

Affine Invariant Detection: Edge Maps, Anisotropic Diffusion, and Active Contours*

Peter J. Olver

Department of Mathematics
University of Minnesota
Minneapolis, MN 55455

Guillermo Sapiro

Department of Electrical Engineering
University of Minnesota
Minneapolis, MN 55455

Allen Tannenbaum

Department of Electrical Engineering
University of Minnesota
Minneapolis, MN 55455

*This work was partially supported by the National Science Foundation ECS 91-22106, DMS 92-04192, DMS 95-00931, by the National Science Foundation Learning and Intelligent Systems Program (LIS), by the Air Force Office of Scientific Research AF/F49620-94-1-0058DEF, by the Army Research Office DAAH04-94-G-0054, DAAH04-93-G-0332, AFOSR-MURI, by the Office of Naval Research ONR-N00014-97-1-0509, by the Office of Naval Research Young Investigator Award, and by the Presidential Early Career Awards for Scientists and Engineers (PECASE).

Abstract

In this paper we undertake a systematic investigation of affine invariant object detection and image denoising. 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 maps are derived as a weighted difference of images at different scales. We then introduce the *affine gradient* as an affine invariant differential function of lowest possible order with qualitative behavior similar to the Euclidean gradient magnitude. These edge detectors are the basis for the extension of 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 gradient path is computed via an algorithm which allows to simultaneously detect any number of objects independently of the initial curve topology. Based on the same theory of affine invariant gradient flows we show that the affine geometric heat flow is minimizing, in an affine invariant form, the area enclosed by the curve.

Key Words: Affine invariant detection, denoising, segmentation, affine scale-space, affine gradient, active contours, gradient flows, geodesics, invariant metrics.

1 Introduction

Despite the extensive activity in recent years on invariant shape recognition algorithms — see [51] 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 [65, 66], 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, 66], a shape can be smoothed in an affine invariant manner before the

computation of invariant descriptors such as those reported in corresponding chapters in [51]; see for example [26]. This work was partially extended for other groups and dimensions in [55, 58, 68]. Motivated by this work, efforts in the derivation of a projective invariant smoothing process was begun in the work of [24, 25]. In [69], the work was extended to the invariant smoothing of shapes without shrinking.

The purpose of this paper is to derive simple geometric object detection and image denoising schemes which incorporate affine invariance. By this we mean that if two images are related by an affine and photometric transformation (see next section for the exact definitions), the images and objects boundaries or edge maps obtained by the algorithms are also related by the same transformation. These invariant edge detection and denoising algorithms should constitute the first step in a fully affine invariant system of object recognition. The second step, if necessary, would be the affine smoothing mentioned above (for curve-based systems), to be followed by the computation of affine invariant descriptors. Incorporating the affine invariant detection and denoising schemes presented here in object recognition systems will reduce the “algorithmic” noise introduced by using non-invariant methods. Indeed, smoothing an image with a non-invariant algorithm before computing its affine curvature, will introduce “artificial noise” to the computation, and when comparing the signatures of different shapes, will be very difficult to know if the differences correspond to the shapes or were introduced by the algorithm. This is illustrated in Figures 1 and 2. Two images related by an affine transformation are used to detect edges, both in an Euclidean (gradient followed by thresholding) and affine form (with the algorithm described in Section 3). Note how in the Euclidean case, entire edge segments are missing from one of the affine related images. In the affine case, some differences appear, due to discretization (the pixel grid is not affine invariant), but the relation between the edges of the two affine related images is much more consistent than in the Euclidean case. Further and more detailed examples are presented later in the paper. The next step in a shape recognition system is to recognize the objects bounded by the edges that have been detected by the segmentation procedure. For this purpose, a new approach, incorporating either Euclidean or affine invariance, based on the general concept of a differential invariant signature curve, was recently proposed in [9, 10]. These papers also derive a new, fully affine-invariant numerical algorithm for computing

the affine invariants required to uniquely characterize the curve up to affine motion, and hence, in conjunction with the algorithms in this paper, form the basis of a fully curve-based affine-invariant object detection and recognition procedure. The schemes here described can in general be used as pre-processing steps for the algorithms described in [51]. For example, the denoising algorithm can be used in combination with the moment-based approach recently introduced in [79]. We should note that the primary goal of this work is to present a framework to derive invariant object detection and image denoising schemes, and not just to improve existing state-of-the-art algorithms. Since those schemes are basically not affine invariant, they are not comparable with the techniques here introduced (although we do demonstrate the importance of affine invariant detection in this paper).

Two different affine edge detectors are presented. The first one is derived by weighted differences of images obtained as solutions of the affine invariant scale-space developed in [1, 66, 67, 70]. The second one is given by the simplest affine invariant function which shows behavior similar to the magnitude of the Euclidean gradient. (By “simplest” we mean the minimal number of spatial derivatives.) This affine gradient is derived from the classification of differential invariants described in [53, 54]. These affine invariant edge maps are then used to define *affine invariant active contours*, extending the work in [12, 13, 36, 37, 83]. (See also the book [6] for a collection of papers on the classical theory of snakes.) The active contours are therefore used to integrate the local information obtained by the affine edge detectors. The boundaries of the scene objects are determined as lying at the bottom of a potential well relative to a geometrically defined energy functional tailored to the affine geometry of the plane weighted by an affine invariant stopping function. (This was the philosophy of [12, 36, 37] in the Euclidean case.) Hence, in contrast with previous approaches, distances in the resulting metric space are affine invariants, and are based on the affine edge maps and classical affine differential geometry [7]. The boundaries of the desired features, are computed via a gradient descent flow based on a new notion of affine invariant curve metric introduced in this paper as well. This curve metric not only makes it possible to define a gradient flow for the active contour computation but also shows that the affine invariant heat flow introduced in [1, 2, 65, 66] is minimizing, in an affine invariant form, the area enclosed by the curve. The affine invariant edge maps are also used to extend the work in [1, 67, 70]




Figure 1: Illustration of the effect of affine invariant edge maps. First row: A disk smoothed by Gaussian filtering and an affine transformation of it. Second row: Results of Euclidean gradient computation followed by thresholding. Note how entire segments are missing. Third row: Results of the affine gradient followed by thresholding. Besides discretization and intrinsic errors (see Section 3 for details), the two edge maps are much closer to each other than the ones in the second row.




Figure 2: Illustration of the effect of affine invariant edge maps. First row: Two images related by an affine transformation. Second row: Results of Euclidean gradient computation followed by thresholding. Note how entire segments are missing. Third row: Results of the affine gradient followed by thresholding. Besides discretization and intrinsic errors (see Section 3 for details), the two edge maps are much closer to each other than the ones in the second row.

to obtain a completely 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.* [5] and by Lindeberg [43]. In [5] the authors presented a very nice affine invariant version of the Mumford-Shah [50] segmentation algorithm. The work of Lindeberg is related to our definition of affine gradient, as will be explained in Section 3. The framework here described for affine invariant edge detectors and active contours can use or be combined with other scale-spaces such as the one proposed in [44].

This paper is organized as follows. In Section 2 we review the affine scale space introduced in [1, 65, 66] and based on it we present a possible affine invariant edge detector. In Section 3 we describe the affine invariant gradient approach following the classification in [53, 54]. Section 4 extends the results of [2, 67, 70] for affine invariant image denoising and simplification. In Section 5, after the affine invariant curve metric is introduced, the affine invariant active contours are presented. Concluding remarks are given in Section 6.

