

Automated Vasculature Extraction from Placenta Images

Nizar Almoussa, Brittany Dutra, Bryce Lampe,
Pascal Getreuer, Todd Wittman, Carolyn Salafia, Luminita Vese

August 8, 2009

Abstract

Recent research in perinatal pathology argues that analyzing properties of the placenta may reveal important information on how certain diseases progress. One important property is the structure of the placental blood vessels, which supply a fetus with all of its oxygen and nutrition. An essential step in the analysis of the vascular network pattern is the extraction of the blood vessels, which has only been done manually through a costly and time-consuming process. There is no existing method to automatically detect placental blood vessels; in addition, the large variation in the shape, color, and texture of the placenta makes it difficult to apply standard edge-detection algorithms. We describe a method to automatically detect and extract blood vessels from a given image by using image processing techniques and neural networks. We evaluate several local features for every pixel, such as intensity, gradient, and variance, in addition to a novel modification to an existing road detector. Pixels belonging to blood vessel regions have recognizable responses; hence, we use an artificial neural network to identify the pattern of blood vessels. A set of images where blood vessels are manually highlighted is used to train the network. We then apply the neural network to recognize blood vessels in new images. The network is effective in capturing the most prominent vascular structures of the placenta.

1 Introduction

Understanding placental vasculature can help doctors identify where prenatal development diverged from normality. This could potentially help lead to earlier diagnoses of significant life-long diseases [21, 23]. However, the process of identifying placental blood vessels is currently costly and time-consuming; an expert must hand-trace the vessels in each placental image. By automating the detection of blood vessels, we hope to increase efficiency and reduce costs for emerging placental research.

1.1 Overview of Methods

To our knowledge, there is no single characteristic that can be used to accurately identify a placental vessel, as there is significant variance between placentas. For example, some placentas are covered in red blood while others appear brown or grey; some have thick, pronounced vessels while others have thinner ones.

Rather than formulating an explicit model for placental vasculature, we account for this variability by computing numerous “features” for the image and then feeding those features to a *neural network*. Such

a neural network must then undergo training, where it is given a ground truth (a hand-traced vascular network) and then iteratively determine weights for these features in order to better detect blood vessels. Once trained, a network can then be *simulated* on additional placenta images in order to identify vessels. For a more thorough introduction to neural networks, see [8].

Unless otherwise specified, all feature computations and network simulations are performed on the green channel of the placenta image. We do this because our initial inspections revealed that blood vessels appeared most distinct and recognizable in the green channel. Preliminary results also suggested that this was desirable over computing features using other channels or RGB-to-gray conversions. Similarly, vessels appeared darker than their surroundings in the green channel, so algorithms described in the following sections assume dark vessels. The validity of this assumption is discussed in Section 7.

2 Pre-Processing

Before being analyzed, all placental images are pre-processed to ultimately improve the performance of



Figure 1: Shown is an example of placenta extraction.

subsequent algorithms.

We first extract the placenta by applying a threshold and some morphological operations on the green channel. We then crop the image, thereby making future calculations more efficient (Fig. 1).

The majority of images we work with feature large patches of glare, so an in-painting approach is used. This method of glare removal is a simplified approach taken by [12] and has comparable results. First, bright spots are identified as those pixels with an intensity above a pre-determined threshold, which we take to be 80% of the maximum intensity. Second, a top-hat filter is applied, and additional thresholding then accurately identifies appropriate glare regions. Third, the regions are dilated by several pixels in order to place the region boundaries on pixels unaffected by glare. Finally, solving Laplace’s equation (1) fills in the regions, which produces satisfactory results as seen in Figure 2.

$$u_{xx} + u_{yy} = 0; u(x, y) = I(x, y) \quad \forall x, y \in \partial\Omega \quad (1)$$

We found that performing glare removal prior to the cropping procedure is preferable, as otherwise some glare regions could be unintentionally cropped.

3 Features

We decided to borrow the neural-net approach of Malik [19], as explained in Section 3.1.4. This meant that numerous features would be computed for a placenta and then later fed to a neural network to detect vessels.

Some features that were computed on placenta images are described in the following subsection. Other features including variance, curvature, eigenvalues of the 2^{nd} moment matrix [19], gradient magnitude, and gradient orientation [19], but we do not describe them here since they are relatively straightforward.

