Phase Field Models and Higher-Order Active Contours

Marie Rochery   Ian Jermyn   Josiane Zerubia [*]
Ariana (INRIA/I3S), Sophia Antipolis, France
FirstName.LastName@inria.fr

Abstract

The representation and modelling of regions is an important topic in computer vision. In this paper, we represent a region via a level set of a 'phase field' function. The function is not constrained, e.g. *to be a distance function; nevertheless, phase field energies equivalent to classical active contour energies can be defined. They represent an advantageous alternative to other methods: a linear representation space; ease of implementation (a PDE with no reinitialization); neutral initialization; greater topological freedom. We extend the basic phase field model with terms that reproduce 'higher-order active contour' energies, a powerful way of including prior geometric knowledge in the active contour framework via nonlocal interactions between contour points. In addition to the above advantages, the phase field greatly simplifies the analysis and implementation of the higher-order terms. We define a phase field model that favours regions composed of thin arms meeting at junctions, combine this with image terms, and apply the model to the extraction of line networks from remote sensing images.*

1. Introduction

The aim of computer vision algorithms is to make statements about the scenes of which images are representations. Many such statements concern the localization of entities in the image (and thereby, implicitly, in the scene). Entities are thus associated with regions in the image domain, and the question arises of how best to represent these regions, both in models and in algorithms, and how to model the prior knowledge we have about the regions of interest.

Many different answers to both these questions have been developed in the literature. Active contours, beginning with Kass *et al.* [6], use energy functionals defined on the space of region boundaries. Early algorithms performed gradient descent on a discretized boundary, but this contour

representation creates a number of difficulties, the most important being that topology changes are hard to handle. This limits the regions that can be reached starting from a given initialization, and thus increases the dependence of the result on the starting point.

Many of the difficulties with discretized boundaries are overcome by using level sets [9]. During gradient descent, the region is represented by a function on the image domain Ω whose zero level set is the region boundary. The function is usually taken to be the signed distance to the boundary, a constraint that is enforced by reinitializing the function during gradient descent. As a consequence, level sets permit changes in the number of connected components of the region only via 'splitting' and 'merging', and changes in the number of handles only via 'wrapping' and 'separating'. Energies may be formulated in terms of the boundary itself, and the descent equations converted to a level set representation, or they may be formulated directly in terms of the level set function. The latter, however, results in singularities in the descent equations, which must then be regularized in an *ad hoc* way to permit computation.

If the image domain is treated as discrete, these singularities are regularized by the discretization; the characteristic function of the discrete region may then be used directly. The best known model in this class is the Ising model. Such models are usually solved using stochastic optimization.

Previous models have incorporated varying degrees of prior geometric knowledge. The standard prior terms to be found in active contour models (and, *e.g.* in region competition [18]) are boundary length and interior area. The Ising model also uses a discrete version of boundary length. This is not surprising: Euclidean invariance is necessary for most problems, and length and area are the only two invariants if higher derivatives are not used. Some models have also used the boundary integral of the magnitude of the boundary curvature or its square. Recently there has been a great deal of interest in the incorporation of more specific prior knowledge. This normally (although see [3]) takes the form of a template shape around which small variations controlled by a Gaussian energy are permitted (see [8] and references therein). Because the template has a particular

position and orientation, repeated pose estimations are required for Euclidean invariance. This means that detecting multiple instances of an entity, whose number must also be known *a priori* [4], is computationally expensive.

An alternative approach to the incorporation of prior geometric (and image) knowledge in active contour models, known as 'higher-order active contours' (HOACs), was introduced by Rochery *et al.* in [11]. The standard length, area, and curvature terms incorporate only local, differential knowledge about the boundary, whereas HOACs include nonlocal interactions between widely separated contour points. The nonlocal terms focus the energy minima on regions with particular geometric characteristics, *e.g.* regions composed of arms with roughly parallel sides that meet at junctions. In [11], such a model was applied to remote sensing images for the extraction of road networks, for which it clearly provides a better prior model than the standard terms alone. HOACs have a number of advantages over other methods for incorporating prior geometric knowledge. First, they are not limited to small variations around a template shape, as the road network example makes clear. Indeed, they provide a general framework for the construction of models of increasing complexity and specificity. Second, they require no pose estimation as the energy is intrinsically invariant. Multiple instances of an entity, whose number need not be known *a priori*, therefore cost no more to detect than a single instance.