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, 55, 58, 65, 66, 67, 70]. We start by a brief review of the relevant results on planar curve evolution, following with the level-sets flows that will lead to the affine edge detector.

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 $[X, Y] := x_1y_2 - x_2y_1$ 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 $\langle X, Y \rangle = x_1y_1 + x_2y_2$ denotes the usual Euclidean inner product.

The geometric transformations here considered are of the form

$$\hat{X} := \mathcal{A}X + V,$$

where X represents the boundary of the (planar) shape (vector in \mathbb{R}^2), V is a translation vector, and \mathcal{A} a 2×2 unimodular matrix (unit determinant). This transformation implies that the camera is relatively far from the object. For general affine transformations of the image plane, the theory here presented gives relative invariants instead of absolute ones.

We also allow for changes in the image intensity (photometric changes, [79]). If $\Phi : [x, y] \rightarrow \mathbb{R}^+$ stands for the image, then we consider *scaling* and *offset* transformations, that is,

$$\widehat{\Phi}(x, y) = s\Phi(x, y) + o,$$

where s and o are constants. Since the theory developed in this paper is based on level-sets curvature flows (which are morphological invariant by definition) and partial derivatives of the image Φ , it is clear that the operations here defined are (at least relative) invariant with respect to this photometric transformation. In this paper, we use the term *affine invariant* to refer to operations that are invariant to the geometric transformation above, as well as to the photometric one when dealing with images. This model of image transformations is the same as the one assumed in the recently published paper [79] for moment-based recognition, and is commonly used in the investigation of pattern recognition systems.

2.1 Planar curve evolution

The theory of planar curve evolution has been considered in a variety of fields such as differential geometry [30, 31, 55, 58, 65], theory of parabolic equations [3], numerical analysis [16, 59], computer vision [25, 26, 38, 39, 41, 42, 60, 62, 64, 66, 76, 82], viscosity solutions [15, 23, 72], phase transitions [34], and image processing [2, 62, 63, 67, 70]. 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.

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 S^1 denotes the unit circle, $p \in S^1$ parametrizes the curve, and $t \in [0, \tau)$ parametrizes the family. Assume that this family of curves evolves

according to the evolution equation

$$\begin{cases} \frac{\partial \mathcal{C}(p, t)}{\partial t} = \frac{\partial^2 \mathcal{C}(p, t)}{\partial v^2} = \kappa(p, t) \mathcal{N}(p, t), \\ \mathcal{C}(p, 0) = \mathcal{C}_0(p). \end{cases} \quad (1)$$

Here

$$v(p) = \int_0^p \| \mathcal{C}_p \| dp$$

is the *Euclidean arc-length*¹ ($\| \mathcal{C}_v \| \equiv 1$), $\kappa = [\mathcal{C}_v, \mathcal{C}_{vv}]$ the *Euclidean curvature*, and \mathcal{N} the *inward unit normal* [33]. The flow given by (1) is called the *Euclidean shortening flow*, since the curve perimeter shrinks as fast as possible when the curve evolves according to it [31]. Gage and Hamilton [30] proved that a simple and smooth convex curve evolving according to (1), converges to a round point. Grayson [31] proved that an embedded planar curve converges to a simple convex one when evolving according to (1), and so any embedded curve in the plane converges to a round point via the flow given in (1).

The flow (1), which is non-linear since v is a time-dependent curve parametrization, is also called the *Euclidean geometric heat flow*. It has been utilized for the definition of a geometric, Euclidean invariant, multiscale representation of planar shapes [1, 38, 39]. As we will show below, this flow is also important for image enhancement applications. Note that in contrast with the classical heat flow given by $\mathcal{C}_t = \mathcal{C}_{pp}$, the Euclidean geometric heat flow is intrinsic to the curve, that is, only depends on the geometry of the curve and not on its parametrization. This flow, as well as the other presented below, can be used also to solve the standard shrinking problem of smoothing processes [69].

Recently, we introduced a new curve evolution equation, the *affine geometric heat flow* [65, 66]:

$$\begin{cases} \frac{\partial \mathcal{C}(p, t)}{\partial t} = \frac{\partial^2 \mathcal{C}(p, t)}{\partial s^2}, \\ \mathcal{C}(p, 0) = \mathcal{C}_0(p), \end{cases} \quad (2)$$

¹We will consistently use v to denote Euclidean arc-length, reserving s for the affine arc length which is the main focus of this paper.

where

$$s(p) = \int_0^p [\mathcal{C}_p, \mathcal{C}_{pp}]^{1/3} dp, \quad (3)$$

is the *affine arc-length* ($[\mathcal{C}_s, \mathcal{C}_{ss}] \equiv 1$), i.e., the simplest² affine invariant parametrization [7]. \mathcal{C}_{ss} is called the *affine normal* [33]. In contrast with the Euclidean version, the affine arc-length is based on area, and not on length (recall that $[\mathcal{C}_p, \mathcal{C}_{pp}]$ is the area between \mathcal{C}_p and \mathcal{C}_{pp}). This is clear since length is not affine invariant, whereas area is the simplest geometric affine invariant. This evolution is the affine analogue of equation (1), and admits affine invariant solutions, i.e., if a family $\mathcal{C}(p, t)$ of curves is a solution of (2), the family obtained from it via unimodular affine mappings, is a solution as well. We have shown that any simple and smooth convex curve evolving according to (2), converges to an ellipse [65]. Since the affine normal \mathcal{C}_{ss} exists just for non-inflection points, we formulated the natural extension of the flow (4) for non-convex initial curves in [66, 68]:

$$\frac{\partial \mathcal{C}(p, t)}{\partial t} = \begin{cases} 0, & p \text{ an inflection point,} \\ \mathcal{C}_{ss}(p, t), & \text{otherwise,} \end{cases} \quad (4)$$

together with the initial condition $\mathcal{C}(p, 0) = \mathcal{C}_0(p)$. The flow (4) defines a geometric, affine invariant, multiscale representation of planar shapes. Indeed, in [66], we proved that this flow satisfies all the required properties of (morphological) scale-space such as causality and order preservation. In this case, we proved (see also [4]) that the curve first becomes convex, as in the Euclidean case, and after that it converges into an ellipse according to the results of [65]. See [66] for a number of explicit examples of planar shape smoothing.

We should also add that in [68], we give a general method for writing down invariant flows with respect to any Lie group action on \mathbb{R}^2 . The idea is to consider the evolution given by $\mathcal{C}_t = \mathcal{C}_{rr}$ where r is the group invariant arc-length. This was formalized, together with uniqueness results, in [55], and extended to surfaces in [58]. Results for the projective group were recently reported in [24, 25].

Recently, algorithms for image smoothing were developed based on the Euclidean and affine geometric heat flows and related equations, and this will be the subject of the next

²Simplest in this context refers to minimal order or minimal number of spatial derivatives.

section. An excellent volume of papers edited by Bart ter Haar Romeny [62] has appeared which is dedicated to such geometry driven diffusion processes. We refer the interested reader to this book for many more details about the subject as well as a rather complete set of references.

