SEGMENTATION OF NETWORKS FROM VHR REMOTE SENSING IMAGES USING A DIRECTED PHASE FIELD HOAC MODEL

Aymen El Ghoul, Ian H. Jermyn and Josiane Zerubia

ARIANA - Joint Team-Project INRIA/CNRS/UNSA
2004 route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex, France
{aymen.el_ghoul,ian.jermyn,josiane.zerubia}@sophia.inria.fr

Commission III, WG III/4

KEY WORDS: shape priors, higher order active contours, phase fields, segmentation, road and hydrographic network.

ABSTRACT:

We propose a new algorithm for network segmentation from very high resolution (VHR) remote sensing images. The algorithm performs this task quasi-automatically, that is, with no human intervention except to fix some parameters. The task is made difficult by the amount of prior knowledge about network region geometry needed to perform the task, knowledge that is usually provided by a human being. To include such prior knowledge, we make use of methodological advances in region modelling: a phase field higher-order active contour of directed networks is used as the prior model for region geometry. By adjoining an approximately conserved flow to a phase field model encouraging network shapes (*i.e.* regions composed of branches meeting at junctions), the model favours network regions in which different branches may have very different widths, but in which width change along a branch is slow; in which branches do not come to an end, hence tending to close gaps in the network; and in which junctions show approximate 'conservation of width'. We also introduce image models for network and background, which are validated using maximum likelihood segmentation against other possibilities. We then test the full model on VHR optical and multispectral satellite images.

1 INTRODUCTION

In this paper, we address the problem of road and hydrographic network segmentation from VHR optical and multispectral images: given an image I, we seek the region R in the image domain Ω that contains the network. We would like to perform this task quasi-automatically, that is, with no human intervention except to fix some parameters. Such segmentation problems remain challenging due to a combination of difficulties. First, the network is usually not distinguishable from the background using image measurements alone. Rather, knowledge of the geometric properties of R (*e.g.* , that it is composed of branches that meet at junctions) is necessary for successful segmentation. Currently, this knowledge is provided, in one way or another, by a human being. Automation of the segmentation process therefore requires models that incorporate this knowledge of region geometry. This is a nontrivial matter, particularly since the regions corresponding to networks have huge variability in their topology as well as their geometry. Second, there is great variability in the appearances of network and background from one image to another. Third, models incorporating the necessary prior knowledge of region geometry are complex, and this leads to efficiency issues when confronted with the large size and number of images to be processed.

The contribution of this paper is a new algorithm for road and hydrographic network segmentation from VHR remote sensing images of rural and non-urban areas which present many occluded parts of the network entity to be extracted. Our new algorithm uses recent advances in shape modelling allowing the incorporation of sophisticated prior knowledge about network region geometry, thereby addressing the first difficulty. A 'phase field higher-order active contour' ('phase field HOAC') model of directed networks (El Ghoul et al., 2009b) is used to favour regions composed of branches that meet at junctions. The network contains a 'flow', which is approximately conserved. This means that the width of each branch changes slowly, while different branches can have very different widths; that branches tend not to end; and that, at junctions, there is approximate 'conservation of width': for example, several small incoming branches combining to form a larger outgoing branch. These characteristic geometric properties are different from those of road networks in VHR images of urban areas (Peng et al., 2010), and from those of networks in medium resolution images (Rochery et al., 2005); the problem therefore requires the use of a new model.

In (El Ghoul et al., 2009b), the model used here was described, but the automatic parameter setting described herein was not used, and the model was not applied to or tested on real images. Real images generate the second difficulty described above. To address it, we also propose generic models for the image in the network region and in the background whose parameters can easily be learned from examples of these two classes. We test these models on several VHR images. The image models are compared to other possibilities using maximum likelihood (ML) classifications. They outperform standard indices, which suggests that their performance when combined with the region geometry model will also be superior. We then test the full model, combining the phase field model of directed networks with the image models, on several satellite images. The segmentation problems involved are very hard, but the results show that the new algorithm is able to ignore confounding factors in the background due to the sophisticated knowledge of region geometry it contains, and is able to complete the network over reasonable gaps.