Neural networks trained using only these more straightforward features were found to be inappropri-

ate for vessel detection since, in most cases, networks that use these features simply detect vessel edges but not the interiors of the vessels themselves. However, these characteristics were still included alongside more intricate features in the event that they might reveal any implicit vessel information.

3.1 Line Detectors

It was noted that in the green channel individual vessels show little variance in intensity after glare has been removed. Hence, when considering additional features to implement we focused on those that could detect thick, uniform, curvilinear structures. We implemented several conventional line detectors, such as Steger’s line detector [25], a phase-coded detector [5], and a slightly modified wide-line detector [15], to be used as additional features for our neural network. We give a brief overview of these detectors as well as any differences between our implementations and the literature.

3.1.1 Phase Coding

For this detector we use a complex-valued linear filter given by [5] with kernel

$$PC(x, y) = e^{2i \tan^{-1}(y-y_0, x-x_0)},$$

where x and y are coordinates within a circle of radius R inscribed within a square window; (x_0, y_0) is the center of the window; and values within this window but outside of the disk are taken to be 0. A potential drawback of this detector is that it will detect both bright and dark lines, and intersections are only weakly identified.

3.1.2 Wide-line Detector

This a more recent detector described by Liu *et al.* [15] that has seen success in similar medical contexts [15, 16]. The detector consists of a nonlinear filter,




Figure 2: A raw image is shown on the left, with glare removed on the right.




Figure 3: The error magnitude for $\text{sech}(x^5)$ and its 20^{th} -order Padé approximant $f(x)$. Note from Equation 2 that $|x| \leq 1/t$, since intensities can differ by at most ± 1 .

where, for each pixel in the image, we tally the number of surrounding pixels within a radius R that differ in intensity by at most t when compared to the center pixel. However, to smooth small intensity changes that occur near the threshold t , Liu actually tallies all pixels within this disk but weights them as

$$s(x, y, x_0, y_0, t) = \text{sech} \left(\frac{(I(x, y) - I(x_0, y_0))^5}{t^5} \right), \quad (2)$$

where $I(x, y)$ is the intensity at (x, y) .

For computational efficiency, we use a 20^{th} -order Padé approximant for $\text{sech}(x^5)$ taken about $x_0 = 0$, namely

$$\text{sech}(x^5) \approx \frac{15120 + (13x^{10} - 660)x^{10}}{15120 + (6900 + 313x^{10})x^{10}},$$

where x^{10} is computed from x using a logarithmic number of multiplications, as MATLAB’s underlying C-based exponentiation function is highly inefficient. This approximant models the desired function very

closely (Fig. 3) and exhibits significant speed improvements.

3.1.3 Steger’s Detector

Steger’s algorithm for thick line detection [25] is available online from the GRASP project [2] and is licensed under the GNU General Public License [10]. We incorporated this code into our MATLAB project by using Mex.

As provided, the Steger detector only gives a response at the center and on the edge of a thick line. Hence, to make this output more appropriate as a feature, these lines were filled in. At each pixel along the center of a line, pixels along the normals (for the appropriate width) were assigned a value of the line’s “response,” or the second-derivative of the line at that point, as described by Steger [25]. This yielded results surprisingly similar to the wide-line detector; the two methods largely agreed on the larger vessels, but differed more in the noisier, vessel-free regions of the placenta.

3.1.4 Modified Road Detection

We also implemented a novel modification to an existing road-detection technique [22] that made the method much more suitable for detecting vessels. Porikli uses a directional line filter to look for elongated rectangular regions in-between two homogeneous regions of different intensity levels on either side. This filter is somewhat limiting because it can only detect lines with a maximum thickness of five or six pixels, whereas blood vessels can be much thicker.

