

A phase field higher-order active contour model of directed networks

Aymen El Ghouli, 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

Abstract

The segmentation of directed networks is an important problem in many domains, e.g. medical imaging (vascular networks) and remote sensing (river networks). Directed networks carry a unidirectional flow in each branch, which leads to characteristic geometric properties. In this paper, we present a nonlocal phase field model of directed networks. In addition to a scalar field representing a region by its smoothed characteristic function and interacting non-locally so as to favour network configurations, the model contains a vector field representing the ‘flow’ through the network branches. The vector field is strongly encouraged to be zero outside, and of unit magnitude inside the region; and to have zero divergence. This prolongs network branches; controls width variation along a branch; and produces asymmetric junctions for which total incoming branch width approximately equals total outgoing branch width. In conjunction with a new interaction function, it also allows a broad range of stable branch widths. We analyse the energy to constrain the parameters, and show geometric experiments confirming the above behaviour. We also show a segmentation result on a synthetic river image.

1. Introduction

The automatic extraction of specific entities from real images is in general a difficult problem. (By ‘extraction’ is meant: ‘find the region R in the image domain that “contains” the entity’, where ‘contains’ means that R is the projection to the image domain of the volume occupied by the entity in the real world.) The difficulty arises because the entity sought is usually not distinguishable from the background using local image measurements alone. Rather, knowledge of the probable ‘shape’ of the region R is necessary in order to extract the entity successfully. In many applications, this knowledge is currently provided, in one way or another, by a human being. Automation of the extraction process therefore requires models that incorporate this

knowledge of region geometry. Mathematically, speaking, we seek to construct the probability distribution $P(R|I, K)$, where I is the image data and K represents prior knowledge about R and the relation between R and I . As usual, this can be written as the product of a likelihood $P(I|R, K)$, and a prior $P(R|K)$ that incorporates knowledge of region ‘shape’. (In practice, we will deal with negative log probabilities, i.e. a total energy $E(R, I)$ that is the sum of a likelihood term $E_l(I, R)$ and a prior term $E_p(R)$.)

The simplest prior models incorporate only local knowledge about ‘shape’, e.g. boundary smoothness. This proves insufficient, however, for all but the simplest images, and as a result, there has been a great deal of work on models incorporating more sophisticated shape knowledge. Most of this work models an ensemble of regions as perturbations of one or more reference regions [1, 3, 6, 7, 8]. This is a very flexible approach, and it works well for many applications. It is not appropriate, however, when the region sought can have arbitrary topology (e.g. if the entity consists of an unknown number of similar objects), since such an ensemble of regions cannot be described as perturbations around a finite number of points in the space of regions. A particular problem that falls into this latter category is the extraction of ‘network’-shaped regions (i.e. regions composed of thin branches that join together at junctions), e.g. road or river networks in satellite images, or vascular networks in medical images. Such network regions may have several connected components, and may be multiply connected.

The case of undirected networks has been addressed in the literature, using the ‘higher-order active contour’ (HOAC) framework [4]. HOACs incorporate non-trivial shape knowledge via long-range interactions between region boundary points, thereby removing the need for reference regions. A reformulation of HOACs as nonlocal ‘phase field’ models [5] avoids certain drawbacks of the contour representation, facilitates model analysis and implementation, allows a ‘neutral’ initialization and complete topological freedom, and results in reduced execution times.

Many of the networks that appear in applications, however, are directed. Each network branch has a ‘flow direc-

tion', and each junction therefore has 'incoming' and 'outgoing' branches. The existence of such a flow typically changes the geometry of the network, because often the flow is in some sense conserved. Thus branches tend not to end, because this would involve the flow stopping, and junctions often consist of small-width incoming branches joining together to form larger-width outgoing branches. This paper proposes a model for such directed networks. The eventual goal is hydrographic network extraction from remote sensing images, but the model is probably relevant to other applications involving networks. More specifically, we would like a prior model $P(R|K)$ that satisfies the following desiderata:

1. a large range of branch widths should be possible, but
2. changes of width should be slow, except
3. at junctions, where the branch widths should be (softly) constrained so that $\sum_i w_i = 0$, where the widths w_i are negative for incoming flow and positive for outgoing flow. Note that this includes the fact that branches should not end, *i.e.* they should be prolonged, since these can be viewed as junctions with only incoming flow.¹

