

Affine Invariant Detection: Edges, Active Contours, and Segments*

Peter J. Olver
University of Minnesota
Minneapolis, MN 55455

Guillermo Sapiro
Hewlett-Packard Labs.
Palo Alto, CA 94304

Allen Tannenbaum
University of Minnesota
Minneapolis, MN 55455

Abstract

In this paper we undertake a systematic investigation of affine invariant object detection. Edge detection is first presented from the point of view of the affine invariant scale-space obtained by curvature based motion of the image level-sets. In this case, affine invariant edges are obtained as a weighted difference of images at different scales. We then introduce the affine gradient as the simplest possible affine invariant differential function which has the same qualitative behavior as the Euclidean gradient magnitude. These edge detectors are the basis both to extend the affine invariant scale-space to a complete affine flow for image denoising and simplification, and to define affine invariant active contours for object detection and edge integration. The active contours are obtained as a gradient flow in a conformally Euclidean space defined by the image on which the object is to be detected. That is, we show that objects can be segmented in an affine invariant manner by computing a path of minimal weighted affine distance, the weight being given by functions of the affine edge detectors. The geodesic path is computed via an algorithm which allows to simultaneously detect any number of objects independently of the initial curve topology.

1 Introduction

Despite the extensive activity in recent years on invariant shape recognition algorithms — see [21] for a representative collection of papers on the topic — the corresponding problem of invariant detection of shapes has received considerably less attention. Some work along these lines has been reported in [30, 31], where the theory of geometric invariant smoothing of planar curves (boundaries of planar shapes) was initiated; see also [1]. In particular, using the methods of [1, 31], a shape can be smoothed in an affine invariant manner before the computation of invariant descriptors such as those reported in corresponding chapters in [21]. This work was partially extended for other groups and dimensions in [10, 24, 25, 32].

The purpose of this paper is to derive simple geometric object detectors which incorporate affine in-

variance. Invariant edge detectors should be the first step in a fully affine invariant system of object recognition. This way, the introduction of non-intrinsic noise to the system is reduced.

Two different affine edge detectors are presented, one based on the affine invariant scale-space [1, 31, 33] and the second one on a novel definition of *affine invariant gradient*. These affine invariant edge maps are then used to define *affine invariant active contours*, extending [8, 9, 16, 17]. The boundary of the scene objects are given by a geodesic or minimal weighted distance path in a Riemannian space. In contrast with previous approaches, distances in this space are affine invariants. These affine invariant edge maps are also used to extend the work in [1, 33] to obtain an affine invariant flow for image denoising and simplification.

To the best of our knowledge, besides the schemes here described, the only works addressing affine invariant detection and segmentation were performed by Ballester *et al.* [4] and by Lindeberg [19]. In [4] the authors presented a very nice affine invariant version of the Mumford-Shah segmentation algorithm. The work of Lindeberg is related to our definition of affine gradient, as will be explained in Section 3.

2 Affine edges from affine scale-space

We begin by deriving the first affine invariant edge detector. It is based on the theory of invariant scale-spaces developed in [1, 24, 25, 30, 31, 33]. We first introduce some preliminary notation. For planar column vectors, $X = (x_1, x_2)^T$, $Y = (y_1, y_2)^T \in \mathbb{R}^2$, we let $[X, Y] := x_1 y_2 - x_2 y_1$ be the area of the parallelogram spanned by X, Y . We also define $Y^\perp := (-y_2, y_1)^T$ by $[X, Y^\perp] = \langle X, Y \rangle$, where $\langle X, Y \rangle = x_1 y_1 + x_2 y_2$ denotes the usual inner product.

