INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Higher-Order Active Contour Energies for Gap Closure

Marie Rochery —— Ian Jermyn —— Josiane Zerubia

N° 5717

octobre 2005

Thème COG

Rapport de recherche

Higher-Order Active Contour Energies for Gap Closure

Marie Rochery , Ian Jermyn , Josiane Zerubia

Abstract: One of the main difficulties in extracting line networks from images, and in particular road networks from remote sensing images, is the existence of interruptions in the data caused, for example, by occlusions. These can lead to gaps in the extracted network that do not correspond to gaps in the real network. In this report, we describe a higher-order active contour energy that in addition to favouring network-like regions composed of thin arms joining at junctions, also includes a prior term that penalizes network configurations containing 'nearby opposing extremities', and thereby makes their appearance in the extracted network less likely. If nearby opposing extremities form during the gradient descent evolution used to minimize the energy, the new energy term causes the extremities to attract one another, and hence to move towards one another and join, thus closing the gap. To minimize the energy, we develop specific techniques to handle the high-order derivatives that appear in the gradient descent equation. We present the results of automatic extraction of networks from real remote-sensing images, showing the ability of the model to overcome interruptions.

Key-words: road extraction, shape, prior, continuity, gap, closure, quadratic, higher-order, active contour, level set

Contours Actifs d'Ordre Supérieur pour l'Extraction de Réseaux Connectés

Résumé : L'un des principaux problèmes lors de l'extraction de réseaux linéiques dans des images, et en particulier l'extraction de réseaux routiers dans des images de télédétection, est l'existence d'interruptions dans les données, causées, par exemple, par des occultations. Ces interruptions peuvent mener à des trous dans le réseau extrait qui n'existent pas dans le réseau réel. Dans ce rapport, nous décrivons une énergie de contour actif d'ordre supérieur qui, en plus de favoriser les régions composées de bras fins et connectés entre eux, inclut un terme d'a priori qui pénalise les configurations du réseau où des extremités proches et se faisant face apparaissent. L'apparition dans le réseau extrait de ces configurations est donc moins probable. Si des extremités proches et se faisant face apparaissent pendant l'évolution par descente de gradient utilisée pour minimiser l'énergie, le nouveau terme dans l'énergie crée une attraction entre ces extremités, qui se rapprochent donc l'une de l'autre et se rejoignent, fermant ainsi le trou entre elles. Pour minimiser l'énergie, nous développons des techniques spécifiques pour traiter les derivées d'ordre élevé qui apparaissent dans l'équation de descente de gradient. Nous présentons des résultats d'extraction automatique de réseaux routiers à partir d'images de télédétection, montrant ainsi la capacité du modèle à surmonter les interruptions.

Mots-clés : extraction de routes, forme, géométrie a priori, continuité, quadratique, contour actif d'ordre supérieur, courbe de niveaux

Contents

1   Introduction

The huge growth in the amount of digital imaging data of various types available in many fields, including remote sensing, medicine, and biology, makes the construction of systems capable of automatically extracting information of semantic value from this data a necessity. While each application comes with its own semantics, prior knowledge, and other specificities that mean it must in principle be treated anew, in practice there are frequently commonalities between different applications that make a somewhat more generic approach feasible. Line networks are an example. They represent information of semantic value in many applications: vascular networks in medicine; filamentary structures in biological images; road and river networks in remote sensing and cartography; and galactic filaments in astronomy; while possessing many properties in common across these applications. The extraction of their properties from imaging data (most often, the identification of the region in the image domain corresponding to the network) is thus of some importance, but semi-automatic extraction remains a time consuming and expensive task. Research concerned with these problems has therefore begun to focus on the development of efficient methods for the automatic extraction of line networks. In this report, we particularly focus on the extraction of line networks from remote-sensing images, but the models described should be equally useful for imagery from other applications, including medical and biological images.

In order to be able to extract networks, we must first be able to model them well, that is, be able implicitly or explicitly to put an image-dependent probability distribution $P(R|I)$ on the space of regions in the image domain whose mass is concentrated on the region corresponding to the network in I. Such a probability can of course be decomposed into two pieces: a prior probability on the space of regions given that the region corresponds to a network, $P(R)$, and a likelihood describing the images to be expected given that the region R in the image domain corresponds to a network, $P(I|R)$. The construction of distributions $P(R)$ and $P(I|R)$ that generate a posterior probability $P(R|I)$ concentrated on the network in a given image is not easy. Generic priors that are not concentrated on network-shaped regions combined with likelihoods based on local image measurements are not sufficient to concentrate $P(R|I)$ on the network sought, because local image measurements typically assume similar values for many regions that do not correspond to the network, and the prior is incapable of distinguishing between them. To improve the situation, one can thus advance in two directions. The first is to design likelihoods that capture some of the dependencies amongst the image values associated with the network, the use of line detection filters being the most common example. Even this does not suffice, though, since the measurements are still relatively local, and may be similar for many structures that do not form part of a network. The second is to design priors that are concentrated on network-shaped regions. Constructing such priors is also non-trivial, however, first because the space of regions is an infinite-dimensional nonlinear space, and second because networks form a subset of this space that is difficult to characterize. Networks possess strongly constrained geometric properties (*e.g.* narrow arms with roughly parallel sides), but cannot be defined, for example, as variations around a mean shape.

Rochery et al. (2003) (see also (Rochery et al., 2005b)) have proposed a method for the quasi-automatic[1] extraction of line networks based on advances in both these directions. These advances make use of a new generation of active contour models, introduced by Rochery et al. (2003), and named 'higher-order active contours' (HOACs). While classical active contours use only boundary length and interior area (and perhaps boundary curvature) as prior knowledge, HOACs allow the incorporation of non-trivial prior knowledge about region geometry, and the relation between region geometry and the data, via nonlocal interactions between tuples (pairs in the case of Rochery et al. (2003)) of contour points. They are also intrinsically Euclidean invariant. They differ from most other methods for incorporating prior geometric knowledge into active contours (Chen et al., 2001; Leventon et al., 2000; Foulonneau et al., 2003; Paragios and Rousson, 2002; Cremers et al., 2003) in not being based upon perturbations of a reference region or regions. Using this new framework, Rochery et al. (2003) proposed a model that goes a long way towards capturing the prior geometric knowledge we have of network regions, as well as the complex dependencies between image values associated with networks. The prior model has as low-energy configurations, regions composed of arms of roughly constant width that join together at junctions. The likelihood model predicts not just high image gradients along the edges of the network, but incorporates longer-range dependencies that predict that image gradients along one side of a network arm will be parallel, while image gradients on opposite sides of a network arm will be anti-parallel.

Thanks to this prior knowledge, the model produces good results using gradient descent to minimize the contour energy, starting from a generic initialization that renders the method quasi-automatic. The primary failure mode of the method is the presence of 'gaps' in the extracted networks caused by 'interruptions' in the image of the road. These interruptions are caused by various types of 'geometric noise': in the case of road networks, for example, trees and buildings close to the network that change its appearance either via occlusion or because of cast shadows. The method fails to close these gaps for three reasons, two related to the model, and one to the algorithm:

1. the prior knowledge concerning the geometry of the network ($P(R)$) does not distinguish between two distant arms that each comes to an end, and two arms that form a gap, once the extremities are more than a few pixels apart. Thus the model does not capture our prior knowledge that road networks, for example, usually do not possess such gaps;

2. the prior knowledge concerning the image to be expected from a given network ($P(I|R)$) does not allow for the possibility that there will be interruptions in the observed road;

3. the gradient descent algorithm may be unable to close the gap even if the configuration with the gap closed has lower energy than the configuration with the gap present, due to the shape of the energy surface between the two configurations.