There are difficulties, however. First, as with other active contour approaches, an initial region is needed for the gradient descent. Although the inclusion of more specific prior knowledge means that a generic, hence automatic initialization can be used, the final result still depends on this choice. Second, a level set representation is used to perform the gradient descent, and this prevents the formation of handles. Road networks, for instance, can have a complicated topology with many handles, and the choice of initialization affects which of these are detected. Finally, the nonlocal terms produce nonlocal forces, whose evaluation requires boundary extraction and integration followed by velocity extension, and this is expensive computationally.

The primary goal of the present paper is to describe how HOACs can be expressed in terms of a fourth region representation and modelling framework: phase fields. Phase field models originated in physics, where they have been much used to model regions and interfaces. They offer the following important advantages:

- the representation space is linear, thus facilitating model building and analysis;

- gradient descent is based solely on the PDE resulting from an energy functional, with no need for *ad hoc* regularization or reinitialization; numerical implementation is simple;

- no initial contour is needed for evolutions; a completely neutral initialization is possible;

- components of a region can be created or destroyed anywhere in Ω; handles can be created and destroyed in the interior of existing regions;

- HOAC terms require no boundary extraction, integration, or velocity extension; indeed prior terms are local in the Fourier domain.

The difficulties mentioned above are thereby overcome, while the ability to include sophisticated prior geometric knowledge is retained. This will be illustrated by applying a particular model to the extraction of road networks from remote sensing images. The first four points apply equally to standard active contour models, which are a special case of HOACs, and a subsidiary goal of this paper is to suggest that phase fields represent an advantageous alternative to other active contour methods in this case also.

In the rest of the paper, we expand on the above. In section 2, we describe the simplest phase field model and its approximate equivalence to classical active contour energies. This equivalence is quite well-known, but does not appear to have been exploited in computer vision. We provide a heuristic but simple derivation. We also illustrate the above advantages as applied to standard active contours. Section 3 is the heart of the paper. We introduce extended phase field models, and show their approximate equivalence to HOACs, as well as their algorithmic simplicity. We illustrate the prior geometric knowledge included in the extended models via gradient descent evolutions using a specific example, and discuss the relation between the descent equations and reaction-diffusion equations in the literature. In section 4, we move from prior models to the construction of image terms for line networks in remote sensing images, and show some examples of network extraction from real images. We summarize and conclude in section 5.

2. Phase Field Models

Phase fields are a level set framework: the phase field ϕ is a function on Ω, which, given a threshold z, defines a region in the space of regions $\mathcal{R}$ via the map $\zeta_z(\phi) \in \mathcal{R} = \{x \in \Omega : \phi(x) > z\}$. However, in contrast to other methods, which use a representation space that is isomorphic to the space of regions $\mathcal{R}$ (*e.g.* distance functions, contours, characteristic functions), no constraint is placed on ϕ: the set of functions considered, Φ, is a linear space. The nonlinear nature of $\mathcal{R}$ is instead taken into account via nonlinear terms in the energy functional defining the model. The sim-

plest phase field energy is

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

When $\alpha = 0$, this is commonly known as the Ginzburg-Landau model. For $\lambda \geq |\alpha|$, the potential on the second line has two minima, at ± 1, and a maximum at α/λ. If $\alpha > 0$, the minimum at -1 has the lowest energy, and is thus preferred. The energy (1) can be minimized subject to the constraint that $\zeta_z(\phi) = R$, which defines a function $\phi_0 : \mathcal{R} \to \Phi$. The form of the potential means that away from the region boundary, ∂R, $\phi_R = \phi_0(R)$ will assume the value 1 inside, and -1 outside R. In the absence of the derivative term, ϕ_R would have a discontinuity on ∂R, where it would jump from -1 to 1. The presence of the derivative term prevents this, however, and the competition between the two parts of the energy results in a smooth transition over an interface region around the boundary. Up to a rescaling and shift, ϕ_R is thus a smoothed version of the characteristic function of R. The function ϕ_0 defines a region energy $E_{C,0} = E_0\phi_0$, which can be described approximately in terms of the boundary length $L(R)$ and interior area $A(R)$ of R. Gradient descent in Φ using E_0 will follow the 'valley' corresponding to the functions $\phi_0(R)$ for a sequence of regions corresponding to gradient descent in $\mathcal{R}$ using $E_{C,0}$. The phase field and the energy E_0 can thus be used to represent and model regions in place of a contour or distance function and $E_{C,0}$. This connection[1] does not appear to have been exploited in computer vision. However, the energy E_0 with ϕ complex and $\alpha = 0$ was used for image inpainting [5], and to detect codimension 2 objects in the image domain [2], while in [14], a model with multiple minima was used with Gamma convergence to construct piecewise constant approximations to an image.