2.1 Planar curve evolution

The theory of planar curve evolution has been considered in a variety of fields; see [18, 29, 31] for references. One of the most important of such flows is derived when a planar curve deforms in the direction of the Euclidean normal, with speed equal to the Euclidean curvature. In [30] we extended this flow to the affine case. Formally, let $\mathcal{C}(p, t) : S^1 \times [0, \tau) \rightarrow \mathbb{R}^2$ be a family of smooth embedded closed curves in the plane (boundaries of planar shapes), where $p \in S^1$ parametrizes the curve, and $t \in [0, \tau)$ the family. This

*This work was partially supported by the National Science Foundation ECS 91-22106, DMS 92-04192 and DMS 95-00931, by the Air Force Office of Scientific Research F49620-94-1-00S8DEF, by the Army Research Office DAAH04-94-G-0054 and DAAH04-93-G-0332.

family of curves evolves according to ($\mathcal{C}(p, 0) = \mathcal{C}_0(p)$)

$$\frac{\partial \mathcal{C}(p, t)}{\partial t} = \frac{\partial^2 \mathcal{C}(p, t)}{\partial s^2}, \quad (1)$$

where $s(p) = \int_0^p [\mathcal{C}_p, \mathcal{C}_{pp}]^{1/3} dp$, is the *affine arc-length* ($[\mathcal{C}_s, \mathcal{C}_{ss}] \equiv 1$), i.e., the simplest¹ affine invariant parametrization [5]. We have shown that any simple and smooth convex curve evolving according to (1), converges to an ellipse [30]. Since the affine normal \mathcal{C}_{ss} exists just for non-inflection points, we formulated the natural extension of the flow (2) for non-convex initial curves in [31, 32]:

$$\frac{\partial \mathcal{C}(p, t)}{\partial t} = \begin{cases} 0, & p \text{ an inflection point,} \\ \mathcal{C}_{ss}(p, t), & \text{otherwise.} \end{cases} \quad (2)$$

This flow defines a geometric, affine invariant, multiscale representation of planar shapes [31]. In this case, we proved (see also [3]) that the curve first becomes convex, as in the Euclidean case, and after that it converges into an ellipse according to [30].

2.2 Euclidean image processing

Algorithms for image smoothing were developed based on the Euclidean and affine heat flows and related equations. In this section, we review a scheme for image processing which is related to the Euclidean heat flow. See [2, 26, 29, 33] for details and references.

Let $\Phi_0 : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ represent a gray-level image, where $\Phi_0(x, y)$ is the gray-level value. Being Φ_0 the initial condition, Alvarez *et al.* [2] proposed the flow

$$\frac{\partial \Phi}{\partial t} = \phi(\|G * \nabla \Phi\|) \|\nabla \Phi\| \operatorname{div} \left(\frac{\nabla \Phi}{\|\nabla \Phi\|} \right), \quad (3)$$

where G is a smoothing kernel, and $\phi(w)$ is a non-increasing function which tends to zero as $w \rightarrow \infty$. Details on the function of each term can be found in [2]. The evolution

$$\Phi_t = \|\nabla \Phi\| \operatorname{div} \left(\frac{\nabla \Phi}{\|\nabla \Phi\|} \right) = \kappa \|\nabla \Phi\| \quad (4)$$

is such that the level-sets of Φ move according to the Euclidean heat flow $\mathcal{C}_t = \kappa \mathcal{N}$ [2, 27], where κ is the Euclidean curvature and \mathcal{N} the unit normal. Equation (3) represents an anisotropic diffusion, extending the ideas first proposed by Perona and Malik [28].

2.3 Affine smoothing and edge detection

It is well-known in the theory of curve evolution, that if the velocity $\mathcal{V} = \mathcal{C}_t$ of the evolution is a geometric function of the curve, then the geometric behavior of the curve is affected only by the normal component of this velocity, i.e., by $\langle \mathcal{V}, \mathcal{N} \rangle$. Therefore, instead of looking at (2), we can consider an Euclidean-type formulation of it. In [30], we proved that the normal component of \mathcal{C}_{ss} is equal to $\kappa^{1/3} \mathcal{N}$. Since $\kappa = 0$ at

¹Simplest in this context refers to minimal order or minimal number of spatial derivatives.

inflection points, and inflection points are affine invariant, we obtain that

$$\mathcal{C}_t = \kappa^{1/3} \mathcal{N} \quad (5)$$

is geometrically equivalent to the affine heat flow (2). The affine invariant property of (5) was also pointed out by Alvarez *et al.* [1], based on a completely different approach. They proved that this flow is unique under certain conditions (uniqueness is obtained also from the results in [24]).