Rochery et al. (2004) made a preliminary attempt to address the gap closure problem. They introduced a 'gap closure' force making nearby opposing network extremities attract one other, thus closing gaps between them. The force was introduced directly to the gradient descent equation. It could not be derived from an energy, thus complicating analysis, and meaning that convergence to

---

[1]By this we mean that no human initialization is required, but that the model possesses parameters that cannot at present be set automatically.

an energy minimum could not be guaranteed. It is the purpose of this report[2] to present a solution to both these problems, and hence to the gap closure problem, by describing a new HOAC energy that penalizes configurations containing gaps, while at the same time changing the shape of the energy surface so that it no longer obstructs the algorithm.

More specifically, based on the geometry of gaps in networks, we design a quadratic HOAC energy for gap closure that penalizes 'nearby opposing extremities'. These extremities are identified by pairs of points that have high positive curvatures, lie outside the contour with respect to one another, and are closer than a certain distance. The effect is that network extremities that are close attract, extend towards one another, and join, thus closing the gap between them.

The new energy leads to a complicated force in the gradient descent equation, a function of third and fourth derivatives of the region boundary. The computation of these terms necessitates careful numerical treatment in order to keep the evolution stable. We use the level set framework to evolve the contour and adapt it to the nature of the gradient descent equation resulting from the HOAC terms in the energy.

Previous work on road extraction has also encountered the problem of interruptions of course, and has dealt with it in different ways, often without addressing it explicitly. Tracking methods (Geman and Jedynak, 1996) and methods minimizing the optimal path between endpoints (Merlet and Zerubia, 1996) generally constrain the topology so that gaps are not possible. The same applies to active contour models requiring endpoints, such as 'ribbon snakes' (Fua and Leclerc, 1990) and 'ziplock snakes' (Neuenschwander et al., 1997). Bicego et al. (2003) use a road tracking method with an 'inertia' term that allows a road extremity to extend a short distance despite lack of support from the data, but do not address gaps as such. A number of methods attempt to close gaps in the extracted network after the fact: Laptev et al. (2000) use ziplock snakes to connect gap endpoints, while Zhang et al. (1999) use morphological operators. Tupin et al. (1998) construct a Markov random field on a graph whose nodes represent line segments, the field labelling the segments as 'road' or 'non-road'. Some of these line segments are extracted from the image by a line detector, while the rest consist of all reasonable potential connections between the extracted segments. A MAP estimate is computed from a model containing the prior geometric knowledge that, for example, roads are long and relatively straight, and that extremities are unlikely. Methods using marked point processes (Stoica et al., 2004; Lacoste et al., 2003) also penalize extremities, not gaps *qua* gaps, but they have the advantage of stochastic algorithms that allow the energetic barrier mentioned above to be overcome. The current method differs from the above methods in two ways. First, it concentrates on gap closure by directly penalizing configurations containing gaps, rather than penalizing isolated extremities. Second, while many of the above methods work with line segments, the method described in the current work deals directly with 2D regions that have an elongated form.

In section 2 we first recall the model proposed by Rochery et al. (2003), and then go on to describe the new energy in detail in section 3. In section 4, we develop the level set method used to evolve the contour. We present results on real aerial images showing the benefits of the new energy in section 5. We conclude in section 6.

---

[2]A shorter article on the work described in this report was published by Rochery et al. (2005a).

2   A model for network extraction

As discussed in section 1, Rochery et al. (2003) have proposed a HOAC energy as a model for automatic line network extraction. In this section, we briefly review this model, and comment on its advantages and deficits.

We will parameterize the space of regions using boundaries, a generic boundary being denoted γ (we will also call it a contour). The model breaks into two pieces, the likelihood energy E_i, and the prior energy E_g, corresponding to $P(I|R)$ and $P(R)$ respectively:

$$E_0(\gamma) = E_i(\gamma; I) + E_g(\gamma) .$$

Here, $I : \Omega \to \mathbb{R}$ is an image, $\Omega \subset \mathbb{R}^2$ being the image domain. The prior energy, E_g, is the sum of three terms: two linear (length and area), and one quadratic HOAC term, which defines an interaction between points:

$$(2.1) \qquad E_g(\gamma) = \lambda \mathcal{L}(\gamma) + \alpha \mathcal{A}(\gamma) - \frac{\beta}{2} \iint dp\, dp'\; \mathbf{t} \cdot \mathbf{t}'\; \Psi(R(p, p')) ,$$

where the integrals are over the contour, parameterized by p; unprimed quantities are supposed evaluated at p or $\gamma(p)$ and primed quantities at p' or $\gamma(p')$; $\mathcal{L}(\gamma)$ is the contour length; $\mathcal{A}(\gamma)$ is the area inside the contour; $\mathbf{t}$ is the tangent vector to the contour; $R(p, p')$ is the Euclidean distance from $\gamma(p)$ to $\gamma(p')$; and Ψ is a function with the form of a smoothed hard-core potential, given by

$$\Psi(x) = \begin{cases} 1 & x - d_{\min} < -\rho, \\ \frac{1}{2}\left(1 - \frac{x - d_{\min}}{\rho} - \frac{1}{\pi}\sin\left(\pi \frac{x - d_{\min}}{\rho}\right)\right) & |x - d| \le \rho, \\ 0 & x - d_{\min} > \rho. \end{cases}$$

The function Ψ is plotted as a dashed line in figure 8 for the values of the parameters we use in the experiments: $d_{\min} = 3$ and $\rho = 1$.

A few comments on the prior energy E_g are necessary. Length and area are classical regularizing terms. The length term suppresses high frequencies in the contour, and thereby enforces contour smoothness, while the area controls the expansion of the contour. The most important part of the model is the HOAC term. This introduces an interaction between pairs of points on the contour. The interaction causes pairs of nearby points with antiparallel tangent vectors to repel each other, and pairs of nearby points with parallel tangent vectors to attract each other. This has two effects: it prevents pairs of points with anti-parallel tangent vectors from coming too close to each other, and it encourages the growth of arm-like structures. As a result, regions with a reticular structure composed of narrow arms with parallel sides have very low energy under this model, and are thus highly favoured. For certain ranges of parameters, such regions are in fact energy minima. (A more detailed description of how this comes about is given by Rochery et al. (2005b), along with examples of purely geometric evolutions.) The energy thus makes a very good prior for networks.

The likelihood energy, E_i, is also composed of three terms:

$$E_i(\gamma; I) = \lambda_i \int dp\, \mathbf{n} \cdot \nabla I + \alpha_i \int_R d^2x\, G[I](x) - \frac{\beta_i}{2} \iint dp\, dp'\; \mathbf{t} \cdot \mathbf{t}'\; (\nabla I \cdot \nabla I')\; \Psi(R(p, p')) .$$

Figure 1: Satellite image (© CNES)and result of line network extraction. Parameter values: $[\lambda = 1, \alpha = -0.2, \beta = 0.2, \lambda_i = 2.4, \alpha_i = 0, \beta_i = 2.4]$.

