SHAPE MODELLING WITH CONTOURS AND FIELDS

Durham University
INRIA

Ian Jermyn

Josiane Zerubia	**INRIA**
Zoltan Kato	University of Szeged
Marie Rochery	INRIA
Peter Horvath	INRIA, Szeged
Ting Peng	INRIA, CASIA
Aymen El Ghoul	INRIA

Overview

- **Why shape and what is shape**?
  - Running example: segmentation.
  - Representation.
- Classical approach:
  - Distances, templates, comparison.
- Nonlocal interactions.
  - (Higher-order) active contours (models, stability);
  - Phase fields (relation to contours, advantages);
  - Binary fields (relation to phase fields, advantages).
- Other models:
  - Low-curvature networks; directed networks; multilayer model; complex shapes.
- Future.

Why shape? Shape is useful

- Many subsets of the world are differentiated from their surroundings.
- The subset has geometry.
- The geometry may be correlated with other properties of the subset.
  - Thus it can be used to make inferences about these properties, and vice-versa.



Example: segmentation

- Find the region R in the data domain that 'corresponds to' a given 'entity'.
- Many applications:
  - Oil strata; roads, rivers, trees; cells; brain fibres; CMEs;...
- Data volume is often huge:
  - EO satellite generates 1TB/day.
  - 62% growth in 2009; 35ZB by 2020.
- Automated inference is crucial for full exploitation of this data.


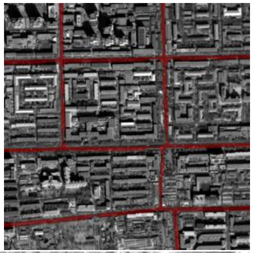
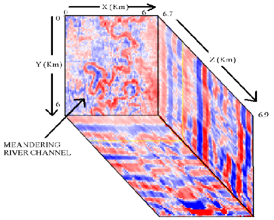

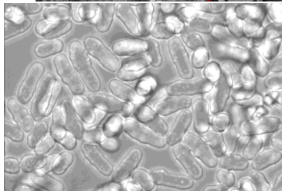
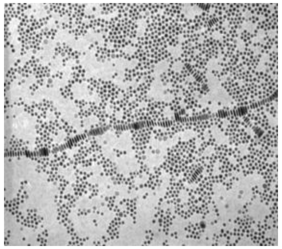
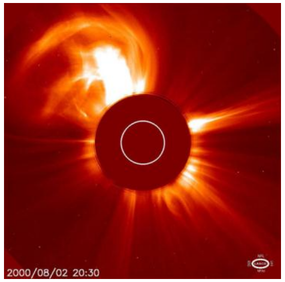

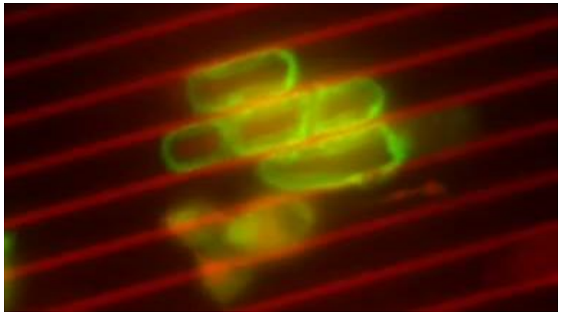

Why shape? Shape is necessary

- Images are complex:
  - Generic techniques are often insufficient for automatic solutions.
- Regions of interest are often distinguished by their shape:
  - A model of likely region shape is crucial.
  - Prior information K.

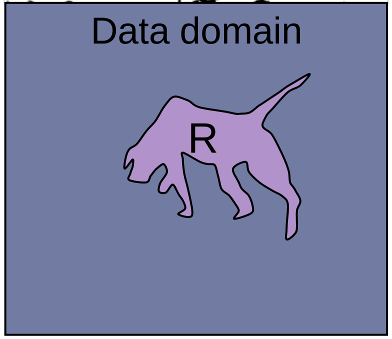

What is shape?

- A subset of a manifold.
  - E.g. codimension-0 submanifold of $\mathbb{R}^n$.
  - Complicated space $\mathcal{R}$.
- Prior information K is not usually enough to specify one subset:
  - Uncertainty.
- Need probability distribution P(R | K):
  - "What we know about R $\in \mathcal{R}$ given prior information K".
    - E.g. that R corresponds to entity X and…
  - Probability distribution often specified by an 'energy':

$$\mathrm{P}(R|K) \propto \exp\{-E(R|K)\}$$

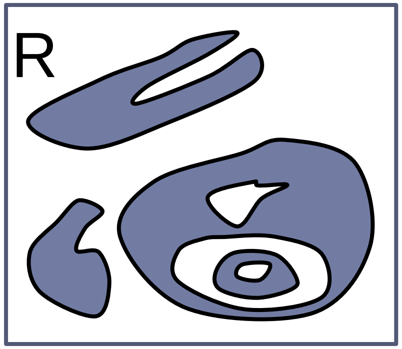

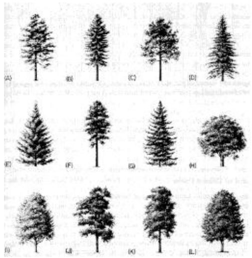

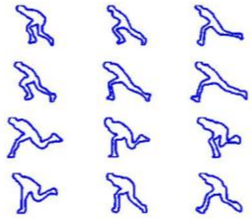

Example: segmentation

- Probability region R corresponds to X, given image I and prior information K.

$$\mathrm{P}(R|I,K) \propto \mathrm{P}(I|R,K)\mathrm{P}(R|K)$$

- P(I | R, K): image model.
  - Probability observed image is I, given region R corresponds to X and prior information K.
- P(R | K): shape model.
  - Probability that R corresponds to X given prior information K.

Which probability distributions?

- Invariance (maybe):
  - $P(g(R)) = P(R)$, $g \in G$.
  - G: Euclidean, similarity, isometry,…
- Low entropy.
- Knowledge of one part of a subset boundary should enable prediction of other parts:
  - $H(X, Y) = H(Y \mid X) + H(X)$.
  - Long-range dependencies.
- How to build such distributions and on which space?

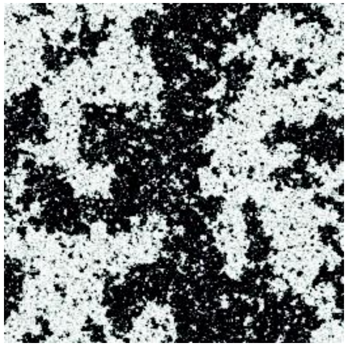

Which space?

- Space $\mathcal{R}$ is hard to deal with directly:
  - Use 'representation space' $\mathcal{S}$.
- One-to-one (ρ: $\mathcal{R} \to \mathcal{S}$ invertible):
  - Characteristic function; distance function.
  - Seems good, but does not simplify space, singular.
- Many-to-one (ρ: $\mathcal{R} \to \mathcal{S}$ not injective):
  - Landmark points; shape space; Fourier descriptors; medial axis; 'canonical form'.
  - Plus: invariance already in representation (ρ(R) = ρ(g(R)));
  - Minus: information lost (R cannot be reconstructed from S = ρ(R)).
- One-to-many (ρ: $\mathcal{S} \to \mathcal{R}$):
  - Parameterized closed curve ('contour'); phase field.
  - Have to deal with redundancy (group H) or induced model.

Overview

- Why shape and what is shape?
  - Running example: segmentation.
  - Representation.
- **Classical approach**:
  - Distances, templates, comparison.
- Nonlocal interactions.
  - (Higher-order) active contours (models, stability);
  - Phase fields (relation to contours, advantages);
  - Binary fields (relation to phase fields, advantages).
- Other models:
  - Low-curvature networks; directed networks; multilayer model; complex shapes.
- Future.

Classical approach

- Define a distance on representation space: $d^2(S, S')$.
  - For one-to-many ρ, minimization is needed to define a distance on $\mathcal{R}$.
  - For many-to-one ρ, distance is pseudometric on $\mathcal{R}$.
- Examples:
  - Length of geodesics in a Riemannian metric on parameterized closed curves.
  - Graph matching metrics, e.g. for medial axis.
  - Gromov-Hausdorff metric on isometry equivalence classes of boundaries.

Classical approach

- Construct P(S) using distance(s) from 'template' shape(s) S_0 (part of K).
  - Energy is distance: 'Gaussian'.

$$\mathrm{P}(S|S_0) \propto \exp\{-d^2(S, S_0)\}$$

- Invariance via mixture model over G:

$$\mathrm{P}(S|[S_0]) \propto \int_G dg \; \exp\{-d^2(S, g(S_0))\} \simeq \exp\{-d^2(S, g^*(S_0))\}$$

  - 'Pose estimation': assumes G acts consistently on $\mathcal{S}$.
  - Gives rise to long-range dependencies.
- Frequently used implicitly in 'shape classification'.

Classical approach

- Distance / template approach is very useful, but…
- Does not apply (or is inefficient) in many important cases:
  - Unknown topology and extent.
  - E.g. multiple objects.
- Need modelling framework allowing strong constraints on shape, but with weak constraints on topology.

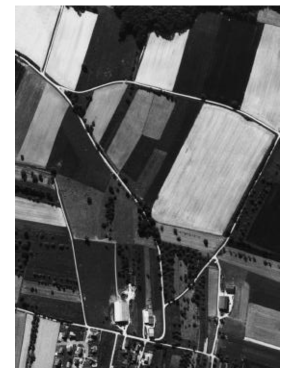

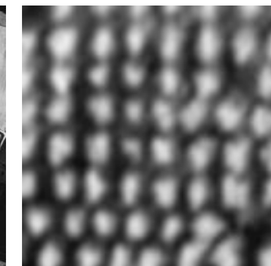

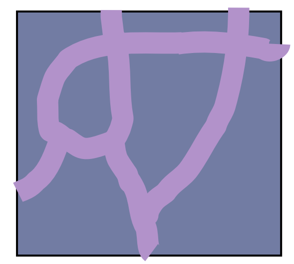

Networks: many loops

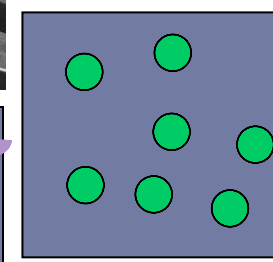

Multiple objects: many components

What to do?

- Explicit nonlocal interactions between boundary points:
  - Create long-range dependencies;
  - Avoid templates: topology / extent need not be constrained.
  - Avoid 'pose estimation': invariance is manifest in model.
- Exploit equivalences between formulations:
  - Contour:
    - Intuitive; facilitates stability analysis.
  - Phase field:
    - Linear space; lower complexity; easy implementation.
  - Binary field:
    - Facilitates sampling, hence learning; allows use of graph cut algorithms.

Overview

- Why shape and what is shape?
  - Running example: segmentation.
  - Representation.
- Classical approach:
  - Distances, templates, comparison.
- **Nonlocal interactions**.
  - (Higher-order) active contours (models, stability);
  - Phase fields (relation to contours, advantages);
  - Binary fields (relation to phase fields, advantages).
- Other models:
  - Low-curvature networks; directed networks; multilayer model; complex shapes.
- Future.

Contours: building E

- R is represented by parameterized closed curve(s), the 'contour' γ .
- Simplest invariant energy:
  - Length of ∂R and area of R:

$$E_{G,0}(\gamma) = \lambda_C L(\gamma) + \alpha_C A(\gamma)$$

    - Cf. Ising model, Brownian motion.
  - Short-range interactions between boundary points.
  - Describes boundary smoothness.
  - Insufficient for all but the simplest problems.

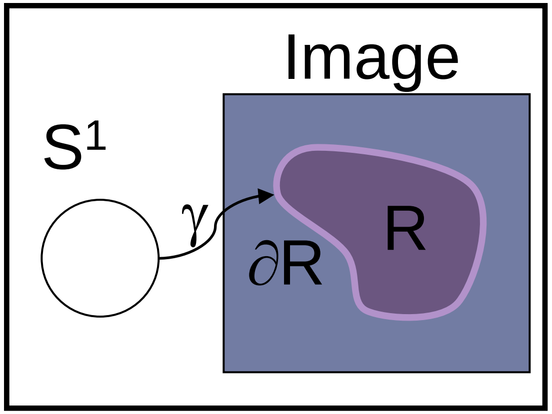

$$L(\gamma) = \int dt |\dot{\gamma}(t)|$$

$$A(\gamma) = \frac{1}{2} \int dt \; [\dot{\gamma}(t) \times \gamma(t)]$$

Expressions invariant to Diff(S^1) hence well-defined on $\mathcal{R}$.

Building E: nonlocal interactions

- Introduce prior information via nonlocal interactions between tuples of points.