Before going on, it is useful to formalize the problem and our approach to solving it, and to introduce some notation. We seek to infer the region R containing the network from the image data I and our prior knowledge K, *e.g.* of image formation, network geometry, and so on. In other words, we wish to construct a probability distribution $\mathrm{P}(R|I,K)$ for the region R containing the network, given the current image data I and our knowledge K. As usual, this can be written as the product of a likelihood $\mathrm{P}(I|R,K)$, which models the images we expect to see given that $R \subset \Omega$ corresponds to a network, that the image is a VHR optical image, etc.; and a prior $\mathrm{P}(R|K)$, which incorporates our

knowledge of network region geometry. We will further factorize the likelihood term into $\mathrm{P}(I_R|R,K)$ and $\mathrm{P}(I_{\bar{R}}|R,K)$, that is into separate models for the image I_R in the network region R and the image $I_{\bar{R}}$ in the background $\bar{R}$. The phase field HOAC model of directed networks then corresponds to $\mathrm{P}(R|K)$, while our image models correspond to $\mathrm{P}(I_R|R,K)$ and $\mathrm{P}(I_{\bar{R}}|R,K)$. In practice, we will deal with negative log probabilities, *i.e.* a total energy $E(R,I) = -\ln \mathrm{P}(R|I,K)$ that is the sum of a likelihood term $E_{\mathrm{I}}(I,R) = -\ln \mathrm{P}(I|R,K)$ and a prior term $E_{\mathrm{P}}(R) = -\ln \mathrm{P}(R|K)$.) We will then extract a maximum *a posteriori* (MAP) estimate for R by minimizing $E(R,I)$ over R. This will be done using gradient descent.

In section 1.1, we justify our choice of shape modelling framework by surveying the alternatives. In section 2, we recall briefly the theoretical background of the models used: sections 2.1 and 2.2 recall the undirected and directed phase field HOAC models respectively. In section 3, we describe the test results obtained in applying the algorithm to VHR images: we compare ML segmentations using different image models, and then test our new algorithm including the full model. We conclude in section 4.

1.1 Previous work

The models used in most previous work on region segmentation do not incorporate any nontrivial knowledge about region geometry. For example, standard active contours, introduced by (Kass et al., 1988), further refined by many authors, and applied in a huge number of other papers, contain only the prior knowledge that the region boundary should be smooth. As we have emphasized, this degree of prior knowledge is almost never enough for the automatic solution of segmentation problems on real images, whatever the domain. As a result, recent work has developed models incorporating more sophisticated knowledge of region geometry. Most of this work models an ensemble of regions as perturbations of one or more reference regions, for example (Cremers et al., 2003, Rousson and Paragios, 2002, Srivastava et al., 2003, Foulonneau et al., 2003). This is an intuitive and useful approach, but it is inappropriate when the region sought can have arbitrary topology, since such an ensemble of regions cannot be described as perturbations around a finite number of reference regions. Since networks can have arbitrary topology (*i.e.* several connected components, each of which can have loops), the above work is not applicable to the network segmentation problem.

To model regions with potentially arbitrary topology, higher-order active contours (HOACs) were introduced by (Rochery et al., 2006). The HOAC prior energy defined by (Rochery et al., 2006) was used to model undirected network regions and to extract road networks from medium resolution optical images. The contour representation used by (Rochery et al., 2006) suffers from many drawbacks, however. To overcome these drawbacks, (Rochery et al., 2005) reformulated HOACs as nonlocal 'phase field' models. This formulation facilitates model analysis and implementation, allows a 'neutral' initialization and complete topological freedom, and results in reduced execution times, sometimes by an order of magnitude. Phase field HOAC models of undirected networks have proved their efficiency for road network segmentation from medium resolution images of rural or semi-rural areas (Rochery et al., 2005), and high resolution images of urban areas (Peng et al., 2010), but the models are not well adapted to non-urban road network segmentation from high resolution images, nor to hydrographic network segmentation. The main problem is that the branch width in these models is very tightly constrained. This works well for medium resolution where the range of visible widths is not large, but at high resolution and in hydrographic networks, the range of widths is much greater. Naive attempts to allow a greater range of widths have the unfortunate side effect of allowing width to vary rapidly along one branch: the branch sides lose their 'rigidity'. The second difficulty is that the early model in (El Ghoul et al., 2009a) had problems closing large gaps in the network. The work of (Peng et al., 2010) solves these problems for urban road networks, but the solution involves favouring long straight branches, which is again not well adapted to non-urban road and hydrographic networks.

To overcome these difficulties, (El Ghoul et al., 2009b) introduced a phase field HOAC model for directed networks, in order to capture some of their distinctive geometric properties. Because these geometric properties are linked to the fact that directed networks carry a conserved flow, the model contains, in addition to the phase field specifying the network region, a vector field representing a 'flow' in the network. The magnitude of the field is approximately constant, and the flow is approximately conserved. As a result, branch width tends to change slowly and branches tend not to end, as both these would change the flow. At junctions, there is approximate 'conservation of width' so that incoming flow be approximately equal to outgoing flow. However, the model was only tested on a synthetic image showing the shape of a river, albeit with success.