The first, linear term favours situations in which the outward normal is opposed to a large image gradient, or in other words, in which the road is brighter than its environment. The second, linear term incorporates a simple line detector filter measurement described by Rochery et al. (2005b). The third, quadratic HOAC term describes the joint behaviour of the data at pairs of points on the contour, given the geometry at the pair of points. It favours situations in which pairs of nearby points with antiparallel tangent vectors (*i.e.* points on opposite sides of the road), lie on large image gradients that point in opposite directions, while pairs of nearby points with parallel tangent vectors (*i.e.* points on the same side of the road), lie on large image gradients that point in the same direction.

Figures 1 and 2 show two results obtained using this model, from a satellite image and an aerial image respectively. Two points are worth noting. First, the region occupied by the network in the image is recovered, and not just its skeleton. Second, a generic initialization consisting of a rounded rectangle slightly smaller than the image was used for both experiments; the amount of prior knowledge included in the model means that no special initialization is necessary.

These results (and in general the results obtained with the energy E_0) are satisfactory: most of the network is extracted in each case. However, they are clearly not completely correct. Consider, for instance, the two images in the top row of figure 3. The corresponding results obtained with the energy E_0 are shown in the bottom row. Clearly there are gaps in the extracted networks that do not correspond to gaps in the real road network. Rather, they correspond to 'interruptions' in the imaged network: places where the luminance of the road changes abruptly and ceases to be different from its immediate surroundings. These interruptions are caused by the presence of trees, buildings, and so on close to the network, which cause the interruptions either by occluding the road, or by casting shadows on the road. Some close-up examples of such interruptions are shown in figure 4.

Figure 2: Aerial image (© IGN) and result of line network extraction. Parameter values: $[\lambda = 1, \alpha = -0.17, \beta = 0.2, \lambda_i = 2, \alpha_i = 1, \beta_i = 2]$.

The presence of gaps in the extracted network caused by interruptions in the imaged network is the main failure mode of the model E_0, and is therefore the first point to address in any attempt to improve the model. This is the subject of the next section.

3   An energy term for gap closure

In order to improve the model E_0, we must first understand in more detail why it fails. There are in fact three reasons for the presence of gaps corresponding to interruptions, two connected to the model (the prior and the likelihood terms), and one connected to the gradient descent algorithm used to minimize the energy.

First, the prior knowledge about geometry described by E_g does not distinguish between two distant arms that each comes to an end and two arms that form an aligned gap, once the extremities are further apart than the range of Ψ, *i.e.* a few pixels: the contribution to the energy is the same. Thus the model as it stands does not capture our prior knowledge that road networks, for example, usually do not possess such gaps; it does not describe what might be called the 'continuity' of roads.

Second, the prior knowledge about the image to be expected from a given network, described by E_i, does not include the possibility that there will be interruptions in the observed road. If there is a road present, E_i says that high gradients are expected normal to its sides; these gradients are expected to be parallel on the same side of the road and antiparallel on opposite sides; and the line detector is expected to respond strongly in the interior of the road. All these expectations are violated by situations such as those shown in figure 4.

Third, the gradient descent algorithm may be unable to close the gap even if the configuration with the gap closed has lower energy than the configuration with the gap present (which can be the case, for example, if the image gradients on either side of the interruption are not too large). This is for two reasons. First, the configuration with a gap may lie at a local energy minimum created

Figure 3: Aerial images (© IGN) with shadows on roads and results of extraction. Parameter values: left, $[\lambda = 1, \alpha = -0.7, \beta = 0.2, \lambda_i = 7.5, \alpha_i = 0, \beta_i = 7.5]$; right, $[\lambda = 1, \alpha = -0.35, \beta = 0.75, \lambda_i = 12.5, \alpha_i = 0, \beta_i = 10]$.

by contributions from both E_i and E_g. The likelihood term E_i contributes because at the edges of an interruption there are image gradients (see figure 4). Moving the extremities of the region off these gradients into the low-gradient area in the interruption means increasing E_i. The prior term E_g contributes because, in order to prevent arms from appearing all over the image domain, the parameters in E_g are adjusted so that the energy per unit length of an arm is slightly positive. This means that if the arms on either side of a gap were to extend towards one another, E_g would increase. Second, a local energy maximum is created by E_g when two extremities are less than a few pixels apart; the same repulsion force that controls the width of the arms in the network causes the extremities to repel one another like two magnetic north poles. The top row of figure 15 illustrates this behaviour. The figure shows the result of a purely geometric evolution using E_g, and starting from the leftmost image. The two arms extend (the parameters are adjusted so that the energy per unit length is negative) but repel one another, resulting in a disconnected network. It should to be

Figure 4: Aerial images (© IGN) with shadows on roads.

stressed that, in contrast with the previous two paragraphs, the points made in this paragraph are all algorithmic issues. They mean that once a gap has formed, it is hard to close it, not that gaps necessarily form for all interruptions; sometimes the data configuration means that an interruption does not produce a gap.

Each of these issues leads to a different approach to the gap closure problem. The first suggests that we should modify the prior term, by penalizing configurations that possess gaps. By raising the energy of such configurations, we decrease the possibility of their occurring in the extracted network (although because the gradient descent algorithm does not find the global minimum, this cannot be guaranteed). The second suggests that we should modify the likelihood term, by allowing for the possibility that interruptions may occur. This means introducing extra variables to model interruptions, and then either estimating these variables or (in a probabilistic framework) marginalizing them away. In principle, both these approaches should be followed, since they are both required by the phenomena we are trying to model. In practice, the second approach increases the complexity of the optimization problem significantly, and consequently we will not pursue it in this work, particularly since a modification of the prior term seems to be sufficient to solve the problem.

The third issue, the algorithm, can be tackled in two ways. One is to use an algorithm with better optimality properties than gradient descent. The other is to attempt to remove the local extrema created by the energy. In conjunction with a modification of the prior term to increase the energy of configurations with gaps, this should allow the gap to close in the course of normal gradient descent. We opt for the second approach here.

The idea then is the following. We will introduce a new term to the prior energy that will penalize gaps, or more specifically 'nearby opposing extremities', a notion that will shortly be made more precise. Since the energy has to take into account the joint geometry at distant points of the contour, it must necessarily be a HOAC energy. The minimal choice is a quadratic energy, and this turns out to be sufficient. The energy will increase with the separation between extremities up to a certain distance, meaning that extremities will attract one another if closer than this distance. This attraction will be large enough to overcome the local minimum produced by the image gradients at the edges of interruptions and the positive energy per unit length of the arms, while the dependence on distance of the new energy term will be designed so that it also removes the local maximum produced by the repulsive effect of the existing quadratic prior term.

3.1 Identification of gaps

In order to identify and penalize configurations containing gaps, we start by defining a gap as consisting of two (or maybe more) 'nearby opposing extremities'. Two points p and p' of the contour will be defined as belonging to nearby opposing extremities if they are close enough together, 'opposing' and if they belong to 'extremities' (both terms to be defined shortly). Below, we define 'switch' functions S_n, S_o, and S_e that measure the extent to which each of these conditions is satisfied, and then combine them to form a function S that measures the extent to which the conjunction of the conditions is satisfied:

$$(3.1) \qquad S(p, p') = S_\mathrm{n}(p, p') S_\mathrm{o}(p, p') S_\mathrm{e}(p, p') \ .$$

We now define the constituent functions of this product.