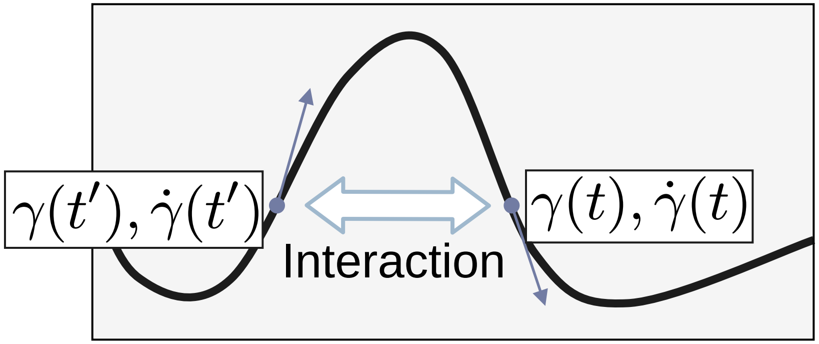

- E.g. Euclidean invariant two-point term:

$$E(\gamma) = -\iint dt\, dt'\; \dot{\gamma}(t) \cdot \dot{\gamma}(t')\; \Psi(|\gamma(t) - \gamma(t')|)$$

Energy for networks

$$E_{\mathrm{G}}(\gamma) = \lambda L(\gamma) + \alpha A(\gamma) - \frac{\beta}{2} \iint dt\, dt'\; \dot{\gamma}(t) \cdot \dot{\gamma}(t')\; \Psi_d(|\gamma(t) - \gamma(t')|)$$

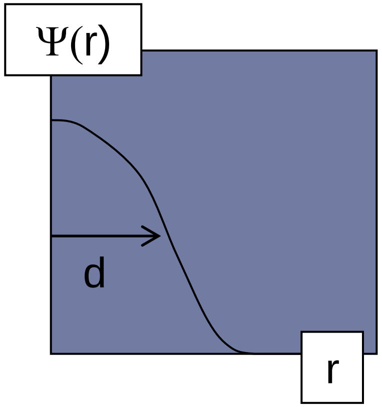

- Gradient descent (with large β):
  - A circle is a saddle point of the energy.

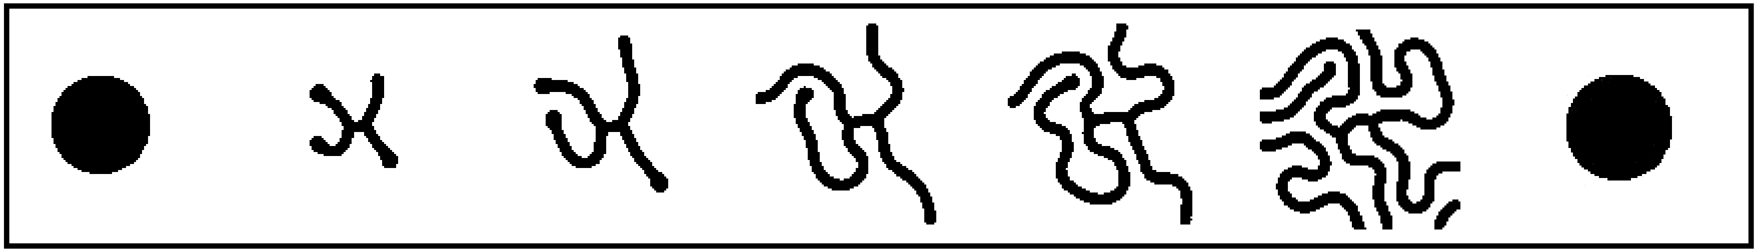

- Network structures have low-energy and are stable:
  - The energy E_G 'models' them.
  - Good for roads, blood vessels, &c.

Energy for a ‘gas of near-circles’

- Experiments show that the same energy E_G can model a ‘gas of near-circles’ :
  - I.e. such configurations are local minima.
- Only true for certain parameter ranges.
  - Which ranges?
- Stability analysis enables fixing of parameters to model different structures.

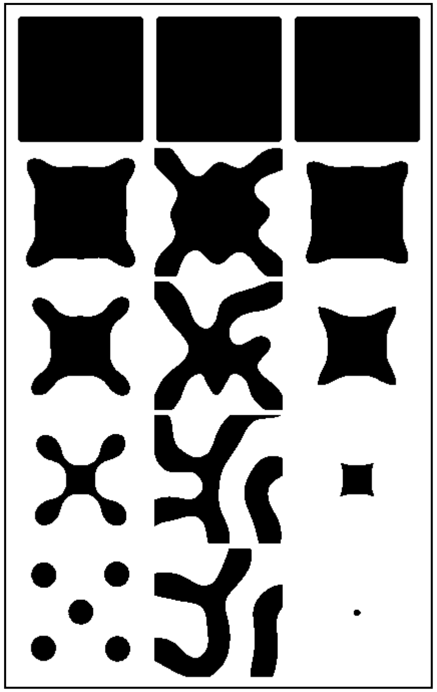

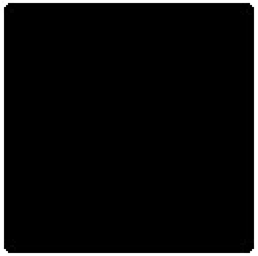

Stability analysis: circle and bar

- Expand $E_G(\gamma_0 + \delta\gamma)$ to second order in $\delta\gamma$:

$$E_G(\gamma_0+\delta\gamma) = E(\gamma_0) + \langle \delta\gamma | \frac{\delta E}{\delta \gamma}(\gamma_0)\rangle + \frac{1}{2}\langle \delta\gamma | \frac{\delta^2 E}{\delta\gamma\delta\gamma'}(\gamma_0)|\delta\gamma'\rangle$$

Must be zero

Must be positive definite

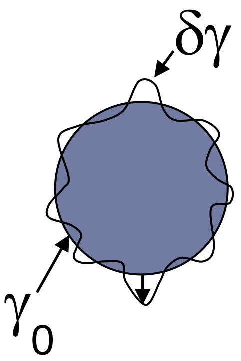

- Rotation invariance means Fourier modes are not coupled.
- Constraints on parameters.

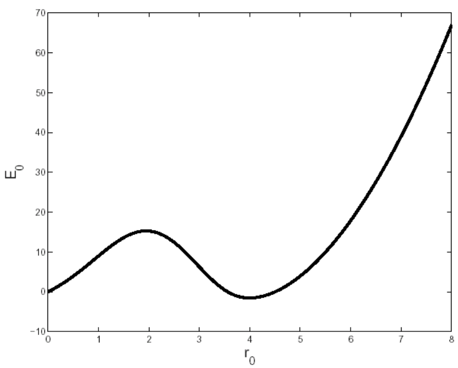

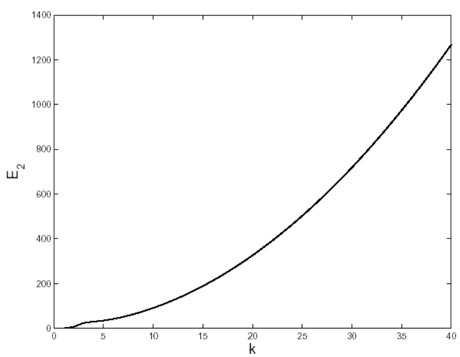

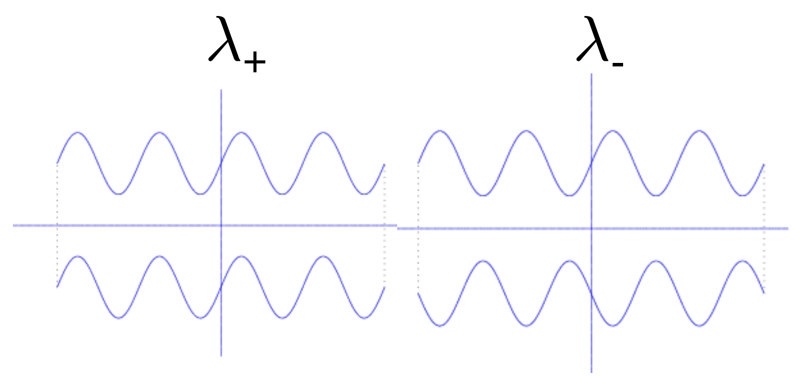

Phase diagram: bars and circles.