Other attempts to solve the hydrographic network segmentation problem include the work of (Dillabaugh et al., 2002). This work uses an interesting multiscale approach, but relies on user input to specify network endpoints, and is limited in the network topologies that it can find. The work of (Lacoste et al., 2004) models the network region using a marked point process of polylines. This model works well when the network has constant width over significant distances, since each polyline has a fixed width. (Lacoste et al., 2005) uses an initial segmentation by Markov random field as a seed from which to build a hierarchical model of the network using a marked point process. This works well when the image is sufficiently clean for the MRF segmentation to capture most of the network, and when the network has a tree structure. For a review of the large number of techniques that have been developed for road network segmentation, see (Mena, 2003). However, none of these solve the problem of network segmentation from high resolution images in a quasi-automatic way.

2 PRIOR MODEL E_{P}

In this section, we present the prior model E_{P}. We begin by recalling the simplest phase field HOAC model of an undirected network (Rochery et al., 2005), since this is the base on which the model of directed networks is built.

2.1 Undirected network model

A phase field ϕ is a real-valued function on the image domain Ω. A phase field determines a region by the thresholding map $\zeta_z(\phi) = \{x \in \Omega : \phi(x) > z\}$ where z is a given threshold. The basic phase field energy is

$$E_0^s(\phi) = \int_\Omega d^2x \left\{ \frac{D}{2} \partial\phi \cdot \partial\phi + \lambda \left(\frac{\phi^4}{4} - \frac{\phi^2}{2} \right) + \alpha \left(\phi - \frac{\phi^3}{3} \right) \right\}. \quad (1)$$

If (1) is minimized subject to $\zeta_z(\phi) = R$, *i.e.* for a fixed region, then away from the boundary, the minimizing function ϕ_R assumes the value 1 inside, and -1 outside R thanks to the ultralocal terms. To guarantee two energy minima (at -1 and 1),

the inequality $\lambda > |\alpha|$ must be satisfied. We choose $\alpha > 0$ so that the energy at -1 is less than that at 1. The derivative term ensures the smoothness of ϕ_R, producing a narrow interface around the boundary ∂R interpolating between -1 and $+1$.

To introduce prior shape information, a nonlocal term is then added to give an energy $E_{\mathrm{P}}^s = E_0^s + E_{\mathrm{NL}}$ (Rochery et al., 2005):

$$E_{\mathrm{NL}}(\phi) = -\frac{\beta}{2}\iint_{\Omega^2} d^2x\, d^2x'\, \partial\phi(x)\cdot\partial\phi(x')\,\Psi\left(\frac{|x-x'|}{d}\right), \quad (2)$$

where d is the interaction range. This term creates long-range interactions between points of ∂R (because $\partial\phi_R$ is zero elsewhere) using an interaction function, Ψ, which decreases as a function of the distance between the points.

In this paper, the interaction function Ψ is taken to be the modified Bessel function of the second kind of order 0, K_0. This choice, as opposed to that used in (Rochery et al., 2005), allows a wide range of stable branch widths. However, it also allows the width of individual branches to fluctuate rapidly, because it reduces branch 'rigidity'. To allow a wide range of branch widths while constraining the rate of change of the widths of individual branches (without imposing straightness of branches as in (Rochery et al., 2005)), and also to favour other important properties of directed networks, it is necessary to augment this model with an extra field representing a 'flow' in the network. This leads to the phase field HOAC model for directed networks, to be described next.

2.2 Directed network model

The phase field HOAC model for directed networks was introduced by (El Ghoul et al., 2009b). To model directed networks, the phase field ϕ is augmented by a new tangent vector field v that loosely speaking represents the 'flow' in the network.[1]

The total prior energy $E_{\mathrm{P}}(\phi, v)$ is then the sum of a local term $E_0(\phi, v)$ and the nonlocal term $E_{\mathrm{NL}}(\phi)$ given by equation (2). E_0 is given by

$$E_0(\phi,v) = \int_\Omega d^2x \Big\{ \frac{D}{2}\,\partial\phi\cdot\partial\phi + \frac{D_v}{2}\,(\partial\cdot v)^2 + \frac{L_v}{2}\,\partial v : \partial v + W(\phi,v)\Big\}. \quad (3)$$