3.1.1 'Extremity'

We measure the extent to which two points belong to extremities by measuring the extent to which they each belong to an extremity:

$$(3.2) \qquad S_\mathrm{e}(p, p') = S_\mathrm{e}(p) S_\mathrm{e}(p') \ ,$$

where $S_\mathrm{e}(p)$ measures the extent to which p belongs to an extremity (we use the same symbol for the two-point and one-point functions). To measure the extent to which a point p of the contour belongs to an extremity, we use the signed curvature $\kappa(p)$ at p. Extremities are the only points in the network that have positive curvature whose magnitude is significant compared to the the reciprocal of the width of the road (all other points have curvature magnitudes small compared to the reciprocal of the width of the road, except for large negative curvatures at junctions), and so we define

$$S_\mathrm{e}(p) = H(\kappa(p)) \ .$$

Here, H is a smoothed Heaviside function given by

$$(3.3) \qquad H(x) = \begin{cases} 0 & x < 0, \\ \frac{1}{2}\left(\frac{x}{\rho_H} - \frac{1}{\pi}\sin\left(\pi \frac{x}{\rho_H} \right) \right) & 0 \leq x \leq 2\rho_H, \\ 1 & x > 2\rho_H, \end{cases}$$

Figure 5: The function H with $\rho_H = 1$.

Figure 6: Identifying points with large, positive curvature

and plotted in figure 5. In practice, we take $\rho_H = 1$. Figure 6 illustrates the idea behind S_{e}.

3.1.2 'Opposing'

We measure the extent to which a pair of points is 'opposing' by measuring the extent to which each is 'external' with respect to the other:

$$(3.4) \qquad S_{\mathrm{o}}(p, p') = S_{\mathrm{ex}}(p, p') S_{\mathrm{ex}}(p', p) \ ,$$

where $S_{\mathrm{ex}}(p, p')$ measures the extent to which p' is 'external' with respect to p. To measure the extent to which point p' is external with respect to point p, we use the dot product between the unit normal at p, $\hat{\mathbf{n}}(p)$, and the unit vector pointing from p to p', $\hat{\mathbf{R}}(p, p') = (\gamma(p') - \gamma(p))/|\gamma(p') - \gamma(p)|$:

$$S_{\mathrm{ex}}(p, p') = H(\hat{\mathbf{R}}(p, p') \cdot \hat{\mathbf{n}}(p)) \ .$$

Figure 7 illustrates the idea behind S_{o}. It is large when p' lies roughly along the outward pointing normal direction from p, and p lies roughly along the outward pointing normal direction from p', which corresponds to the fact that 'opposing' extremities lie outside the contour with respect to one another and are roughly aligned.

Figure 7: Identifying external points

3.1.3 'Nearby'

To measure extent of closeness, we use a function of the distance $R(p, p')$ between the points:

$$S_n(p, p') = \Psi_A(R(p, p')) \, .$$

The function Ψ_A is given by

(3.5)
$$\Psi_A(x) = \begin{cases} \frac{x}{\rho_A} + \frac{1}{\pi} \sin\left(\frac{\pi x}{\rho_A}\right) - 1 & x \le \rho_A \\ 0 & x > \rho_A. \end{cases}$$

Ψ_A tends to zero at distance ρ_A, and is zero thereafter. In combination with the other functions, as in equation (3.1), this means that if the two points are further apart than ρ_A, they are assumed to belong to extremities that do not form a gap in a continuous piece of the network. In practice, we take $\rho_A = 40$.

3.2   Defining the energy for gap closure

Using the function S defined in the last subsection, we can now form a quadratic HOAC energy term by integrating this function (twice) over the contour, thus effectively 'counting' the number of pairs of points corresponding to nearby opposing extremities. This energy term thus penalizes configurations containing nearby opposing extremities:

$$E_{pen}(\gamma) = -\frac{\beta_A}{2} \iint dp \, dp' \, \mathbf{t} \cdot \mathbf{t}' \, S(p, p')$$
$$= -\frac{\beta_A}{2} \iint dp \, dp' \, \mathbf{t} \cdot \mathbf{t}' \, \Psi_A(R(p, p')) \, S_o(p, p') S_e(p, p') \, .$$

This term could have been constructed using $|\mathbf{t}||\mathbf{t}'|$ instead of $\mathbf{t} \cdot \mathbf{t}'$ (Rochery et al., 2005b), but the use of the tangent vectors reinforces the condition that the extremities should be opposing. Note the similarity between E_{pen} and the quadratic term in equation (2.1). The differences are threefold: first, the presence of the functions Ψ_A, S_o, and S_e mean that this energy 'switches on' only when the two points in the integrand belong to nearby opposing extremities; second, the function Ψ_A

Figure 8: The interaction functions. The short range Ψ from the original model is shown dashed ($d_{\min} = 3$, $\rho = 1$), while the new attractive interaction function Ψ_A is shown solid ($\rho_A = 40$). Note the difference in sign and range.

increases with increasing distance, rather than decreasing as does Ψ, thus producing an attractive force between antiparallel tangent vectors rather than a repulsive one; third, the range over which Ψ_A and its gradient are non-zero is set to about ten times the width of a road, and is thus much greater than the range of Ψ, which is about equal to the width of a road. Ψ and Ψ_A are plotted in figure 8, illustrating the difference in range and behaviour.

While the energy E_{pen} succeeds in overcoming the local minimum caused by the image gradients at the edges of interruptions and the positive energy per unit length of the arms, it is not sufficient to overcome the local maximum produced by the repulsive effect of the existing quadratic prior term. To achieve this, we define a second new energy term, E_{can}, whose function is to cancel the effect of this repulsion between pairs of points that belong to nearby opposing extremities. It is given by

$$(3.6) \qquad E_{\text{can}}(\gamma) = \frac{\beta}{2} \iint dp\, dp'\, \mathbf{t} \cdot \mathbf{t}'\, \Psi(R(p, p'))S_{\text{o}}(p, p')S_{\text{e}}(p, p') \ .$$

Equation (3.6) is the negative of the quadratic term in equation (2.1), except that the integrand now includes the functions that identify opposing extremities.

We sum these energies to form a gap closure energy, E_{gap}, given by

$$(3.7) \qquad E_{\text{gap}} = \frac{1}{2} \iint dp\, dp'\, \mathbf{t} \cdot \mathbf{t}' \left(\beta \Psi(R(p, p')) - \beta_A \Psi_A(R(p, p')) \right) S_{\text{o}}(p, p') S_{\text{e}}(p, p') \ .$$

The functions Ψ_A and Ψ thus combine to form one positive function $\Psi_C = \beta\Psi - \beta_A\Psi_A$ that controls the dependence on distance of the new energy.

Finally, the new prior energy we propose for the extraction of road networks, including gap closure, is given by

$$E = E_0 + E_{\text{gap}} \ .$$

3.2.1 Functional derivative of E_{gap}

We will minimize the energy E using gradient descent:

$$(3.8) \qquad \frac{\partial \gamma}{\partial \tau}(s) = -\frac{\delta E}{\delta \gamma(s)} = -\left(\frac{\delta E_0}{\delta \gamma(s)} + \frac{\delta E_{\text{gap}}}{\delta \gamma(s)} \right) \ .$$

The functional derivative of E_0 is given by Rochery et al. (2003). What remains is to calculate the functional derivative of E_{gap}. The presence of the switches means that this calculation is complicated. The result is given in equation (3.9) overleaf.