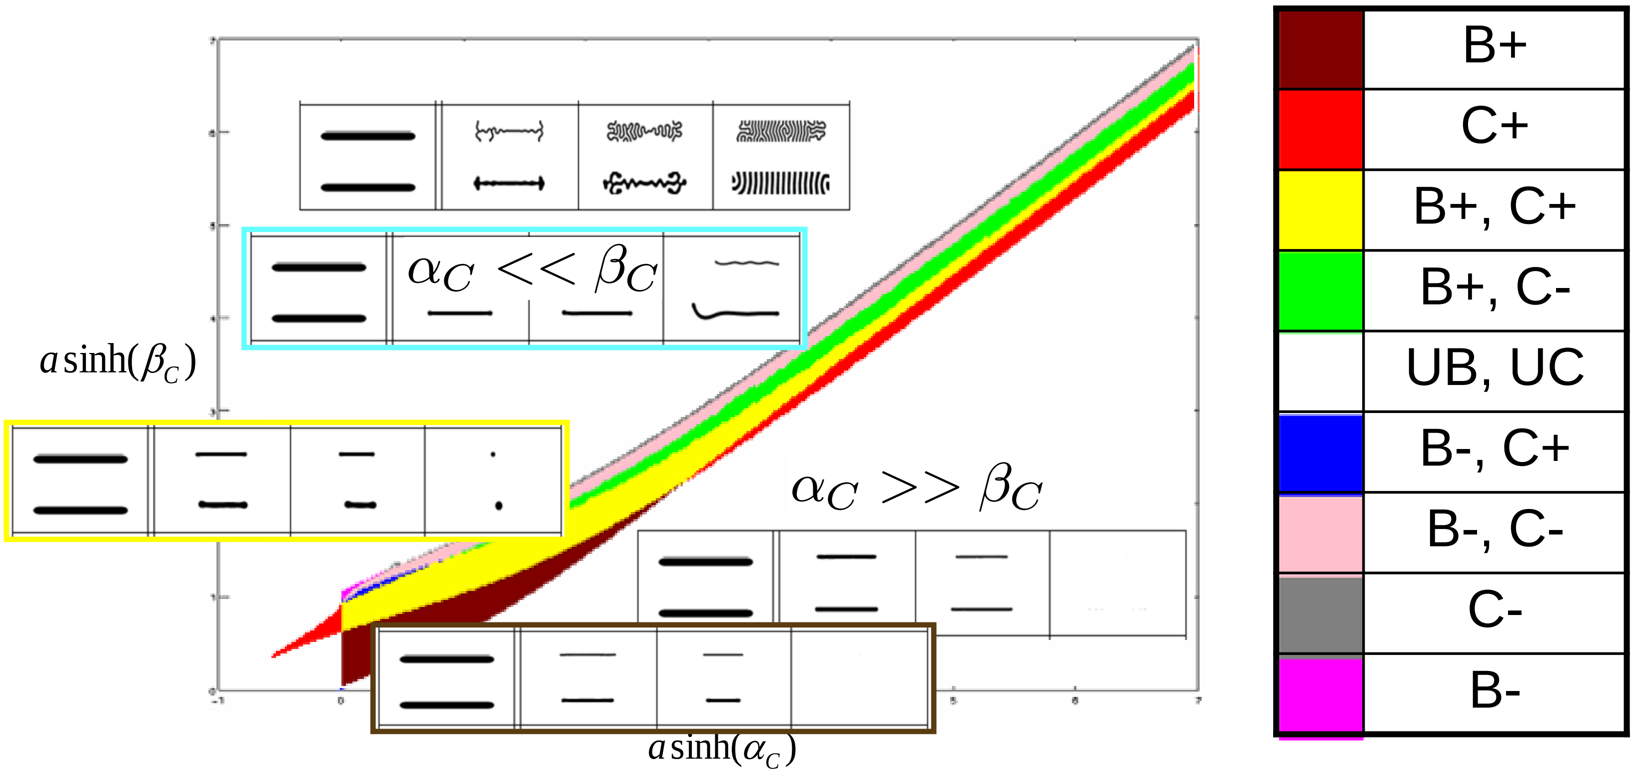

Numerical check: circles

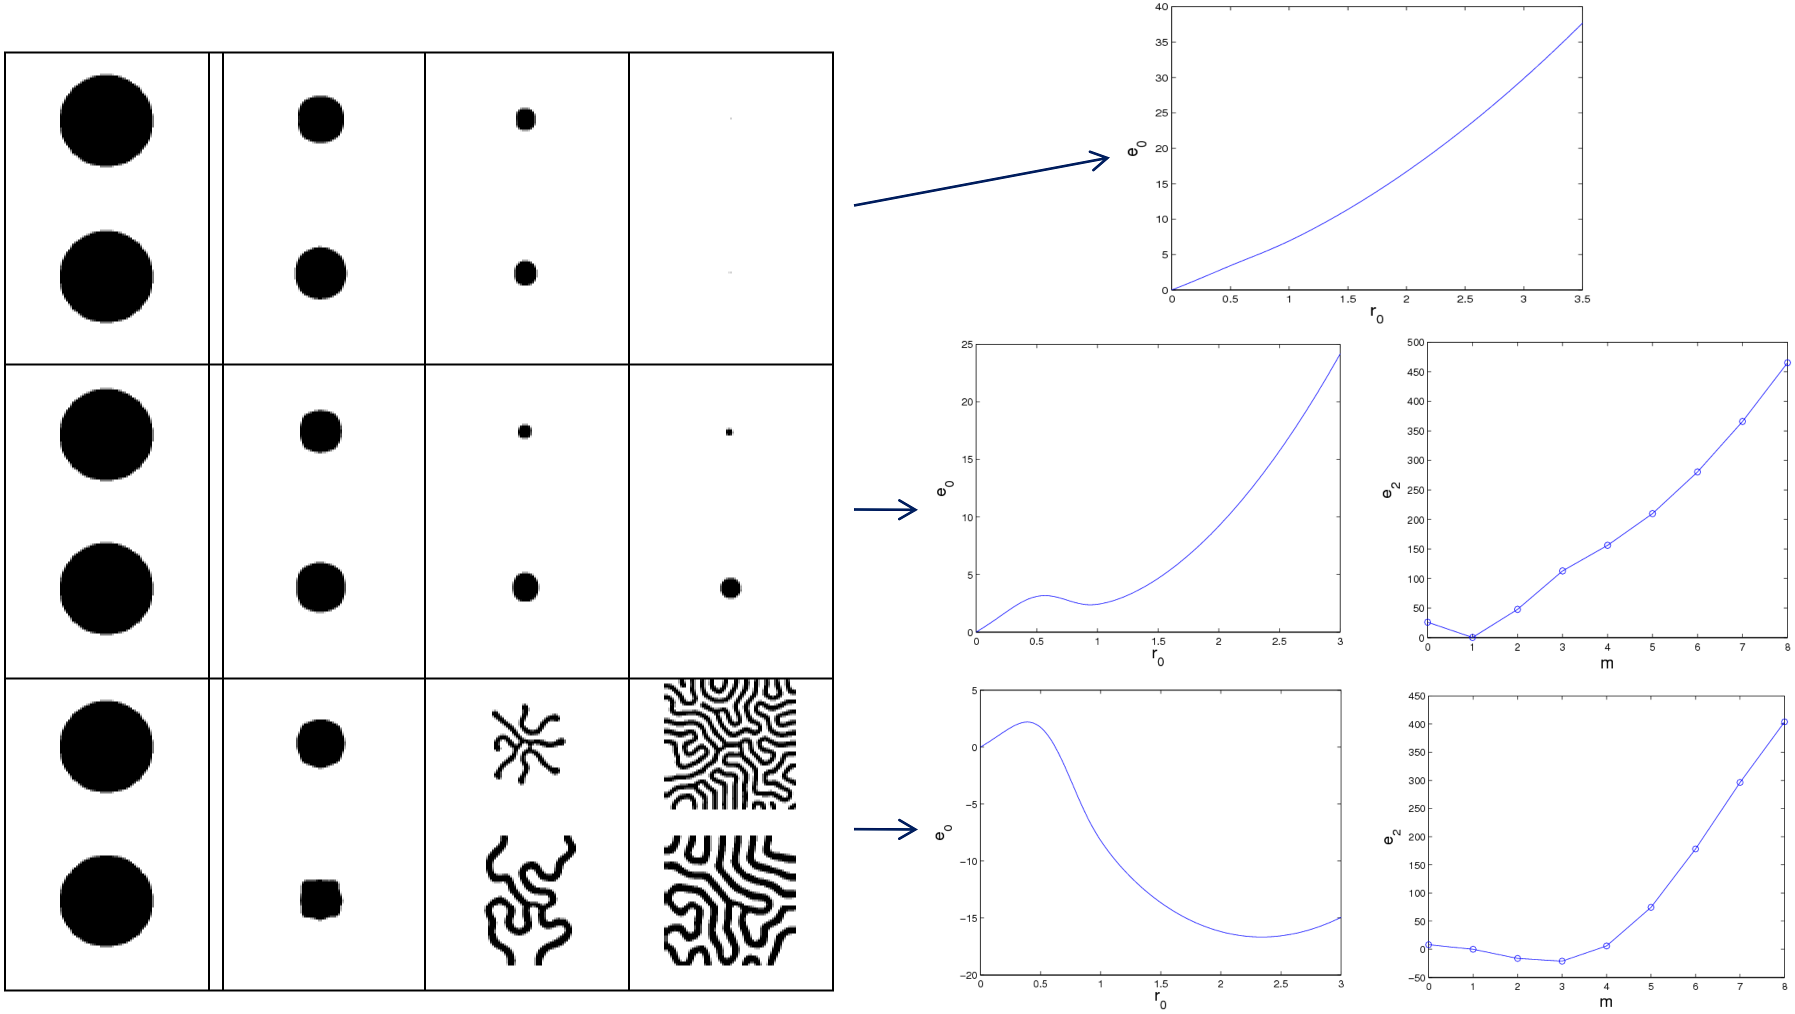

Numerical check: bars

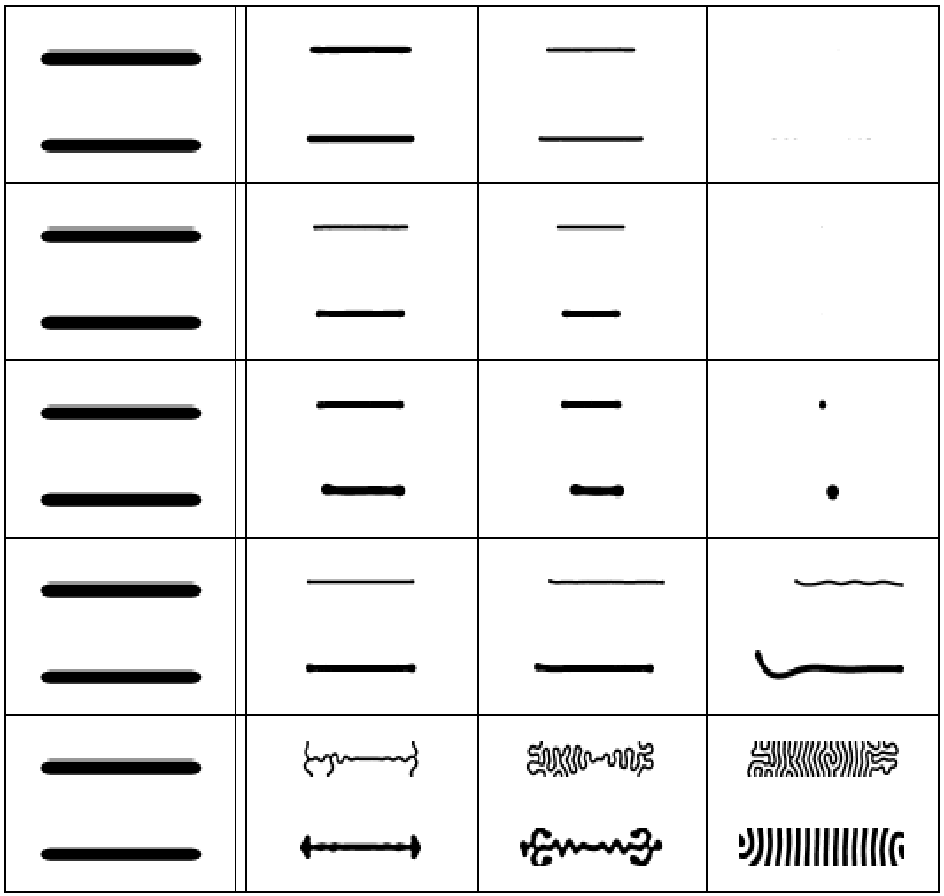

$$E_G(\gamma) = \ell \, e_G(\gamma)$$

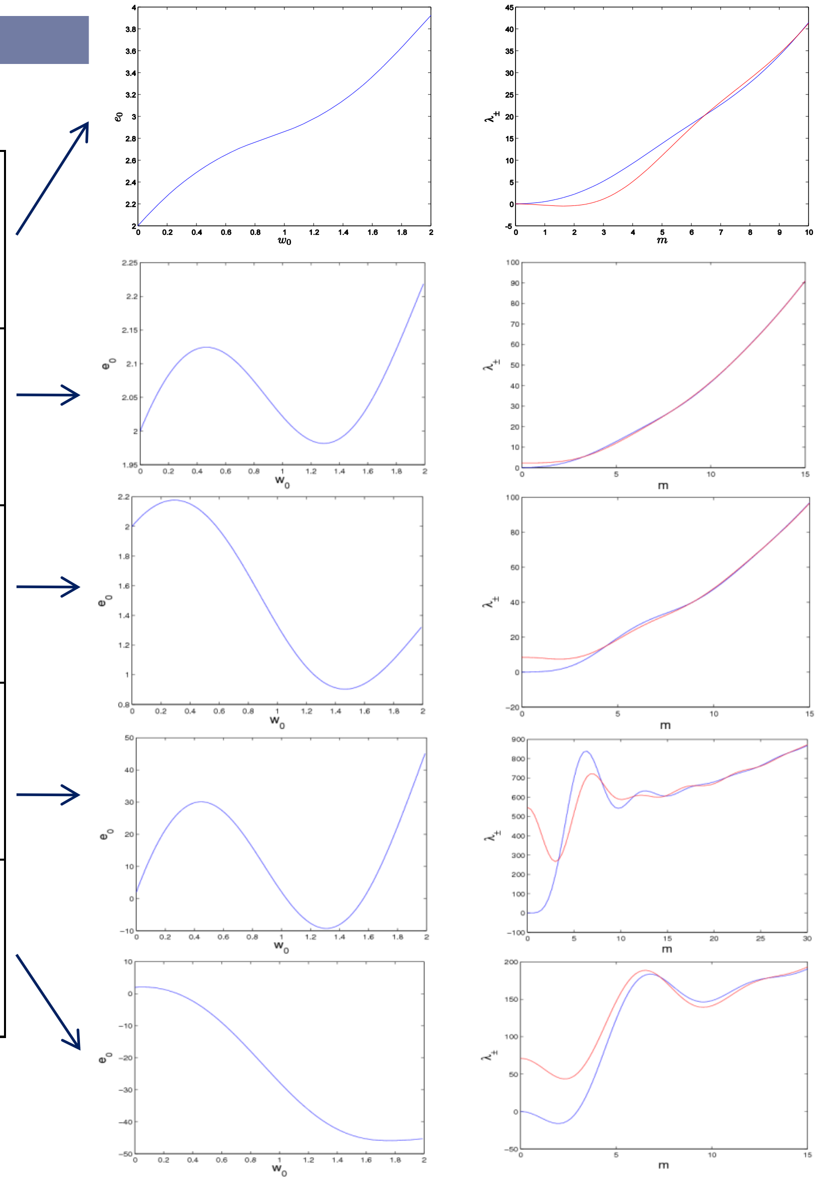

Example: segmentation

- Extract information from P(R | I, K) via a MAP estimate:

$$\hat{R} = \arg\max_R \mathrm{P}(R|I,K) = \arg\min_R E(R,I)$$

  - Where

$$\begin{aligned} E(R,I) &= -\ln \mathrm{P}(R|I,K) = -\ln \mathrm{P}(I|R,K) - \ln \mathrm{P}(R|K) + \text{const} \\ &= E_{\mathrm{I}}(I,\gamma) + E_{\mathrm{G}}(\gamma) + \text{const} \end{aligned}$$

- Algorithm: gradient descent using distance function level sets.
- But nonlocal term requires:
  - Extracting the contour;
  - Multiple integrations around the contour;
  - 'Velocity extension'.