We multiply Porikli’s directional line filter [22] by a Gaussian function in order to allow the filter to identify wider structures, with more weight towards the center. Our new filter is given as

$$g_{\theta}(i, j) = \cos\left(\frac{\pi i'}{2M}\right) \cos\left(\frac{\pi j'}{M}\right) e^{\frac{-(i^2+j^2)}{2(0.9M)^2}} \quad (3)$$

where M is fixed, $i, j = -M, \dots, M$, and i', j' are coordinates rotated by θ . We tried different values for the coefficient of M in the Gaussian function and found $M = 0.9$ to be the most effective for reducing noise while still retaining all the vessel information. We convolve this new filter over the entire image. Porikli’s filter assumes that roads are generally brighter than their surrounding areas, but we found that most of the placentas we worked with had blood vessels that were darker than the neighboring tissue, so we resolve this by inverting the convoluted image. We repeat this process for values of θ between 0 and π and combine the outputs for an overall result of

$$L = \left| \sum_{\theta \in \Theta} e^{2\pi i} (g_{\theta} \star I)^+ \right| \quad (4)$$

where $\Theta = \{0, \pi/8, \dots, 7\pi/8\}$ and I is the input image.

We found that using differently sized filters on the image could help detect small as well as large vessels, so we applied the filter with four different sizes of M , including 4, 7, 11, and 15. We created a simple noise reduction technique that labels all the connected segments and eliminates those with fewer pixels than a certain threshold, usually between 200 and 400 pixels. To combine the results from the small and large line filters, we fuse together the four images. Again we can threshold the image to increase accuracy and reduce noise. Since there will only be four values in the normalized image, it’s very simple to show points where a certain number of the four filters found a vessel. For example, by setting all points with a value of 0.25 to 0, the resulting image will show only points where two or more filters found a vessel, eliminating any accidental noise from one filter.

We found that our enhanced road detection method was superior to the other line detectors when used on placental images. However, this method fails on blood vessels that are softer and brighter than their surrounding tissue. No single detector could sufficiently and reliably identify all blood vessels; and simple arithmetic and morphological combinations of the detectors were also inadequate. Thus, we chose to borrow the neural-net approach taken by Malik [19]; the purpose of our neural network being to resolve the output from all of the features above, particularly the line detectors, while minimizing false-positives in the final result.

4 Neural Network Training

We use a set of manually traced images to train a neural network to detect blood vessels. Ideally, we would choose all pixels belonging to the placenta region to form the training set; however, our machines have limited memory, and the training time is preferred to be minimal. We overcome these problems by applying the following steps before training:

- Randomly pick n pixels from each traced image.
- For each pixel, compute its features and a binary value indicating whether it is a blood vessel.
- Randomly permute the sequence of pixels so that no single image can bias the training.

Manually tracing a vasculature network is subject to human interpretation. When we examined the traced images, we noticed that traced lines sometimes do not accurately cover vessels. This imprecision may affect the accuracy of the neural network output. To reduce the impact of these outliers, we added an option to transform the binary traced data to grayscale by convolving it with a Gaussian kernel, thereby giving greater weight to regions that had been traced while still allowing positive responses outside the traced areas. We also considered using mean-absolute error instead of mean-squared error when evaluating a network’s performance, as it is more robust to outliers in the training data.

To determine the optimal combination of the features described in Section 3, numerous neural networks were trained with all possible combinations of 3-or-more features from all available features. This exhaustive search of the feature space was necessary because the processes of neural network training is not sufficient to determine which features are unneeded; determining those features to assign zero




Figure 4: The original piece of the placenta (left) compared with the results from the smallest and largest road detection filters (middle). The rightmost image shows the modified road-detection feature on a whole placenta.

weight requires exponential time (since the greedy approach is not guaranteed to be optimal) but training occurs in polynomial time.

In addition, various parameters for the networks, such as the performance function (mean-squared error or mean-absolute error), the number of hidden nodes (5 or 15), how to normalize features (between $[0, 1]$ or $[-1, 1]$), and whether to apply a Gaussian blur to training data were also tested.

5 Post-Processing

Our neural networks largely produced soft classifications of blood vessels, so further processing of these results was necessary. Grayscale neural-net outputs were thresholded to obtain a binary classification. These black-and-white images were then filtered for size; components smaller than usually 400 pixels were discarded as noise. Further filtering this by eccentricity (e.g., taking only highly eccentric, line-line objects), as suggested by [16], was found to be unreliable. Due the relatively high connectivity of our identified vessels, many components have low eccentricity, which is uncharacteristic of lines. Additionally, the edge is eroded since many of our detectors detected the edge of the placenta as a vessel.

In a few cases the neural network would produce an *almost* binary classification, with very few gray pixels. Thresholding was also applied in those cases.

