}, s) \quad (10)$$

Clearly, this pixel-wise choice ignores the desired regularization in the scale field. In the next section, we formulate vessel scale-selection as a graph labeling problem and solve it using MRF multi-label optimization.

Figure 1 shows an example retinal image from the DRIVE database [36] and the vessel analysis results according to the method of Frangi et al. [12].

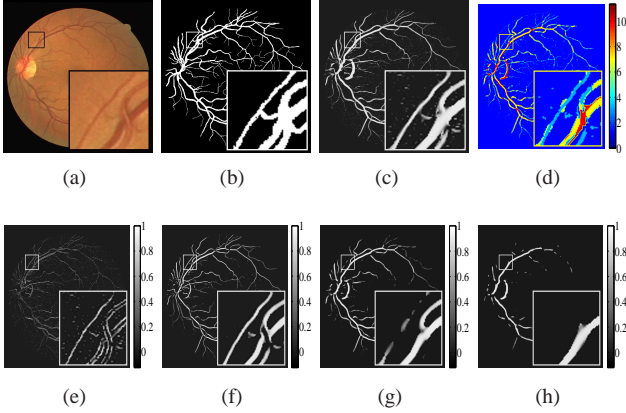


Figure 1. Frangi et al.’s approach to vessel image analysis. (a) Retinal image from the DRIVE database [36]. (b) Expert vessel segmentation. (c) The vessel enhanced image with pixel intensity set to the maximum-vesselness across all scales, calculated according to (9). (d) The selected scales that give maximal vesselness response found using (10) (different scales represented in different colors). (e-h) Frangi et al. [12] vesselness images calculated at different scales, according to (6).

2.3. Vessel Scale-Selection using MRF Labeling

The goal is to assign to each pixel a suitable scale $s \in \{s^1, s^2, \dots, s^K\}$. To this end, we construct a graph, as described in section 2.1, representing the 2D image I , with graph edges representing the 4-connectedness of pixels. We adopt Frangi et al.’s vesselness response at different scales, i.e. $\{\nu(\mathbf{x}, s^1), \nu(\mathbf{x}, s^2), \dots, \nu(\mathbf{x}, s^K)\}$, as the likelihood of choosing the corresponding scale at a particular pixel \mathbf{x} . However, since background pixels do not contain any vessels, it is not meaningful to assign any of these vessel scale labels $\{s^1, \dots, s^K\}$ to background pixels. Therefore, we define an additional label for the background, denoted s^{bg} , with the following likelihood term:

$$\nu(\mathbf{x}, s^{bg}) = \xi \exp\left(-\frac{\bar{\nu}(\mathbf{x})}{\sigma^2}\right) \quad (11)$$

where ξ is a normalization constant ensuring fair comparison between the newly defined data term for s^{bg} and those for the previous scales $\{s_i\}_{i=1}^K$. We set ξ to:

$$\xi = \max_{s \in \{s^1, \dots, s^K\}} \nu(\mathbf{x}, s) \quad (12)$$

and $\bar{\nu}(\mathbf{x})$ is the average vesselness across scales

$$\bar{\nu}(\mathbf{x}) = \frac{1}{K} \sum_{i=1}^K \nu(\mathbf{x}, s^i) \quad (13)$$

Consequently, pixels with low vesselness responses will have a high likelihood of belonging to the background, i.e. high $\nu(\mathbf{x}, s^{bg})$, and vice versa.

The unary penalty term in (5) is then set as:

$$\varphi_p(f_p = s^i) = \max_{s^j \in \{s^1, \dots, s^K, s^{bg}\}} \nu(p, s^j) - \nu(p, s^i) \quad (14)$$

We adopt $\varphi_{pq}(f_p, f_q) = \min\{1, |f_p - f_q|\}$, i.e. Pott’s model [1], as the regularization penalty. We use α -expansion graph cuts [1] to find the optimal MRF labeling. This will assign a scale $s \in \{s^1, s^2, \dots, s^K, s^{bg}\}$ to every pixel in the image.

3. Results

3.1. Evaluation Approaches

Let $\hat{s}_{MVR}(p)$, $\hat{s}_{MRF}(p)$, and $s^*(p)$ denote the estimated scale using Frangi et al.’s maximum vesselness response (MVR) (10), our MRF-based method, and the correct (ground truth) scale, respectively, at pixel p of image I .