Roads

$$E_I(I,\gamma) = \int_{\partial R} n \cdot \partial I(\gamma) - \int_{\partial R \times \partial R} \partial I \cdot \partial I' \; \dot{\gamma} \cdot \dot{\gamma}' \; \Psi(|\gamma - \gamma'|)$$

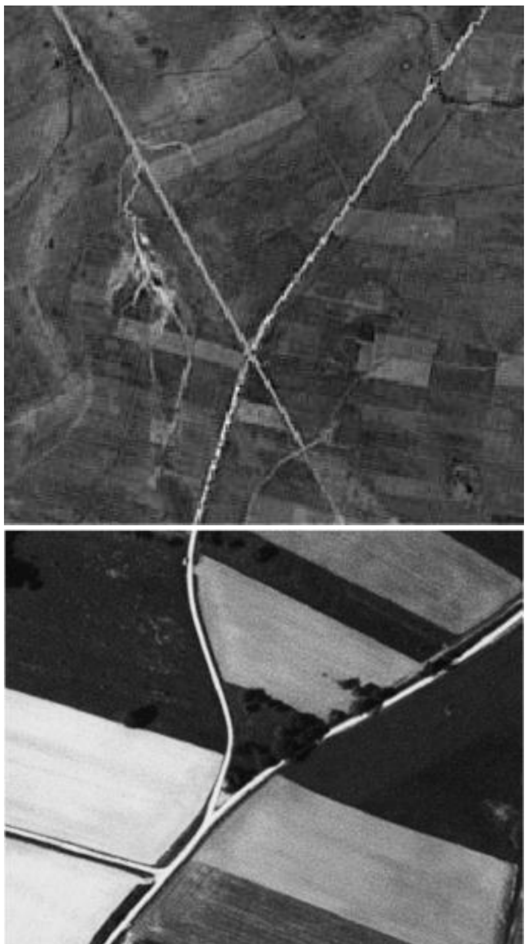

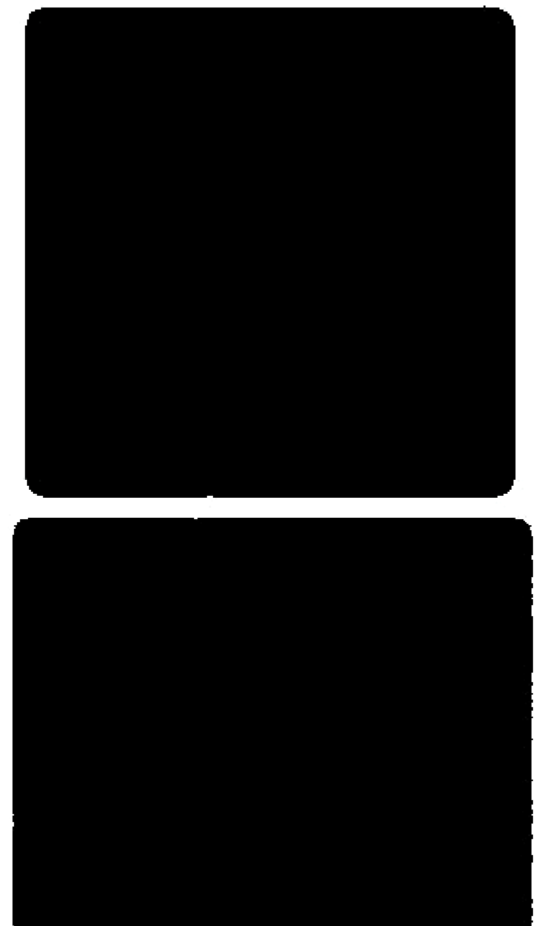

Tree crowns

$$E_I(I,\gamma) = \int_{\partial R} n \cdot \partial I(\gamma) + \int_R \frac{(I-\mu)^2}{2\sigma^2} + \int_{\bar{R}} \frac{(I-\bar{\mu})^2}{2\bar{\sigma}^2}$$

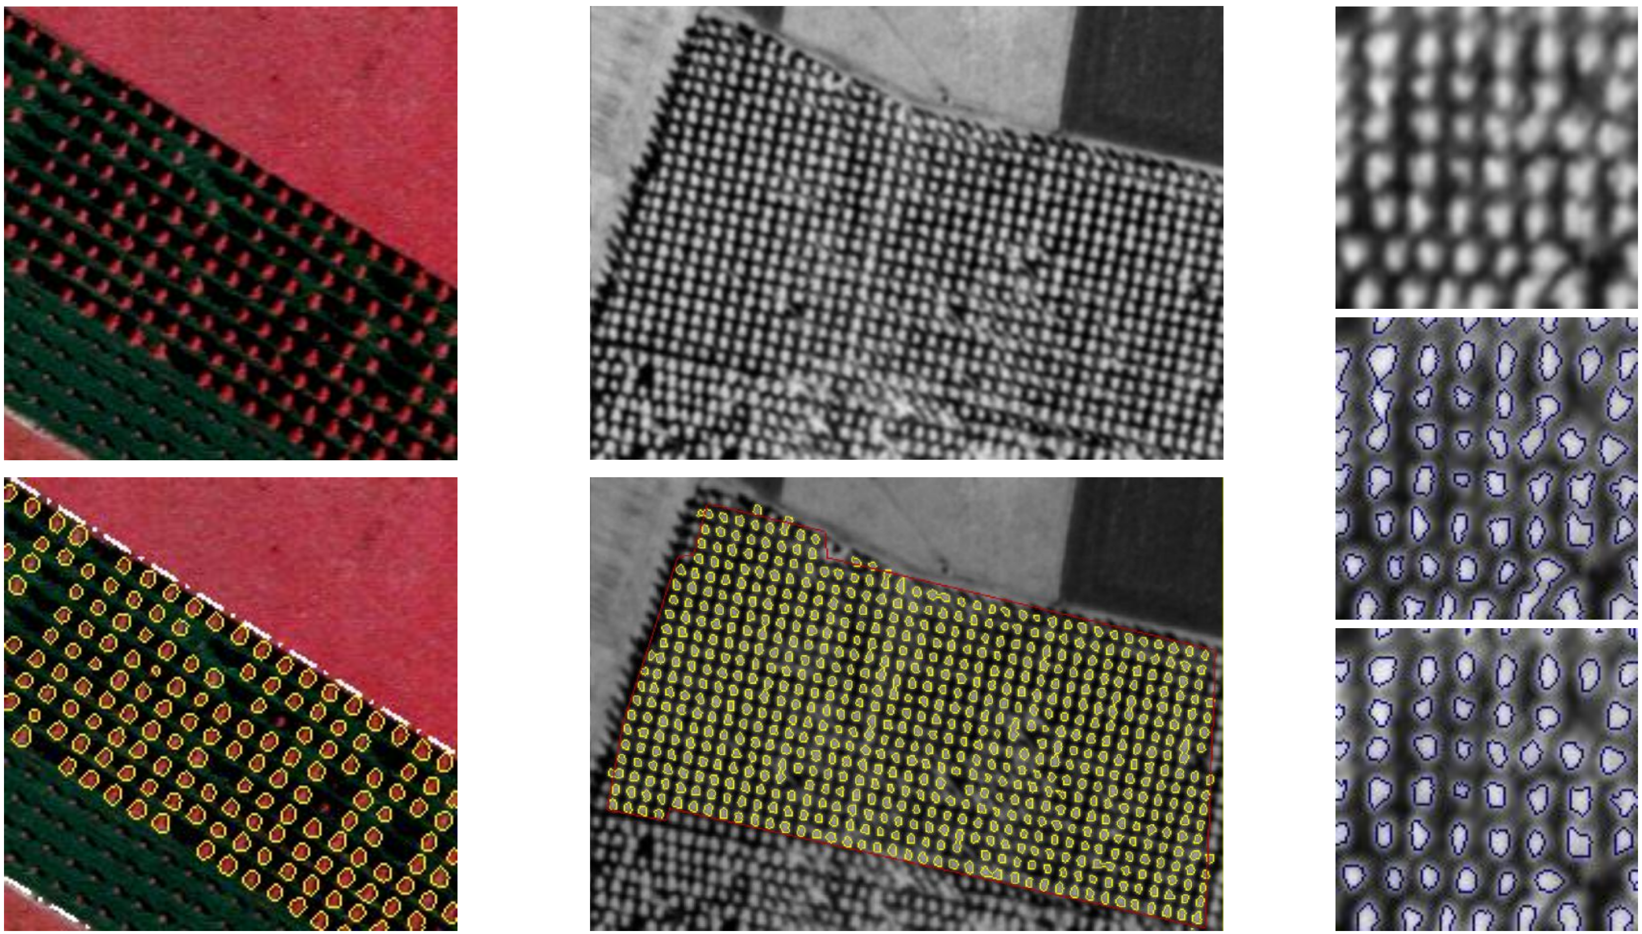

Problems with contours

- Modelling:
  - Space of regions is complicated in terms of contours:
    - No (self)-intersections; relative orientations; not connected; not linear space.
  - Probabilistic formulation is difficult.
- Algorithm (distance function level sets):
  - Topology change is limited:
    - Not robust to initial conditions.
  - Gradient descent is complicated to implement.
    - Slow.
- Solution: phase fields.

Phase fields

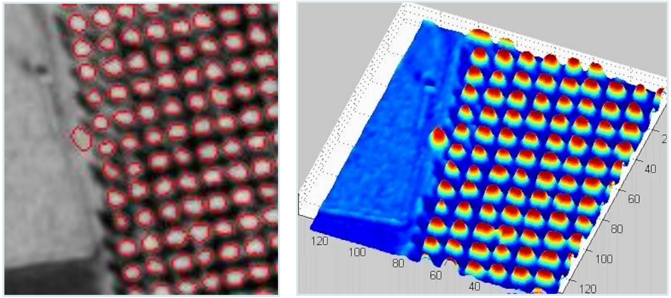

- Phase fields are a level set representation:
  - $\zeta_z(\phi) = \{x : \phi(x) > z\}$.
- Basic model: Ginzburg-Landau energy:

$$E_0(\phi) = \int \left\{ \frac{D}{2} \partial\phi \cdot \partial\phi + \lambda\left(\frac{1}{4}\phi^4 - \frac{1}{2}\phi^2\right) \right\}$$

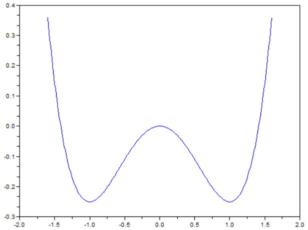

- Define ϕ_R:

$$\phi_R = \arg \min_{\phi:\, \zeta(\phi) = R} E_0(R)$$

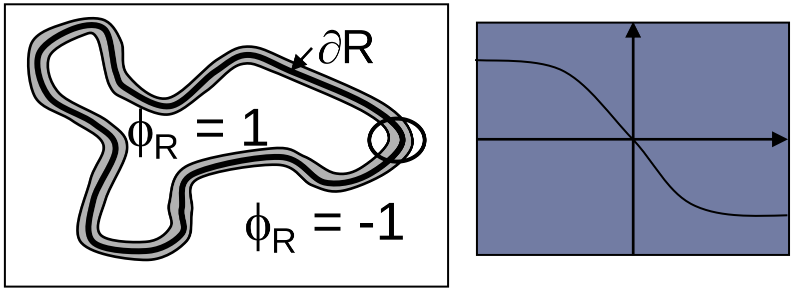

Relation to contours

- So induced model is approximately

$$E_0(\phi_R) \simeq \lambda_C L(\partial R)$$

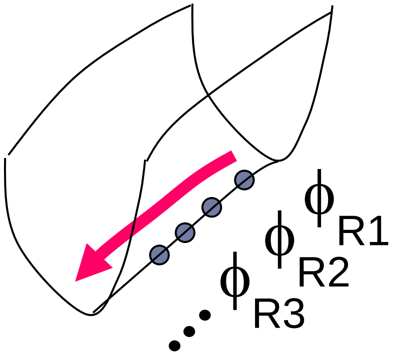

- Because ϕ_R is a minimum for fixed R, descent with E_0 mimics descent with L.
- Also true to first order in fluctuations:

$$\int_{\zeta_z(\phi)=R} D\phi\, e^{-E_0(\phi)} \approx e^{-\lambda_C L(\partial R)}$$

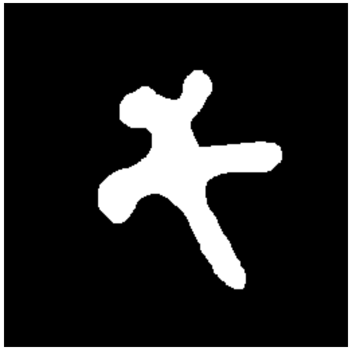

Adding nonlocal interactions

- Use that $\partial\phi_R$ is zero except near ∂R, where it is proportional to the normal vector.

$$E_Q(\gamma) = -\frac{\beta_C}{2} \iint_{\partial R^2} dt\, dt'\, \boxed{\dot{\gamma}(t)} \cdot \mathbf{G}_C(\gamma(t), \gamma(t')) \cdot \boxed{\dot{\gamma}(t')}$$

$$E_{NL}(\phi) = -\frac{\beta}{2} \iint_{\;{}^2} d^2x\, d^2x'\, \boxed{\partial\phi(x)} \cdot \mathbf{G}(x, x') \cdot \boxed{\partial\phi(x')}$$

  - Can show that $E_{NL}(\phi_R;\, \beta,\, G) \simeq E_Q(\gamma;\, \beta_C,\, G_C)$; induced model is higher-order active contour.
- Stability analysis constraints translate accurately from contour to phase field.

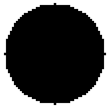

Advantages: model

- Complex topologies are easily represented:
  - No constraints on ϕ.
- Representation space is linear:
  - ϕ can be expressed in any basis, e.g., in wavelet basis for multiscale analysis of shape.
  - Probabilistic formulation (relatively) simple (continuum Markov random field).
- Nonlocal terms are quadratic.

Advantages: algorithm

- Complex topologies and multiple instances come at no extra cost.
- Descent is based solely on gradient:
  - Implementation is simple: no reinitialization.
- Neutral initialization and topological freedom:
  - No initial region; no bias.
  - Number of connected components and handles changes easily.
  - More robust to choice of initial condition.
- Nonlocal terms are linear:
  - Pointwise evaluation in Fourier domain.
  - Computation time improved.
  - Simple implementation.