The process of embedding a curve in a 3D surface, and looking at the evolution of the level-sets, is frequently used for the digital implementation of curve evolution flows [27]. It is easy to show that the level-sets evolution equation corresponding to (5) is

$$\Phi_t = \kappa^{1/3} \|\nabla \Phi\| = (\Phi_y^2 \Phi_{xx} - 2\Phi_x \Phi_y \Phi_{xy} + \Phi_x^2 \Phi_{yy})^{1/3}. \quad (6)$$

This equation was used in [31] for the implementation of the novel affine invariant scale-space for planar curves mentioned before and in [1, 33] for image denoising. See [26] for details on the advantages of (6) over (4).

From the results in [1, 30, 31] the general behavior of a curve (or level-set) evolving according to the Euclidean or affine heat flows are very similar. The affine based flow will perform edge preserving anisotropic diffusion as well. Based on this, we obtain our first affine invariant edge detection scheme:

Definition 1 Let $\mathcal{S}_{edge}(t_0, t_1) := a\Phi(t_1) - b\Phi(t_0)$, such that $\Phi(\cdot)$ is the solution of (6) with initial datum $\Phi(0)$, $a, b \in \mathbb{R}^+$ and $t_1 > t_0 \geq 0$. $\mathcal{S}_{edge}(t_0, t_1)$ is denoted as the scale-space affine invariant edge detector.

Note that if $t_0 > 0$, noise is (efficiently and affine) removed before edges are computed. Varying t_0 and t_1 gives affine edges at different scales. An example of this flow is presented in Figure 1. The function \mathcal{S}_{edge} can be thresholded without affecting the affine invariance.



Figure 1: Example of \mathcal{S} .

3 Affine invariant gradient

Let $\Phi : \mathbb{R}^2 \rightarrow \mathbb{R}^+$ be a given image in the continuous domain. In order to detect edges in an affine invariant form, a possible approach is to replace the classical gradient magnitude $\|\nabla \Phi\|$, which is only Euclidean invariant, by an *affine invariant gradient*. For doing this, we have to look if we can use basic affine invariant descriptors that can be computed from Φ to find an expression that behaves like

$\|\nabla\Phi\|$. Using the classification developed in [22, 23], we found that the two basic independent affine invariant descriptors are $H := \Phi_{xx}\Phi_{yy} - \Phi_{xy}^2$, and $J := \Phi_{xx}\Phi_y^2 - 2\Phi_x\Phi_y\Phi_{xy} + \Phi_x^2\Phi_{yy}$. We should point out that there is no (non-trivial) first order affine invariant descriptor, and that all other second order differential invariants are functions of H and J . Therefore, the simplest possible affine gradient must be expressible as a function $\mathcal{F} = \mathcal{F}(H, J)$ of these two invariant descriptors.

If a curve \mathcal{C} is defined as the level-set of Φ , then the curvature of \mathcal{C} is given by $\kappa = \frac{J}{\|\nabla\Phi\|^3}$. Lindeberg [19] used J to compute edges in an affine invariant form, that is, $\mathcal{F} = J = \kappa \|\nabla\Phi\|^3$, which singles out edges as a combination of high gradient and high curvature of the level sets. Note that in general edges do not have to lie on a unique level-set. Here, by combining both H and J , we present a more general affine gradient approach. Since both H and J are second order derivatives of the image, the order of the affine gradient is not increased while using both invariants.

Definition 2 *The (basic) affine invariant gradient of a function Φ is defined by $\hat{\nabla}_{aff}\Phi := \left|\frac{H}{J}\right|$.*

Technically, since $\hat{\nabla}_{aff}\Phi$ is a scalar, it measures just the magnitude of the affine gradient, so our definition may be slightly misleading. However, an affine invariant gradient direction does not exist, since directions (angles) are not affine invariant, and so we are justified in omitting “magnitude” for simplicity. The justification for our definition is based on a (simplified) analysis of the behavior of $\hat{\nabla}_{aff}\Phi$ near edges in the image defined by Φ [26].

In order to avoid possible difficulties when the affine invariants H or J are zero, we replace $\hat{\nabla}_{aff}$ by a slight modification. Indeed, other combinations of H and J can provide similar behavior, and hence be used to define affine gradients. Here we present the general technique as well as a few examples.