To construct such a model, we extend the nonlocal phase field model of undirected networks described in [5]. The phase field function ϕ still represents the region R corresponding to the network, and still interacts nonlocally so as to favour network configurations. The novel element is a tangent vector field v representing the 'flow' through the network branches. The vector field is coupled to ϕ in such a way that it is strongly encouraged to be zero outside R ; to have unit magnitude inside R ; to have zero divergence; and, more weakly, to be smooth. The transition from unit magnitude inside the region to zero magnitude outside, coupled with small divergence, encourages the vector field to be parallel to the region boundary. Both the divergence and smoothness constraints then tend to propagate this parallelism to the interior of network branches. (It is possible to encourage such parallelism directly, but this is less physical, and in practice we have not found it necessary.) Small divergence and parallelism, coupled with the constraint on the magnitude, aids prolongation of network branches; allows a larger range of stable widths; controls rate of change of width along a branch; and encourages asymmetric junctions

¹Such a linear constraint arises if 'flow volume' is proportional to branch width in the image. This will be true if flow speed is roughly constant, and if 'channel volume' is proportional to branch width in the image, which is true for river networks if channel depth is roughly constant. On the other hand, for tubular networks in three dimensions, one would rather expect the sum of the (appropriately signed) squares of the widths to be zero. In any case, our aim is not to model the detailed physics of each situation, but to model networks that possess certain qualitative geometric properties.

for which total incoming branch width equals total outgoing branch width.

Section 2 recalls the model in [5] and then describes the new directed network model. Section 3 describes numerical experiments showing that the above behaviour is indeed realized, and a segmentation experiment showing the improvement brought by using the directed network model.

2. Phase field model of directed networks

We first recall the undirected network model in [5]. A phase field ϕ is a real-valued function on the image domain Ω . A phase field determines a region by the map $\zeta_z(\phi) = \{x \in \Omega : \phi(x) > z\}$ where z is a given threshold. The basic phase field energy is [5]

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

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. The derivative term ensures the smoothness of ϕ_R , producing a narrow interface around the boundary ∂R interpolating between -1 and $+1$. The energy of ϕ_R is given approximately by a linear combination of the length of ∂R and the area of R [5]. E_0^s is thus equivalent to a classical active contour model.

To introduce prior shape information, a nonlocal term is then added to give a total energy $E_P^s = E_0^s + E_{NL}$, where [5]

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

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, G , which decreases as a function of the distance between the points. In this paper, the interaction function G will be taken to be either the interaction function described in [5], or the modified Bessel function of the second kind of order 0, K_0 . For certain parameter ranges, E_P^s favours undirected network-shaped regions where the network branches have approximately the same width. The parameter ranges can be found by a stability analysis [2].

2.1. Directed networks

Directed networks by definition possess a sense of direction in each branch, usually due to a unidirectional flow

through a channel represented by the branch. Conservation of flow then leads to geometric constraints on the network, particularly at junctions, meaning that directed networks possess different characteristic geometric properties to undirected networks. To model such networks, we introduce a phase field prior, $E_P(\phi, v)$, that in addition to ϕ , is a functional of a tangent vector phase field v which ‘represents’ the flow through the network. We note immediately that v is not supposed to be equal to the physical flux through the network. This would require a much more complicated model than the one we are proposing. Rather, v is an auxiliary quantity (probabilistically speaking, a hidden variable) that introduces interactions that constrain the geometry of the network. We will see, however, that in many ways it does behave like a physical flux.

Since v is zero outside the network, we design the ultralocal potential of the model to have only two local minima: the background, where $(\phi(x), |v(x)|) = (-1, 0)$; and the foreground where $(\phi(x), |v(x)|) = (1, 1)$. This control of the magnitude of v is in one sense unphysical: for example, the channel may widen while the flow speed decreases, thereby conserving the flow. However, in another sense, it represents real physical effects. Often rigidity in the physical nature of the channel (*e.g.* stiffness of the channel wall, resistance to widening in the substrate in which the channel is embedded) means that such widening is not possible. Some rigidity is already built into the model via ϕ , but control of the magnitude of v , coupled with the divergence term to be described in a moment, reinforces this: it will control the rapidity of width variations. Control of the magnitude of v also represents the fact that in directed networks there is a force that pushes the flow through the network (*e.g.* gravity, pressure), which, in conjunction with viscosity and frictional forces, produces a preferred speed. Again, we do not pretend to be modelling the physics in detail: the constraint on the magnitude of v is a stand-in for these effects, designed to realize certain constraints on the geometry.