In order to calculate the ground truth vessel scale $s^*(p)$ at each pixel p , we first calculate the Euclidean distance $d(b)$ between each vessel boundary pixel b and the closest vessel centreline pixel (utilizing the distance transform of the vessels’ skeleton image). Then, the pixel p is assigned $s^*(p)$ equal to $d(b)$ of the nearest boundary pixel to p (see Figure 2 for an illustration of this procedure). This approach assumes the same ground truth scale along all pixels perpendicular to the vessel centerline. An alternative approach is to restrict the vessel’s ground truth scale and the subsequent error analysis to only the vessel centrelines. This ground truth scale calculation approach is used for both synthetic images (Section 3.2) and real retinal images (Section 3.3).

Given the ground truth scale s^* , we propose and calculate two types of scale-selection error at a pixel p :

$$e_\delta(p) = 1 - \delta(\hat{s}(p) - s^*(p)) \quad (15)$$

$$e_\Delta(p) = (\hat{s}(p) - s^*(p))^2 \quad (16)$$

The error e_δ returns 0 if the estimated scale \hat{s} is equal to the ground truth scale s^* and 1 if the scales are different, regardless of how much different they are. The error e_Δ , on the other hand, considers the difference between the estimated and ground truth scales. We further distinguish between two sets of pixels: pixels that reside inside the linear structures of interest i.e. foreground or vessel pixels, and those that do not, i.e. background pixels. We denote these two sets of pixels as P_{fg} and P_{bg} . $P = P_{fg} \cup P_{bg}$ is the set of all pixels in the image I . The following four error measures are used to evaluate the scale-selection in the subsequent experiments:

$$e_\delta = \frac{1}{|P|} \sum_{p \in P} e_\delta(p) \quad (17)$$

$$e_\delta^{fg} = \frac{1}{|P_{fg}|} \sum_{p \in P_{fg}} e_\delta(p) \quad (18)$$




Figure 2. Calculating the ground truth vessel thickness. (a) Binary vessel mask. (b) Skeleton of (a). (c) Distance transform of (b) denoted by I_{dt} . (d) and (e) are close-ups of (c) and (a), respectively. (f) The scale of each pixel $p \in P_{fg}$ inside the vessel is set to the value of the $I_{dt}(q)$, where q is the closest edge pixel to p , i.e. $s^*(p) = I_{dt}(q)$. (g) Quantized version of s^* .




Figure 3. (a) Generated synthetic image (tubular structures with the different scales). (b) The ground-truth for the thickness (different thickness values are represented by different colors).

$$e_{\Delta} = \frac{1}{|P|} \sum_{p \in P} e_{\Delta}(p) \quad (19)$$

$$e_{\Delta}^{fg} = \frac{1}{|P_{fg}|} \sum_{p \in P_{fg}} e_{\Delta}(p) \quad (20)$$

Note that e_{Δ} and e_{Δ}^{fg} are defined for s^* (Figure 2(f)), whereas calculating e_{δ} or e_{δ}^{fg} , requires that we first quantize s^* into $K + 1$ different values (Figure 2(g)).

3.2. Synthetic Data

To evaluate our proposed MRF-based scale-selection, we generate a synthetic binary image containing several lines with different thicknesses (Figure 3) and measure the scale-selection errors by (17)-(20) for different Gaussian noise levels added to the image (Figure 4). Figure 5 shows e_{Δ}^{fg} and e_{δ}^{fg} for both $\hat{s}_{MVR}(p)$ and $\hat{s}_{MRF}(p)$ for the K labeling problem ($s \in \{s^1, s^2, \dots, s^K\}$). The different colored curves correspond to different spatial regularization values $\eta \in [0, 1]$ in (5). It can be noted that we obtain improved scale estimates compared to MVR. We also note the local minima in the error curves of the MRF approach (Figures 5(c)-5(d)). The reason for these minima is that the amount of spatial regularization should ideally be related the level of noise in the image [5, 29], i.e. for a given noise level, a specific choice of η will give the least error.