2.2 Euclidean image processing

In this section, we review a number of algorithms for image processing which are related to the Euclidean shortening flow (1). The algorithms were developed in continuous spaces, and tested on digital computers by very accurate and stable numerical implementations. These numerical implementations were developed by the various authors for their specific algorithm. Only the basic concepts of the algorithms are given here. For more details, see the appropriate references given below.

In general, $\Phi_0 : \mathbb{R}^2 \rightarrow \mathbb{R}^+$ represents a gray-level image, where $\Phi_0(x, y)$ is the gray-level value. The algorithms that we describe are based on the formulation of partial differential equations, with Φ_0 as initial condition. The solution $\Phi(x, y, t)$ of the differential equation gives the processed image.

Rudin *et al.* [63] presented an algorithm for noise removal, based on the minimization of the total first variation of Φ , i.e.,

$$\int_{Image} \|\nabla\Phi\| \, dx dy. \tag{5}$$

The minimization is performed under certain constraints and boundary conditions (zero flow on the boundary). The constraints they employed are zero mean value and given variance σ^2 of the noise, but other constraints clearly can be considered as well. More precisely, if the noise is additive, the constraints are given by

$$\int_{Image} \Phi \, dx dy = \int_{Image} \Phi_0 \, dx dy, \quad \int_{Image} (\Phi - \Phi_0)^2 \, dx dy = 2\sigma^2. \tag{6}$$

Note that κ , the Euclidean curvature of the level-sets, is exactly the Euler-Lagrange derivative of this total variation. Then, for the minimization of (5) with the constraints given by

(6), the following gradient-descent flow is obtained: ³

$$\Phi_t = \kappa - \lambda(\Phi - \Phi_0), \quad (7)$$

and the solution to the variational problem is given when Φ achieves steady state. The level-sets curvature κ may be computed via standard formulas for curves defined by implicit functions. The quantity σ is used in the computation of λ . The authors computed λ from the steady state solution ($\Phi_t = 0$).

Alvarez *et al.* [2] described an algorithm for image selective smoothing and edge detection. In this case, the image evolves according to

$$\Phi_t = \phi(\| G * \nabla \Phi \|) \|\nabla \Phi\| \operatorname{div} \left(\frac{\nabla \Phi}{\|\nabla \Phi\|} \right), \quad (8)$$

where G is a smoothing kernel (for example, a Gaussian), and $\phi(w)$ is a nonincreasing function which tends to zero as $w \rightarrow \infty$. Note that

$$\|\nabla \Phi\| \operatorname{div} \left(\frac{\nabla \Phi}{\|\nabla \Phi\|} \right)$$

is equal to $\Phi_{\xi\xi}$, where ξ is the direction normal to $\nabla \Phi$. Thus it diffuses Φ in the direction orthogonal to the gradient $\nabla \Phi$, and does not diffuse in the direction of $\nabla \Phi$. This means that the image is being smoothed on both sides of the edge, with minimal smoothing at the edge itself. Note that the evolution

$$\Phi_t = \|\nabla \Phi\| \operatorname{div} \left(\frac{\nabla \Phi}{\|\nabla \Phi\|} \right) = \kappa \|\nabla \Phi\| \quad (9)$$

is such that the level-sets of Φ move according to the Euclidean shortening flow given by equation (1) [2, 59]. Finally, the term

$$\phi(\| G * \nabla \Phi \|)$$

is used for the enhancement of the edges. If $\|\nabla \Phi\|$ is “small”, then the diffusion is strong. If $\|\nabla \Phi\|$ is “large” at a certain point (x, y) , this point is considered as an edge point, and the diffusion is weak.

³Note that the gradient-descent flow is not the only possible technique to minimize the total-variation. We concentrate on this approach here since the flow obtained is related to those relevant to our algorithms.

In summary, equation (8) gives anisotropic or edge preserving (and enhancement) diffusion, extending the ideas first proposed by Perona and Malik [61]. The equation looks like the level-sets of Φ are moving according to (1), with the velocity value “altered” by the function $\phi(w)$.

2.3 Affine smoothing and edge detection

As we saw in the previous section, there is a close relationship between the curve evolution flow (1), and recently developed image enhancement and smoothing algorithms (see equation (9)). In this section we employ the affine shortening flow (4) for a similar purpose.

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 $\langle \mathcal{V}, \mathcal{N} \rangle$. The tangential velocity component only affects the parametrization of the evolving curve [22, 66]. Therefore, instead of looking at (4), we can consider a Euclidean-type formulation of it. In [65], we proved that the normal component of \mathcal{C}_{ss} is equal to $\kappa^{1/3}\mathcal{N}$. This is very easy to prove, since

$$\mathcal{C}_{ss} = \frac{\mathcal{C}_{vv}}{(ds/dv)^2} + f(\kappa, \kappa_v)\mathcal{C}_v, \quad \mathcal{C}_{vv} = \kappa\mathcal{N}, \quad \mathcal{C}_v = \mathcal{T},$$

and

$$ds = [\mathcal{C}_v, \mathcal{C}_{vv}]^{1/3} dv = \kappa^{1/3} dv. \tag{10}$$

Since $\kappa = 0$ at inflection points, and inflection points are affine invariant, we obtain that the evolution given by

$$\mathcal{C}_t = \kappa^{1/3}\mathcal{N}, \tag{11}$$

is geometrically equivalent to the affine shortening flow (4). Then the trace (or image) of the solution to (11) is affine invariant.

It is important to note that the affine invariant property of (11) 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 [55]). Moreover, they give an extensive characterization of PDE based multiscale analysis, and

remarked that the flows (1) and (11) are well-defined also for non-smooth curves, using the theory of viscosity solutions [20]. This is also true for the corresponding image flows, where the level-sets deform according to the geometric heat flows [15, 23] (see below). The existence of the Euclidean and affine geometric heat flows for Lipschitz functions is obtained from the results in [3, 4] as well. These results on extensions of the flows to non-smooth data are fundamental for all image processing applications, since images, being discrete, are non-smooth. The results prove that the flows are mathematically correct (well-defined and stable).

We proceed now to show how the affine curve flow (11) can be extended to process images. The technique 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 [59]. Let us consider now what occurs when the level-sets of Φ evolve according to (11). It is easy to show that the corresponding evolution equation for Φ is given by

$$\Phi_t = \kappa^{1/3} \|\nabla\Phi\| = (\Phi_y^2\Phi_{xx} - 2\Phi_x\Phi_y\Phi_{xy} + \Phi_x^2\Phi_{yy})^{1/3}. \quad (12)$$

This equation was used in [66] for the implementation of the novel affine invariant scale-space for planar curves mentioned in the Introduction. It was also used in [1, 67, 70] for image denoising (here we extend those flows; see Section 4). Note again that, based on the theory of viscosity solutions, equations (9) and (12) can be analyzed even if the level-sets (or the image itself), are non-smooth; see [1, 15, 20, 23]. This flow is well-posed and stable. The maximum principle holds, meaning that the flow is smoothing the image.