$$\frac{\delta E_{\text{gap}}}{\delta \gamma(s)} =$$

$$- \int ds' \ \Big\{ \hat{\mathbf{n}} \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{n}}') \, \dot{\Psi}_C \, H_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{t}}' \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{n}})_s \, \Psi_C \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{t}}' \, (\hat{\mathbf{R}}' \cdot \hat{\mathbf{n}}')_s \, \Psi_C \, H_{\text{ex}} \, \dot{H}'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- 2\hat{\mathbf{t}}' \, \kappa_s \, \Psi_C \, H_{\text{ex}} \, H'_{\text{ex}} \, \dot{H}_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{n}} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \frac{1}{R} \, \Psi_C \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$+ \hat{\mathbf{n}}' \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \frac{1}{R} \, \Psi_C \, H_{\text{ex}} \, \dot{H}'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$+ \hat{\mathbf{R}} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \frac{(\hat{\mathbf{R}} \cdot \hat{\mathbf{n}})}{R} \, \Psi_C \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{R}}' \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \frac{(\hat{\mathbf{R}}' \cdot \hat{\mathbf{n}}')}{R} \, \Psi_C \, H_{\text{ex}} \, \dot{H}'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{R}}^{\perp} \, (\boldsymbol{\kappa}_s \cdot \mathbf{t}') \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$(3.9) \qquad + \hat{\mathbf{R}}^{\perp} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{t}}) \, \Psi_C \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{R}}^{\perp} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{n}})_s \, \Psi_C \, \ddot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{R}}^{\perp} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, (\hat{\mathbf{R}}' \cdot \hat{\mathbf{n}}')_s \, \Psi_C \, \dot{H}_{\text{ex}} \, \dot{H}'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{R}}^{\perp} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \kappa_s \, \Psi_C \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, \dot{H}_{\text{e}} \, H'_{\text{e}}$$

$$- \hat{\mathbf{R}}^{\perp} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \frac{(\hat{\mathbf{R}} \cdot \hat{\mathbf{t}})}{R} \, \Psi_C \, \dot{H}_{\text{ex}} \, H'_{\text{ex}} \, H_{\text{e}} \, H'_{\text{e}}$$

$$+ \hat{\mathbf{n}} \, (\boldsymbol{\kappa}_{ss} \cdot \mathbf{t}') \, \Psi_C \, H_{\text{ex}} \, H'_{\text{ex}} \, \dot{H}_{\text{e}} \, H_{\text{e}}$$

$$- \hat{\mathbf{n}} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \left(\ddot{\Psi}_C \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{t}}) + \dot{\Psi}_C \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{t}})_s \right) \, H_{\text{ex}} \, H'_{\text{ex}} \, \dot{H}_{\text{e}} \, H_{\text{e}}$$

$$+ \hat{\mathbf{n}} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \Psi_C \, \left(\ddot{H}_{\text{ex}} \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{n}})_s + \dot{H}_{\text{ex}} \, (\hat{\mathbf{R}} \cdot \hat{\mathbf{n}})_{ss} \right) \, H'_{\text{ex}} \, \dot{H}_{\text{e}} \, H_{\text{e}}$$

$$+ \hat{\mathbf{n}} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \Psi_C \, H_{\text{ex}} \, \left(\ddot{H}'_{\text{ex}} \, (\hat{\mathbf{R}}' \cdot \hat{\mathbf{n}}')_s + \dot{H}_{\text{ex}} \, (\hat{\mathbf{R}}' \cdot \hat{\mathbf{n}}')_{ss} \right) \, \dot{H}_{\text{e}} \, H'_{\text{e}}$$

$$+ \hat{\mathbf{n}} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \Psi_C \, H_{\text{ex}} \, H'_{\text{ex}} \, \left(\dddot{H}_{\text{e}} \, \kappa_s + \ddot{H}_{\text{e}} \, \kappa_{ss} \right) \, H'_{\text{e}}$$

$$+ \boldsymbol{\kappa}_s^{\perp} \, (\hat{\mathbf{t}} \cdot \hat{\mathbf{t}}') \, \Psi_C \, H_{\text{ex}} \, H'_{\text{ex}} \, \dot{H}_{\text{e}} \, H'_{\text{e}} \Big\}$$

$$+ \text{ tangential terms that do not contribute to the motion of the contour.}$$

The notation $\cdot^{\perp}$ indicates rotation by $\pi/2$ anticlockwise. As before, primed quantities are evaluated at s' or $\gamma(s')$, while unprimed quantities are evaluated at s or $\gamma(s)$. For functions of both

Figure 9: Geometric configuration of two arms, with various geometric quantities marked.

variables, a prime indicates the exchange of s and s' with respect to their definition. Dots indicate derivatives. Subscript s indicates a derivative with respect to s. The arguments of Ψ_C and its derivatives are understood to be R. The functions H_{ex} and H_{e} are both the function H, but their arguments and those of their derivatives are understood to be $\hat{\mathbf{R}} \cdot \hat{\mathbf{n}}$ and κ respectively.

3.3 Analysis of the thin road case

In order better to understand the new energy term E_{gap}, it is useful to consider it in the limit when the width of the arms becomes very small compared to the distance between them: that is, they effectively become line elements. We consider the case in which the arms consist of two parallel lines separated by a distance $2a$, with 'caps' at the ends. The two caps are identical up to rotation. Figure 9 illustrates the situation and defines the various geometric quantities involved.

The caps are parameterized by the angles $-\pi/2 \leq \theta_i \leq \pi/2$, $i \in \{1, 2\}$, and described by their 'profile' $r(\theta_i)$ relative to the centre of each cap. The equations for the caps are therefore

$$\mathbf{r}_i(\theta_i) = (x_i(\theta_i), y_i(\theta_i)) = (x_{i,0}, y_{i,0}) + r(\theta_i)(\cos(\theta_i + \alpha_i), \sin(\theta_i + \alpha_i)) \, .$$

Without loss of generality, we take $x_{1,0} = y_{1,0} = \alpha_1 = 0$. Since $\kappa = 0$ everywhere except the caps, S_{e} means that the only parts of the contour entering into the integral in equation (3.7) will be the caps. We further assume that the caps are convex, so that for two points on the same cap, $\hat{\mathbf{R}} \cdot \hat{\mathbf{n}} \leq 0$. S_{o} then means that only pairs of points on different caps interact. Since we can reasonably assume that $r \sim a$, the fact that the width of the arms is much less than the distance between them, $a \ll R(\theta_1, \theta_2)$, implies that the vector from any point on the first cap to any point on the second cap

is given to a first approximation by $\mathbf{R} = (x_{2,0} - x_{1,0}, y_{2,0} - y_{1,0})$, which reduces to $\mathbf{R}_0 = (x_{2,0}, y_{2,0})$ in our coordinate system. It does not depend on the θ_i. The distance $R = |\mathbf{R}_0|$ between points on different extremities is then also constant to a first approximation. All this information results in the following factorized expression for the new energy term:

$$(3.10) \qquad E_{\text{gap}} = \Psi_C(R) \int\limits_{-\pi/2}^{\pi/2} d\theta_1 \, \mathbf{t}_1 \, H(\hat{\mathbf{R}}_0 \cdot \hat{\mathbf{n}}_1) \; \cdot \int\limits_{-\pi/2}^{\pi/2} d\theta_2 \, \mathbf{t}_2 \, H(-\hat{\mathbf{R}}_0 \cdot \hat{\mathbf{n}}_2) \; .$$