Figure 6 illustrates e_{Δ} and e_{δ} , the scale-selection error for the set of all pixels in the image ($p \in P_{fg} \cup P_{bg}$), for the $K + 1$ labeling problem ($s \in \{s^1, s^2, \dots, s^K, s^{bg}\}$), i.e. the case in which we include the additional background label mentioned in section 2.3. It can be seen that our proposed MRF approach yields lower error (e_{Δ} and e_{δ}) compared

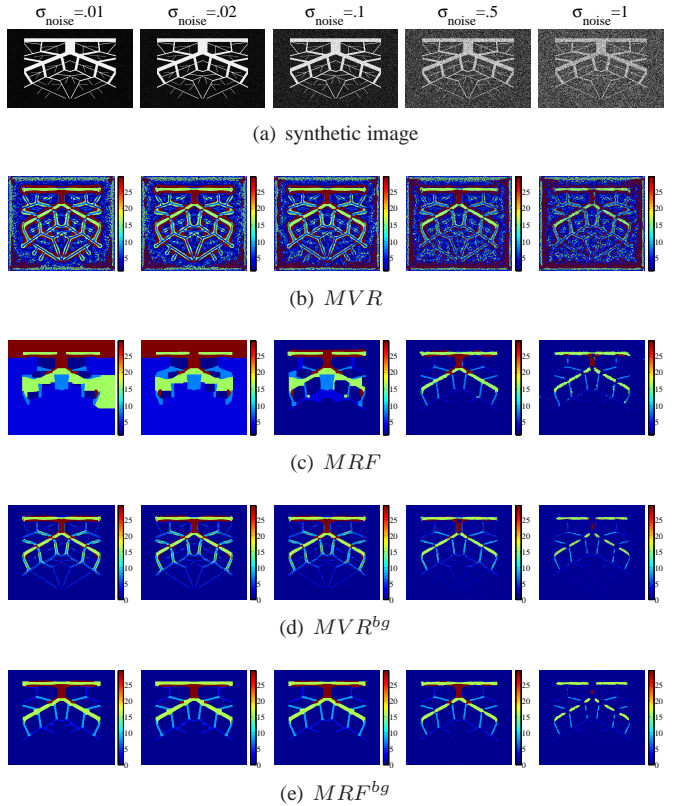


Figure 4. Results on synthetic data. The different columns (from left to right) correspond to increasing levels of noise added to the image. The rows (from top to bottom) show: The noisy image, estimated scale based on MVR, MRF, MVR with a background (BG) label, and MRF with a BG label

to MVR. Also, we obtain lower e_{Δ}^{fg} and e_{δ}^{fg} (error measured for the pixels inside the vessels) compared to MVR. These results are omitted as they are similar to those in Figures 5(a) and 5(b). In Figure 6(a), and in particular for $\sigma_n \in [0, 0.4]$, we observe a decrease in error as the noise level increases. Surprising at first, this is explained again by the fact that a specific spatial regularization level η is ideal for a certain noise level. So, for a fixed η , if the noise is negligibly small, then the error will be high since it is favorable to avoid regularization in this case. However, as the noise increases, the regularization matches the noise and the error will decrease.




Figure 5. Scale-selection error based on MVR and MRF (K labeling problem) for increasing levels of additive Gaussian noise σ_n and different regularization η . (a-b) show e_{Δ}^{fg} and e_{δ}^{fg} , respectively. The dashed curve corresponds to MVR error and the other colored curves correspond to MRF error for different η . (c-d) show e_{Δ}^{fg} and e_{δ}^{fg} versus η with MVR shown as dashed lines. Different colors correspond to different noise levels.

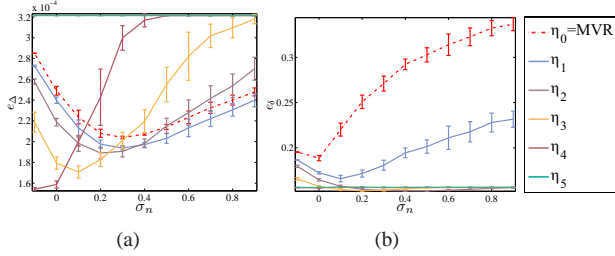


Figure 6. (a-b) Show e_{Δ} and e_{δ} (similar to Figure 5) when we define a label for the background ($K + 1$ labeling problem). Different colors correspond to different regularization levels.