If we compare (9) with (12), we observe that *the denominator is eliminated*. This not only makes the evolution (11) affine invariant [1, 66], it also makes the numerical implementation more stable [59]. The $1/3$ power is the *unique one* which eliminates this denominator. This is of course an important advantage of the affine flow over the Euclidean one in image processing. Moreover, for high curvatures, $\kappa^{1/3}$ is smaller than κ , which further prevents sharp regions from moving. Finally, since the symmetry group (the affine group) of (12) is much larger than that of equation (9) (the Euclidean heat flow), more structure is preserved up to a higher degree of smoothing. This phenomenon has been observed, for example, in Niessen *et al.* [52] in which elliptical structures of MRI images of the brain were preserved

up to a very high degree of smoothing using equation (12).

We are now ready to define the first affine invariant edge map. Note that from the results in [1, 65, 66], the affine based flow derived from (11) will perform edge preserving anisotropic diffusion as in the Euclidean case. Based on this, we obtain the first affine invariant edge detection scheme, from the following function (see also [19]):

Definition 1 *Let*

$$\mathcal{S}_{edge}(t_0, t_1) := a\Phi(t_1) - b\Phi(t_0), \tag{13}$$

*such that $\Phi(\cdot)$ is the solution of (12) 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.*⁴

From the results above, we first of all obtain that \mathcal{S}_{edge} is affine-invariant. (If $a \geq b$, then \mathcal{S}_{edge} is also positive, since the flow satisfies the maximum principle [1, 66].) \mathcal{S}_{edge} then gives an affine invariant edge map of the original image $\Phi(0)$. Note that if $t_0 > 0$, noise is (efficiently) removed before edges are computed. Varying t_0 and t_1 gives affine edges at different scales. Examples of this flow are presented in Figure 3 for different values of t_1 (t_0 is left fixed in these experiments). The implementation of the affine invariant scale-space in this example, as well as all the other examples in this paper, is based on the popular level-sets numerical scheme [59] and details can be found for example in [59, 66]. The function \mathcal{S}_{edge} can be followed by a pre-defined threshold without affecting the affine invariance. Note that we use thresholding here only for ease of visualization. When using \mathcal{S}_{edge} for further computations, as in the affine active contours described in following sections, this thresholding is not used, since it can remove important information. The affine invariant edge detection in Figure 3 gives similar results to those obtained with basic edge detection algorithms reported in the literature, with the additional property of being affine invariant. As can be observed from the figures, edges are not perfectly detected (compare with Figure

⁴After this paper was concluded [56, 57], we received a copy of [40], where the function \mathcal{S}_{edge} is also defined. As pointed out by the author when defining this function, his work was inspired by our remarks in [70], where we reported the possibility of defining an affine invariant edge map from the affine invariant geometric heat flow. The interested reader is referred to [40] for applications of this edge map as well as nice examples further supporting the importance of affine invariance in this type of computation.

4). In particular, at clean areas, where the curvature at the edge is exactly zero, it will take a number of iterations for the edge to “move,” allowing only then its detection via \mathcal{S}_{edge} . In spite of this, \mathcal{S}_{edge} can be used, as we see later, as a basic edge map for invariant active contours or other edge integration algorithms, as well as for affine invariant image denoising.




Figure 3: Examples of the scale-space based affine invariant edge detector. The original image is presented on the left, followed by results of $\mathcal{S}(t_0, t_1)$ for t_0 fix and different values of t_1 . For the time step $\Delta t = 0.05$, the (non-optimized) parameters used in this Figure are: $t_0 = 10\Delta t$, $t_1 = 15\Delta t$, $20\Delta t$, $25\Delta t$, $a = b = 12$.

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 $\|\nabla\Phi\| = \sqrt{\Phi_x^2 + \Phi_y^2}$, which is only Euclidean invariant, by an *affine invariant gradient*. By this we mean that we search for an affine invariant function from \mathbb{R}^2 to \mathbb{R} that has, at image edges, values significantly different from those at flat areas, and such that this values are preserved, at corresponding image points, under affine transformations. In order

to accomplish this, we have to verify that we can use basic affine invariant descriptors which can be computed from Φ in order to find an expression that (qualitatively) behaves like $\|\nabla\Phi\|$. Using the classification developed in [53, 54], we found that the two basic independent affine invariant descriptors are ⁵

$$H := \Phi_{xx}\Phi_{yy} - \Phi_{xy}^2, \quad J := \Phi_{xx}\Phi_y^2 - 2\Phi_x\Phi_y\Phi_{xy} + \Phi_x^2\Phi_{yy}.$$

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.

The differential invariant J is related to the Euclidean curvature of the level-sets of the image. Indeed, 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 [43] used J to compute corners and edges, in an affine invariant form, that is,

$$\mathcal{F} := J = \kappa \|\nabla\Phi\|^3.$$

This singles out image structures with a combination of high gradient (edges) and high curvature of the level-sets (corners). Note that in general edges and corners 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 the equation⁶*

$$\widehat{\nabla}_{\text{aff}} \Phi := \left| \frac{H}{J} \right|. \tag{14}$$

Technically, since $\widehat{\nabla}_{\text{aff}} \Phi$ is a scalar (a map from \mathbb{R}^2 to \mathbb{R}), it measures just the magnitude of the affine gradient, so our definition may be slightly misleading. However, an affine

⁵Note that the simplest Euclidean invariant differential descriptor is exactly $\|\nabla\Phi\|$, which is enough to formulate a basic Euclidean invariant edge detector.

⁶This functions is not well defined at perfect straight lines, that is, straight lines that have exactly the same gray value. Invariants of additional complexity need to be investigated to avoid this problem.

invariant gradient direction does not exist, since directions (angles) are not affine invariant, and so we are justified in omitting “magnitude” for simplicity.

Note also that if photometric transformations are allowed, then $\widehat{\nabla}_{\text{aff}} \Phi$ becomes only a relative invariant. In order to obtain an absolute invariant, we can use, for example, the combination $\frac{H^{3/2}}{J}$. Since in this case, going from relative to absolute invariants is straightforward, we proceed with the development of the simpler function $\widehat{\nabla}_{\text{aff}} \Phi$.

The justification for our definition is based on a (simplified) analysis of the behavior of $\widehat{\nabla}_{\text{aff}} \Phi$ near edges in the image defined by Φ . Near the edge of an object, the gray-level values of the image can be (ideally) represented via $\Phi(x, y) = f(y - h(x))$, where $y = h(x)$ is the edge, and $f(t)$ is a slightly smoothed step function with a jump near $t = 0$. Straightforward computations show that, in this case,

$$H = -h'' f' f'', \quad J = -h'' f'^3.$$

Therefore

$$H/J = f''/f'^2 = (-1/f')'.$$

Clearly H/J is large (positive or negative) on either side of the object $y = f(x)$, creating an approximation of a zero crossing at the edge.⁷ This is due to the fact that $f(x) = \text{step}(x)$, $f'(x) = \delta(x)$, and $f''(x) = \delta'(x)$. (We are omitting the points where $f' = 0$). Therefore, $\widehat{\nabla}_{\text{aff}} \Phi$ behaves like the classical Euclidean gradient magnitude.

In order to avoid possible difficulties when the affine invariants H or J are zero, we replace $\widehat{\nabla}_{\text{aff}}$ 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.

For further applications described in this paper, we will be more interested in edge stopping functions than in edge maps. These are functions which are as close as possible to zero at edges, and close to the maximal possible value at flat regions. We then proceed to make

⁷Note that the Euclidean gradient is the opposite, high at the ideal edge and zero everywhere else. Of course, this does not make any fundamental difference, since the important part is to differentiate between edges and flat regions. In the affine case, edges are given by doublets.

modifications on $\widehat{\nabla}_{aff}$ which allow us to compute well-defined stopping functions instead of just edge maps.