$W(\phi, v)$ is a potential with two degenerate sets of minima, at $(\phi, |v|) = (-1, 0)$ and $(\phi, |v|) = (1, 1)$. These minima define the stable phases corresponding to $\bar{R}$ and R respectively, replacing $\phi = \pm 1$ in the undirected model. The potential W is a fourth order polynomial in ϕ and $|v|$, constrained to be differentiable:

$$W(\phi,v) = \frac{|v|^4}{4} + (\lambda_{22}\frac{\phi^2}{2} + \lambda_{21}\phi + \lambda_{20})\frac{|v|^2}{2} + \lambda_{04}\frac{\phi^4}{4} + \lambda_{03}\frac{\phi^3}{3} + \lambda_{02}\frac{\phi^2}{2} + \lambda_{01}\phi\,. \quad (4)$$

[1] To avoid misunderstanding, we stress that v is not intended to represent the physical flow of, *e.g.* water, in the network, nor is the model intended to model the physical behaviour of the flow. Rather, v is an auxiliary variable that acts to favour certain geometric properties of the network. (Since it is not coupled to the image, it is, probabilistically speaking, a 'hidden variable'.) At the same time, it is not a coincidence that v shares many of the properties of physical flows, such as smoothness and conservation, nor that the resulting stable configurations resemble physical flows in the network.

The second term in equation (3) penalizes the divergence of v. This represents a soft version of flow conservation, but in practice the parameter multiplying this term will be large, so that in general the divergence will be small. The third term is a small overall smoothing term on v ($\partial v : \partial v = \sum_{m,n}(\partial_m v^n)^2$, where $m, n \in \{1, 2\}$ label the two Euclidean coordinates), since constraining the divergence is not sufficient to ensure smoothness. Because of the transition from $|v| = 1$ to $|v| = 0$ across the boundary of the region, the divergence term tends to make v parallel to the region boundary, since this results in zero divergence. The smoothness and divergence terms then propagate this parallelism to the interior of the branch, with the result that the flow tends to be along the branch. This fact, when coupled with the constraint on $|v|$ inside the channel, means that width variations are constrained to be slow along a channel, since total flow is directly related to branch width. At the same time, the use of $\Psi = K_0$, means that different branches may have very different widths. At junctions, the conserved flow along each branch favours 'conservation of width': the (soft) constraint that total incoming flow be approximately equal to total outgoing flow translates to the sum of the incoming widths being approximately equal to the sum of the outgoing widths. Thus the introduction of the new field v can favour network regions with geometric properties characteristic of directed networks.

2.2.1 Parameter settings Requiring $(-1, 0)$ and $(1, 1)$ to be extrema of the potential w reduces the number of free parameters of W from seven to four, while requiring these points to be minima (*i.e.* the Hessian at these two points should be positive-definite) generates further lower and upper bounds on the remaining parameter values.

We fix further relations between the parameters by requiring that the two minima described above be the only local minima; that W be bounded below; and that the potential energy of the network region R be greater than of the background $\bar{R}$, *i.e.* $w(1, 1) > w(-1, 0)$. The resulting potential has a saddle point lying between the two minima at a point (ϕ_s, v_s). This point plays an important role: the 'neutral' initialization of the gradient descent algorithm is given by $(\phi, |v|) = (\phi_s, v_s)$, the direction of v being random. In addition, we constrain the parameter β in E_{NL} so that the part of E_{P} containing derivatives, *i.e.* everything except W, be positive definite (it is a quadratic form). Since constant values of ϕ and v produce zero in these derivative terms, which is the global minimum value of these terms, and since constant values of ϕ and v equal to those at the global minimum of W, which is $W(-1, 0)$, produce the global minimum of W, the global minimum of E_{P} is at $(\phi, |v|) = (-1, 0)$.

The energy E_{P} can favour different stable geometric configurations depending on the values of the parameters remaining after the above constraints have been imposed. Since we are interested in modelling networks, we need to choose parameter values that favour networks as stable structures. Such values can be found using a stability analysis of the model. We assume that network branches are long enough and straight enough that their stability can be analysed by considering the limit of a long, straight bar, whose symmetry facilitates the analysis. We do not detail here the stability calculations because they are lengthy: they will be reported elsewhere. The stability analysis of a network branch places constraints on the parameter values of the model. In all the experiments, we use these constraints to fix further parameter values, and to replace others with physical parameters such as average branch width.

3 EXPERIMENTAL RESULTS

Having presented the prior model E_{P}, we are now nearly in a position to describe experiments testing the properties of the model and its performance in network segmentation. First, however, we have to describe the likelihood energy $E_{\mathrm{I}}(R, I) = -\ln \mathrm{P}(I|R, K)$ that we will use to create the complete model $E = E_{\mathrm{P}} + \theta E_{\mathrm{I}}$ where the parameter θ balances the two terms. Since we test the model on 0.61m resolution multispectral VHR Quickbird images with four channels (red (R), green (G), blue (B) and infrared (I)), and optical colour images, we need a model of such images. As already stated, we will assume that the likelihood can be factorized as $\mathrm{P}(I|R, K) = \mathrm{P}(I_R|R, K)\mathrm{P}(I_{\bar{R}}|R, K)$, and we thus need models for the image in R and $\bar{R}$.