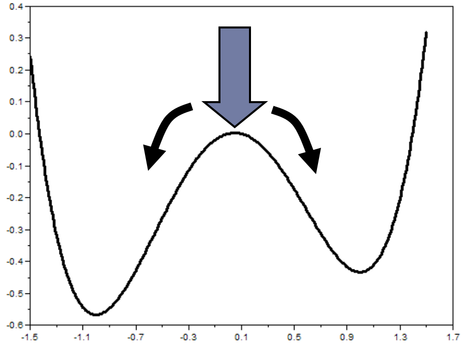

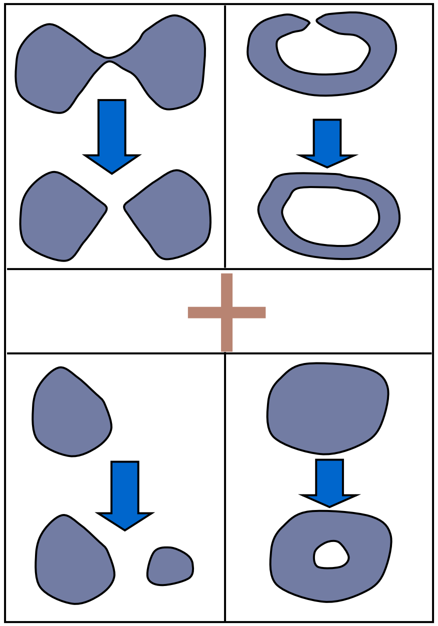

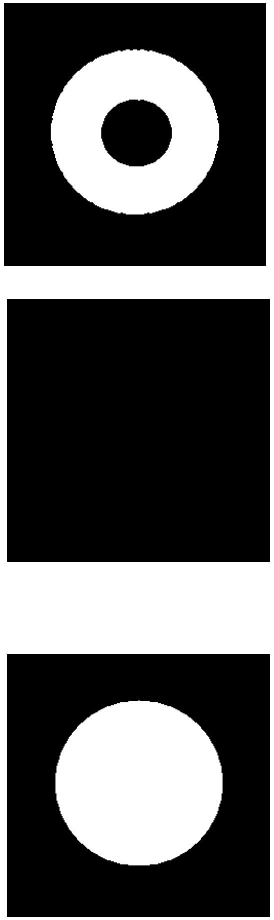

Example: segmentation

- One can build equivalents of contour likelihood energies $E_I(I, \gamma)$:
  - $\partial\phi$ the normal vector;
  - $(1 + \phi)/2$ is the characteristic function.
  - For example:

$$E_I(I,\gamma) = \int_{\partial R} n \cdot \partial I(\gamma) \simeq \int \partial\phi_R \cdot \partial I = E_I(I,\phi)$$

Example: segmentation

- Calculating the MAP estimate:

$$\hat{R} = \arg\min_{R} E(R, I) = \zeta_z(\hat{\phi})$$

$$\hat{\phi} = \arg\min_{\phi} E(\phi, I) = E_{\mathrm{I}}(I, \phi) + E_0(\phi) + E_{NL}(\phi)$$

- Algorithm: gradient descent, but, whereas…
  - Nonlocal contour terms are complex to evaluate,
  - Nonlocal phase field terms require only convolution:

$$\frac{\delta E_{NL}}{\delta \phi(x)} = \beta \int d^2x' \; \partial^2 G(x - x')\phi(x')$$

Comparison to contours

- ‘Evolution everywhere’ as opposed to ‘shrink-wrapping’ of contour.
  - Potential for parallelization and hardware implementation.

Comparison to contours

- Execution time reduced with respect to contour implementation by factor of ~10-100.

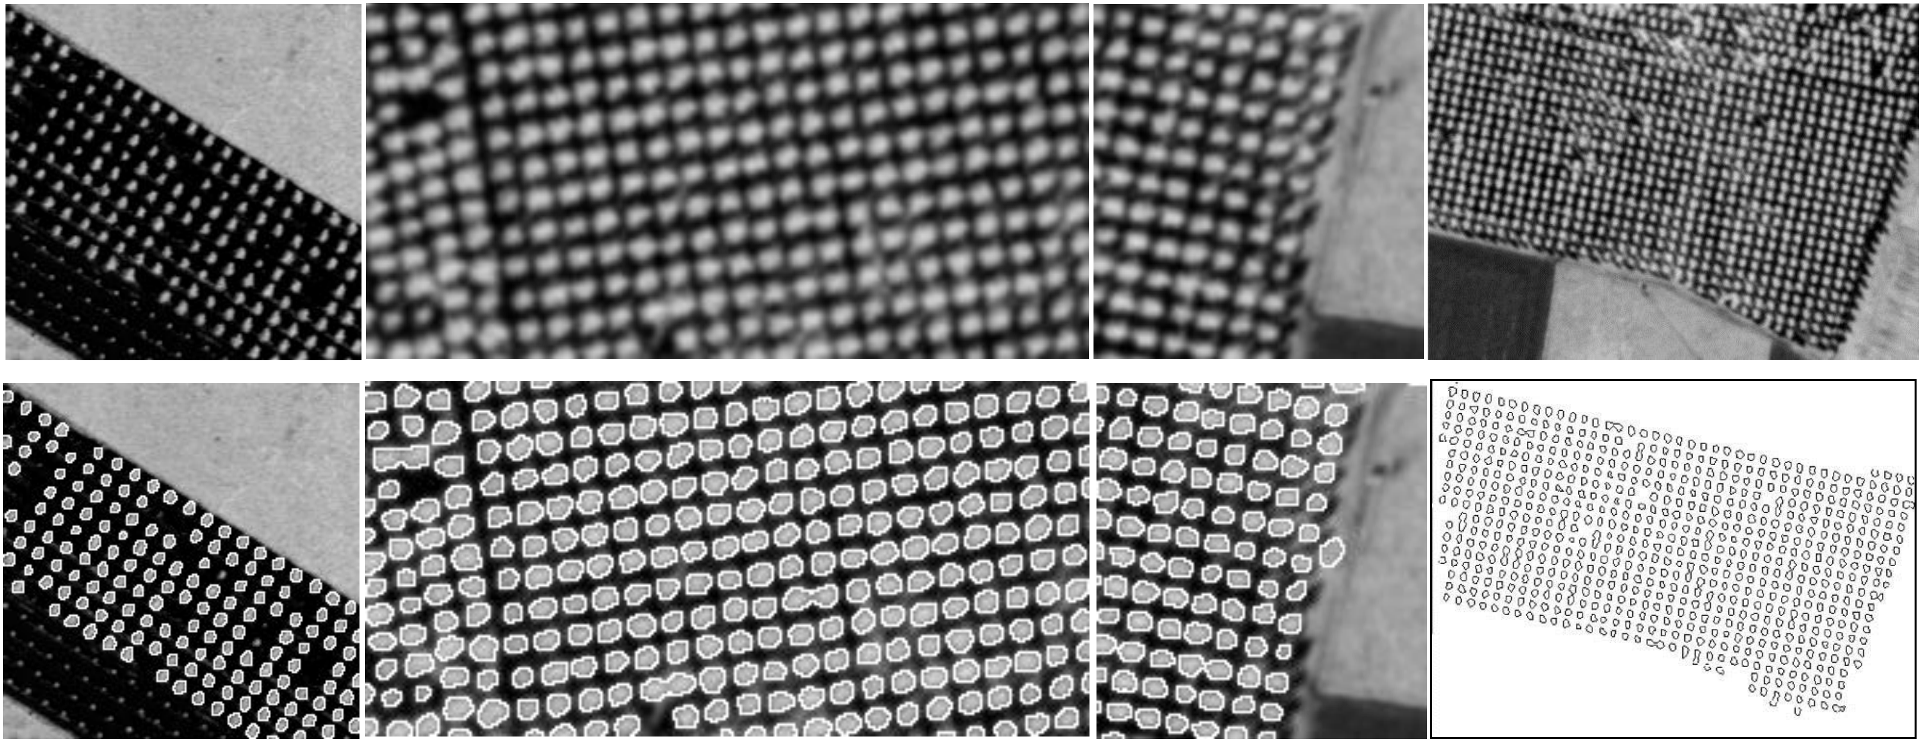

Binary field

- Because $\phi \simeq \pm 1$, one can binarize (and discretize space):
  - Ising model (i.e. boundary length) plus long-range interactions.
- Advantages:
  - MCMC sampling is more efficient:
    - Can use simulated annealing.
  - Can use QPBO algorithm:
    - Gives access to most of global minimum.
    - Conclusion: gradient descent works well.
- Potential: facilitates parameter and model learning.

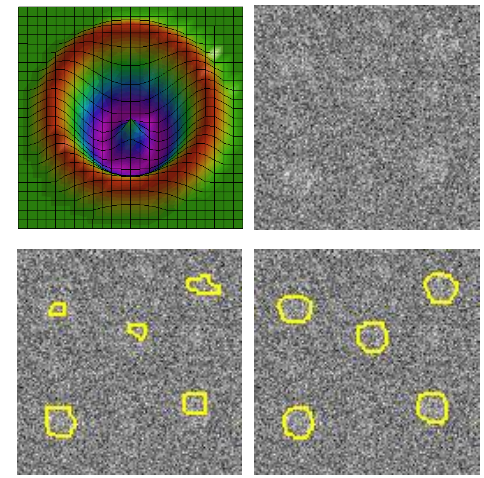

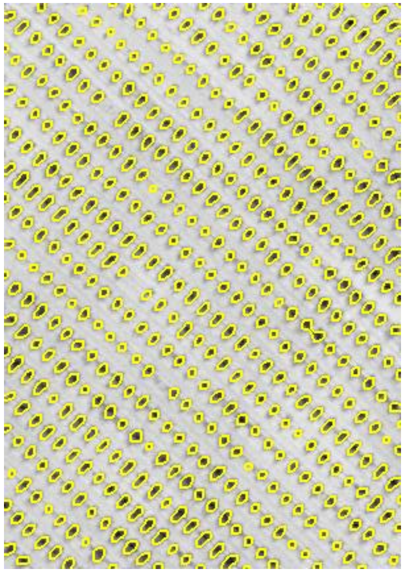

Overview

- Why shape and what is shape?
  - Running example: segmentation.
  - Representation.
- Classical approach:
  - Distances, templates, comparison.
- Nonlocal interactions.
  - (Higher-order) active contours (models, stability);
  - Phase fields (relation to contours, advantages);
  - Binary fields (relation to phase fields, advantages).
- **Other models**:
  - Low-curvature networks; directed networks; multilayer model; complex shapes.
- Future.

Low-curvature networks

- Use a more complex interaction to achieve long, straight network branches.

90 pixels

QuickBird image of Beijing (0.6m)

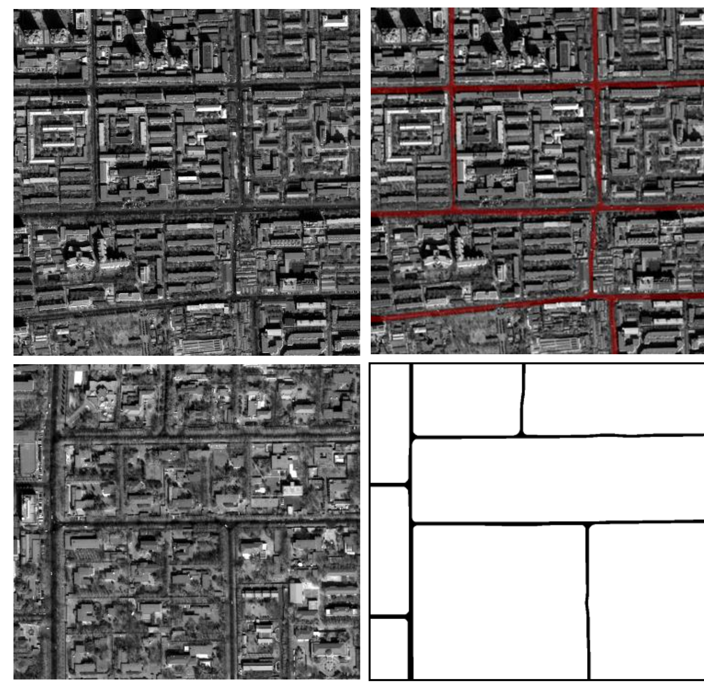

Low-curvature networks: model

- Decouple interactions along and across bar.
- Make 'along-bar' interactions longer-range and stronger.
- In terms of the contour:

$$E_Q(\gamma) = -\frac{1}{2}\iint_{\partial R \times \partial R'} \left\{ \beta_1 \left(\dot{\gamma}\cdot\hat{r}\right)\left(\dot{\gamma}'\cdot\hat{r}\right)\ \Psi\left(\tfrac{|r|}{d_1}\right) + \beta_2\left(\dot{\gamma}\cdot\hat{r}^{\perp}\right)\left(\dot{\gamma}'\cdot\hat{r}^{\perp}\right)\ \Psi\left(\tfrac{|r|}{d_2}\right) \right\}$$

  - Where r = γ - γ'
- Phase field version analogous.

