Visualizing level lines and curvature in images

Pascal Monasse¹ (joint work with Adina Ciomaga² and Jean-Michel Morel²)

 ¹ IMAGINE/LIGM, École des Ponts ParisTech/Univ. Paris Est, France
 ² CMLA, ENS de Cachan, France



IPAM, Computational Photography and Intelligent Cameras Febuary 4-6, 2015



2 Motions by curvature in image processing

3 Numerical schemes





Plan

- Introduction
- 2 Motions by curvature in image processing
- 3 Numerical schemes
- 4 Tree of level lines
- 5 Experiments

Introduction: What is an image?



Computer scientist's point of view: array of pixel values: 0=black, 255=white

58	61	67	58	52
77	56	62	54	62
71	66	63	44	53
78	60	37	55	62
82	86	51	31	70

Mathematician's point of view.

 $u: \mathbb{R}^2 \to \mathbb{R}$ (but only samples are known). Which regularity?

Introduction: description of a shape by its curvatures

Idea: a shape is well described by its extrema of curvature:



Question: how to compute these curvatures in the presence of pixelization? Answer: We need to smooth the lines



Introduction

Introduction: images, contrast, level lines





P. Monasse (IMAGINE)

Introduction: topographic map



Introduction: importance of invariance

- There is a necessity to smooth the image
- The curvature is a Euclidean invariant, so our smoothing process should not disturb that invariance
- Commutation of the smoothing process with a rotation is a nice property and will be satisfied
- Actually, we will get even stronger invariance: commutation with special affine transforms

Plan

Introduction

2 Motions by curvature in image processing

- 3 Numerical schemes
- 4 Tree of level lines

5 Experiments

Fundamental geometric flows in image processing

Theorem (Alvarez, Guichard, Lions, Morel 93)

Causal, local, Euclidean and contrast covariant scale-spaces are all governed by a curvature equation of type:

$$\frac{\partial u}{\partial t} = |Du|G(\operatorname{curv}(u), t).$$

Two of these are the most interesting for image processing:

• the simplest one, the mean curvature motion:

$$\frac{\partial u}{\partial t} = |Du|\operatorname{curv}(u) \tag{MCM}$$

• the unique special affine covariant one, the affine curvature motion:

$$\frac{\partial u}{\partial t} = |Du|\operatorname{curv}(u)^{1/3} \tag{ACM}$$

Curvatures

We can define the curvature in two different ways:

• If u is C^2 and $Du(x_0) \neq 0$, the scalar curvature at x_0 is

$$\operatorname{curv}(u)(x_0) = \frac{u_{xx}u_y^2 - 2u_{xy}u_xu_y + u_{yy}u_x^2}{(u_x^2 + u_y^2)^{3/2}}(x_0)$$
(1)

• If x(s) is a C^2 curve parameterized by length s (|x'(s)| = 1), the vector curvature at $x_0 = x(s_0)$ is

$$\kappa(x_0) = x''(s_0) \tag{2}$$

Link: denote by x(s) the level line of u passing by x_0 ($u(x(s)) = u(x_0)$, $x(s_0) = x_0$), then

$$\kappa(x_0) = -\operatorname{curv}(u)(x_0)\frac{Du}{|Du|}(x_0).$$

This suggests two ways to compute curvature: 2D differential operator (1) or curvature of level line (2).

P. Monasse (IMAGINE)

Curve shortening

[Mackworth Mokhtarian 92]: The curvature motion $\frac{\partial x}{\partial t} = \kappa(x)$ (CS) can be implemented by applying the heat equation to each coordinate independently:

Theorem (Grayson 87)

If x(s,0) is a C^2 Jordan curve, then applying the intrinsic heat equation:

- For t > 0, x(s, t) is C^{∞} and satisfies (CS).
- For t > 0, x(., t) has a finite and non-increasing number of inflection points and curvature extrema.
- For $t \ge t_o$, x(.,t) is convex, and for $t \ge t_1$, x(.,t) is a point.

Algorithm: Input: polygon P. Output: smoothed polygon P. Iteratively:

- Sample P uniformly by length.
- Convolve coordinates of P by Gaussian kernel G_{σ} .