3.1 Likelihood energy

In (El Ghoul et al., 2009a), the road network segmentation performance of a phase field HOAC model for undirected networks was tested using two classes of likelihoods (the same class was used for both R and $\bar{R}$): a multivariate Gaussian (MG) and a mixture of two multivariate Gaussians (MMG). In ML segmentations, the performance was mixed, but when combined with the prior energy, the MMG model was found to outperform the MG model, with the improvement being most significant when the image was very heterogeneous. Here, we test the ML performance of these two likelihood classes on the images used in this paper, and compare them to segmentations obtained using the normalized difference vegetation index ($NDVI = (I-R)/(I+R)$) (Rouse et al., 1973, Tucker, 1979) and the normalized difference water index ($NDWI = -(I-G)/(I+G)$) (McFeeters, 1996). We apply the former to images of road networks in which the background is mostly vegetation, and the latter to an image of a hydrographic network. The 4^{th}, 5^{th}, and 6^{th} rows of figure 1 show ML segmentations using NDVI/NDWI, MG, and MMG respectively.

Table 1 shows quantitative evaluations of the quality of the ML segmentations using NDVI, MG, and MMG. The bold numbers show the best ML segmentation method. In all experiments, the NDVI results show lower performance, according to the quality measure, compared to the MG and MMG results. The NDVI results on the first and second images show that most of the hidden parts of the network are not retrieved because they resemble vegetation more than network. The result is the presence of many lengthy gaps in the ML segmentation. Because these gaps are so long, it is very unlikely that the prior term would close them. In contrast, the MG and MMG segmentations include most of the network, but also many points of the background, which the prior model should be able to eliminate. When coupled with the results of (El Ghoul et al., 2009a) showing that MMG outperforms MG, these results lead us to choose the MMG model to construct the likelihood energy E_{I}.

3.2 Segmentation results

In order to compute a MAP estimate for $\mathrm{P}(R|I, K)$, we use gradient descent to seek minima of $E = E_{\mathrm{P}} + \theta E_{\mathrm{I}}$. To reduce the computational complexity, we primarily test the model on small images of size 256×256, some of which were obtained by reducing the resolution of the original images. We discuss this point further in section 4.

Figure 2 shows segmentations of the images shown in figure 1 obtained using the new algorithm.[2] The results in the first and second rows show that most of the gaps present in the original images are indeed closed thanks to the contribution of the divergence term: when divergence of the vector field is heavily penalized, network branch extremities prefer to meet and close gaps because in this way flow can be conserved. This does not occur using the undirected phase field HOAC model described in section 2.1, as shown in (El Ghoul et al., 2009a).

		Completeness	Correctness	Quality
a	NDVI	0.4296	0.3965	0.2598
	MG	0.6920	0.3423	**0.2970**
	MMG	0.7510	0.3000	0.2728
b	NDVI	0.4745	0.6536	0.3791
	MG	0.7166	0.4958	**0.4145**
	MMG	0.7983	0.3521	0.3233
c	NDWI	0.6280	0.9446	0.6057
	MG	0.7835	0.8468	**0.6862**
	MMG	0.8485	0.7424	0.6555
d	NDVI	0.6776	0.4517	0.3718
	MG	0.9060	0.4099	0.3932
	MMG	0.8634	0.4641	**0.4323**

Table 1: Quantitative evaluations of the ML segmentations given in figure 1. a, b, c, and d correspond to the four images in figure 1, from left to right. Completeness= TP/(TP + FN), correctness= TP/(TP + FP) and quality = TP/(TP + FP + FN). T, F, P, and N correspond to true, false, positive, and negative respectively.

The result in the third row emphasize the role of the divergence term at junctions. The divergence-free property of the vector field favours junctions where total incoming branch width equals total outgoing branch width. Figure 3 shows streamline plots of the final vector field configuration superimposed on the thresholded ϕ corresponding to the last two segmentation results in figure 2. The vector field is indeed of constant (unit) magnitude inside the network, parallel to the network boundaries, and smooth; the flow is approximately conserved along network branches and in particular at junctions, where the total incoming flow is approximately equal to total outgoing flow.

Figure 4 shows the segmentation of a river network from a colour optical image.[3] The likelihood model used was the same, but with one less band. The flow conservation property and its geometric consequences enable the algorithm to avoid confounding factors in the background and segment the network to a good accuracy.