We take the function H in equation (3.3) to be a Heaviside function ($\lim_{\rho_H \to 0}$), which together with the assumption of convexity means that the effect of the Heaviside functions in equation (3.10) is simply to impose limits on the integrations over the θ_i. For θ_i, these limits are functions of β_i only, being also functionals of r. (We define $\beta_2 = \alpha_2 - \beta_1 - \pi$. See figure 9.) For θ_1, we will name the upper limit $B(\beta_1)$ and the lower limit $A(\beta_1)$. The upper and lower limits for θ_2 are then $-A(\beta_2)$ and $-B(\beta_2)$ respectively, due to the difference in definition of β_1 and β_2. Once these limits are imposed, the integrals can be performed, because they are just integrals of the tangent vectors, which are exact. The result is that

$$E_{\text{gap}} = \Psi_C(R)[\mathbf{r}_1(B(\beta_1)) - \mathbf{r}_1(A(\beta_1))] \cdot [\mathbf{r}_2(-A(\beta_2)) - \mathbf{r}_2(-B(\beta_2))] \; .$$

A particularly simple case is when the caps are semi-circular: $r(\theta_i) = a$. In this case

$$B(\beta) = \min(\frac{\pi}{2}, \frac{\pi}{2} + \beta) \qquad \text{and} \qquad A(\beta) = \max(-\frac{\pi}{2}, -\frac{\pi}{2} + \beta) \; .$$

The resulting value of E_{gap} is

$$(3.11) \qquad E_{\text{gap}} = -a^2 \Psi_C(r)[1 + \cos(\beta_1) + \cos(\beta_2) + \cos(\beta_1 + \beta_2)] \; .$$

Note that despite the local definition of E_{gap}, the final energy depends on the overall directions of the arms. Figure 10 shows a contour plot of the angular term in the brackets. The units are multiples of π. As can be seen from the figure, for fixed r, the energy is a minimum for $\beta_1 = \beta_2 = 0$. For fixed $\beta_1 + \beta_2$, *i.e.* for fixed angle between the arms, the energy is minimum when $\beta_1 = \beta_2$, that is, when the arms are co-circular.

The radial force, $-\partial E_{\text{gap}}/\partial r$, is given by

$$(3.12) \qquad F_r = a^2 \dot{\Psi}_C(r)[1 + \cos(\beta_1) + \cos(\beta_2) + \cos(\beta_1 + \beta_2)] \; ,$$

where it should be noted that $\dot{\Psi}_C < 0$. This force is thus maximal when the two line segments are anti-parallel and opposing, *i.e.* $\beta_1 = \beta_2 = 0$. It reduces to zero when $|\beta_1 + \beta_2| = \pi$, *i.e.* when the two segments are parallel. There is thus no attraction in this case. When $|\beta_1 + \beta_2| > \pi$, *i.e.* the arms 'diverge', the force becomes repulsive. The force is zero whenever β_1 or $\beta_2 = \pi$, *i.e.* whenever one arm points radially away from the position of the other extremity. Note that for a fixed angle between the line segments ($\beta_1 + \beta_2$ fixed), the maximum force occurs when $\beta_1 = \beta_2$, that is when

Figure 10: A contour plot of the angular factor in equation (3.11).

the two line segments are co-circular, and that this force increases as the radius of the inscribed circle increases, *i.e.* as β_1 and β_2 decrease.

The angular forces, $-\partial E_{\text{gap}}/\partial\beta_1$ and $-\partial E_{\text{gap}}/\partial\beta_2$ are plotted as a vector field in the β_1, β_2 plane in figure 11. Note that the force tends to anti-align the segments ($\beta_1 = -\beta_2$), while also tending to make them oppose one another ($\beta_1 = \beta_2 = 0$).

Figure 12 shows the direction that minimizes the energy of (and maximizes the radial force on) a second segment as a function of its position, r and β_1, with respect to a first segment situated at the origin and pointing along the positive x-axis. The magnitude of the vectors is proportional to the magnitude of the radial force F_r acting at this optimal direction, but note that the arrows do not represent the force *direction*, which in this simple model is always radial.

It is interesting to compare this behaviour with that of stochastic completion fields (Williams and Jacobs, 1997), extension fields (Guy and Medioni, 1996), and other variants (Williams and Thornber, 1999). The comparison is not completely clear, since the energy E_{gap} in equation (3.11) does not govern the gap completion in the same way that completion energies and affinities do. The behaviour is similar but not identical. For example, for small β_i, equation (3.11) becomes

$$E_{\text{gap}} = a^2 \Psi_C(r)[\frac{3}{4}(\beta_1 + \beta_2)^2 + \frac{1}{4}(\beta_1 - \beta_2)^2 - 4] \, ,$$

which should be compared to the approximate expression for the scale-invariant elastica energy given by Sharon et al. (2000). Although the pertinence of perceptual completion criteria is not obvious in

Figure 11: The angular force resulting from equation (3.11).

an application to road networks, further improvements in E_{gap} may be possible using such criteria as a guide.

If we consider the scaling with arm width a of E_{gap} and F_r in equations (3.11) and (3.12), and assume that typical interruption lengths do not depend on a (although arguably such interruptions will be longer for narrower roads, since they are easier to occlude), so that $\Psi_C(r)$ typically assumes the same value for arms of different widths, it appears that narrow arms will attract more weakly than wider arms. To a large extent, however, this is a result of the form of the function H that we have assumed, which is scale-invariant. In practice the value of S_{e} depends on the curvature of the extremities, and for a range of curvature values this dependence is approximately linear. For $\rho_H = 1$, the linear region is approximately $1/2 < \kappa < 3/2$, which translates into $4/3 < w < 4$, where $w = 2a$ is the width of the road. For the image resolution we are considering here, this covers the majority of road widths. The value of $S_{\mathrm{e}}(p, p')$ will thus scale as $\simeq a^{-2}$, thereby removing the strong dependence on arm width.

We can also consider the scaling of E_{gap} with image resolution. Taking into account the behaviour of S_{e} discussed above, we see, by re-expressing all lengths in physical units rather than pixels, that E_{gap} is invariant to changes in image resolution provided we keep ρ_A, and ρ_H as it appears in S_{e} (but not S_{o}) constant in physical units ($\rho_A = b^{-1}\tilde{\rho_A}$ and $\rho_H = b\tilde{\rho_H}$, where b is the image resolution in metres per pixel, and $\tilde{\rho_A}$ and $\tilde{\rho_H}$ are constant), and make the maximum value of S_{e} proportional to b.

Figure 12: The lowest energy direction of a second segment as a function of position (r and β_1) with respect to a first segment situated at the origin and pointing along the positive x-axis. The length of the arrows is proportional to the magnitude of the energy associated with the optimal direction (the longer the arrow, the lower the energy), or equivalently to the magnitude of radial force at that point (but note that the arrows do not represent the force *direction*, which in this simple model is always radial).

4   Energy minimization

We now move on to the numerical implementation of the gradient descent equation (3.8). HOAC energies lead to non-local forces in the evolution equation given by integrals over the contour, as can be seen in equation (3.9). These integrals require specific treatment in the contour evolution. We do not discuss this further here, but refer the reader to Rochery et al. (2005b).

In addition, the energy E_{gap} introduces its own complications. Note that amongst the terms in equation (3.9) are some which contain first and second derivatives of the contour curvature, which translate into third and fourth derivatives of the level set function, and first, second, and third derivatives of the smoothed Heaviside function H. These terms, which either have high-order or large derivatives, can cause instabilities numerically. We adopt specific measures, described below, to ameliorate the difficulties that these terms could cause.