Definition 3 *The normalized affine invariant gradient is given by $\nabla_{aff}\Phi := \sqrt{\frac{H^2}{J^2+1}}$*

The motivation comes from the form of the *affine invariant stopping term* analogue to ϕ in (3), which is now given by $\frac{1}{1+(\nabla_{aff}\Phi)^2} = \frac{J^2+1}{H^2+J^2+1}$, avoiding all difficulties where either H or J vanishes, and hence is the proper candidate for affine invariant edge detection. The analysis of this edge detector is given in [26]. An example is given in Figure 2, after thresholding.

4 Affine invariant image denoising

According to (3), a stopping term ϕ should be added to the directional derivative to stop diffusion across edges. Following the work in [33] (see also [1]), where the affine flow (6) is used as “directional diffusion,” we can replace the function ϕ in (3) by an




Figure 2: Example of the affine invariant gradient.

affine invariant edge stopping function ϕ_{aff} . Assume that $\phi_{aff} = \phi(w_{aff})$ where, as before, $\phi(w) \rightarrow 0$ when $w \rightarrow \infty$. We let now $w = w_{aff}$ be either one of the affine edge detectors defined above, i.e., \mathcal{S}_{edge} or $\nabla_{aff}(\Phi)$. This results in a completely affine invariant flow,

$$\Phi_t = \phi_{aff} \kappa^{1/3} \|\nabla\Phi\|. \quad (7)$$

This flow is tested in Figure 3. Note that since this type of flow, as well as the one proposed in [1], moves an image towards piecewise constant, its results can be used to simplify (segment) an image in an affine invariant fashion.



Figure 3: Example of the affine invariant image flow for image denoising and simplification. The original image is presented on the left, noisy on the middle, and the result of the affine invariant flow on the right.

5 Affine invariant active contours

In this section we derive the affine invariant active contours, based on the treatments of [8, 9, 16, 17]. It is important to note that after affine edges are computed locally based on the scale-space or affine gradient derive above, affine invariant fitting can be performed (see [26, 36] for the relevant references). In this work, the affine invariant integration is done by means of active contours.

5.1 Euclidean geodesic active contours

We present now the geodesic active contours derived in [8, 16, 17]. These schemes are based on previous work reported in [7, 15, 20, 35].

Let $\mathcal{C} = \mathcal{C}(p, t)$ be as before, with $\mathcal{C}(0, t) = \mathcal{C}(1, t)$, $\mathcal{C}'(0, t) = \mathcal{C}'(1, t)$. Consider the length functional $L(t) := \int_0^1 \|\mathcal{C}_p\| dp$. Then differentiating (i.e., taking the “first variation”), and integrating by parts, we find $L'(t) = -\int_0^1 \langle \frac{\partial \mathcal{C}}{\partial t}, \kappa \mathcal{N} \rangle dv$, where dv is the Euclidean arc-length. Now, in the standard way, we can define a norm $\|\cdot\|_{euc}$ on the (Fréchet) space of twice-differentiable closed curves in the plane. Indeed, the norm is given by the length of the curve,

$\|\mathcal{C}\|_{\text{euc}} := \int_0^1 \|\mathcal{C}_p\| dp = \int_0^L dv = L$. Thus the direction in which $L(t)$ is decreasing most rapidly is when \mathcal{C} satisfies the gradient flow $\mathcal{C}_t = \kappa\mathcal{N}$. The Euclidean heat flow is precisely a gradient flow.

We should note that this flow has arisen in the finding of closed geodesics on Riemannian manifolds (it can be defined with respect to any Riemannian metric), and the basic idea is that as long as it remains regular it will converge to a closed geodesic. The deep part is the regularity; for details see [12, 13, 14]. The active contours models which we are about to give are completely straightforward consequences of these principles.