The results we have shown here still have false positives and false negatives. The main raison for the former is that the gradient descent algorithm becomes locally stuck in a local minimum, so that some of the background remains classified as network even if this is globally energetically unfavourable. The main reason for the latter are long gaps in the visible network caused by occlusions.

4 CONCLUSION

We have proposed a new algorithm for the segmentation of networks from VHR satellite images. Such networks have characteristic geometric properties: network branch widths change slowly although different branches may have very different widths; while, at junctions, there is approximate 'conservation of width'. To segment such networks (quasi-)automatically, requires models that

[2] Parameter values were, for the 4 images in figure 1 from left to right: $(\lambda_{04}, \lambda_{03}, \lambda_{22}, \lambda_{21}, D, \beta, d, L_v, D_v, \theta) = (0.3375, -0.1767, 0.2712, -0.6, 0.2645, 0.0629, 1.68, 0.2649, 100, 0.03)$, $(0.1, 0.0164, 0.1162, -0.8, 0.0512, 0.0205, 1.45, 0.0227, 200, 0.01)$, $(0.412, -0.0008, 0.0022, -0.6, 0.257, 0.0083, 8.33, 0.275, 50, 0.04)$, and $(0.4, -0.018, 0.15, -0.8, 0.548, 0.0316, 3.45, 0.150, 50, 0.04)$.

[3] Parameter values were: $(\lambda_{04}, \lambda_{03}, \lambda_{22}, \lambda_{21}, D, \beta, d, L_v, D_v, \theta) = (0.25, 0.0323, 0.1138, -0.8, 0.1903, 0.0176, 2.56, 0.0644, 100, 0.07)$.

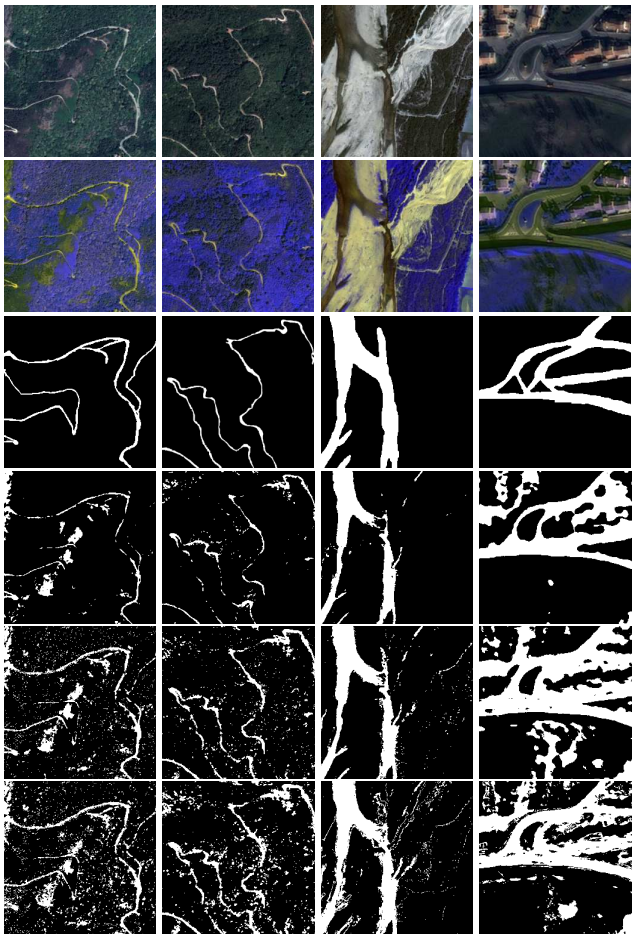

Figure 1: Each column corresponds to a multispectral Quickbird image. The resolution of the first three images is $1/4$ the original resolution (2.44m) while the fourth image is at full resolution (0.61m). From top to bottom: RGB channels of the images; GBI (mapped to RGB) channels of the images; reference segmentation; segmentations obtained using NDVI (columns 1, 2, and 4) and NDWI (column 3) and optimal thresholding; ML segmentations using the MG model; and ML segmentations using the MMG model. (Original images ©DigitalGlobe, CNES processing, images acquired via ORFEO Accompaniment Program).

incorporate this prior knowledge of network region geometry. We use a phase field HOAC model of directed networks to model region geometry and incorporate these properties. In addition to terms favouring network regions, the model contains a vector field representing a 'flow' in the network. This flow has approximately constant speed and is approximately conserved, which leads the model to favour network regions possessing the above geometric properties. Coupled with suitable image models, the model results in a quasi-automatic network segmentation algorithm with good performance, capable of closing many gaps, and of avoiding confounding factors in the background.