In Euclidean invariant edge detection algorithms based on active contours, as well as in anisotropic diffusion, the stopping term is usually taken in the form $(1 + \|\nabla\Phi\|^2)^{-1}$, the extra 1 being taken to avoid singularities where the Euclidean gradient vanishes. Thus, in analogy, the corresponding affine invariant stopping term should have the form

$$\frac{1}{1 + (\widehat{\nabla}_{aff}\Phi)^2} = \frac{J^2}{H^2 + J^2}$$

However, this can still present difficulties when both H and J vanish, so we propose a second modification.

Definition 3 *The normalized affine invariant gradient is given by:*

$$\nabla_{aff}\Phi = \sqrt{\frac{H^2}{J^2 + 1}} \tag{15}$$

The motivation comes from the form of the *affine invariant stopping term*, which is now given by

$$\frac{1}{1 + (\nabla_{aff}\Phi)^2} = \frac{J^2 + 1}{H^2 + J^2 + 1}. \tag{16}$$

Formula (16) avoids all difficulties where either H or J vanishes, and hence is a proper candidate for affine invariant edge detection. Indeed, in the neighborhood of an edge we obtain

$$\frac{J^2 + 1}{H^2 + J^2 + 1} = \frac{f'^6 h''^2 + 1}{h''^2 f'^2 (f'^4 + f''^2) + 1},$$

which, assuming h'' is moderate, gives an explanation of why it serves as a barrier for the edge. Barriers, that is, functions that go to zero at (salient) edges, will be important for the affine active contours presented in the following sections.

Examples of the affine invariant edge detector (16) are given in Figure 4. As with the affine invariant edge detection scheme introduced in previous section, this algorithm might produce gaps in the objects boundaries, as a result of the existence of perfectly straight segments with the same gray value. In this case, an edge integration algorithm is needed to complete the object boundary. The affine invariant active contours presented later is a possible remedy of this problem.



Figure 4: Examples of the affine invariant edge detector (after thresholding).

4 Affine invariant image denoising and simplification

We now present an additional application of the affine invariant edge maps introduced above. Indeed, we will describe a new model of anisotropic diffusion for image denoising and simplification, which is motivated by the ideas in [1, 2, 70]. According to the anisotropic diffusion flow of Alvarez *et al.* [2] given by (8), a stopping term ϕ should be added to the directional derivative to stop diffusion across edges. Following the work in [67, 70] (see also [1]), where the affine flow (12) is used as “directional diffusion,” we can further replace the Euclidean function ϕ in (8) by an affine invariant edge stopping function ϕ_{aff} . Accordingly, assume that $\phi_{aff} = \phi(w_{aff})$ where, as before, $\phi(w) \rightarrow 0$ when $w \rightarrow \infty$. We let $w = w_{aff}$ be either one of the affine edge detectors defined above, i.e., \mathcal{S}_{edge} as in (13) or $\nabla_{aff}(\Phi)$ as in (15). This results in a completely affine invariant flow,

$$\Phi_t = \phi_{aff} \kappa^{1/3} \|\nabla\Phi\| = \phi_{aff} (\Phi_y^2 \Phi_{xx} - 2\Phi_x \Phi_y \Phi_{xy} + \Phi_x^2 \Phi_{yy})^{1/3}. \quad (17)$$

This flow is analogous to the flow in [2], while having the additional property of being affine invariant, and therefore similar qualitative behavior is to be expected. The edge preserving denoising property of the flow is tested in Figure 5 (the affine gradient is used for

ϕ_{aff} in this and following examples). Note that since this type of flow (as well as the ones proposed in [2, 63]) tends to make an image piecewise constant, its results can be used to simplify (segment) an image in an affine invariant fashion.




Figure 5: Examples of the affine invariant image flow for image denoising and simplification. The original image is presented on the top row. Two different noise levels are given on the left at the second and last row, and the corresponding results of the affine invariant flow on the right.

In Figure 6, we compute the affine curvature for two images related by an affine transformation. Both images are first altered by random Gaussian noise, and after that, are processed by the denoising algorithm described above. The same number of steps was used for processing both images ($20\Delta t$). The affine curvature (the simplest non-trivial differential

affine invariant), given by [7]

$$\mu = [\mathcal{C}_{ss}, \mathcal{C}_{sss}],$$

is then computed for the mid-range level-set of both cleaned images. This was done using the implicit functions formula for the affine curvature (see for example [40] for the specific expression). Although this is not the best possible way to compute the affine curvature in discrete curves [9, 10], it is sufficient to show the qualitative behavior of the algorithms here described. The affine arc-length was computed using the relation between affine and Euclidean arc-lengths given by (10). Since we do not use the same starting point to compute the affine curvature, the plots are shifted one with respect to the other. Apart from this shift, the affine signatures are very close. The affine invariant edge map \mathcal{S}_{edge} is also shown for both images.

In Figure 7, we test the importance of affine invariant denoising. The noisy image is processed with the affine flow introduced above and also with its Euclidean analogue introduced in [2] (or [63]). After denoising, the affine curvature of the mid-range level-set is computed and compared to that of the original image. We note how this signature is much better preserved with the affine processing than with the Euclidean one.

5 Affine invariant active contours

In this section, we derive our affine invariant active contour models. We start with a brief description of classical energy “snakes” and curve evolution based snakes, followed by the presentation of Euclidean geodesic active contours, following the treatments of [12, 13, 36, 37]. For complete comparisons between these models, the interested reader is referred to [12, 37, 47]. We then proceed to derive the affine active contours based on the Euclidean version by defining the proper gradient flow. 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 [8, 27, 28, 80]. In this work, the affine invariant integration is done by means of active contours.




Figure 6: Affine denoising, edge maps, and curvature computation for two affine related images. The first row shows the two original noisy images (two left figures). Gaussian random noise was added independently to each one of the images after the affine transformation was performed. The subsequent images on the same row show the results of the affine denoising algorithm. The second row shows, on the left, \mathcal{S}_{edge} for the images after affine denoising. The last two figures on the second row show plots of the affine curvature vs. affine arc-length of the mid-range level-set for these images. The affine curvature was computed using implicit functions. Although this is not the best possible way to compute the affine curvature in discrete curves, it is sufficient to show the qualitative behavior of the algorithm. The affine arc-length was computed using the relation between affine and Euclidean arc-lengths described in the text. The curve was smoothed with the affine geometric heat flow for a small number of steps to avoid large noise in the discrete computations. Note that different starting points were used for both images, and therefore the corresponding plots are shifted.