Figure 13: Contour with a gap

4.1 Gaussian smoothing and derivative computations

Before computing geometric quantities such curvature and its derivatives that are part of the functional derivative of E_{gap} and thus present in the gradient descent equation, we apply Gaussian smoothing to the level set function ϕ. We use an isotropic Gaussian kernel with $\sigma = 1$. This smoothing corresponds to using a larger stencil for the computation of these geometric quantities, and produces smoother results. In particular, the contour curvature is smoothed, meaning that extremities are more clearly recognized. For example, consider the contour depicted in figure 13.

In figure 14, curvatures computed before and after smoothing the level set function are depicted. The curvature computed directly on the level set function is very irregular. In contrast, the curvature computed after smoothing the level set function is smoother and reveals two clear peaks corresponding to the two extremities in figure 13. This coincides well with the purpose of E_{gap}, which is to identify extremities that should be connected.

In addition to the above, we use finite differences of order four to calculate numerical derivatives. This provides a reasonable degree of smoothness in the curvature derivatives, which are otherwise very noisy even after Gaussian smoothing.

4.2 Computation of contour integrals

In order to compute the non-local terms of the force, we approximate the integrals over the contour by sums over the extracted contour segments (Rochery et al., 2005b). When the contour is closed, we can improve the precision of this approximation and obtain accuracy of order four if we use equispaced points, as follows. Consider a function f that we want to integrate over the interval $[a, b]$. If we construct approximations $\hat{f}_i$ to f over a partition of $[a, b]$ into intervals $[x_i, x_{i+1}]$, $i \in \{0, \dots, n-1\}$, then the integral can be approximated by

$$\int_a^b dx\, f(x) \simeq I = \sum_{i=0}^{n-1} I_i = \sum_{i=0}^{n-1} \int_{x_i}^{x_{i+1}} dx\, \hat{f}_i(x) \,.$$

Using linear approximations on each interval leads to

$$I = h \left(\tfrac{1}{2}(f(a) + f(b)) + \sum_{i=1}^{n-1} f_i \right) \,.$$

Figure 14: Curvatures computed before (top) and after (bottom) smoothing.

where for any function g, $g_i = g(x_i)$, $\bar{g}_i = \frac{1}{2}(g_i + g_{i+1})$ and $\delta g_i = g_{i+1} - g_i$, and we have taken $\delta x_i = h$, constant.

On the other hand, if we approximate f on each interval by a third-order polynomial,

$$\hat{f}_i(x) = a_{i,3}(x - x_i)^3 + a_{i,2}(x - x_i)^2 + a_{i,1}(x - x_i) + a_{i,0} ,$$

and require that

$$\hat{f}_i(x_i) = f_i , \qquad \hat{f}_i(x_{i+1}) = f_{i+1} , \qquad \dot{\hat{f}}_i(x_i) = \dot{f}_i , \qquad \dot{\hat{f}}_{i+1}(x_{i+1}) = \dot{f}_{i+1} ,$$

where as before the dot indicates derivative, we find, on substituting the resulting values for the a_i, that

$$I_i = \int\limits_0^h dx \, \hat{f}_i(x + x_i) = a_{i,3}\frac{h^4}{4} + a_{i,2}\frac{h^3}{3} + a_{i,1}\frac{h^2}{2} + a_{i,0}h = h\bar{f}_i - \frac{h^2}{12}\delta\dot{f}_i ,$$

and hence that

$$I = h \left(\frac{1}{2}(f(a) + f(b)) + \sum_{i=1}^{n-1} f_i \right) - \frac{h^2}{12}(\dot{f}_n - \dot{f}_0) \,.$$

In the case of a function defined on a closed contour, we have $\dot{f}_0 = \dot{f}_n$, whence all the derivatives disappear and the fourth order accurate approximation is the same as the second order accurate approximation. Thus, using equispaced points allows us to have fourth order accuracy at no extra cost. Therefore, in order to compute the non-local force at each point p of the extracted contour, we first redistribute the extracted contour points around p so that they are equispaced (details of the redistribution procedure are given in); we then compute the necessary geometric quantities at these points, and finally perform the numerical integration. The redistribution of the points in itself may introduce errors, so that the above analysis is suggestive rather than directly applicable; nevertheless, empirically we find that this procedure improves the stability of the evolution.

5   Experimental results

In order to give an understanding of the behaviour of the new energy term E_{gap}, we first show the result of an experiment using only the prior energies E_{g} and E_{gap}. We then show the results of road network extraction on real images using the new model $E = E_0 + E_{\text{gap}}$, illustrating the performance of the new model in the presence of occlusions. Parameter values in the experiments are shown in the corresponding figure captions. In common with most variational methods, these parameters are fixed by hand.

5.1   Prior energy

Figure 15 shows the results of contour evolution using E_{g} and $E_{\text{g}} + E_{\text{gap}}$, given an initial condition containing a gap, shown on the left. The top row is the evolution given by E_{g} alone. The two arms repel one another due to the antiparallel tangent vectors at the extremities, and develop into a network with two connected components. This effect is exactly what the new energy term E_{gap} is designed to avoid. The second row shows the evolution using the full prior energy $E_{\text{g}} + E_{\text{gap}}$. Now the arms extend towards one another and join, resulting in a connected network with the gap closed.

5.2   Extraction of road networks from real images

Figure 16 shows a real image in which a tree and its shadow obscure a road. The lower image in the first column shows the result of applying the previous model E_0 to this image. There is a gap in the result for all the reasons discussed at the beginning of section 3.

Moving to the second column, the upper image shows the previous solution, complete with gap. We then apply the full model, $E = E_0 + E_{\text{gap}}$, using this as an initial condition. The lower image in the second column shows the result. As hoped, the gap is closed.

Figure 15: Evolutions based on prior energies only (time runs from left to right). Top row: E_{g} only, without gap closure energy E_{gap}. Bottom row: $E_{\mathrm{g}} + E_{\mathrm{gap}}$. Note the closure of the gap. Parameter values: top, $[\lambda = 1, \alpha = -0.1, \beta = 0.4]$; bottom, $[\lambda = 1, \alpha = -0.1, \beta = 0.4, \beta_A = 0.4]$.

Figure 16: First column: aerial image (© IGN) and the result of extraction using E_0. Second column: initial contour and final closed contour using the new model $E = E_0 + E_{\mathrm{gap}}$. Third column: generic initialization and result of extraction with the new model $E = E_0 + E_{\mathrm{gap}}$. Parameter values: left, $[\lambda = 1, \alpha = -0.14, \beta = 0.03, \lambda_i = 2.9, \alpha_i = 0.3, \beta_i = 2.9]$; middle and right, $[\lambda = 1, \alpha = -0.14, \beta = 0.06, \lambda_i = 2.9, \alpha_i = 0.3, \beta_i = 2.9, \beta_A = 6]$.

In the third column, the lower image shows the result of applying the full model $E = E_0 + E_{\mathrm{gap}}$ starting from the generic initialization shown in the upper image. Again the gap is closed.

Figure 17: Top: force resulting from E_{gap} for the contour shown in the leftmost image of figure 16. Bottom: graph of the function given in equation (5.1), showing identification of opposing extremities.

To illustrate what is happening, consider the contour shown in the leftmost image of the second row of figure 16. For this contour, the top row of figure 17 shows the value of the force resulting from the new term E_{gap}. Positive values indicate forces along the outward pointing normal. Clearly the forces generated by the new term are pulling both extremities outwards, and hence towards one another. The rest of the contour is unaffected.

The second row of figure 17 shows the graph of the following quantity:

$$(5.1) \qquad\qquad S(p) = \int dp' \, S_{\text{o}}(p, p') \, S_{\text{e}}(p, p') \ .$$

It is clear that the extremities are very well identified; $S(p)$ is zero expect for points that lie on nearby opposing extremities.