Affine shortening

Algorithm [Moisan 98]

- Break P into convex and concave parts
- Replace each part by the middle points of σ -chords originating from vertices of P.
- Concatenate the pieces of curves



Properties:

- Fast and simple algorithm (but numerically delicate)
- Special affine covariance (only inflection points, areas and middle points are involved)

P. Monasse (IMAGINE)

Plan

Introduction

2 Motions by curvature in image processing

3 Numerical schemes

4 Tree of level lines

5 Experiments

Algorithms on sets

Algorithm [Koenderink and van Doorn 86] Input: a closed subset X of \mathbb{R}^N

- Compute $u(x,t) = G_t \star \mathbf{1}_X(x)$
- Threshold at 1/2: $X_t = [u(., t) \ge 1/2].$

Problem: fusion of shapes that are close.

An improvement is the threshold dynamic shape: Algorithm [Merriman Bence Osher 92] Iteratively:

- Convolve $\mathbf{1}_X$ with G_σ
- Threshold at 1/2: $X \leftarrow [\mathcal{G}_{\sigma} \star \mathbf{1}_X \ge 1/2]$

This is a Gaussian-weighted median filter applied to a binary image.

Median filters

Algorithm: Iterated median filter. Iteratively, for every point x:

- Gather points y in a discrete neighborhood of x
- Put at x the median value of the discrete neighborhood

52	49	50		52	49	50
47	20	53	ightarrow 20,47,48,49,50,50,51,52,53 $ ightarrow$	47	50	53
51	48	50		51	48	50

Problem: the algorithm is blind to small curvatures, due to its discrete nature.

Theorem (Ishii 95)

Iterated weighted median filters on images converge to the MCM.

Finite difference schemes

 [Guichard Morel 97] uses the second derivative in the direction orthogonal to the gradient: |Du| curv(u) = u_{ξξ} with ξ = Du[⊥]/|Du|

$$(u_{\xi\xi})_{ij} = \begin{pmatrix} \lambda_3 & \lambda_2 & \lambda_4 \\ \lambda_1 & -4\lambda_0 & \lambda_1 \\ \lambda_4 & \lambda_2 & \lambda_3 \end{pmatrix} (\theta) \star u_{i,j} \text{ with } Du = |Du| \begin{pmatrix} \cos \theta \\ \\ \sin \theta \end{pmatrix}$$

and all $\lambda_i = 1/2$ when $|D_u| < \epsilon$ (half Laplacian)

• [Crandall Lions 96] Explicit scheme

$$u^{n+1} = u^n + \frac{dt}{h^2} \sum_{i=1}^N (u^n (x + ha(Du^n)e_i) + u^n (x - ha(Du^n)e_i) - 2u^n(x))$$

with (e_i) basis of \mathbb{R}^N and $a(p) = I - \frac{pp^T}{|p|^2 + \epsilon}$

Stack filters

Finite difference schemes do not commute with contrast changes. To recover the contrast covariance, we can "stack" the results on level sets: Algorithm [Stack filter]

- Extract upper level sets $X_{\lambda} = [u \ge \lambda]$
- Apply FDS to $\mathbf{1}_{X_{\lambda}}$
- Set X'_{λ} by threshold at 1/2 of the results
- Reconstruct by superposition: $u(x) \leftarrow \max\{\lambda : x \in X'_{\lambda}\}$

Anyway, other problems of FDSs persist, including lack of:

- Monotonicity, we always have slightly oscillatory solutions
- Euclidean (or affine) covariance, since they are grid dependent

Plan

Introduction

- 2 Motions by curvature in image processing
- 3 Numerical schemes
- 4 Tree of level lines

5 Experiments

Tree of level lines

Inclusion tree



Morse functions and critical points



Tree of inclusion of Jordan curves



Level lines at different levels:

Left: pixels still visible. Center: Jordan curves. Right: saddle points

20 / 39

Bilinear interpolated images

• Convolution of a lattice of Dirac masses by separable triangle function

$$arphi(x)arphi(y)$$
 with $arphi(x)=(1-|x|)^+$

• Between a square ABCD of data values, we have the equation

$$u(x, y) = axy + bx + cy + d$$

• Level line is intersection of the square with hyperbola

$$a(x-x_s)(y-y_s) = \lambda - \lambda_s$$

 If λ_s ∈ [min_{P=ABCD} u(P), max_{P=ABCD} u(P)], we have a saddle point and the piece of level line at λ_s is two orthogonal segments