We can 'derive' the form of $E_{C,0}$ heuristically using a simple ansatz $\tilde{\phi}_R$ for ϕ_R. We divide Ω into three pieces: a narrow interface region R_C containing ∂R (the width w of R_C can vary with position on ∂R); an interior region $R_+ = R\backslash R_C$; and an exterior region $R_- = \bar{R}\backslash R_C$, where $\bar{R} = \Omega\backslash R$. The situation is illustrated in figure 1. We take $\tilde{\phi}_R = \pm 1$ on $R_\pm$, while on R_C it changes linearly from -1 to 1 along the normal to ∂R. For a segment of ∂R small enough that the curvature can be considered constant, in polar coordinates based on the centre of curvature, ϕ in R_C is thus

$$\tilde{\phi}_R(r, \theta) = -\frac{2}{w}(r - r_0) + z \quad (2)$$

---

[1]It is interesting to note that when (1) is discretized on a lattice using standard finite differences, and $\lambda \to \infty$, the Ising model is recovered.
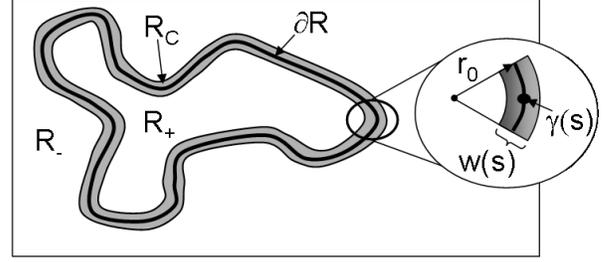

Figure 1. Ansatz used to derive $E_{C,0}$.

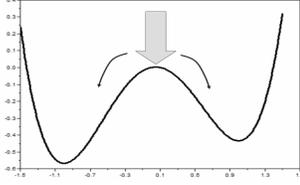Substituting this form into E_0, one finds, up to an additive constant, that

$$E_0(\tilde{\phi}_R) = \int_0^{L(R)} ds \left\{ \frac{2D}{w} + w \left(\frac{2\lambda K}{15} + \frac{3\alpha}{40}w\kappa \right) \right\}$$
$$+ \frac{4\alpha}{3}A(R) , \quad (3)$$

where s is an arc length parameterization of ∂R, $K = 1 + 5(\alpha/\lambda)^2$, and κ is the boundary curvature. Minimizing E_0 with respect to w gives the optimal width w_* as a function of the parameters. If $w\kappa \ll 1$, w_* is constant: $w_*^2 = 15D(\lambda K)^{-1}$. Thus we find that

$$E_{C,0}(R) \cong \tilde{E}_{C,0}(R) = \lambda_C L(R) + \alpha_C A(R) , \quad (4)$$

where $\alpha_C = 4\alpha/3$ and $\lambda_C^2 = 16D\lambda K/15$. These values are close to those obtained by more sophisticated methods.

Of course there are curvature dependent corrections to $E_{C,0}$ stemming from the $w\kappa$ term in (3), meaning that E_0 does not mimic exactly the behaviour of $\tilde{E}_{C,0}$. To achieve this, the interface width must be as small as possible, which requires very fine and/or adaptive discretization, and this is computationally expensive. This dilemma constitutes a commonly voiced objection to phase field methods. The objection has force, however, only when it is known that $\tilde{E}_{C,0}$ is the appropriate region energy, i.e. that the correction terms should be very small at the scale of interest. In physical systems, this may arise from more fundamental principles. In computer vision, however, region energies are not derived from fundamental principles. Rather, they are designed to represent our prior knowledge about region geometry. If the maximum curvature expected of the regions of interest is $\kappa_{\max}$, then w must be $\lesssim \kappa_{\max}^{-1}$ to assign them significant probability. Then if $\kappa_{\max} \gtrsim 1$ in pixel units, w must be $\lesssim 1$. The need to reduce computation may then outweigh the advantages of phase field methods. If this is not the case, however, then there is no good reason to assert that one must use $\tilde{E}_{C,0}$ rather than $E_{C,0}$, especially since the lowest-order term in the difference is, anyway, of