Figure 18: An aerial image (© IGN) with occlusions and the result of applying $E = E_0 + E_{\text{gap}}$ to this image. Parameter values: $[\lambda = 1, \alpha = -0.14, \beta = 0.06, \lambda_i = 2.9, \alpha_i = 0.3, \beta_i = 2.9, \beta_A = 6]$.


Figure 19: An aerial image (© IGN) and the result of applying $E = E_0 + E_{\text{gap}}$ to this image. Parameter values: $[\lambda = 1, \alpha = -0.14, \beta = 0.05, \lambda_i = 2.9, \alpha_i = 0.3, \beta_i = 2.9, \beta_A = 10]$.

Figure 18 shows another result of extraction using $E = E_0 + E_{\text{gap}}$. Despite the trees obscuring the network, the road is perfectly reconstructed with the new model. Note that the initial contour is again a rounded rectangle covering the image. Despite this generic initialization a long way from the solution, the model does not become stuck in local minima, for example, those due to the borders of fields.

Figure 19 shows the result obtained on a larger image. The extracted network is again connected, the model closing the gaps caused by occlusions.

6   Summary

When attempting to extract line networks from images, and in particular road networks from remote sensing images, one of the key difficulties is created by the presence of interruptions in the imaged

network due to occlusions, cast shadows, and other effects. Such interruptions can lead to gaps in the extracted network that do not correspond to gaps in the real network. In this report, we have described a solution to this problem within the framework of higher-order active contours (Rochery et al., 2005b). A previously proposed higher-order active contour model for the extraction of line networks (Rochery et al., 2003) was in general successful, but it was unable to surmount the problem of interruptions. Building on this model, we have defined a quadratic gap closure energy that penalizes network configurations containing nearby opposing extremities. In the gradient descent algorithm used to minimize the energy, the effect of this new energy is to cause such extremities to attract one another, to move together, and to join, thereby closing the gap. We note that the new energy is inherently higher-order: it involves the long-range interaction of two different extremities, *i.e.* of widely separated points on the contour; it thereby demonstrates the ability of higher-order active contours to include sophisticated prior morphological knowledge.

Gradient descent using this new energy is a delicate matter due to the presence of numerous force terms containing higher derivatives; these require special attention if instabilities are to be avoided. Working within the level set framework, we have developed techniques to ameliorate the numerical difficulties caused by these force terms. Experiments on real remote sensing images demonstrate that, with the exception of very long interruptions, the new energy succeeds in overcoming interruptions in imaged networks to produce networks without gaps.

Although we have focused on the extraction of road networks from remote sensing images, as emphasized in section 1, diverse line networks in different imagery types have much in common. The prior knowledge captured by the model described in this report is thus also relevant to other line network extraction problems, for example, the extraction of hydrographic networks from remote sensing images, or vascular and other networks from medical and biological images.

References

M. Bicego, S. Dalfini, G. Vernazza, and V. Murino. Automatic road extraction from aerial images by probabilistic contour tracking. In *Proc. IEEE International Conference on Image Processing (ICIP)*, volume 3, pages 585–588, Barcelona, Spain, September 2003.

Y. Chen, S. Thiruvenkadam, H.D. Tagare, F. Huang, D. Wilson, and E.A. Geiser. On the incorporation of shape priors into geometric active contours. *Proc. IEEE Workshop Variational, Geometric and Level Set Methods in Computer Vision*, pages 145–152, 2001.

D. Cremers, T. Kohlberger, and C. Schnörr. Shape statistics in kernel space for variational image segmentation. *Pattern Recognition*, 36(9):1929–1943, September 2003.

A. Foulonneau, P. Charbonnier, and F. Heitz. Geometric shape priors for region-based active contours. *Proc. IEEE International Conference on Image Processing (ICIP)*, 3:413– 416, 2003.

P. Fua and Y. G. Leclerc. Model driven edge detection. *Machine Vision and Applications*, 3:45–56, 1990.

D. Geman and B. Jedynak. An active testing model for tracking roads in satellite images. 18:1–14, 1996.

G. Guy and G. Medioni. Inferring global perceptual contours from local features. *International Journal of Computer Vision*, 20(1-2):113–133, October 1996.

C. Lacoste, X. Descombes, and J. Zerubia. Road network extraction in remote sensing by a Markov object process. In *Proc. IEEE International Conference on Image Processing (ICIP)*, volume 3, pages 1017–1020, Barcelona, Spain, September 2003.

I. Laptev, T. Lindeberg, W. Eckstein, C. Steger, and A. Baumgartner. Automatic extraction of roads from aerial images based on scale space and snakes. *Machine Vision and Applications*, 12:23–31, 2000.

M.E. Leventon, W.E.L. Grimson, and O. Faugeras. Statistical shape influence in geodesic active contours. *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, 1:316–322, 2000.

N. Merlet and J. Zerubia. New prospects in line detection by dynamic programming. 18(4):426–431, 1996.

W. M. Neuenschwander, P. Fua, L. Iverson, G. Székely, and O. Kubler. Ziplock snakes. *International Journal of Computer Vision*, 25(3):191–201, 1997.

N. Paragios and M. Rousson. Shape priors for level set representations. *Proc. European Conference on Computer Vision (ECCV)*, pages 78–92, 2002.

M. Rochery, I. H. Jermyn, and J. Zerubia. Higher order active contours. *International Journal of Computer Vision*. To appear.

M. Rochery, I. H. Jermyn, and J. Zerubia. Higher order active contours and their application to the detection of line networks in satellite imagery. In *Proc. IEEE Workshop Variational, Geometric and Level Set Methods in Computer Vision*, at ICCV, Nice, France, October 2003.

M. Rochery, I. H. Jermyn, and J. Zerubia. Gap closure in (road) networks using higher-order active contours. In *Proc. IEEE International Conference on Image Processing (ICIP)*, Singapore, October 2004.

M. Rochery, I. H. Jermyn, and J. Zerubia. New higher-order active contour energies for network extraction. In *Proc. IEEE International Conference on Image Processing (ICIP)*, Genoa, Italy, September 2005a.

M. Rochery, I. H. Jermyn, and J. Zerubia. Higher order active contours. Research Report 5656, INRIA, France, August 2005b.

E. Sharon, A. Brandt, and R. Basri. Completion energies and scale. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(10):1117–1131, 2000.

R. Stoica, X. Descombes, and J. Zerubia. A Gibbs point process for road extraction from remotely sensed images. *International Journal of Computer Vision*, 57(2):121–136, 2004.

F. Tupin, H. Maitre, J-F. Mangin, J-M. Nicolas, and E. Pechersky. Detection of linear features in SAR images: Application to road network extraction. *IEEE Trans. Geoscience and Remote Sensing*, 36(2):434–453, 1998.

L. R. Williams and D. W. Jacobs. Stochastic completion fields: A neural model of contour shape and salience. *Neural Computation*, 9:849–870, 1997.

L. R. Williams and K. K. Thornber. A comparison of measures for detecting natural shapes in cluttered backgrounds. *International Journal of Computer Vision*, 34(2-3):81–96, August 1999.

C. Zhang, S. Murai, and E. Baltsavias. Road network detection by mathematical morphology. In *Proceedings of ISPRS Workshop "3D Geospatial Data Production: Meeting Application Requirements"*, pages 185–200, Paris, France, April 1999.