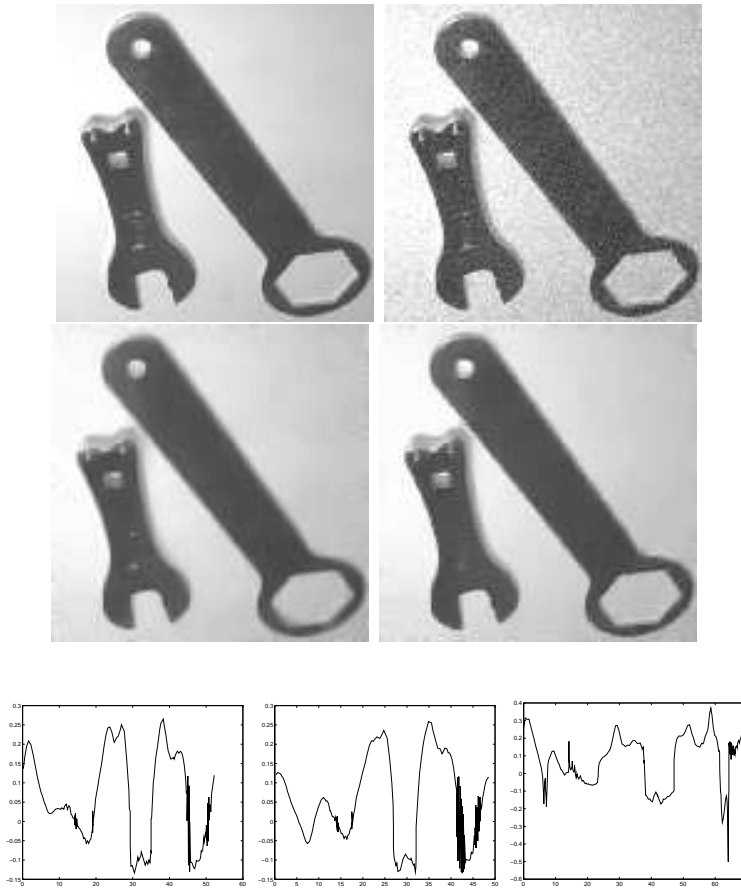


Figure 7: Test of the importance of affine invariant denoising. The first row shows the original image followed by the noisy one. The second row shows the denoising results for the affine and Euclidean flows. The third row shows plots of the affine curvature vs. affine arc-length of the mid-range level-set for the original image and those obtained from the affine and Euclidean denoising algorithms (second row images). The affine curvature was computed using implicit functions. The affine arc-length was computed using the relation between affine and Euclidean arc-lengths described in the text. In both cases, the curve was smoothed with the corresponding affine and Euclidean geometric heat flows for a small number of steps to avoid large noise in the discrete computations.

5.1 Classical snakes

Let $\mathcal{C}(p) : [0, 1] \rightarrow \mathbb{R}^2$ be a parametrized planar curve, and $\Phi : [0, a] \times [0, b] \rightarrow \mathbb{R}^+$ a given image where we want to detect the objects boundaries. The classical snakes approach [35, 77] associates to the curve \mathcal{C} an energy given by

$$E(\mathcal{C}) = \alpha \int_0^1 \| \mathcal{C}_p \|^2 dp + \beta \int_0^1 \| \mathcal{C}_{pp} \|^2 dp - \lambda \int_0^1 \| \nabla \Phi(\mathcal{C}(p)) \| dp, \quad (18)$$

where α , β , and λ are real positive constants. Here α and β determine the elasticity and rigidity of the curve, so that first two terms represent internal energy, and essentially control the smoothness⁸ of the contours to be detected, while the third term represents external energy, and is responsible for attracting the contour towards the desired object in the image. Solving the problem of snakes amounts to finding, for a given set of constants α , β , and λ , the curve \mathcal{C} that minimizes E . The classical (energy) approach of snakes can not deal with changes in topology, unless special topology handling procedures are added [49, 73]. This is the basic formulations of 2D active contours. Other related and 3D formulations have been proposed in the literature (e.g., [17, 18]). Reviewing all of them is beyond the scope of this paper.

5.2 Deformable models based on curve shortening

Our approach to geometric based active contours is strongly motivated by the papers [11, 46, 47, 48], and we now present the basic results reported there. Assume in the 2D case that the deforming curve \mathcal{C} is given as a level-set of a function $u : \mathbb{R}^2 \rightarrow \mathbb{R}$, so that we can represent the deformation of \mathcal{C} via that of u . In this case, the proposed 2D deformation is obtained modifying the edge detection algorithm (8) by including an inflationary force in the normal direction governed by a positive real constant ν .⁹ The evolution equation takes the form

$$\frac{\partial u}{\partial t} = \phi \| \nabla u \| \operatorname{div} \left(\frac{\nabla u}{\| \nabla u \|} \right) + \nu \phi \| \nabla u \|, \quad (t, x) \in [0, \infty[\times \mathbb{R}^2, \quad (19)$$

$$u(0, x) = u_0(x), \quad x \in \mathbb{R}^2, \quad (20)$$

⁸Other smoothing constraints can be used, but this is the most common one.

⁹Note that in (8), Φ is the image, while here u is an auxiliary embedding function.

where the stopping term typically has the form

$$\phi = \frac{1}{1 + \|\nabla \hat{\Phi}\|^m}, \quad (21)$$

where $m = 1$ or 2 , and $\hat{\Phi}$ is a regularized version of the original image Φ . We are looking for the contour of an object O , so, in the case of outer snakes (curves evolving towards the boundary of O from the exterior of O) the initial condition $u(0, x) = u_0(x)$ is typically taken as a regularized version of $1 - \chi_{\mathcal{C}}$ where $\chi_{\mathcal{C}}$ is the characteristic function of a curve \mathcal{C} containing O . Using once again the fact that $\operatorname{div}\left(\frac{\nabla u}{\|\nabla u\|}\right) = \kappa$, where κ is the Euclidean curvature of the level-sets \mathcal{C} of u , equation (19) can be written in the form

$$u_t = \phi \cdot (\nu + \kappa)\|\nabla u\|.$$

Equation (19) may then be interpreted as follows: Suppose that we are interested in following a certain level-set of u , which we can take to be the zero level-set. Suppose also that this level-set is a smooth curve. Then the flow

$$u_t = (\nu + \kappa)\|\nabla u\|,$$

means that the the level-set \mathcal{C} of u we are considering is evolving according to

$$\mathcal{C}_t = (\nu + \kappa)\mathcal{N}, \quad (22)$$

where \mathcal{N} is the inward normal to the curve. This equation was first proposed in [59, 74, 75], where extensive numerical research on it was performed. It was introduced in computer vision in [38, 39], where deep research on its importance for shape analysis is presented. The motion

$$C_t = \kappa\mathcal{N}, \quad (23)$$

is the Euclidean heat flow presented before, which is very well-known for its excellent geometric smoothing properties [3, 30, 31]. As was pointed out before, this flow is also called the ‘‘Euclidean shortening flow,’’ since it moves the curve in the gradient direction of the length functional given by

$$L := \oint_{\mathcal{C}} dv, \quad (24)$$

where $dv = \| \mathcal{C}_p \| dp$ is the Euclidean arc-length element, and therefore this flow decreases the length of the curve as fast as possible using only local information. This idea is the key for the snake models of [12, 13, 36, 37], as well as their extension to the affine case, as we shall soon see.