Figure 2. Neutral initialization

Figure 3. Top: evolution starting from the neutral initialization. Bottom: evolution starting from a circle.

the κ^2 form that has been used in computer vision. In what follows, including the HOAC case, we will therefore simply define the phase field energy to be our region model, thereby implicitly defining an energy on regions.

The functional derivative of E_0 is given by

$$\frac{\delta E_0}{\delta \phi} = -D\nabla^2 \phi + \lambda(\phi^3 - \phi) + \alpha(1 - \phi^2) , \quad (5)$$

and gradient descent thus follows the Allen-Cahn equation [1]. Provided the parameters are such that the interface width is not small compared to the discretization, this can be implemented using simple finite differences. No reinitialization or *ad hoc* regularization are required.

Phase fields have a major advantage in terms of initialization. We can initialize gradient descent using the constant function $\phi_{\text{init}} \equiv \alpha/\lambda$. If we then choose the map $\zeta_{\alpha/\lambda}$ to map functions to regions, this is a completely neutral initialization, in two senses. First, it corresponds to the maximum of the potential, and hence is not biased towards one minimum or the other. Figure 2 illustrates this. Second, ϕ_{init} is not biased towards interior or exterior, in the sense that both $\{x \in \Omega : \phi_{\text{init}} \gtrless \alpha/\lambda\}$ are empty.

Such an initialization means that ϕ can produce regions 'where required' in the image domain at the beginning of the evolution, rather than having to evolve towards them from some initial configuration. A simple example is shown in the top row of figure 3. The underlying data is a white annulus on a black background. The prior energy is E_0 with $\alpha = 0$. The image energy is the first term of (10). The figure shows the evolution of the region $\zeta(\phi)$ starting from the neutral initialization (all-white image at left). Note the formation of regions near the high gradient zones followed by their extension to the full annulus. This behaviour is particularly important when multiple instances of an entity must be found. In addition, it seems intuitively clear that there will be less problems with local minima starting from this initialization, although making this precise is hard.

As advertised, the new framework also has much greater topological freedom. In the standard level set framework, the number of connected components of a region can only change by splitting and merging. Now, as figure 3 makes clear, new components can appear or disappear where needed. In addition, in the standard level set framework,
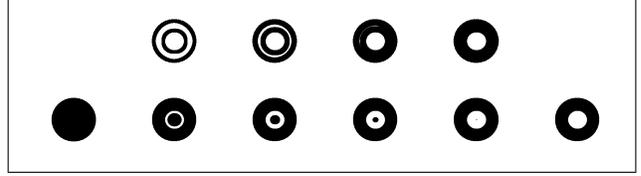
the number of handles of a region can only change by wrapping or separating. Phase fields again show more topological freedom: handles can appear and disappear within an existing region. This is illustrated in the bottom row of figure 3. The data and energy are the same as for the top row, but now the initialization is a circle. Note the formation of the hole in the centre.

The greater topological freedom offered by phase field models means that the evolving region can reach more of $\mathcal{R}$ starting from a given initialization than other frameworks, which renders the method more independent of the initialization. These freedoms are critical for applications dealing with topologically non-trivial regions, *e.g.* road network extraction. If topological restrictions are required, then they must be imposed at the level of the model, as terms in the energy, which is where they should appear anyway, rather than appearing as side-effects of the algorithm.

We conclude that phase fields offer a distinctly advantageous alternative to standard active contours, whether the latter are used with contour or level set representations. All these advantages, and more, apply equally to the phase field models corresponding to HOACs, which we now discuss.