Low-curvature networks: results

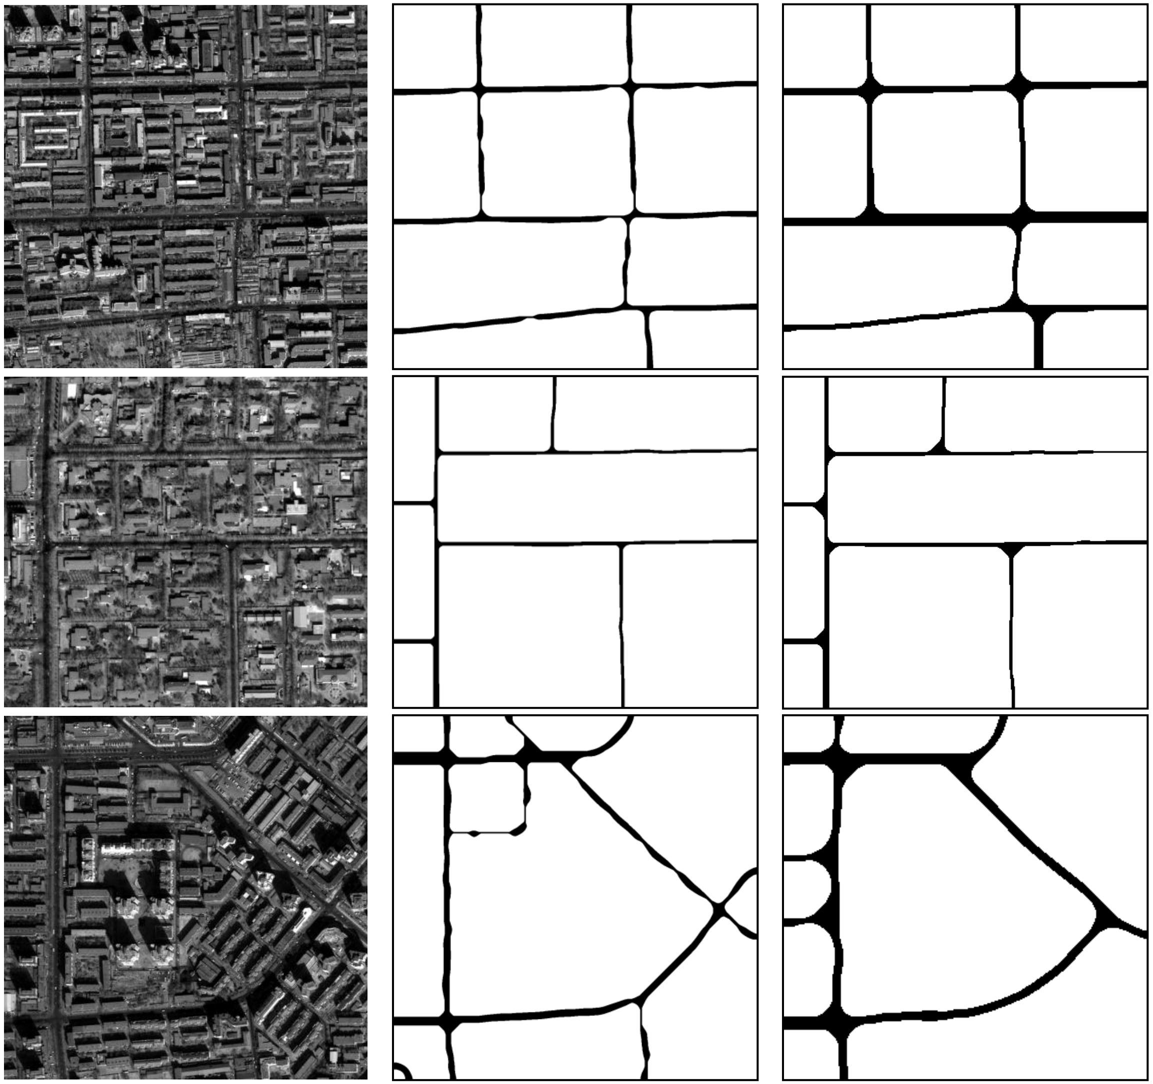

Q= TP / (TP + FP + FN)

(Quality measure based on GIS maps of Beijing.)

L0	L1	L2
0.8644		0.8086
0.9779	0.9543	
0.8237		0.6233

Low-curvature networks: results

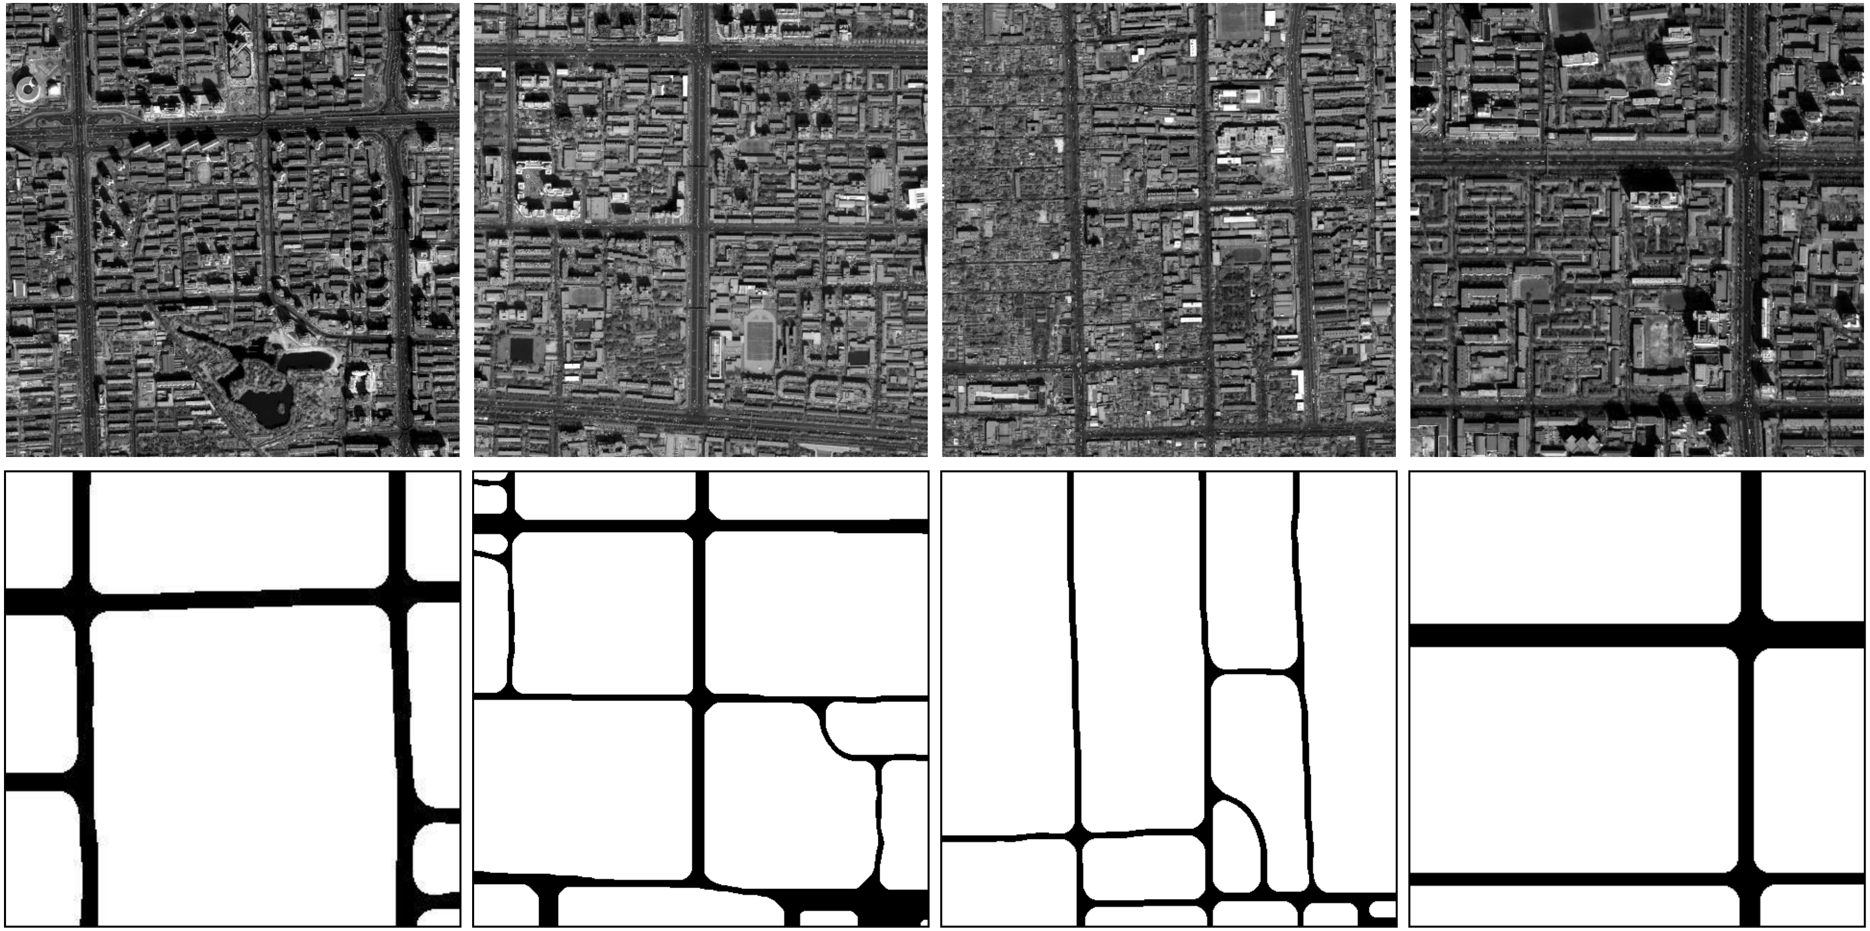

Directed networks

- Directed networks carry 'flow'.
  - Adding a conserved, 'fixed magnitude' vector field prolongs branches; stabilizes a range of widths; produces asymmetric junctions.

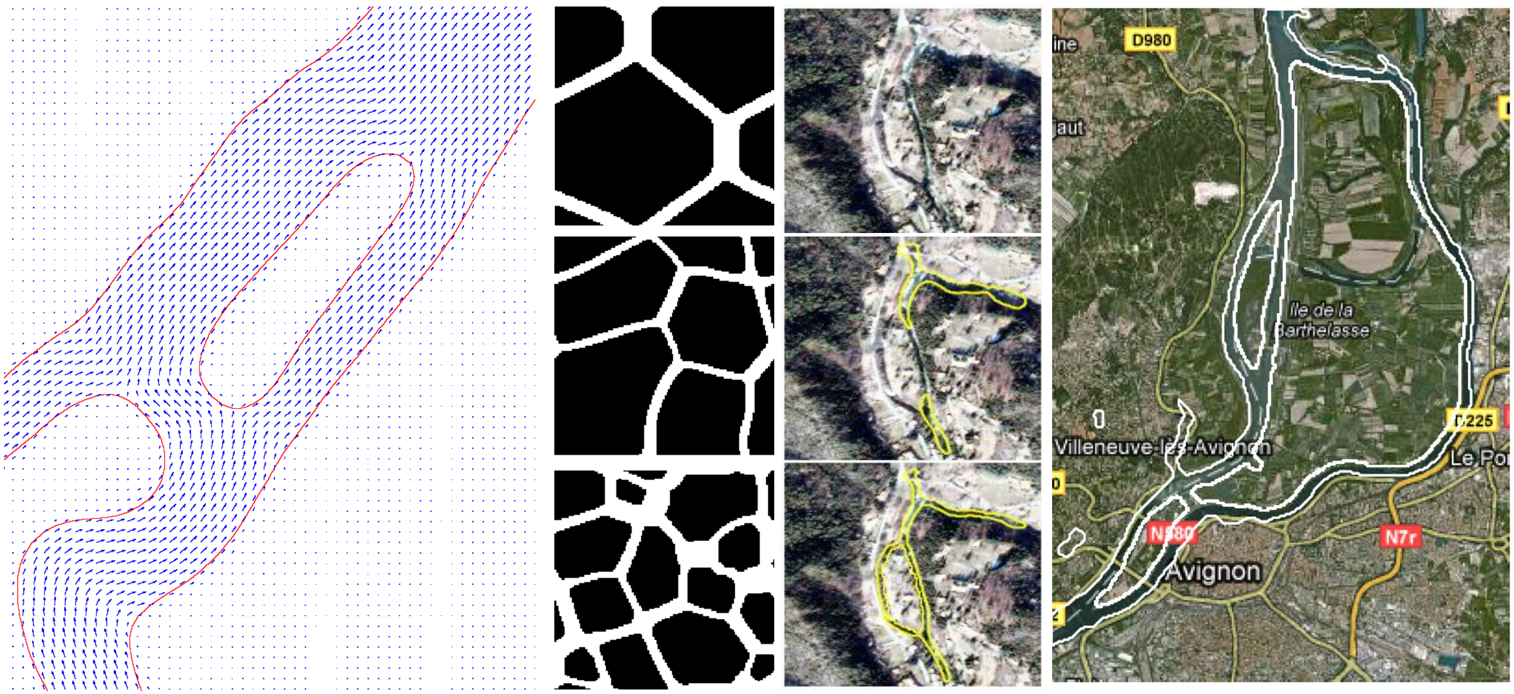

Directed networks: idea

- Directed networks carry 'flow'.
- Typical properties:
  - a large range of branch widths, but
  - changes of width are slow, except
  - at junctions, where flow is conserved.
- Model should reproduce these properties.
- Solution: use vector field to represent 'conserved flow'.