Figure 7 shows lower scale-selection error for pixels inside the vessels ($p \in P_{fg}$) for the K labeling problem compared to the $K + 1$ labeling (when we add a background label). This behavior is expected because this additional background label would decrease the chances of assigning the correct scale to a vessel pixel. In other words, adding the background label increases the accuracy of assigning scale s^{bg} to background pixels (i.e. identifying the background pixels) at the expense of decreasing the scale-selection accuracy of the foreground vessel pixels.

In summary, the results for synthetic data show an average reduction in error measures e_{Δ}^{fg} and e_{δ}^{fg} of about 14% and 17%, respectively, over all noise levels. We also evaluated the scale-selection for only the centreline pixels and obtained, respectively, 14% and 18% reduction in error.




Figure 7. Red (respectively blue) represent e^{fg} when we use (respectively, do not use) the background label. (a) e_{Δ}^{fg} resulting from MVR. (b) e_{δ}^{fg} resulting from MVR.

3.3. Real Data

We evaluate our method on 40 2D-retinal images from the DRIVE database [36], in which expert binary segmentations of the images are available. Using the same method shown in Figure 2, we calculate the ground truth scale (vessel radii) $s^*(p)$ at each pixel of each image. Figure 8 shows the values of s^* for a set of sample retinal images.

A comparison between the scale-selection error based on MVR and based on our MRF method for the retinal images is shown in Figure 9. Our approach decreases the scale-selection errors e_{Δ}^{fg} and e_{δ}^{fg} on average by 18% and 15%, respectively. Limiting the evaluation to the centerline pixels only results in 4% and 11% reduction in error. Note that for the DRIVE database, the best results are achieved for $\eta = 0.2$.

4. Conclusion

We formulated scale-selection as a graph labeling problem and employed Markov random field multi-label optimization. Due to the clinical importance of measuring vessel thickness, e.g. for detecting stenoses, we focused on scale-selection in vascular images. Nevertheless, we foresee MRF optimization being utilized for other scale-dependent features. Although we only showed the development of our technique for 2D data, this algorithm can be easily generalized to handle 3D images (e.g. magnetic resonance or computed tomography angiography), since 3D vesselness functions [12] and MRF optimization methods are readily available for 3D.

In this work, even with the regularization weight set to a constant, we achieved improved results over Frangi et al.'s scale-selection method. However, a single value may not be optimal for all regions of the image. Spatially adapting the regularization weights may improve the scale-selection even further, similar to improvements obtained in spatially adaptive regularization for image segmentation and denoising applications [11, 16, 29].

Although we adopted the Pott's penalty and α -expansion graph cuts in our implementation, other label interaction

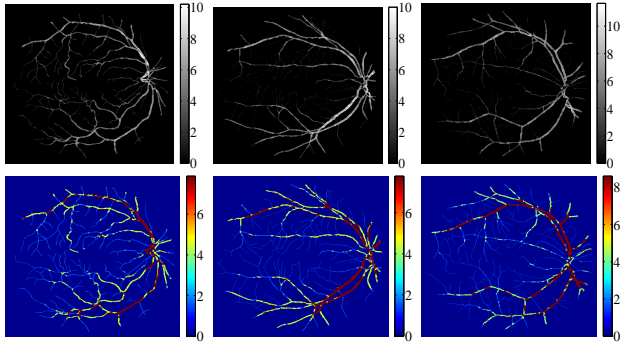


Figure 8. Some examples of s^* (ground truth vessel thickness) for images from the DRIVE database [36]. (top) fuzzy and (bottom) quantized s^* .

penalties (e.g. those mentioned in section 2.1) and a variety of alternative optimization methods can be used instead [37].