We are now ready to formulate the geodesic active contours model from [8, 16, 17]. In [8], the model is derived from the principle of least action in physics, showing the mathematical relation between energy and curve evolution based snakes. In [16, 17], the model is derived immediately from curve shortening, and is compared to similar flows in continuum mechanics, in particular, phase transitions. Of course, the two obtained flows are mathematically identical and present active contours as geodesic computations. The basic idea is to change the ordinary Euclidean arc-length function $dv = \|\mathcal{C}_p\| dp$ along a curve $\mathcal{C}(p)$ by multiplying by a conformal factor $\phi(x, y) > 0$, which is assumed to be a positive, differentiable function. The resulting *conformal Euclidean metric* on \mathbb{R}^2 is given by $\phi dx dy$, and its associated arc length element is $dv_\phi = \phi dv = \phi \|\mathcal{C}_p\| dp$. As in ordinary curve shortening, we compute the corresponding gradient flow for $L_\phi(t) := \int_0^L \phi dv$. Taking the derivative and integrating by parts, we find that [8, 16, 17] $-L'_\phi(t) = \int_0^{L_\phi(t)} \langle \mathcal{C}_t, \phi\kappa\mathcal{N} - (\nabla\phi \cdot \mathcal{N})\mathcal{N} \rangle dv$, which means that the direction in which the L_ϕ perimeter is shrinking as fast as possible is given by

$$\frac{\partial \mathcal{C}}{\partial t} = \phi\kappa\mathcal{N} - (\nabla\phi \cdot \mathcal{N})\mathcal{N}. \quad (8)$$

As long as the flow remains regular, we will get convergence to a closed geodesic in the plane relative to the conformal Euclidean metric $\phi dx dy$. Regularity may be deduced from the classical curve shortening case.

To introduce the level-set formulation [27], let us assume that a curve \mathcal{C} is parametrized as a level-set of a function $u : [0, a] \times [0, b] \rightarrow \mathbb{R}$. Then, the level-set formulation of the steepest descent method says that solving the above geodesic problem starting from \mathcal{C}_0 amounts to searching for the steady state ($u_t = 0$) of the evolution equation $\frac{\partial u}{\partial t} = \|\nabla u\| \operatorname{div} \left(\phi \frac{\nabla u}{\|\nabla u\|} \right)$, with initial datum $u(0, x) = u_0(x)$. As in [7, 20], we may add an inflationary constant, to derive

$$\frac{\partial u}{\partial t} = \|\nabla u\| \operatorname{div} \left(\phi \frac{\nabla u}{\|\nabla u\|} \right) + \nu\phi \|\nabla u\|. \quad (9)$$

In the context of image processing, we take ϕ to be a stopping term depending on the image. In this case, notice that $\nabla\phi$ will look like a doublet near an edge.

The new gradient term directs the curve towards the boundary of the objects since $-\nabla\phi$ points toward the center of the boundary. The advantages of this model over previous ones is reported in [8, 9, 16, 17]. Existence, uniqueness and stability results for the gradient active contour model (9) were studied in [8, 9, 16, 17]. See [34, 37] for related approaches.

5.2 Affine invariant active contours

Based on the geodesic active contours and affine invariant edge detectors above, it is almost straightforward to define affine invariant gradient active contours. In order to carry this program out, we will first have to define the proper norm. Since affine geometry is defined only for convex curves [5], we will initially have to restrict ourselves to the (Fréchet) space \mathbf{C}_0 of thrice-differentiable convex closed curves in the plane. As above, let ds denote the affine arc-length, and $L_{\text{aff}} := \oint ds$ is the *affine length* [5]. On \mathbf{C}_0 , we define the affine metric $\|\mathcal{C}\|_{\text{aff}} := \int_0^1 \|\mathcal{C}(p)\|_a dp = \int_0^{L_{\text{aff}}} \|\mathcal{C}(s)\|_a ds$, where $\|\mathcal{C}(p)\|_a := [\mathcal{C}(p), \mathcal{C}_p(p)]$. Note that the area enclosed by \mathcal{C} is $A = \frac{1}{2} \|\mathcal{C}\|_{\text{aff}}$. Observe that $\|\mathcal{C}_s\|_a = [\mathcal{C}_s, \mathcal{C}_{ss}] = 1$, $\|\mathcal{C}_{ss}\|_a = [\mathcal{C}_{ss}, \mathcal{C}_{sss}] = \mu$, where μ is the *affine curvature*, i.e., the simplest non-trivial differential affine invariant. This makes the affine norm $\|\cdot\|_{\text{aff}}$ consistent with the properties of the Euclidean norm on curves relative to the Euclidean arc-length dv .