In addition to the potential, we introduce a term that penalizes the divergence of v . This represents a soft version of flow conservation, but the parameter multiplying this term will be large so that in general the divergence will be small. We also add a small overall smoothing term on v , 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 boundary, *i.e.* the flow is along the channel. Coupled with the constraint on $|v|$ inside the channel, this means that width variations are constrained to be slow along a channel, while at junctions, it tends to produce configurations where total incoming flow is approximately equal to total outgoing flow, which translates to the sum of the incoming widths being approximately equal to the sum of the outgoing widths.

The total prior energy, $E_P(\phi, v)$, is then the sum of a local term E_0 and the nonlocal term E_{NL} . E_0 is

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

The third term is the smoothing term: $\partial v : \partial v = \sum_{m,n} (\partial_m v^n)^2$, where $m, n \in \{1, 2\}$ label the two Euclidean coordinates. $W(\phi, v)$ is a potential which defines the stable phases $(\phi, |v|) = (-1, 0)$ and $(\phi, |v|) = (1, 1)$. The generic form of the potential W we use is a fourth order polynomial in ϕ and $|v|$, constrained to be differentiable. We define $W(\phi, v) = w(\phi, |v|)$, where W is given by

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

We fix many relations between the parameters by requiring that the two minima described above are the only local minima, by requiring that W be bounded below, and by requiring that the potential energy of the foreground is greater than of the background, *i.e.* $w(1, 1) > w(-1, 0)$. Figure 1 shows a contour plot of w satisfying these requirements, showing local minima at the desired points. The saddle point (ϕ_s, v_s) between the two minima plays an important role in initializing the gradient descent algorithm: the ‘neutral’ initialization 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 , is 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 $(-1, 0)$. Thus the background is globally stable, not only metastable.

Just as in the case of the the undirected network model, we expect that E_P has local minima corresponding to network shapes. This was directly verified for the undirected network model via a stability analysis of a long bar [2]. We intend to perform such an analysis for the current model in the future. However, the numerical experiments show that such an expectation is indeed correct.

3. Experiments

In sections 3.1 and 3.2, we study gradient descent evolutions using E_P^S and E_P . In section 3.3, we show a segmentation result on a synthetic image of a ‘river configuration’.




Figure 1. The potential w .

Derivatives are discretized using the Fourier basis, and we employ a forward Euler scheme. We omit the expressions for the functional derivatives, as these are standard calculations.

3.1. Geometric evolutions of v for fixed ϕ

Figure 2 shows gradient descent evolutions of v with ϕ fixed to a ‘junction’ configuration.² The first experiment uses the divergence term but no smoothing. Initialized with constant $v = (0, -v_s)$, it shows the effects of the divergence term and the potential, which align the field with the network while conserving the flow. In the second experiment, the initial v was given a random direction at each point. In this case, the divergence of v is small and v is parallel to the boundary near the boundary, but it is not smooth. The condition of small divergence still allows a great deal of freedom: we can add to v any u with $\partial \cdot u = 0$ that preserves the magnitude. The third experiment uses the smoothing term instead of the divergence term. The result is a smooth field, but the field does not run along the network. The fourth experiment uses both divergence and smoothing terms. The vector field evolves from an initially random configuration towards a smooth and divergence-free configuration that runs along the network.

3.2. Geometric evolutions of v and ϕ

Figure 3 shows gradient descent evolutions using the new energy E_p .³ For comparison, figure 4 shows gradient descent evolutions using the undirected phase field model

²The parameter values were $(\lambda_{04}, \lambda_{03}, \lambda_{22}, \lambda_{21}) = (2.275, -0.467, 1.34, -3)$ for all evolutions, while from top to bottom: $(L_v, D_v) = (0, 0.1); (0, 0.1); (0.1, 0)$ and $(0.1, 0.1)$.

³From top to bottom, parameter values were: $(\lambda_{04}, \lambda_{03}, \lambda_{22}, \lambda_{21}, D, \beta, d, L_v, D_v) = (3.13, -0.99, 0.131, -2, 0.7, 0.4, 4, 7, 240); (3.13, -0.99, 0.131, -2, 1, 2, 4, 7, 240); (1, 0.072, 0.207, -1, 1, 2, 3, 2, 20)$ and $(1.25, -0.325, 0.368, -1, 1, 2, 3, 7, 240)$.