3. HOACs as Phase Field Models.

In the previous section, we described the basic phase field model corresponding to standard active contour prior energies. In this section, we show how to extend E_0 to $E = E_0 + E_{NL}$ such that $E_C = E\phi_0$ is a HOAC energy. We can thus use phase field models to incorporate sophisticated prior geometric knowledge about the regions being modelled, thus accruing all the advantages mentioned previously. In addition, phase field models greatly simplify the analysis and algorithmic treatment of HOAC terms. We begin by briefly discussing HOACs [11].

Standard active contour energies are defined by single integrals of local, differential geometric properties of the region boundary. They thus incorporate prior knowledge only about these local quantities. HOACs, in contrast, incorporate long-range interactions between boundary points,

and this enables the inclusion of complex prior geometric knowledge. The long-range interactions are expressed via multiple integrals over the boundary. For instance, the simplest form involves two integrals:

$$E_Q(R) =$$
$$- \frac{\beta_C}{2} \int_{S^1 \times S^1} dt \, dt' \, \dot{\gamma}(t) \cdot \mathbf{G}_C(\gamma(t), \gamma(t')) \cdot \dot{\gamma}(t') \, , \quad (6)$$

where $\gamma : S^1 \to \Omega$ is a parameterization of ∂R (for notational simplicity, we consider only single connected component, simply connected regions), a dot indicates d/dt, and $\mathbf{G}_C$ is a map from Ω^2 to 2×2 matrices.

Now define the phase field energy

$$E_{NL}(\phi) =$$
$$- \frac{\beta}{2} \int_{\Omega^2} d^2x \, d^2x' \, \nabla\phi(x) \cdot \mathbf{G}(x, x') \cdot \nabla\phi(x') \, , \quad (7)$$

where $\mathbf{G}$ is the same type of operator as $\mathbf{G}_C$. Consider inserting the ansatz (2) into this energy. The regions R_+ and R_- contribute nothing, so all the contribution comes from the double integral over R_C. The gradient of ϕ in R_C is given by $\nabla\phi = -(2/w)\hat{n}$, where $\hat{n}$ is the outward normal to the boundary, extended to R_C. The integrals across the interface in (7) are thus easy to perform if we assume that $\mathbf{G}$ is roughly constant over distances $\sim w$. Each integral simply contributes a factor w, which cancels that in the corresponding gradient. We are left with two boundary integrals in the form (6), with $\beta_C = 4\beta$, and $\mathbf{G}_C = \epsilon^\dagger \mathbf{G} \epsilon$, where ϵ is the 2×2 Levi-Civita tensor and $\dagger$ is transposition. (This conclusion is confirmed by a more sophisticated analysis [7].) The phase field model $E = E_0 + E_{NL}$ is thus approximately equivalent to the HOAC model $E_C = \tilde{E}_{C,0} + E_Q$, and can be used in its place, thus profiting from all the advantages of the phase field framework.

The functional derivative of E_{NL} is given by

$$\frac{\delta E_{NL}}{\delta\phi(x)} = \beta \int_\Omega d^2x' \, \nabla \cdot \mathbf{G}(x, x') \cdot \nabla\phi(x') \, . \quad (8)$$

Note that this is a nonlocal term. However, for prior terms, E_{NL} will normally be Euclidean invariant, which means that $\mathbf{G}$ is rotationally and translationally invariant:

$$\mathbf{G}(x, x') = G(x - x')\delta = \Psi(|x - x'|)\delta \, , \quad (9)$$

where δ is the 2×2 unit matrix. Translation invariance means that $\mathbf{G}$ is diagonalized by the Fourier basis; its action on ϕ can thus be computed trivially in the Fourier domain. This shows the particular advantage of phase field methods for HOACs. Whereas the standard level set method requires boundary extraction, integration, and velocity extension at

each iteration to compute the nonlocal force, here it is simple to compute. Implementation is thus much easier and execution faster.