6 Measuring Network Performance

Once the neural-net output was post-processed, it remained to be shown which resulting binary classification was “best.” Several metrics were used to compare the classifications to the traced data.

Between a given pixel P_1 from the target image and the corresponding pixel P_2 from the predicted image, four different conditions can occur:

- True Positive: Both pixels P_1 and P_2 have a value of one.
- True Negative: Both pixels P_1 and P_2 have a value of zero.
- False Positive: P_1 has a value of zero while P_2 has a value of one.
- False Negative: P_1 has a value of one while P_2 has a value of zero.

From these four conditions, we form the following binary classifications:

- TP: the total number of true positive pixels
- TN: the total number of true negative pixels
- FP: the total number of false positive pixels
- FN: the total number of false negative pixels.

These binary classifications can be used to form different metrics for a network’s performance. The ROC curve is one of these metrics. It is usually used to visualize the performance of different algorithms, but the ROC curves are not a good tool to choose the optimal algorithm [9]. To plot the ROC curves, two quantities are needed: the false positive rate

$$FPR = \frac{FP}{FP + TN}$$

and the true positive rate

$$TPR = \frac{TP}{TP + FN}.$$

The ROC curve is the set of all such pairs (FPR, TPR) .

We can use the ROC curve to visually compare the performance of the neural-network technique versus individual features such as the wide-line and the road detection algorithms. Since the output of each algorithm is a grayscale image, a threshold is applied as described in Section 5 to obtain a binary image. This binary image and the image of manually traced blood vessels work as the predicted and target images, respectively. The threshold value is varied increasingly between 0 and 1. For each threshold value, FPR and TPR are computed from the binary classifications. The result is a ROC curve, as shown in Figure 6.

While the ROC curve is a good tool to visualize the performance of a network, it is not efficient in choosing the optimal algorithm, nor the best threshold to use for a particular network. A better metric is the Matthew Correlation Coefficient, or MCC, which puts more emphasis on the correlation between the two images rather than the actual accuracy in the prediction (see Fig. 6). The MCC uses the binary classification as following

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

The area under the MCC curve was measured for each neural network, and the network with the largest area was chosen to be the optimal network. From the MCC curve, one can also compute the optimal threshold value. This value was used as the initial threshold value for the graphical user interface.

7 Results

A result of our method is shown in Figure 6. While not perfect, the neural network does identify many prominent vessels. We note that the width of the detected vessels is more accurate than in the manual tracing. Noise is present in the classifications, which is why we provide a feature in our GUI frontend for the user to erase classified vessels.

While the neural nets can perform well for some placentas, such as the one shown in Figure 6, there are many others where the classification is very unreliable. This is because, contrary to our assumption, many placentas feature vessels that appear *brighter* in the green channel. Hence, identifying dark lines in the green channel is not always appropriate; this is discussed further in Section 8.

A side-effect of the randomization of training data, described in Section 4, is that several networks trained using the same features can have variable performance. To account for this, we train several (usually 3) networks using the same features and the same training images. We then simulate *all* of these networks and average their results; this produces more

connected vasculature than just taking a single network on its own. This is demonstrated in Figure 7.

With 5 binary options and 8 features in all, there were $2^5 \sum_{k=3}^8 \binom{8}{k} = 7008$ possible nets to generate. This number was reduced to approximately 2000 by eliminating the binary options that did not seem immediately beneficial after preliminary results (e.g., blurring training data and normalizing to $[-1, 1]$). However, due to hardware limitations we were only able to generate roughly 1292 of those 2000. Hence our results are somewhat incomplete, but still promising.

With that incompleteness in mind, we found that the neural net with the highest area under its MCC curve (0.22) used: the gradient magnitude and gradient angle, the wide-line detector, the Steger detector, and our modified road detector.

We found that, in-general, nets that used the mean-absolute error for their performance function performed slightly better than those that used mean-squared error. Blurring training data to reduce the impact of outliers had little effect on performance. Somewhat surprisingly, networks with only five hidden nodes performed better than networks with 15-or-more hidden nodes; they were also faster to train and simulate. We were unable to determine whether normalizing features to $[0, 1]$ or $[-1, -1]$ was preferable, although we use the range $[0, 1]$ in our GUI.