Directed networks: model

- Use two field variables:
  - The phase field, ϕ, to describe the region;
  - A vector field, v, to describe the flow.
- Desiderata:
  - $(\phi, |\mathrm{v}|) = (1, 1)$ and $(-1, 0)$ should be stable, corresponding to the region and its complement;
  - v should be smooth;
  - v should be parallel to the boundary;
  - v should have small divergence.

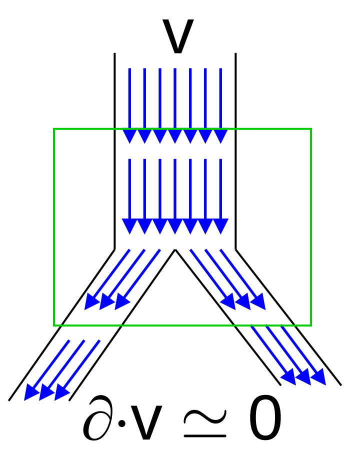

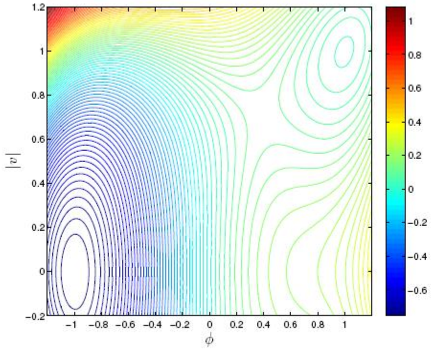

$$E_P(\phi, v) = E_{NL}(\phi) + \int \frac{D}{2}|\partial\phi|^2 + \frac{D_v}{2}(\partial \cdot v)^2 + \frac{L_v}{2}|\partial v|^2 + W(\phi, |v|)$$