To illustrate the prior knowledge contained in (7), we will show the results of gradient descent using the energy E, and compare with the results of gradient descent using E_C. We choose Ψ to be a smoothed step function:

$$\Psi(z) = \begin{cases} \frac{1}{2}\left(1 - \frac{z-d}{\epsilon} - \frac{1}{\pi}\sin(\frac{\pi(z-d)}{\epsilon})\right) & \text{if } |z - d| < \epsilon \, , \\ 1 - H(z - d) & \text{else} \, . \end{cases}$$

where H is the Heaviside function. With this choice of interaction, $E = E_0 + E_{NL}$, favours regions composed of arms with roughly parallel sides that meet at junctions. In figure 4, the results of gradient descent starting from a circle are shown. The upper evolution is based on E, the lower on E_C, with parameters chosen to be equivalent according to the relations given earlier. As can be seen the evolutions are rather similar, with eight arms with a width of ~ 5 pixels being formed from the circle. The differences are also notable. The phase field evolution preserves the symmetry of the initial condition, whereas the contour evolution breaks it. The symmetry breaking is due to the contour extraction and integration steps in the latter, which introduce small numerical errors that are then amplified by the energy E_C. While this has little effect in computer vision applications, where the effect of the image terms dominates, the experiments nevertheless show that the phase field implementation is numerically more stable.
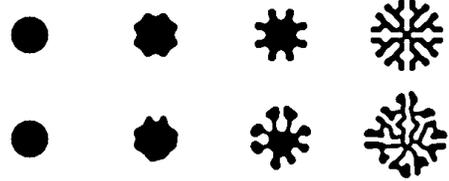
Figure 4. Geometrical evolutions based on E (top) and E_C (bottom).

The energy E clearly provides a better model for the geometry of road networks than E_0 alone. We will thus use it as the prior energy in a model for the extraction of road networks in remote sensing images.

3.1. Relation to coupled reaction-diffusion

The gradient descent equation derived from E is a nonlocal reaction-diffusion equation. For very special forms of $\mathbf{G}$, E_{NL} can be rewritten in terms of local operators using an auxiliary field. Gradient descent using E is then equivalent to two coupled reaction-diffusion equations in

the 'slaving limit', *i.e.* when the auxiliary field reacts instantaneously to changes in ϕ. Coupled reaction-diffusion equations have been used in graphics for texture synthesis [15, 16] and in computer vision for image restoration [10]. With appropriate initial conditions, the models used here can produce patterns similar to those generated in these works, but there are several important differences. First, the equations therein were far from the slaving limit. Second, the equations here are based on an energy, which greatly simplifies their analysis. This is not true of many of the equations used, *e.g.* for texture synthesis, and brings the current work closer to [17], in which reaction-diffusion equations arise from the Gibbs energies of images learned using minimax entropy. Although the energies used here fit within the framework of [17], which is rather general, [17] did not address region representation and modelling.

The third and most important difference between the work on texture synthesis and the present work is that pattern generation in the former relies on the Turing instability of homogeneous solutions: such solutions 'decay' into periodic patterns, resulting in spontaneous region creation in homogeneous areas, even in the absence of data. This might be useful in some applications, but it is of no use in the present context, since we wish to be able to describe localized regions such as those in figure 4. This requires that the homogeneous solutions be Turing stable (*i.e.* the second functional derivative of E evaluated at $\phi_\pm$ must be positive definite), to prevent the decay of the background or region, which in turn constrains the possible parameter values. Another way to say the same thing is that ϕ_0 is not well-defined in the unstable case. However, as figure 4 shows, stability does not mean that complex patterns cannot arise [7].

4. Line network detection.

In order to extract line networks from remote sensing images, image terms must be defined in terms of the phase field function ϕ. We define two local terms and one nonlocal term similar to those introduced by Rochery *et al.* in [11]:

$$E_I(\phi) = -\int_\Omega d^2x \left\{ \lambda_I \nabla I \cdot \nabla \phi + \alpha_I \left(\frac{\phi + 1}{2} \right) F[I] \right\}$$
$$- \frac{\beta_I}{2} \int_{\Omega^2} d^2x \, d^2x' \, \nabla I \cdot \nabla I' \, \nabla \phi \cdot \mathbf{G} \cdot \nabla \phi' \,, \quad (10)$$

where (un)primed quantities are evaluated at x' (x), and $\mathbf{G}$ is given by (9). Note that $\nabla \phi$ is zero everywhere except for the interface, where it points inwards. The first term thus predicts large inwardly pointing image gradients on the road boundary. Note that $(\phi + 1)/2$ is approximately the characteristic function of the region. F is a simple line detection filter; the second term thus predicts large values