Figure 2. Geometric evolutions of v keeping ϕ fixed. First column: initial configuration. Second column: intermediate configuration. Third column: final configuration. From top to bottom: result with the divergence term using a vertical initialization; result with the divergence term using a random initialization; result with the smoothing term; result with the smoothing and divergence terms.

E_p^s .⁴ The initial regions are shown in the leftmost column; time runs from left to right. The binary images are obtained by thresholding ϕ at ϕ_s . The initial configuration for v had $(0, -v_s)$ everywhere, while ϕ had the value -1 outside the region and ϕ_s inside.

Both models produce stable network configurations, but it is illuminating to examine the differences in detail. The first two experiments in each set used the interaction function in [4, 5]. This interaction function constrains the possible stable widths quite severely. This renders moot desiderata 1 and 3 in section 1, and as a result, in these four experiments, the widths of the branches are all more or less the same. The directed model, though, tends to produce straighter branches with even less width variation than those in the undirected case. This corresponds to desideratum 2.

The last two experiments in each group use K_0 as the

⁴From top to bottom, parameter values were: $(\lambda, \alpha, D, \beta, d) = (1.24, 0.038, 0.75, 0.137, 4); (1.24, 0.056, 0.75, 0.125, 4); (1.23, 0.076, 0.75, 0.83, 5.65)$ and $(1.23, 0.076, 0.75, 0.917, 5.65)$.

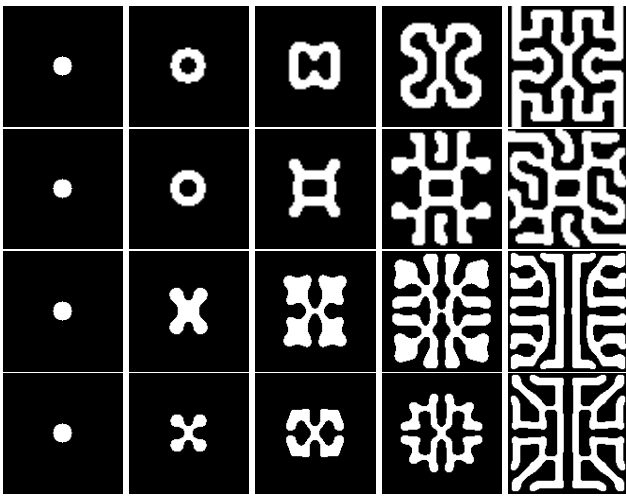


Figure 3. Gradient descent evolutions using the new, directed network model E_p . The initial regions are shown in the leftmost column; time runs from left to right.




Figure 4. Gradient descent evolutions using the undirected network model E_p^s . The initial regions are shown in the leftmost column; time runs from left to right.

interaction function. This constrains the stable width far less severely, and as a result we see a large range of widths in both groups. We have thus satisfied desideratum 1 for our directed model. However, the spatial distribution of this range is very different in the undirected and the directed cases. In the undirected case, each branch varies in width along its length, tending to bulge out away from junctions. In the directed case, in contrast, each branch tends to preserve the same width, although the extent to which it does this depends on the value of the divergence term. This can be seen by comparing the last two experiments in the directed case. The parameter D_v was 20 in the third experiment and 240 in the fourth. As a result, there are some width variations along each branch in the third experiment, whereas in the fourth, they are absent. This satisfies desideratum 2. In addition, the fourth experiment produces long straight branches, while at junctions the sum of incoming and outgoing widths tend to be similar. This can be seen in figure 3, but it is more clearly illustrated in figure 5, which shows a zoom on the bottom-left quarter of the final configuration in the fourth experiment with the directed model.⁵ Thus desideratum 2 is satisfied too.

3.3. Segmentation

Figure 6 shows an experiment using a synthetic image of a ‘river’, consisting of three regions each with a differ-

⁵Although not shown here, the evolution of the vector field in this experiment showed an interesting behaviour. At a certain point, it ran from right to left across the short, narrow, horizontal branch (call it B) in figure 5. The branch joining B from the lower left widened during the gradient descent, and the flow in B gradually reversed to accommodate the extra flow. This in turn increased the width of the wide, vertical branch flowing down from the right-hand end of B.