Input images were 1,200 by 1,600 pixels, and the computations to compute image features were completed within two minutes for each image using a 2.13GHz Core 2 Duo with 2GB of RAM without down sampling. Most of this time was spent computing the wide-line detector and our modified road detection feature. The time to train a neural network usually took 5 minutes, however only a subset (8 images) of all available images were used to train. Training with more images had little impact, possibly implying that the sampled points taken from those 8 images were abundant enough to be reflective of an “average” placenta’s vascular network.

8 Future Work

We found our method is effective at identifying blood vessels, and we are confident it can eventually be used to help facilitate placental research. Future work will involve testing additional features for use with the neural network; evaluating performance metrics for the effectiveness of a network; and implementing a post-processing technique to improve the results of the network. A more long-term goal is to eventually be able to compute network statistics for identified




Figure 5: An example to motivate the MCC metric. A hypothetical neural-net prediction is shown in red, with a desired target in red. Although the prediction largely does not agree with the target, there is a strong correlation between the two shapes.




Figure 6: The original placenta (left) compared with the hand-traced data (right) and the result of our neural network. The particular features used are given in Section 7. The placenta is the same as the one shown in Fig. 3.1.4.




Figure 7: An ROC curve for a neural network is shown (left) and compared with individual features, and an MCC profile is also shown (right) for a different neural network. Neither of these networks are the same as the one used in Figure 6, but almost all networks exhibited similar performance.




Figure 8: On the left is the result from one neural network. On the right is the average (and the post-processed) result from three neural nets trained using the same features and training data. The randomization of the training process seems to be mitigated by using several nets.

vasculature and then correlate those statistics with medical data. These correlations might then help identify placental characteristics that could be used to identify at-risk children.

Currently, inputs to our neural networks are features computed for individual pixels. To make the results of these networks more context-aware, we could borrow the technique used by Boggess [4] where the network is fed features computed at the pixel in question as well as the features for all neighboring pixels. Although he found this approach to be ineffective for road detection in satellite imagery, it may still be useful for placentas, especially considering that Boggess’ result is from 1993. Other learning methods such as k -nearest neighbors could also be investigated.

A C implementation of the wide-line detector would provide tremendous speed improvements wherever that feature is used.

We have adapted Steger’s detector to return a grayscale image with entire curvilinear structures filled in. However, the process used to fill in these structures is naïve and often leaves 1-pixel “holes” in the lines. This should be improved, as it could potentially lead to problems if morphological operations are ever performed on results from the detector.

The algorithm presented in [18] might be useful for filling in gaps, where neural network output would be given as the necessary seed values. This algorithm could also be used to more easily distinguish between arteries and veins, because it seems capable of correctly handling vessel crossings instead of treating all crossings as intersections.

We considered blurring as part of the pre-processing stage but did not have time to implement it. Ideally, this would blur away smaller vessels and wrinkles, leaving the larger, more important vessel

structures intact. While this could potentially reduce the number of wrinkles we classify as vessels, in reality we would expect this to, at most, reduce noise in some of our features; most of our detectors already fail to identify smaller vessels.

The scale of images used is currently not considered at all, which is problematic when handling images of different resolutions. All of our parameters have been tuned to best detect vessels in 1,200-by-1,600 images, meaning our methods would almost certainly not work as well on larger or smaller pictures. Luckily, most placental images contain a ruler or penny, which could be extracted and then used for scale normalization.

Finally, our assumption that vessels appear dark in the green channel does not hold for all placentas. Hence, detecting brighter lines should also be investigated. Sometimes a placenta’s vessels will appear as dark *and* bright lines in the green channel, so bright-line detection should also be investigated, as well as how to accurately resolve the two line types (e.g., a dark vessel is often sandwiched between bright areas, which could be incorrectly identified as bright vessels). The wide-line detector as well as Steger’s detector describe techniques for detecting bright lines [15, 25]; these require only small changes in the code. We had difficulty modifying our road-detection method to detect bright lines.