when the filter output is integrated over the road interior. The third term is nonlocal, and corresponds to a quadratic HOAC term (6) in which $\mathbf{G}_C$ depends on the image. If $\nabla\phi(x)$ and $\nabla\phi(x')$ are parallel and close (*i.e.* x and x' are on the same side of the road), it predicts large, parallel values of $\nabla I(x)$ and $\nabla I(x')$. If $\nabla\phi(x)$ and $\nabla\phi(x')$ are antiparallel and close (*i.e.* x and x' are on opposite sides of the road), it predicts large, antiparallel values of $\nabla I(x)$ and $\nabla I(x')$. Note that unlike the prior term E_{NL}, this term is not translation invariant, and so is not diagonalized by the Fourier basis. For the moment, we treat this term by brute force, convolving a varying filter with ϕ.

We used $E_I + E$ as a model for road networks in the experiments shown in figures 5 and 6, performing gradient descent from the neutral initialization ϕ_{init}. In figure 5, two steps of the evolution are shown, in addition to the data and the result, to illustrate the very different nature of the evolution found with the phase field model as compared to a contour model. Figure 6 shows two further results. The parameter values were chosen by hand, but they are constrained by the relation with w and Turing stability. For the experiments shown here, they were $(D, \lambda, \alpha, \beta, \lambda_I, \alpha_I, \beta_I) = (1, 3.5, 0.55, 0.16, 1.8, 0, 3)$ and $(1.1, 5, 1.7, 0.3, 2, 0.4, 1.5)$ and $(0.35, 0.6, 0.05, 0.01, 2, 0.2, 0.1)$. In both cases, despite the neutral initialization, the network is largely extracted, while confounding features such as image gradients at the edges of fields are ignored. The main failure mode involves 'gaps' in the network. This was addressed within the HOAC framework by Rochery *et al.* [12, 13]. The same idea can be translated into the phase field framework, and would certainly lead to reduction in the number of gaps. Note that if we run the model on these images with $\beta = \beta_I = 0$ (*i.e.* with a standard active contour), even the best results are not at all close to the network.

5. Conclusion

Phase field models offer an advantageous alternative to classical active contour methods: a linear representation space; ease of implementation (a PDE with no reinitialization); neutral initialization; and greater topological freedom. Further, extended models that are equivalent to higher-order active contour prior energies can be constructed, meaning that sophisticated prior geometric knowledge can be included while retaining these advantages. In addition, the phase field framework greatly simplifies the analytic and algorithmic treatment of the HOAC terms. Phase fields can also be used to build standard image terms, and again, more sophisticated image models can be constructed using terms equivalent to image-dependent HOAC energies. Road network extraction results on aerial and satellite images confirm the efficacy of the models and algorithms: starting from a neutral initialization, the network is
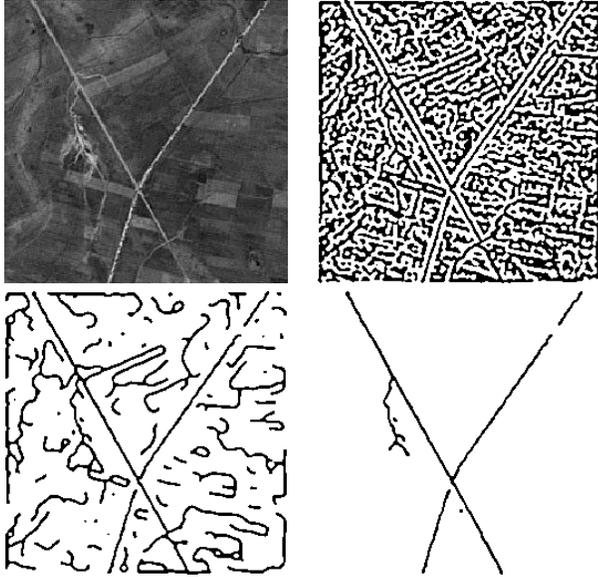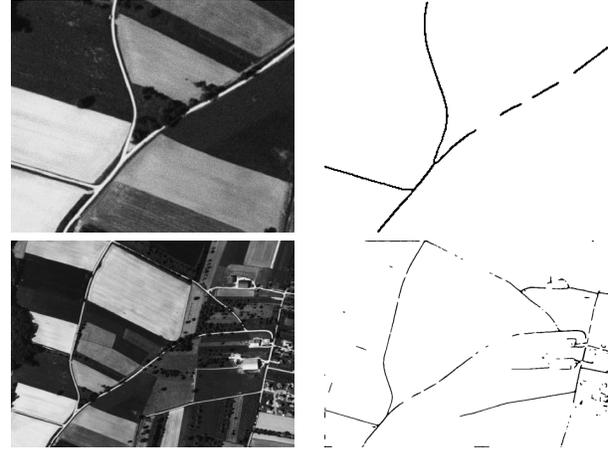
Figure 5. From left to right, top to bottom: satellite image (©CNES); two steps of the evolution; result.)