Assume now that $\phi_{\text{aff}} = \phi(w_{\text{aff}})$ is an affine invariant stopping term, based on the affine invariant edge detectors as above. Therefore, ϕ_{aff} behaves as the weight ϕ in L_ϕ , being now affine invariant. As in the Euclidean case, we regard ϕ_{aff} as an affine invariant conformal factor, and replace the affine arc length element ds by a conformal counterpart $ds_{\phi_{\text{aff}}} = \phi_{\text{aff}} ds$ to obtain the first possible functional for the affine active contours

$$L_{\phi_{\text{aff}}} := \int_0^{L_{\text{aff}}(t)} \phi_{\text{aff}} ds, \quad (10)$$

where $L_{\text{aff}} = \oint ds$ is now the affine length of our curve. The obvious next step is to compute the gradient flow corresponding to $L_{\phi_{\text{aff}}}$ in order to produce the affine invariant model. Unfortunately, as we will see, this will lead to an impractically complicated geometric contour model which involves four spatial derivatives.

The snake model which we will use comes from another (special) affine invariant, namely *area*. Let $\mathcal{C}(p, t)$ be a family of curves in \mathbf{C}_0 . The first variation of the area functional $A(t)$ is given by $A'(t) = -\int_0^{L_{\text{aff}}(t)} [\mathcal{C}_t, \mathcal{C}_s] ds$. Therefore the gradient flow which will decrease the area as quickly as possible relative to $\|\cdot\|_{\text{aff}}$ is exactly $\mathcal{C}_t = \mathcal{C}_{ss}$, which, modulo tangential terms, is equivalent to $\mathcal{C}_t = \kappa^{1/3}\mathcal{N}$, precisely the affine invariant heat equation studied in [30]! We define now the conformal area functional to be

$$A_{\phi_{\text{aff}}} := \int_0^1 [\mathcal{C}, \mathcal{C}_p] \phi_{\text{aff}} dp = \int_0^{L_{\text{aff}}(t)} [\mathcal{C}, \mathcal{C}_s] \phi_{\text{aff}} ds.$$

The first variation of this will turn out to be much simpler than that of $L_{\phi_{\text{aff}}}$ and will lead to an implementable geometric snake model.

Lemma 1 *Let $L_{\phi_{\text{aff}}}$ and $A_{\phi_{\text{aff}}}$ denote the conformal affine length and area functionals respectively.*

1. *The first variation of $L_{\phi_{\text{aff}}}$ is given by*

$$-\int_0^{L_{\text{aff}}(t)} [\mathcal{C}_t, (\nabla \phi_{\text{aff}})^\perp] ds + \int_0^{L_a(t)} \phi_{\text{aff}} \mu [\mathcal{C}_t, \mathcal{C}_s] ds. \quad (11)$$

2. *The first variation of $A_{\phi_{\text{aff}}}$ is given by*

$$-\int_0^{L_{\text{aff}}(t)} [\mathcal{C}_t, (\phi_{\text{aff}} \mathcal{C}_s + \frac{1}{2}[\mathcal{C}, (\nabla \phi)^\perp \mathcal{C}_s])] ds. \quad (12)$$

The affine invariance of the resulting variational derivatives follows from a general result governing invariant variational problems having volume preserving symmetry groups [25].

We now consider the corresponding gradient flows computed with respect to $\|\cdot\|_{\text{aff}}$. First, the flow corresponding to the functional $L_{\phi_{\text{aff}}}$ is $\mathcal{C}_t = \{(\nabla \phi_{\text{aff}})^\perp + \phi_{\text{aff}} \mu \mathcal{C}_s\}_s$. As before, we ignore the tangential components, which do not affect the geometry of the evolving curve, to obtain

$$\mathcal{C}_t = \phi_{\text{aff}} \mu \kappa^{1/3} \mathcal{N} + \langle ((\nabla \phi_{\text{aff}})^\perp)_s, \mathcal{N} \rangle \mathcal{N}. \quad (13)$$

The geometric interpretation of the affine gradient flow (13) minimizing $L_{\phi_{\text{aff}}}$ is analogue to that of the corresponding Euclidean geodesic active contours [26]. Unfortunately, this flow involves μ which makes it difficult to implement for affine invariant segmentation. (Possible techniques to compute μ numerically were recently reported in [6, 11].)