Directed networks: geometry

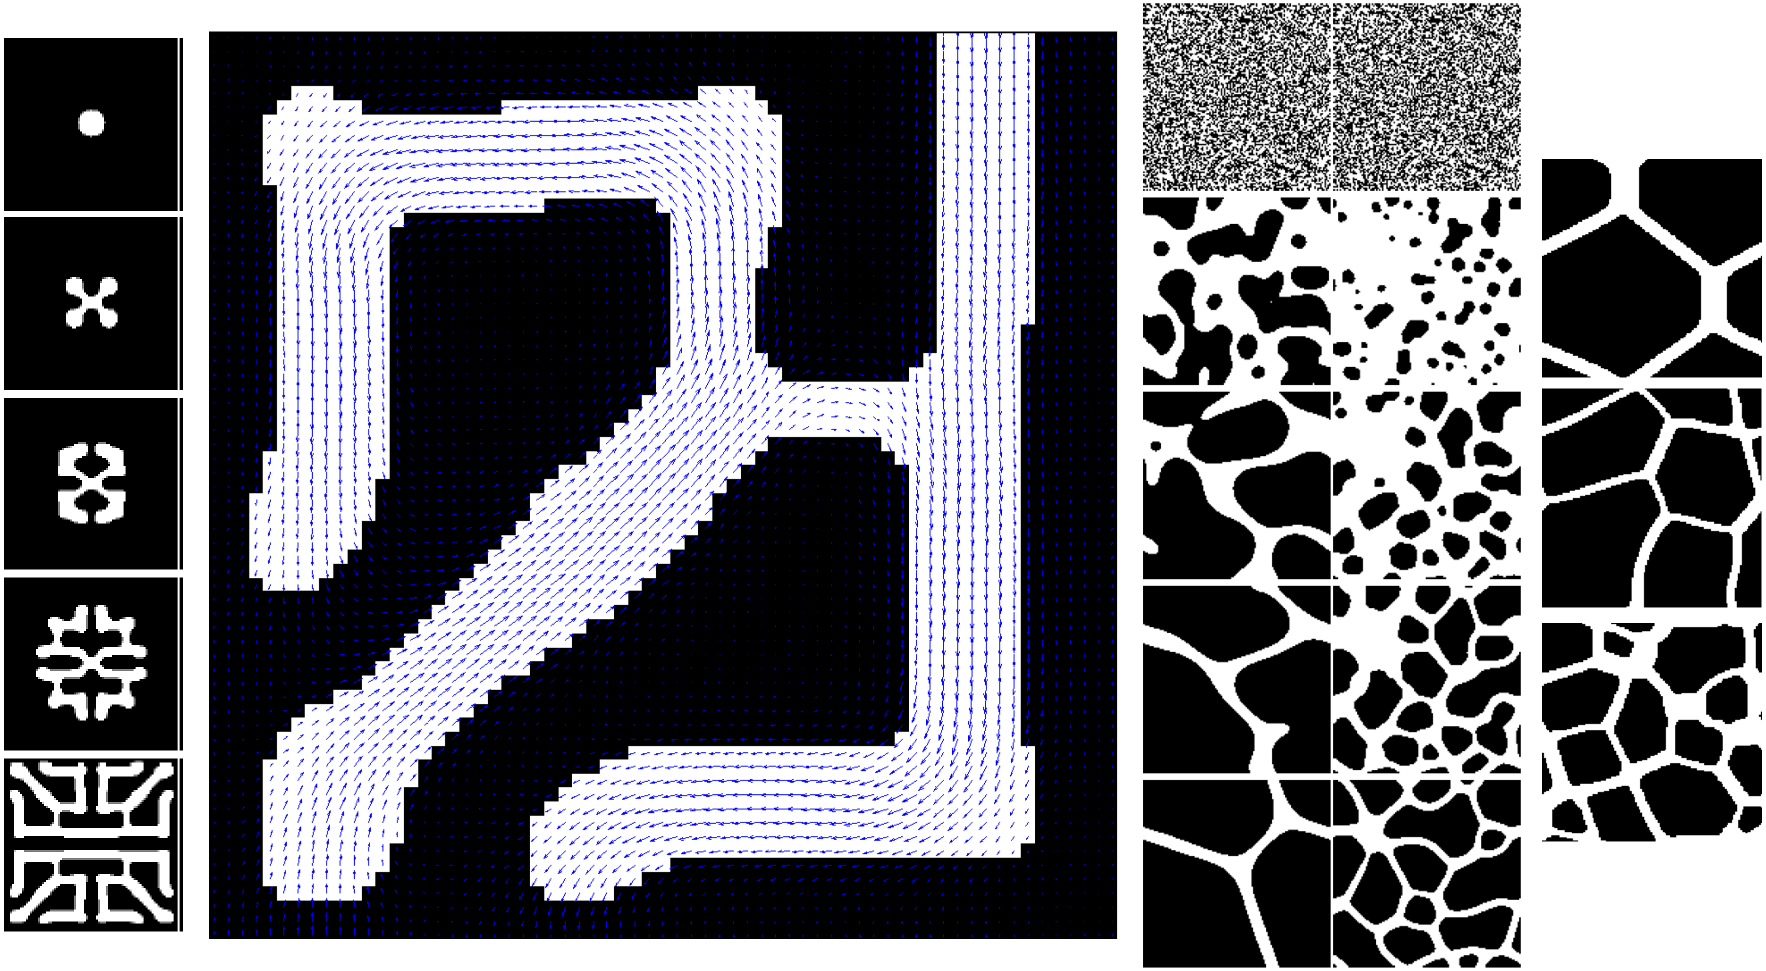

Evolution of v for gap closure

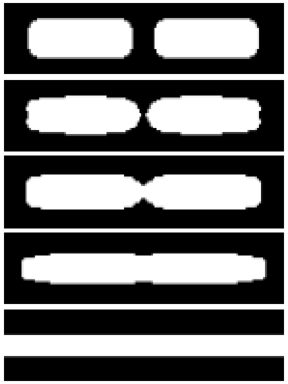

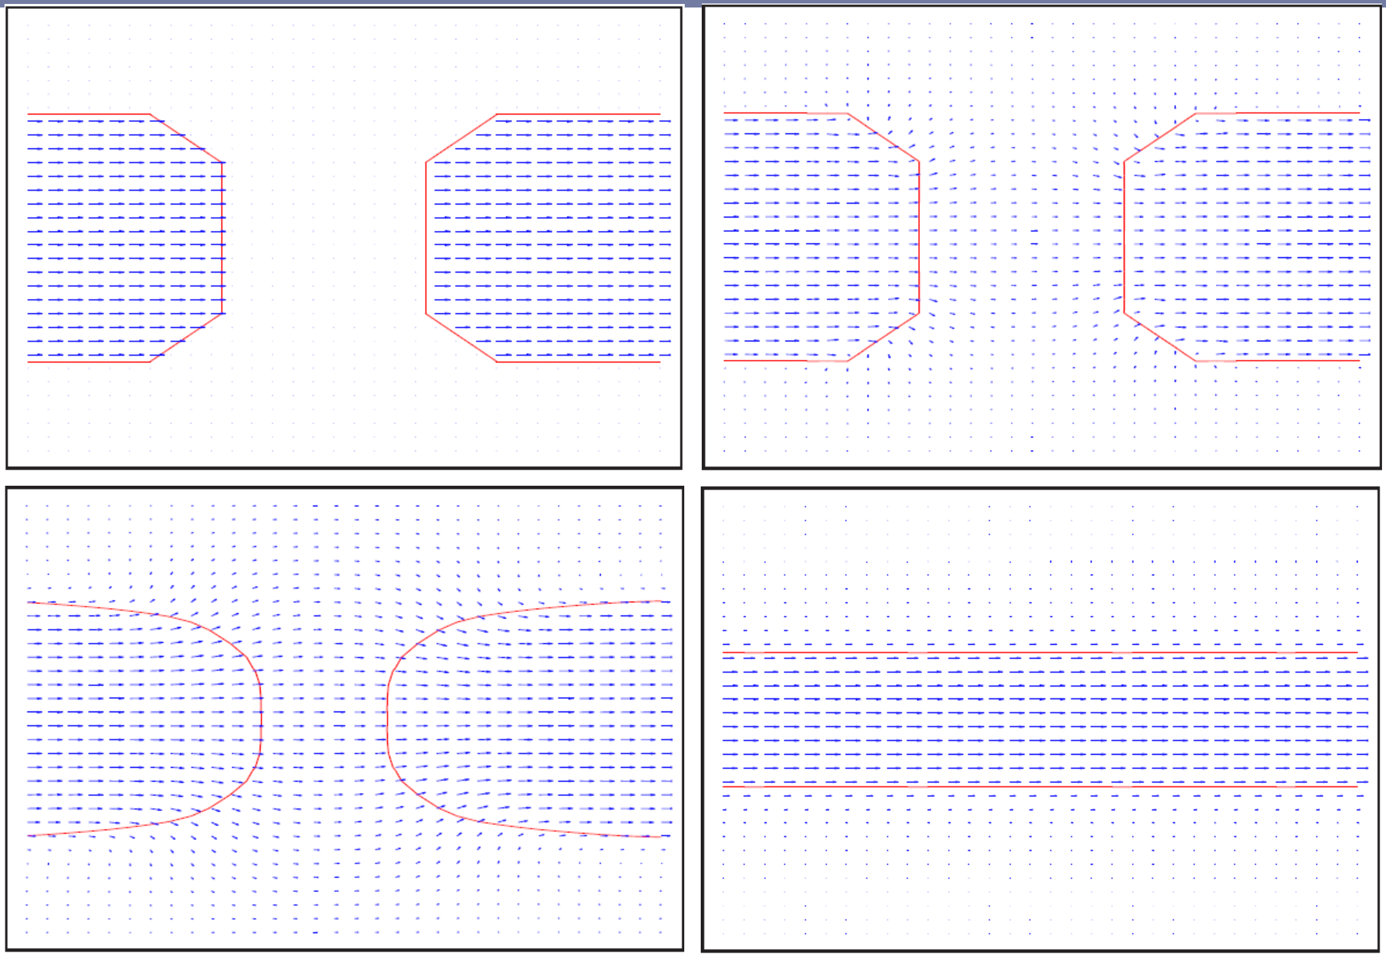

Directed networks: synthetic result

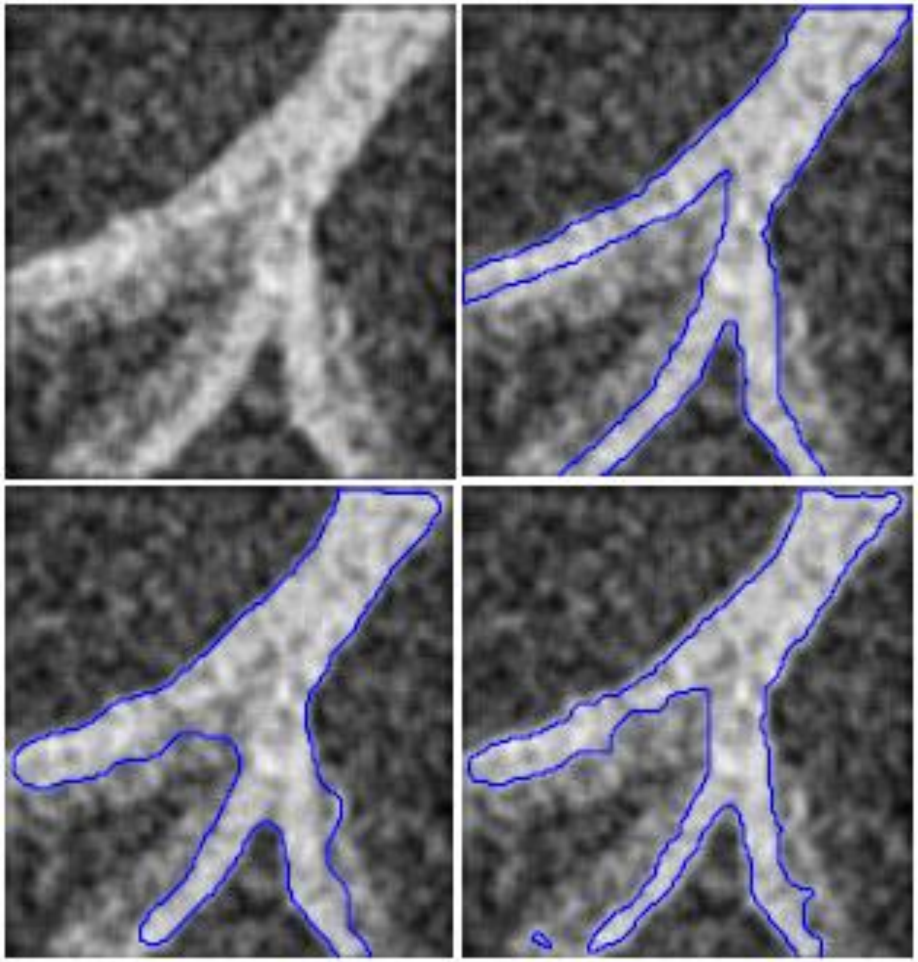

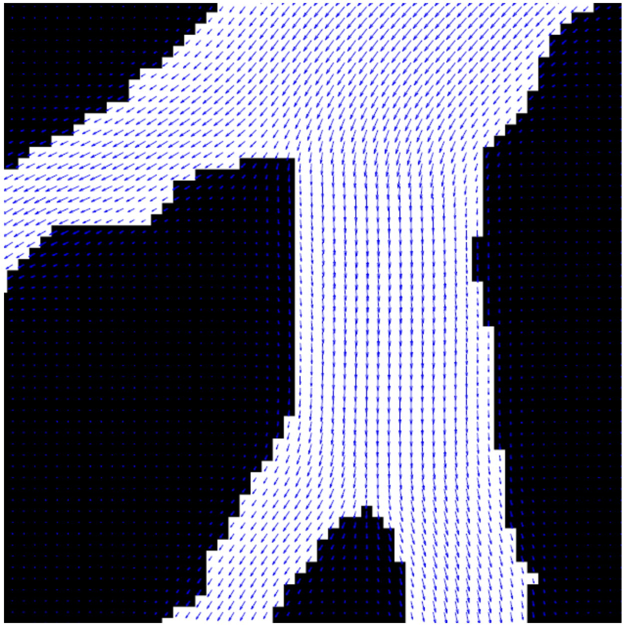

Directed networks: results

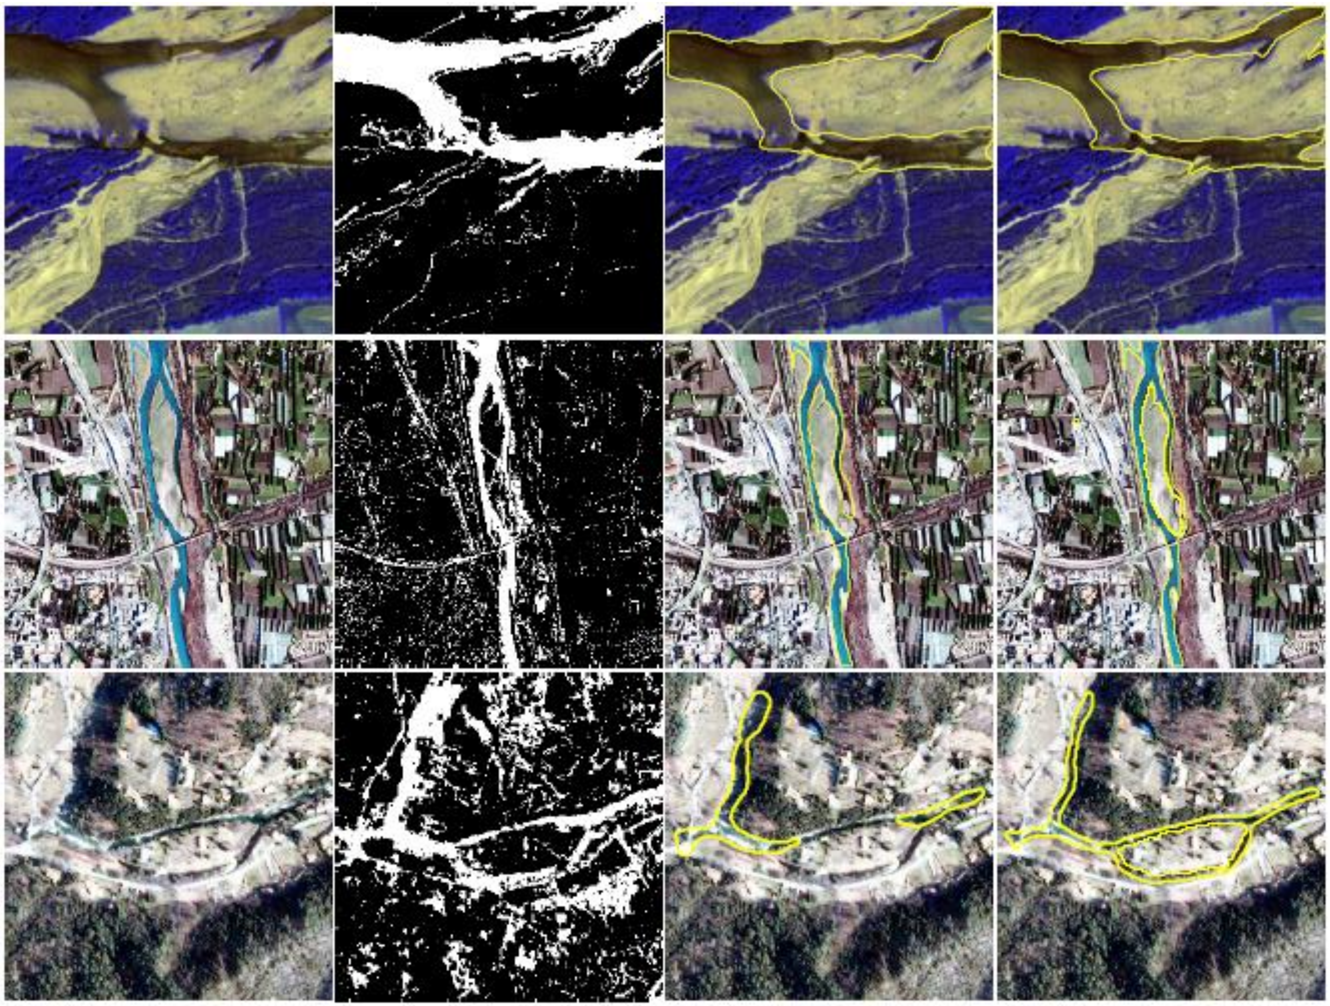

Multilayer binary field

- Removes a limitation by representing overlapping objects on different layers.
- Enables control of 'inter-object' interactions.

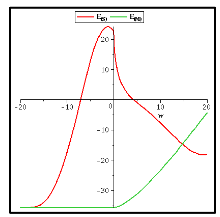

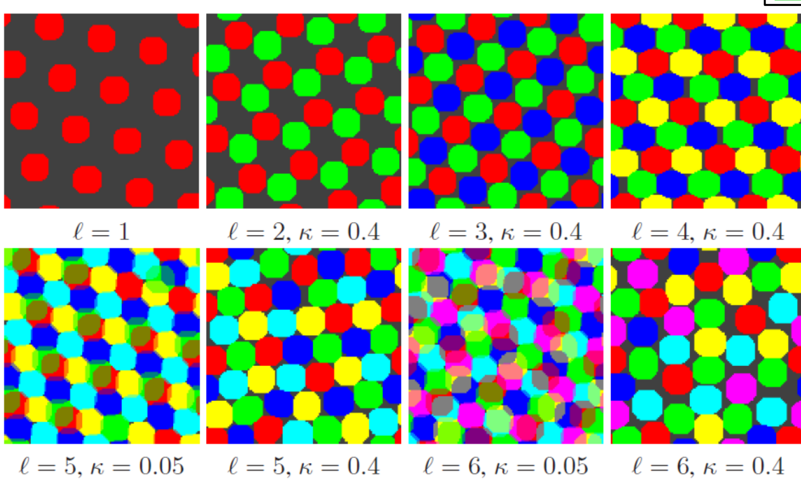

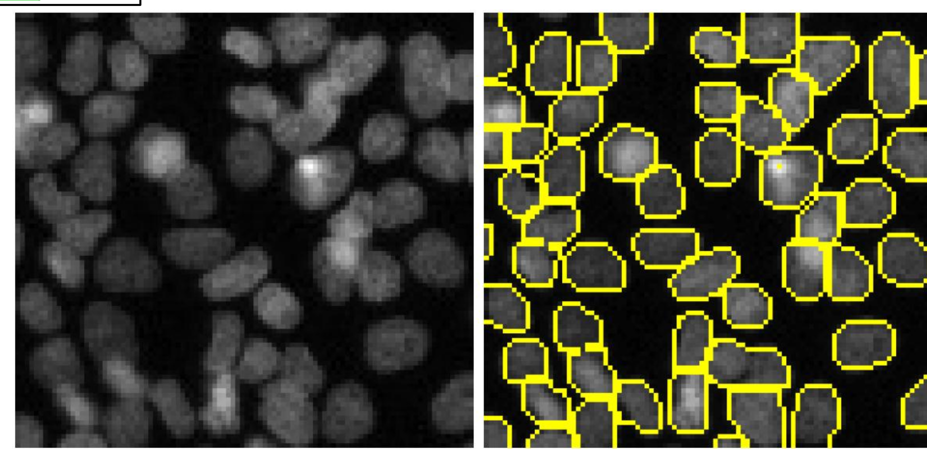

Complex shapes

- Parameter changes can render one or more Fourier perturbations of a stable circle unstable.
- Can higher-order effects stabilize them?
  - If so, it can lead to new stable shapes.
  - In principle, could model all star domains.

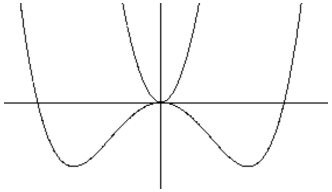

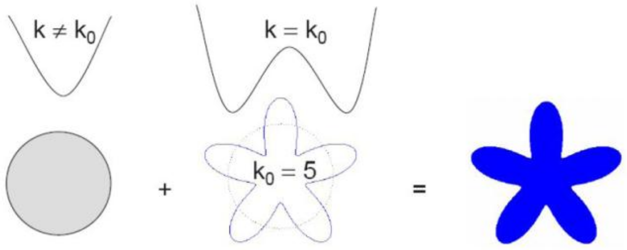

- Hierarchy of shapes of increasing complexity bifurcating from circle?

Summary

- Use explicit nonlocal interactions between boundary points to model shape while avoiding templates.
- Exploit equivalences between formulations:
  - Contour:
    - Intuitive; stability analysis for parameter estimation.
  - Phase field:
    - Linear space; lower complexity; easy implementation.
  - Binary field:
    - Facilitates sampling, hence learning; allows use of graph cut algorithms.

Future directions

- Models:
  - Complex shapes.
  - Learning models from examples.
  - Higher dimensions.
  - Multiscale: wavelets.
  - Analysis of binary field model.
  - Connection to point processes.
- Algorithms:
  - Efficient sampling (via wavelets?).
  - Analysis of the behaviour of graph cut algorithms.
- Many new applications:
  - Segmentation of cells; oil strata; CME; solar and ionospheric electron density reconstruction;...

Thank you

Very Vague Big Picture