Tree extraction algorithm

- Choose a number of levels, avoiding:
 - Initial values
 - Saddle values
- Two steps: follow level lines, then recover inclusion structure
- Step 1: Follow the level line inside square: given entry point, find exit point and sample the curve in between



- Store intersections with horizontal edges
- Step 2: Order intersection points at each horizontal line
- Set inclusion by parity argument:
 - Odd: get inside
 - 2 Even: go outside

Sampling in a dual pixel

• We can write

$$f(x,y) = a(x-x_S)(y-y_S) + \lambda_S$$
 $(a = u_{00} + u_{11} - u_{01} - u_{10}).$

When $a \neq 0$, level lines are equilateral hyperbola branches.

- Maximum curvature point is at $|x x_S| = |y y_S|$, we add it if inside the dual pixel.
- We can sample the hyperbola branch by writing y(x) (if $|y'| \le 1$) or x(y) (if $|x'| \le 1$) and sampling uniformly along x or y.



- \bullet Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.

- \bullet Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



- \bullet Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



- $\bullet\,$ Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



- $\bullet\,$ Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



- \bullet Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



- $\bullet\,$ Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



- \bullet Between two adjacent pixels, linear variation, thus a single point at level λ
- We follow the level line from dual pixel to dual pixel.



Particular cases

• Saddle point inside the dual pixel



Particular cases

• Saddle point inside the dual pixel





Particular cases

• Saddle point inside the dual pixel





Defect at saddle point



B goes through two saddle points: by a local decision, it is not possible to treat both according to formal definition.

Reconstruction from the tree

- Walk of the tree in preorder, and for each node representing a level line do (noting λ its level):
- Two steps: intersection with horizontal lines, then filling
- Step 1: similar to step 2 of extraction, intersection with each horizontal line and ordering inside each horizontal line, (xⁱ₁...xⁱ_{2Ni})
- Step 2: for all i and all $k = 0, \ldots, N_i 1$ do

• for all
$$j \in \mathbb{N} \cap [x_{2k+1}^i, x_{2k+2}^i]$$
:

pixel $(j, i) \leftarrow \lambda$

Geometric scheme curvature motion

Proceed in 3 steps:

- Extract tree of bilinear level lines
- 2 Let each level line evolve by curve shortening or affine shortening
- ③ Reconstruct image from the tree of (shortened) level lines



Curvature map

- Each level line is represented by a polygon.
- For each vertex P_i of the polygon, compute curvature as the inverse of the radius of the circumscribed circle of $P_{i-1}P_iP_{i+1}$
- Inside a pixel, average the curvatures of level lines passing through

Plan

- Introduction
- 2 Motions by curvature in image processing
- 3 Numerical schemes
- 4) Tree of level lines



Fattening effect



(a) Original and interpolation(c) Finite difference scheme

(b) Level lines shortening(d) Same with stack filter

P. Monasse (IMAGINE)

Level lines & curvature

Experiments

JPEG artifacts



Original image



LLS image



Original level lines



Shortened LL



Curvature map



After filtering

Experiments

Accurate mean curvature evolution



Left: detail of painting. Center: diff with LLS. Right: diff with stack FDS

P. Monasse (IMAGINE)

Level lines & curvature

Junctions



Curvature microscope



Experiments

Topographic map of a digital elevation model







Experiment

Textures







Conclusion

- The best scheme for motion by curvature of images proceeds in 3 steps:
 - Decompose the image into its level lines
 - Smooth each line independently
 - ③ Reconstruct from shortened level lines
- Such a scheme satisfies the covariance requirements (geometric and contrast)
- This is necessary to estimate reliably the curvature and avoid pixel artifacts
- This can be used as a microscope on the image: look at fine structures at any scale

- On inclusion tree and applications: Vicent Caselles and PM, *Geometric description of images as topographic maps* (Springer Lecture Notes in Mathematics) 2010
- FDSs in IPOL: http://www.ipol.im/pub/algo/cm_fds_mcm_amss/
- http://dev.ipol.im/~monasse/ipol_demo/cmmm_image_ curvature_microscope/