Finally, the constant velocity $\nu \mathcal{N}$ in (22), acts as the balloon force in [17] and is related to classical mathematical morphology [38, 64]. If $\nu > 0$, this velocity pushes the curve inwards and it is crucial in the model in order to allow convex initial curves to become non-convex, and thereby detect non-convex objects. Of course, the ν parameter must be specified a priori in order to make the object detection algorithm automatic, being this a non trivial issue. Recapping, the “force” $\nu + \kappa$ acts as the internal force in the classical energy based snakes model. The external force is given by ϕ , which is supposed to prevent the propagating curve from penetrating into the objects in the image. In [11, 46, 47, 48], the authors choose ϕ given by (21). $\hat{\Phi}$ may be obtained by Gaussian filtering, but more effective geometric smoothers, as those in [2, 70], can be used as well [45]. For an ideal edge, $\nabla \hat{\Phi} = \delta$, and the curve stops at the edge since $u_t = 0$ there. The boundary is then given by the set $u = 0$.

This curve evolution model given by (19) automatically handles different topologies. This is achieved with the help of an efficient numerical algorithm for curve evolution, developed by Osher and Sethian [59, 74, 75], used by many others for different image analysis problems, and analyzed, for example, in [15, 23].

5.3 Euclidean gradient snakes

We present now the geodesic active contours derived in [12, 36, 37]. Because of the central role played by Euclidean curve shortening in these models as well as the affine extension to be given below, we would like to explain in some detail now the relationship between curve shortening, gradient flows, and closed geodesics.

Let $\mathcal{C} = \mathcal{C}(p, t)$ be a smooth family of closed curves where t parametrizes the family and p the given curve, say $0 \leq p \leq 1$. (We assume that $\mathcal{C}(0, t) = \mathcal{C}(1, t)$ and similarly for the first derivatives.) 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

$$L'(t) = - \int_0^{L(t)} \left\langle \frac{\partial \mathcal{C}}{\partial t}, \kappa \mathcal{N} \right\rangle dv,$$

where dv is the Euclidean arc-length. Now, in the standard way, the length functional

$$\| \mathcal{C} \|_{euc} := \int_0^1 \| \mathcal{C}_p \| dp = \int_0^L dv = L,$$

will define a norm on the (Fréchet) space of twice-differentiable closed curves in the plane

$$\mathbf{C} := \{ \mathcal{C} : [0, 1] \rightarrow \mathbb{R}^2 : \mathcal{C} \text{ is closed and } C^2 \}.$$

Consequently, the direction in which $L(t)$ is decreasing most rapidly is when \mathcal{C} satisfies the gradient flow $\mathcal{C}_t = \kappa \mathcal{N}$, proving that the Euclidean curve shortening flow (23) is precisely a gradient flow. This analysis will be essential when we discuss the affine versions of this 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 [3, 30, 31, 32]. The active contours models which we are about to define are completely straightforward consequences of these general principles.

We are now ready to formulate the gradient active contours model from [12, 36, 37]. In [12], the model is derived from the principle of least action in physics [21], showing the mathematical relation between energy and curve evolution based snakes. In [36, 37], the model is derived immediately from curve shortening, and is compared to similar flows in continuum mechanics, in particular, phase transitions [34]. Of course, the two obtained flows are mathematically identical and present active contours as geodesic computations. We prefer to use here the simple curve shortening argument since it is the easier to generalize to the affine case. 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

$$dv_\phi = \phi dv = \phi \| \mathcal{C}_p \| dp. \tag{25}$$

As in ordinary curve shortening, we want to compute the corresponding gradient flow for the modified length functional

$$L_\phi(t) := \int_0^L \phi \, dv = \int_0^1 \|\mathcal{C}_p\| \phi \, dp. \quad (26)$$

Taking the derivative and integrating by parts, we find that [12, 36, 37]

$$-L'_\phi(t) = \int_0^{L_\phi(t)} \langle \mathcal{C}_t, \phi \kappa \mathcal{N} - (\nabla \phi \cdot \mathcal{N}) \mathcal{N} \rangle \, dv$$

which means that the direction in which the L_ϕ perimeter is shrinking as fast as possible is given by

$$\frac{\partial \mathcal{C}}{\partial t} = \phi \kappa \mathcal{N} - (\nabla \phi \cdot \mathcal{N}) \mathcal{N}. \quad (27)$$

This is precisely the gradient flow corresponding to the minimization of the length functional L_ϕ . 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 [59, 74, 75], let us assume now 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 following evolution equation:

$$\frac{\partial u}{\partial t} = \|\nabla u\| \operatorname{div} \left(\phi \frac{\nabla u}{\|\nabla u\|} \right) = \phi \|\nabla u\| \operatorname{div} \left(\frac{\nabla u}{\|\nabla u\|} \right) + \nabla \phi \cdot \nabla u, \quad (28)$$

with initial datum $u(0, x) = u_0(x)$. As in [11, 47], we may add an inflationary constant, to derive

$$\frac{\partial u}{\partial t} = \|\nabla u\| \operatorname{div} \left(\phi \frac{\nabla u}{\|\nabla u\|} \right) + \nu \phi \|\nabla u\| = \phi(\nu + \kappa) \|\nabla u\| + \nabla u \cdot \nabla \phi. \quad (29)$$

In the context of image processing, we take ϕ to be a stopping term depending on the image as in (21). In this case, $\nabla \phi$ 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. This new force then increases the attraction of the deforming contour towards the boundary, being of special help when the boundary includes high variations of its

gradient values. Note that in contrast with this gradient flow (29), in the model of [11, 47], the curve stops only when $\phi = 0$, which happens only along an ideal edge. The second advantage of this new term is that it allows the detection of non-convex objects as well, thus removing the necessity to *a priori* estimate the inflationary constant given by ν . In case we wish to add this constant velocity, in order for example to increase the speed of convergence, we can just consider the term $\nu\phi|\nabla u|$ of (29) as an extra speed in the gradient problem (26), minimizing the enclosed area, [17, 85]. Existence, uniqueness and stability results for the gradient active contour model (29) were studied in [12, 13, 36, 37].

We should point out it is trivial to write down the 3D extensions of such active contour models, as was done in [13, 14, 36, 37, 83, 84]; see also [81, 82]. We should also note that Shah [71] recently presented an active contours formulation using a weighted length formulation as in (26) as starting point. In his case, ϕ is obtained from an elaborated segmentation procedure obtained from the Mumford-Shah approach [50]. Extensions of the model in [11, 46] are studied also in [76], motivated in part by the work in [38, 39]. Related work may also be found in [29].

5.4 Affine invariant gradient snakes

Based on the gradient 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 [7], we will initially have to restrict ourselves to the (Fréchet) space of thrice-differentiable convex closed curves in the plane, i.e.,

$$\mathbf{C}_0 := \{\mathcal{C} : [0, 1] \rightarrow \mathbb{R}^2 : \mathcal{C} \text{ is convex, closed and } C^3\}.$$

As above, let ds denote the affine arc-length; see (3). Then, letting $L_{aff} := \oint ds$ be the *affine length* [7], we proceed to define the *affine norm* on the space \mathbf{C}_0

$$\|\mathcal{C}\|_{aff} := \int_0^1 \|\mathcal{C}(p)\|_a dp = \int_0^{L_{aff}} \|\mathcal{C}(s)\|_a ds,$$

where

$$\|\mathcal{C}(p)\|_a := [\mathcal{C}(p), \mathcal{C}_p(p)].$$

Note that the area enclosed by \mathcal{C} is just

$$A = \frac{1}{2} \int_0^1 \|\mathcal{C}(p)\|_a dp = \frac{1}{2} \int_0^1 [\mathcal{C}, \mathcal{C}_p] dp = \frac{1}{2} \|\mathcal{C}\|_{aff} . \quad (30)$$

Observe that

$$\|\mathcal{C}_s\|_a = [\mathcal{C}_s, \mathcal{C}_{ss}] = 1, \quad \|\mathcal{C}_{ss}\|_a = [\mathcal{C}_{ss}, \mathcal{C}_{sss}] = \mu,$$

where μ is the *affine curvature*, i.e., the simplest non-trivial differential affine invariant. This makes the affine norm $\|\cdot\|_{aff}$ consistent with the properties of the Euclidean norm on curves relative to the Euclidean arc-length dv . (Here we have that $\|\mathcal{C}_v\| = 1$, $\|\mathcal{C}_{vv}\| = \kappa$.)