The gradient flow coming from the first variation of the modified area functional on the other hand is much simpler, $\mathcal{C}_t = (\phi_{\text{aff}} \mathcal{C}_s + \frac{1}{2}[\mathcal{C}, (\nabla \phi_{\text{aff}})^\perp] \mathcal{C}_s)_s$. Ignoring tangential terms (those involving \mathcal{C}_s) this flow is equivalent to $\mathcal{C}_t = \phi_{\text{aff}} \mathcal{C}_{ss} + \frac{1}{2}[\mathcal{C}, (\nabla \phi_{\text{aff}})^\perp] \mathcal{C}_{ss}$, which in Euclidean form gives the second possible affine snakes model:

$$\mathcal{C}_t = \phi_{\text{aff}} \kappa^{1/3} \mathcal{N} + 1/2 \langle \mathcal{C}, \nabla \phi_{\text{aff}} \rangle \kappa^{1/3} \mathcal{N}. \quad (14)$$

Notice that both models (13) and (14) were derived for *convex curves*, even though the flow (14) makes sense in the non-convex case. Formal results regarding existence of (14) can be derived following [1, 7, 8, 9, 16, 17].

Figure 4 illustrates simulations of these active contour models (the implementation is as in [8, 9, 16, 17, 20], based on the level-sets formulation [27]).




Figure 4: Example of the affine invariant snakes.

6 Concluding remarks

The problem of affine invariant detection was addressed in this paper. Two different affine invariant edge detectors were first discussed. One is obtained from weighted difference of images at different scales obtained from the affine invariant scale-space developed in [1, 30, 31]. The second one is obtained from a function which behaves like the Euclidean gradient magnitude, having, in addition, the affine invariance property. From the classification of invariants developed in [22, 23], this function is the simplest possible with this characteristic.

We then presented two models for affine invariant active contours, extending the results presented in [8, 9, 16, 17] for the Euclidean group. We showed that objects can be obtained as gradient flows relative to modified area and affine arc-length functionals. The induced metric is a function of the affine invariant edge maps. Therefore, objects are modeled as paths of minimal weighted affine distance. The same affine maps were used to extend the image flows in [1, 33], obtaining a complete affine invariant flow for image denoising and simplification.

Acknowledgments

Images in this paper are courtesy of Prof. Andrew Zisserman from Oxford University.

References

- [1] L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel, "Axioms and fundamental equations of image processing," *Arch. Rational Mechanics* **123**, 1993.
- [2] L. Alvarez, P. L. Lions, and J. M. Morel, "Image selective smoothing and edge detection by nonlinear diffusion," *SIAM J. Numer. Anal.* **29**, 1992.
- [3] S. Angenent, G. Sapiro, and A. Tannenbaum, "On the affine heat flow for non-convex curves," preprint, 1995.
- [4] C. Ballester, V. Caselles, and M. Gonzalez, "Affine invariant segmentation by variational method," *Technical Report, U. of Illes Balears*, 1994.
- [5] W. Blaschke, *Vorlesungen über Differentialgeometrie II*, Verlag Von Julius Springer, Berlin, 1923.
- [6] E. Calabi, P. J. Olver, and A. Tannenbaum, "Affine geometry, curve flows, and invariant numerical approximations," TR, Department of EE, University of Minnesota, June 1995.