References

- [1] *Modeling placental growth and structure*, <http://www.mathinstitutes.org/nuggets/placental.html>.
- [2] Bruce Allen, *Gravitational radiation*

- analysis and simulation package*, <http://www.lsc-group.phys.uwm.edu/~ballen/grasp-distribution/GRASP/include/>, January 1999.
- [3] Pablo Arbelaez, Charless Fowlkes, and David Martin, *The Berkeley segmentation dataset and benchmark*, <http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/>, June 2007.
- [4] Julian E. Bogges, *Identification of roads in satellite imagery using artificial neural networks: A contextual approach*, Tech. report, Mississippi State University, 1993.
- [5] Simon Clode, Emanuel Zelniker, Peter Kootsookos, and Vaughan Clarkson, *A phase coded disk approach to thick curvilinear line detection*, Tech. report, ARROW Discovery Service, 2004.
- [6] L.D. Cohen, *Minimal paths and fast marching methods for image analysis*, MMCV05, 2005.
- [7] Howard Demuth, Mark Beagle, and Martin Hagan, *Neural Networks Toolbox 6 user guide*, http://www.mathworks.com/access/helpdesk/help/pdf_doc/nnet/nnet.pdf, March 2009.
- [8] Laurence Fausett, *Fundamentals of neural networks*, Prentice-Hall, 1994.
- [9] Tom Fawcett, *ROC graphs: Notes and practical considerations for researchers*, (2004), 1–38.
- [10] Free Software Foundation Inc., *GNU general public license*, <http://www.gnu.org/copyleft/gpl.html>, June 2007.
- [11] V. Lacroix and M. Acheroy, *Feature-extraction using the constrained gradient*, *PandRS* **53** (1998), 85–94.
- [12] Holger Lange, *Automatic glare removal in reflectance imagery of the uterine cervix*, *Proc. SPIE Medical Imaging* **5747** (2005), 2183–2191.
- [13] Shu-Xiao Li, Hong-Xing Chang, and Cheng-Fei Zhu, *Fast curvilinear structure extraction and delineation using density estimation*, *Comput. Vis. Image Underst.* **113** (2009), no. 6, 763–775.
- [14] Wenjing Li, Marcelo Soto-Thompson, and Ulf Gustafsson, *A new image calibration system in digital colposcopy*, *Opt. Express* **14** (2006), no. 26, 12887–12901.
- [15] L. Liu, D. Zhang, and J. You, *Detecting wide lines using isotropic nonlinear filtering*, *IEEE Transactions on Image Processing* **16** (2007), 1584–1595.
- [16] Laura Li Liu and David Zhang, *Extracting tongue cracks using the wide line detector*, *ICMB*, 2008, pp. 49–56.
- [17] Rong-Chin Lo and Wen-Hsiang Tsai, *Gray-scale hough transform for thick line detection in gray-scale images*, *Pattern Recognition* **28** (1995), no. 5, 647–661.
- [18] G. Peyré M. Péchaud and R. Keriven, *Extraction of vessels networks over an orientation domain*, *CVPR* (2009).
- [19] David R. Martin, Charless C. Fowlkes, and Jitendra Malik, *Learning to detect natural image boundaries using local brightness, color, and texture cues*, *PAMI* **26** (2004), 530–549.
- [20] R. Matijevic and A. Kurjak, *The assessment of placental blood vessels by three-dimensional power doppler ultrasound*, *J Perinat Med* (2002).
- [21] Oleksandr Shlakhter Danielle Haas Barbara Eucker Michael Yampolsky, Carolyn M. Salafia and John Thorp, *Modeling the variability of shapes of a human placenta*, *Placenta* **29** (2008), 790–797.
- [22] Fatih Porikli, *Road extraction by point-wise Gaussian models*, *SPIE AeroSense Technologies and Systems for Defense and Security*, 2003, pp. 758–764.
- [23] Carolyn M. Salafia and Michael Yampolsky, *Metabolic scaling law for fetus and placenta*, *Placenta* **30** (2008), 468–471.
- [24] Kai Sheng, Keisuke Kameyama, Kazuki Katagishi, and Kazuo Toraichi, *Matched filter design by fluency analysis for a more accurate fundus image blood-vessel extraction*, *Proceedings of the 24th IASTED international conference on Biomedical engineering (Anaheim, CA, USA)*, ACTA Press, 2006, pp. 350–355.
- [25] Carsten Steger, *An unbiased detector of curvilinear structures*, *IEEE TPAMI* **20** (1998), 113–125.
- [26] Radu Stoica, Xavier Descombes, and Josiane Zerubia, *A gibbs point process for road extraction from remotely sensed images*, *Int. J. Comput. Vision* **57** (2004), no. 2, 121–136.