Figure 5. A zoom on the bottom-left quarter of the final configuration in the fourth experiment with the directed model, shown in figure 3, showing v as well as ϕ .

ent (constant) intensity value, plus added Gaussian noise. The highest intensity I_1 corresponds to the ‘river’, while the lowest I_{-1} and intermediate $I_0 = (I_{-1} + I_1)/2$ intensities correspond to the background. The zone with intensity I_0 is designed to resemble a network, but one that does not respect ‘flow’ conservation. The likelihood $P(I|R, K)$ is a product of Gaussian distributions for the intensity at each pixel, with means I_1 for points in R and I_{-1} for points in its complement \bar{R} . The variances for R and \bar{R} are the same. Maximum likelihood classification is thus unable to classify points with intensity I_0 ; the prior that decides whether such points are part of the estimated river region or not.




Figure 6. From left to right, top to bottom: synthetic image with three grey levels and added noise; ground truth; segmentation using undirected network model; segmentation using directed network model. Note how the constraint on branch width in the directed network model avoids including parts of the background that have similar intensity to the ‘river’.

As can be seen in the bottom row of figure 6, the undirected network model, although it finds a network-shaped region, includes a significant amount of the confounding region with intensity I_0 . The directed network model on the other hand, is considerably more accurate because of the geometric constraints arising from ‘low’ conservation.⁶ Figure 7 shows a zoom on the central part of the bottom-right result in figure 6, showing v as well as ϕ .

4. Conclusion

We have introduced a nonlocal phase field model for directed network-shaped regions. The model contains two field variables: a scalar field that describes the region by its smoothed characteristic function, with a long-range interaction that tends to produce network-shaped regions, and a vector field that represents the ‘flow’ through the network. The vector field is strongly encouraged to be divergence-free, and of unit magnitude inside and zero magnitude outside the region. This forces the field in the region to be parallel to the region boundary, and to conserve flow. This results in slow width variations along a network branch, ex-

⁶The parameters used for the undirected and directed network models were: $(\lambda, \alpha, D, \beta, d) = (25, 0.053, 15, 0.01)$ and $(\lambda_{04}, \lambda_{03}, \lambda_{22}, \lambda_{21}, D, \beta, d, L_v, D_v) = (1.025, -0.211, 0.564, -1.5, 0.2, 0.2, 5, 8, 1000)$.




Figure 7. A zoom on the central part of the result in figure 6, showing v as well as ϕ .

cept at junctions, where total incoming flow/width is encouraged to be equal to total outgoing flow/width. We have confirmed the expected behaviour of the model via gradient descent evolutions, and via an extraction experiment on a synthetic image representing part of a hydrographic network, showing that the new model can avoid errors that arise if the undirected network model is used.

Acknowledgments

The authors thank the French Space Agency (CNES) for providing the satellite images, and CNES, the PACA Region, the MUSCLE Network of Excellence (FP6-507752), and INRIA Associated Team SHAPES for partial financial support.

References

- [1] D. Cremers, T. Kohlberger, and C. Schnörr. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36:1929–1943, 2003.
- [2] A. El Ghouli, I. H. Jermyn, and J. Zerubia. Phase diagram of a long bar under a higher-order active contour energy: application to hydrographic network extraction from VHR satellite images. In *International Conference on Pattern Recognition*, Tampa, Florida, Dec. 2008.
- [3] N. Paragios and M. Rousson. Shape priors for level set representations. In *Proceedings of the European Conference on Computer Vision*, Copenhagen, Denmark, May 2002.
- [4] M. Rochery, I. Jermyn, and J. Zerubia. Higher order active contours. *International Journal of Computer Vision*, 69(1):27–42, 2006.
- [5] M. Rochery, I. H. Jermyn, and J. Zerubia. Phase field models and higher-order active contours. In *Proceedings of the IEEE*

International Conference on Computer Vision, Beijing, China, Oct. 2005.

- [6] A. Srivastava, S. Joshi, W. Mio, and X. Liu. Statistical shape analysis: Clustering, learning, and testing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(4):590–602, 2003.
- [7] M. Taron, N. Paragios, and M.-P. Jolly. Registration with uncertainties and statistical modeling of shapes with variable metric kernels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31:99–113, 2009.
- [8] M. Vaillant, M. I. Miller, A. L. Younes, and A. T. D. Statistics on diffeomorphisms via tangent space representations. *NeuroImage*, 23:161–169, 2004.