- [7] V. Caselles, F. Catta, T. Coll, F. Dibos, "A geometric model for active contours," *Numerische Mathematik* **66**, pp. 1-31, 1993.
- [8] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International Journal of Computer Vision*, to appear. Also in *Proc. ICCV*, Cambridge, MA, June 1995.
- [9] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert, "Minimal surfaces: A three dimensional segmentation approach," *Technion-Israel EE Pub.* **973**, June 1995. Also in *Proc. ECCV '96*.
- [10] O. Faugeras, "On the evolution of simple curves of the real projective plane," *Comptes rendus de l'Acad. des Sciences de Paris* **317**, pp. 565-570, September 1993.
- [11] O. Faugeras and R. Keriven, "Scale-spaces and affine curvature," *Proc. Europe-China Workshop on Geometrical Modeling and Invariants for Computer Vision*, R. Mohr and C. Wu (Eds.), pp. 17-24, 1995.
- [12] M. Gage and R. S. Hamilton, "The heat equation shrinking convex plane curves," *J. Differential Geometry* **23**, pp. 69-96, 1986.
- [13] M. Grayson, "The heat equation shrinks embedded plane curves to round points," *J. Differential Geometry* **26**, pp. 285-314, 1987.
- [14] M. Grayson, "Shortening embedded curves," *Annals of Mathematics* **129**, pp. 285-314, 1989.
- [15] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: Active contour models," *International Journal of Computer Vision* **1**, pp. 321-331, 1988.
- [16] S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. Yezzi, "Gradient flows and geometric active contour models," *Proc. ICCV*, Cambridge, June 1995.
- [17] S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. Yezzi, "Conformal curvature flows: From phase transitions to active vision," to appear *Arch. Rational Mechanics and Analysis*.
- [18] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, "Shapes, shocks, and deformations, I," *International Journal of Computer Vision* **15**, pp. 189-224, 1995.
- [19] T. Lindeberg, *Scale-Space Theory in Computer Vision*, Kluwer, 1994.
- [20] R. Malladi, J. A. Sethian and B. C. Vemuri, "Shape modeling with front propagation: A level set approach," *IEEE Trans. on PAMI* **17**, 1995.
- [21] J. L. Mundy and A. Zisserman (Eds.), *Geometric Invariance in Computer Vision*, MIT Press, 1992.
- [22] P. J. Olver, *Applications of Lie Groups to Differential Equations*, second ed., Springer-Verlag, New York, 1993.
- [23] P. J. Olver, *Equivalence, Invariants, and Symmetry*, Cambridge University Press, Cambridge-UK, 1995.
- [24] P. J. Olver, G. Sapiro, and A. Tannenbaum, "Differential invariant signatures and flows in computer vision: A symmetry group approach," in [29].
- [25] P. J. Olver, G. Sapiro, and A. Tannenbaum, "Invariant geometric evolutions of surfaces and volumetric smoothing," *SIAM J. of Appl. Math.*, to appear.
- [26] P. J. Olver, G. Sapiro, and A. Tannenbaum, "Affine invariant edge maps and active contours," *Geometry Center Technical Report 90*, University of Minnesota, October 1995, submitted.
- [27] S. J. Osher and J. A. Sethian, "Fronts propagation with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics* **79**, pp. 12-49, 1988.
- [28] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Machine Intell.* **12**, 1990.
- [29] B. Romeny, Editor, *Geometry Driven Diffusion in Computer Vision*, Kluwer, 1994.
- [30] G. Sapiro and A. Tannenbaum, "On affine plane curve evolution," *Journal of Functional Analysis* **119:1**, pp. 79-120, 1994.
- [31] G. Sapiro and A. Tannenbaum, "Affine invariant scale-space," *International Journal of Computer Vision* **11:1**, pp. 25-44, 1993.
- [32] G. Sapiro and A. Tannenbaum, "On invariant curve evolution and image analysis," *Indiana University Mathematics Journal* **42:3**, 1993.
- [33] G. Sapiro, A. Tannenbaum, Y. L. You, and M. Kaveh, "Experiments on geometric image enhancement," *First IEEE-International Conf. on Image Proc.*, Austin-Texas, November 1994.
- [34] J. Shah, "Recovery of shapes by evolution of zero-crossings," *Technical Report, Math. Dept. Northeastern Univ*, Boston MA, 1995.
- [35] D. Terzopoulos, A. Witkin, and M. Kass, "Constraints on deformable models: Recovering 3D shape and nonrigid motions," *Artificial Intelligence* **36**, pp. 91-123, 1988.
- [36] I. Weiss, "Geometric invariants and object recognition," *International Journal of Computer Vision*, pp. 207-231, 1993.
- [37] R. T. Whitaker, "Algorithms for implicit deformable models," *Proc. ICCV'95*, Cambridge, June 1995.