The main difficulties with the current method are computational. To ensure the necessary geometric properties, the divergence term must be strong. This implies a small time step in the gradient descent algorithm to avoid the algorithm diverging, which means long computation times. A second difficulty is that despite the stability analysis, the model still has several free parameters that need to be tuned by hand. The use of graph cut algorithms and parameter learning are possible ways to overcome these difficulties, and we are currently studying them.

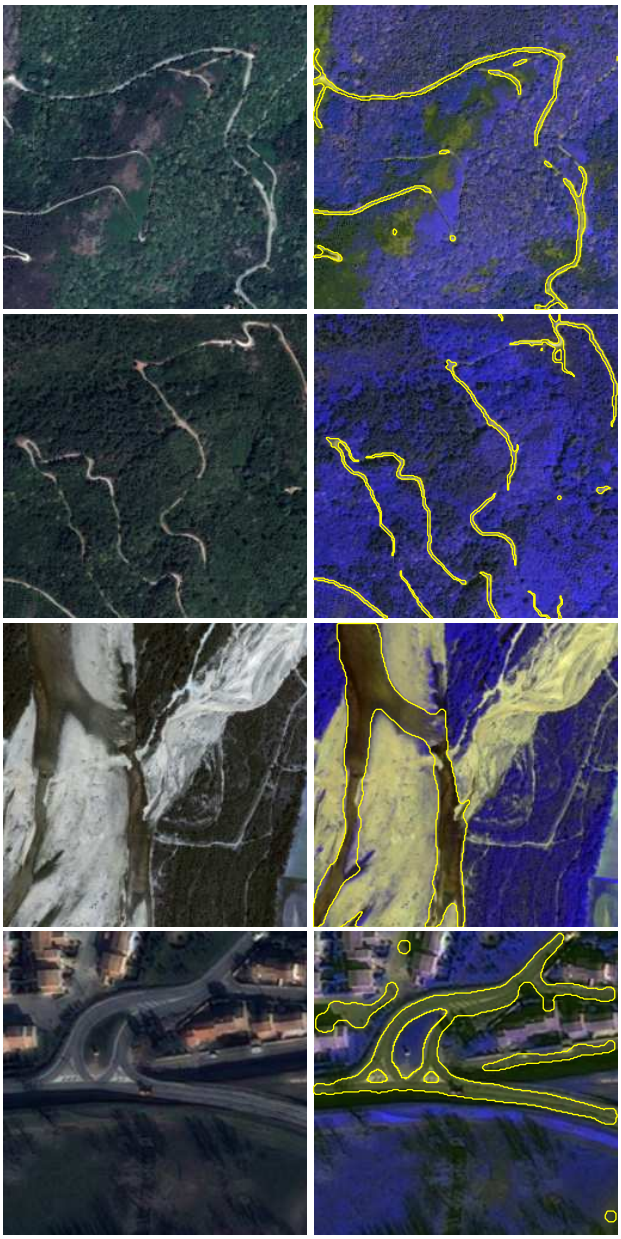

Figure 2: Left: RGB channels of multispectral QuickBird images; right: segmentation results, using the new algorithm, superimposed on GBI (mapped to RGB) channels.

ACKNOWLEDGMENTS

The authors thank the French Space Agency (CNES) for the satellite images, and CNES, the PACA Region, and INRIA Associated Team SHAPES for partial financial support.

REFERENCES

Cremers, D., Kohlberger, T. and Schnörr, C., 2003. Shape statistics in kernel space for variational image segmentation. Pattern Recognition 36, pp. 1929–1943.

Dillabaugh, C. R., Niemann, K. O. and Richardson, D. E., 2002. Semi-automated extraction of rivers from digital imagery. Geoinformatica 6(3), pp. 263–284.

El Ghoul, A., Jermyn, I. H. and Zerubia, J., 2009a. Inflection point model under phase field higher-order active contours for network extraction from vhr satellite images. In: Proc. European Signal Processing Conference, Glasgow, Scotland.