References

- [1] 3
- [2] S. Aylward, S. Pizer, D. Eberly, and E. Bullitt. Intensity ridge and widths for tubular object segmentation and description. pages 131–138, 1996. 2
- [3] C.A. Bouman and M. Shapiro. A multiscale random field model for bayesian image segmentation. *IEEE TIP*, 3(2):162–177, 1994. 2
- [4] Y. Boykov and G. Funka-Lea. Graph cuts and efficient n-d image segmentation. *IJCV*, 70(2):109–131, 2006. 2
- [5] Y. Boykov and M. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in n-D images. *IEEE ICCV*, 1:105–112, 2001. 2, 4
- [6] Y. Boykov, O. Veksler, and R. Zabih. Markov random fields with efficient approximations. In *IEEE CVPR*, pages 648–655, 1998. 2
- [7] E. Bullitt, G. Gerig, S.R. Aylward, S.C. Joshi, K. Smith, M. Ewend, and W. Lin. Vascular attributes and malignant brain tumors. *MICCAI*, 1:671–679, 2003. 2
- [8] H. Chen and J. Hale. An algorithm for MR angiography image enhancement. *MRM*, 33(4):534–540, 1995. 1
- [9] G. DeAngelis, I. Ohzawa, and R. Freeman. Receptive-field dynamics in the central visual pathways. *Trends in Neurosciences*, 18(10):451–458, 1995. 1
- [10] F. Dell’Acqua, P. Gamba, and G. Trianni. Semi-automatic choice of scale-dependent features for satellite sar image classification. *Pattern Recognition Letters*, 27(4):244–251, 2006. 1
- [11] E. Erdem and S. Tari. Mumford-shah regularizer with contextual feedback. *J. Math. Imaging Vis.*, 33(1):67–84, 2009. 5
- [12] Ro F. Frangi, Wiro J. Niessen, Koen L. Vincken, and Max A. Viergever. Multiscale vessel enhancement filtering. In *MICCAI*, pages 130–137, 1998. 2, 3, 5
- [13] S. Geman, D. Geman, K. Abend, T. J. Harley, and L. N. Kanal. Stochastic relaxation, gibbs distributions and the bayesian restoration of images. *Journal of Applied Statistics*, 20(5):25–62, 1993. 1
- [14] H. Ishikawa. Exact optimization for markov random fields with convex priors. *IEEE TPAMI*, 25(10):1333–1336, 2003. 2
- [15] P. Kohli, A. Shekhovtsov, C. Rother, V. Kolmogorov, and P. Torr. On partial optimality in multi-label mrfs. In *ICML*, pages 480–487, 2008. 2
- [16] I. Kokkinos, G. Evangelopoulos, and P. Maragos. Texture analysis and segmentation using modulation features, generative models, and weighted curve evolution. *IEEE TPAMI*, 31(1):142–157, 2009. 5
- [17] T.M. Koller, G. Gerig, G. Szekely, and D. Dettwiler. Multiscale detection of curvilinear structures in 2-D and 3-D image data. *IEEE ICCV*, page 864, 1995. 2
- [18] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE TPAMI*, 28(10):1568–1583, 2006. 2
- [19] V. Lempitsky, C. Rother, and A. Blake. Logcut - efficient graph cut optimization for markov random fields. pages 1–8, 2007. 2
- [20] Q. Li, J. You, L. Zhang, and P. Bhattacharya. Automated retinal vessel segmentation using multiscale analysis and adaptive thresholding. *IEEE SSIAI*, 1:139–143, 2006. 2
- [21] T. Lindeberg. Principles for automatic scale selection. *Handbook on Computer Vision and Applications*, 2:239–274, 1998. 1
- [22] T. Liu, K. Luk Chan, and S. Li. Model-based active object recognition using mrf matching and sensor planning. pages 225–232, 1997. 2
- [23] H. Lombaert, Y. Sun, and F. Chriet. Landmark-based non-rigid registration via graph cuts. *LNCS*, pages 166–175, 2007. 2