Figure 6. Top row: piece of aerial image (©IGN); result. Bottom row: aerial image (©IGN); result.

largely extracted while ignoring confounding features that confuse models which do not include prior geometric information or nonlocal image terms. The results obtained are at least equal in quality to those obtained in [11], while computation times are reduced. The gaps that remain in the networks were a problem for [11] also. Addressing them along the lines of [12, 13] is future work.

References

[1] S. M. Allen and J. W. Cahn. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsing. *Acta Metallurgica*, 27:1085–1095, 1979.

[2] G. Aubert, J. F. Aujol, and L. Blanc-Féraud. Detecting codimension-two objects in an image with Ginzburg-Landau models. Research Report 5254, INRIA, July 2004.

[3] D. Cremers, S. Osher, and S. Soatto. Kernel density estimation and intrinsic alignment for knowledge-driven segmentation: Teaching level sets to walk. In *Proc. Patt. Rec.*, volume 3175 of *LNCS*, pages 36–44. Springer, 2004.

[4] D. Cremers, N. Sochen, and C. Schnörr. Multiphase dynamic labeling for variational recognition-driven image segmentation. In *Proc. ECCV*, volume 3024 of *LNCS*, pages 74–86. Springer, 2004.

[5] H. Grossauer and O. Scherzer. Using the complex Ginzburg-Landau equation for digital inpainting in 2D and 3D. In *Proc. Scale-Space*, volume 2695 of *LNCS*, pages 225–236. Springer, 2003.

[6] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *IJCV*, 1(4):321–331, 1988.

[7] S. A. Langer, R. E. Goldstein, and D. P. Jackson. Dynamics of labyrinthine pattern formation in magnetic fluids. *Phys. Rev. A*, 46(8):4894–4904, 1992.

[8] S. Osher and N. Paragios, editors. *Geometric Level Set Methods in Imaging, Vision, and Graphics*. Springer, New York, 2003.

[9] S. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations. *J. Comp. Phys.*, 79:12–49, 1988.

[10] C. B. Price, P. Wambacq, and A. Oosterlinck. Image enhancement and analysis with reaction-diffusion paradigm. *IEE Proceedings I*, 137(3):136–145, June 1990.

[11] 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. VLSM*, at ICCV, Nice, France, Oct. 2003.

[12] M. Rochery, I. H. Jermyn, and J. Zerubia. Gap closure in (road) networks using higher-order active contours. In *Proc. ICIP*, Singapore, Oct. 2004.

[13] M. Rochery, I. H. Jermyn, and J. Zerubia. New higher-order active contour energies for network extraction. In *Proc. ICIP*, Genova, Italy, Sept. 2005.

[14] C. Samson, L. Blanc-Féraud, G. Aubert, and J. Zerubia. A variational model for image classification and restoration. *Trans. PAMI*, 22(5):460–472, May 2000.

[15] G. Turk. Generating textures on arbitrary surfaces using reaction-diffusion. *Comp. Graph.*, 25(3), July 1991.

[16] A. Witkin and M. Kass. Reaction-diffusion textures. *Comp. Graph.*, 25(3), July 1991.

[17] S. C. Zhu and D. Mumford. Prior learning and Gibbs reaction-diffusion. *Trans. PAMI*, 19(11):1236–1250, Nov. 1997.

[18] S. C. Zhu and A. Yuille. Region competition: Unifying snakes, region growing, and Bayes/MDL for multi-band image segmentation. *Trans. PAMI*, 18(9):884–900, 1996.