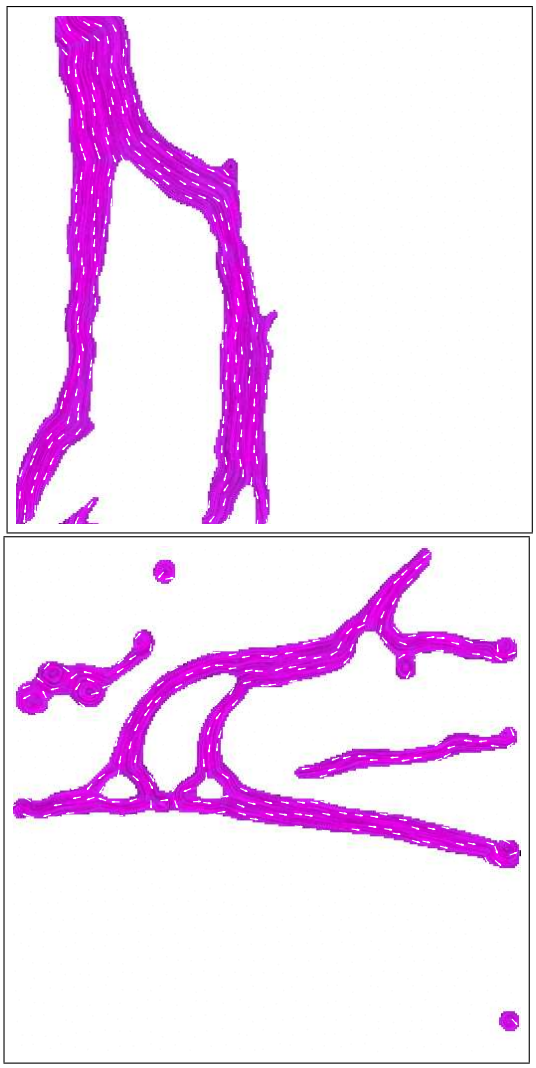

Figure 3: The final configuration of ϕ (thresholded) and v corresponding to the images in the third and fourth rows in figure 2. The vector field is indeed zero outside the network and of constant (unit) magnitude inside, smooth, parallel to branch boundaries, and conserved along network branches and at junctions.

Figure 4: River network segmentation from a colour image using the new algorithm. Flow conservation is satisfied at junctions, illustrated by a significant change of width, and along branches, as illustrated by a slow change of width. (©2010 Google - Imagery ©2010 TerraMetrics, Map data ©2010 Tele Atlas.)

El Ghoul, A., Jermyn, I. H. and Zerubia, J., 2009b. A phase field higher-order active contour model of directed networks. In: 2nd IEEE Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment, at ICCV, Kyoto, Japan.

Foulonneau, A., Charbonnier, P. and Heitz, F., 2003. Geometric shape priors for region-based active contours. In: Proceedings of the IEEE International Conference on Image Processing, Barcelona, Spain.

Kass, M., Witkin, A. and Terzopoulos, D., 1988. Snakes: Active contour models. Int. J. Comput. Vis. 1(4), pp. 321–331.

Lacoste, C., Descombes, X., Zerubia, J. and Baghdadi, N., 2004. Unsupervised line network extraction from remotely sensed images by polyline process. In: Proc. European Signal Processing Conference, University of Technology, Vienna, Austria.

Lacoste, C., Descombes, X., Zerubia, J. and Baghdadi, N., 2005. Extraction of hydrographic networks from satellite images using a hierarchical model within a stochastic geometry framework. In: Proc. European Signal Processing Conference, Antalya, Turkey.

McFeeters, S., 1996. The use of the normalized difference water index (NDWI) in the delineation of open water features. International Journal of Remote Sensing 17(7), pp. 1425–1432.

Mena, J. B., 2003. State of the art on automatic road extraction for GIS update: a novel classification. Pattern Recogn. Lett. 24(16), pp. 3037–3058.

Peng, T., Jermyn, I. H., Prinet, V. and Zerubia, J., 2010. Extended phase field higher-order active contour models for networks. Int. J. Comput. Vis. 88(1), pp. 111–128.

Rochery, M., Jermyn, I. and Zerubia, J., 2006. Higher order active contours. Int. J. Comput. Vis. 69(1), pp. 27–42.

Rochery, M., Jermyn, I. H. and Zerubia, J., 2005. Phase field models and higher-order active contours. In: Proceedings of the IEEE International Conference on Computer Vision, Beijing, China.

Rouse, J., Haas, R., Schell, J. and Deering, D., 1973. Monitoring vegetation systems in the great plains with ERTS. In: ERTS Symposium, NASA SP-351, pp. 309–317.

Rousson, M. and Paragios, N., 2002. Shape priors for level set representations. In: Proceedings of the European Conference on Computer Vision, Copenhagen, Denmark.

Srivastava, A., Joshi, S., Mio, W. and Liu, X., 2003. Statistical shape analysis: Clustering, learning, and testing. IEEE Trans. Pattern Anal. Mach. Intell. 27(4), pp. 590–602.

Tucker, C., 1979. Red and photographic infrared linear combinations for monitoring vegetation. Remote Sensing of Environment 8(2), pp. 127–150.