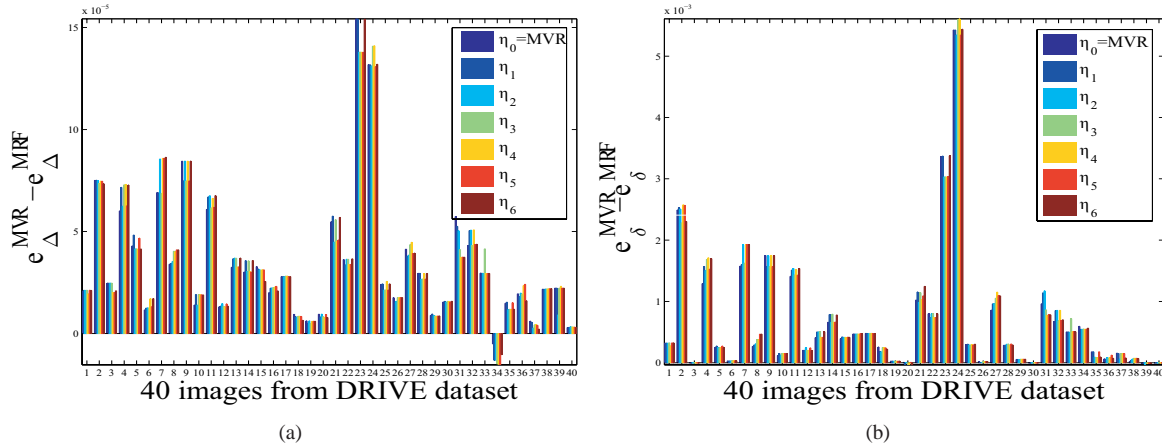


Figure 9. Scale-selection error based on MVR and MRF on all the images from the DRIVE database [36]. The y-axis represent $e^{MVR} - e^{MRF}$, i.e. positive values imply better performance of our propose method. Different colors correspond to different levels of spatial regularization (η) in (5). It can be seen that in almost all the 40 cases (except case 34), our approach achieved lower scale-selection error compared to MVR.

- [24] C. Lorenz, I. Carlsen, T. Buzug, C. Fassnacht, and J. Weese. Multi-scale line segmentation with automatic estimation of width, contrast and tangential direction in 2D and 3D medical images. *CVRMed-MRCAS*, pages 233–242, 1997. 2
- [25] C. McIntosh and G. Hamarneh. Vessel crawlers: 3D physically-based deformable organisms for vasculature segmentation and analysis. In *IEEE CVPR*, pages 1084–1091, 2006. 2
- [26] T. Nguyen, J. Wang, A. Sharrett, F. Islam, R. Klein, B. Klein, M. Cotch, and T. Wong. Relationship of retinal vascular caliber with diabetes and retinopathy. *Diabetes Care*, 31:544–549, 2008. 2
- [27] J. Novatnack and K. Nishino. Scale-dependent 3D geometric features. *IEEE ICCV*, pages 1–8, 2007. 1
- [28] M. Poon, G. Hamarneh, and R. Abugharbieh. Livevessel: Extending livewire for simultaneous extraction of optimal medial and boundary paths in vascular images. In *MICCAI*, pages 444–451, 2007. 2
- [29] J. Rao, G. Hamarneh, and R. Abugharbieh. Adaptive contextual energy parameterization for automated image segmentation. 5875-I:1089–1100, 2009. 4, 5
- [30] H. Rehrauer. Selecting scales for texture models. In *Machine Perception and Artificial Intelligence*, 2000. 1
- [31] E.M. Rikxoort, B. Ginneken, M. Klik, and M. Prokop. Supervised enhancement filters: Application to fissure detection in chest ct scans. *IEEE TMI*, 27(1):1–10, 2008. 2
- [32] S. Roy and I. Cox. A maximum-flow formulation of the n-camera stereo correspondence problem. pages 492–499, 1998. 2
- [33] Y. Sato, S. Nakajima, H. Atsumi, T. Koller, G. Gerig, S. Yoshida, and R. Kikinis. 3D multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. In *Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, pages 213–222, 1997. 1, 2
- [34] D. Schlesinger and B. Flach. Transforming an arbitrary min-sum problem into a binary one. In *Technical Report TUD-FI06-01, Dresden University of Technology*, 2006. 2
- [35] J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE TPAMI*, 22(8):888–905, 2000. 2
- [36] J.J. Staal, M.D. Abramoff, M. Niemeijer, M.A. Viergever, and B. van Ginneken. Ridge based vessel segmentation in color images of the retina. *IEEE TMI*, 23(4):501–509, 2004. 2, 3, 5, 6, 7
- [37] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *IEEE TPAMI*, 30(6):1068–1080, 2008. 2, 6
- [38] T. Tang and A. Chung. Non-rigid image registration using graph-cuts. *MICCAI*, pages 916–924, 2007. 2
- [39] Y. Yuan and A. Chung. Multi-scale model-based vessel enhancement using local line integrals. pages 2225–2228, 2008. 2