We can now formulate the functionals that will be used to define the affine invariant snakes. Accordingly, assume that $\phi_{aff} = \phi(w_{aff})$ is an affine invariant stopping term, based on the affine invariant edge detectors considered before. Therefore, ϕ_{aff} plays the role of the weight ϕ in L_ϕ . 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 $ds_{\phi_{aff}} = \phi_{aff} ds$ to obtain the first possible functional for the affine active contours

$$L_{\phi_{aff}} := \int_0^{L_{aff}(t)} \phi_{aff} ds, \quad (31)$$

where as above L_{aff} is the affine length. The obvious next step is to compute the gradient flow corresponding to $L_{\phi_{aff}}$ 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. In the meantime, using the connection (10) between the affine and Euclidean arc lengths, note that the above equation can be re-written in Euclidean space as

$$L_{\phi_{aff}} = \int_0^{L(t)} \phi_{aff} \kappa^{1/3} dv, \quad (32)$$

where $L(t)$ denotes the ordinary Euclidean length of the curve $\mathcal{C}(t)$.

The snake model which we will use comes from another (special) affine invariant, namely *area*, cf. (30). Let $\mathcal{C}(p, t)$ be a family of curves in \mathbf{C}_0 . A straightforward computation reveals that the first variation of the area functional

$$A(t) = \frac{1}{2} \int_0^1 [\mathcal{C}, \mathcal{C}_p] dp$$

is

$$A'(t) = - \int_0^{L_{\text{aff}}(t)} [\mathcal{C}_t, \mathcal{C}_s] ds.$$

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

$$\mathcal{C}_t = \kappa^{1/3} \mathcal{N},$$

which is precisely the affine invariant heat equation studied in [65]! It is this functional that we will proceed to modify with the conformal factor ϕ_{aff} . Therefore, we define the conformal area functional to be ¹⁰

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

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

The precise formulas for the variations of these two functionals are given in the following result. They use the definition of Y^\perp given in (2). The proof follows by an integration by parts argument and some manipulations as in [12, 13, 36, 37].

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*

$$\frac{dL_{\phi_{\text{aff}}}(t)}{dt} = - \int_0^{L_{\text{aff}}(t)} [\mathcal{C}_t, (\nabla \phi_{\text{aff}})^\perp] ds + \int_0^{L_{\text{aff}}(t)} \phi_{\text{aff}} \mu[\mathcal{C}_t, \mathcal{C}_s] ds. \quad (33)$$

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

$$\frac{dA_{\phi_{\text{aff}}}(t)}{dt} = - \int_0^{L_{\text{aff}}(t)} [\mathcal{C}_t, (\phi_{\text{aff}} \mathcal{C}_s + \frac{1}{2} [\mathcal{C}, (\nabla \phi)^\perp \mathcal{C}_s])] ds. \quad (34)$$

¹⁰An alternative definition would be to use $A_{\phi_{\text{aff}}} := \int \int \phi_{\text{aff}} dx dy$, which is actually the affine analogue of the weighted area that produces the Euclidean weighted constant motion $\mathcal{C}_t = \phi \vec{\mathcal{N}}$.

The affine invariance of the resulting variational derivatives follows from a general result governing invariant variational problems having volume preserving symmetry groups [58]:

Theorem 1 *Suppose G is a connected transformation group, and $\mathcal{I}[\mathcal{C}]$ is a G -invariant variational problem. Then the variational derivative (or gradient) $\delta\mathcal{I}$ of \mathcal{I} is a differential invariant if and only if G is a group of volume-preserving transformations.*

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

$$\mathcal{C}_t = \{(\nabla\phi_{\text{aff}})^\perp + \phi_{\text{aff}}\mu\mathcal{C}_s\}_s = ((\nabla\phi_{\text{aff}})^\perp)_s + (\phi_{\text{aff}}\mu)_s\mathcal{C}_s + \phi_{\text{aff}}\mu\mathcal{C}_{ss}.$$

As before, we ignore the tangential components, which do not affect the geometry of the evolving curve, and so obtain the following possible model for geometric affine invariant active contours:

$$\mathcal{C}_t = \phi_{\text{aff}}\mu\kappa^{1/3}\mathcal{N} + \langle((\nabla\phi_{\text{aff}})^\perp)_s, \mathcal{N}\rangle\mathcal{N}. \quad (35)$$

The geometric interpretation of the affine gradient flow (35) minimizing $L_{\phi_{\text{aff}}}$ is analogous to that of the corresponding Euclidean geodesic active contours. The term $\phi_{\text{aff}}\mu\kappa^{1/3}$ minimizes the affine length L_{aff} while smoothing the curve according to the results in [65, 66], being stopped by the affine invariant stopping function ϕ_{aff} . The term associated with $((\nabla\phi_{\text{aff}})^\perp)_s$ creates a potential valley, attracting the evolving curve to the affine edges. Unfortunately, this flow involves μ which makes it difficult to implement. (Possible techniques to compute μ numerically were recently reported in [9, 10, 26].)

The gradient flow coming from the first variation of the modified area functional on the other hand is much simpler:

$$\mathcal{C}_t = (\phi_{\text{aff}}\mathcal{C}_s + \frac{1}{2}[\mathcal{C}, (\nabla\phi_{\text{aff}})^\perp]\mathcal{C}_s)_s \quad (36)$$

Ignoring tangential terms (those involving \mathcal{C}_s) this flow is equivalent to

$$\mathcal{C}_t = \phi_{\text{aff}}\mathcal{C}_{ss} + \frac{1}{2}[\mathcal{C}, (\nabla\phi_{\text{aff}})^\perp]\mathcal{C}_{ss}, \quad (37)$$

which in Euclidean form gives the second possible affine contour snake model:

$$\mathcal{C}_t = \phi_{\text{aff}} \kappa^{1/3} \mathcal{N} + \frac{1}{2} \langle \mathcal{C}, \nabla \phi_{\text{aff}} \rangle \kappa^{1/3} \mathcal{N}. \quad (38)$$

Notice that although both models (35) and (38) were derived for *convex curves*, the flow (38) makes sense in the non-convex case as well, which makes this the only candidate for a practical affine invariant geometric contour method. Thus we will concentrate on (38) from now on, and just consider (35) as a model with some theoretical interest.

In order to better capture concavities, to speed up the evolution, as well as to be able to define outward evolutions, a constant inflationary force of the type $\nu \phi \mathcal{N}$ may be added to (38). This can be obtained from the affine gradient descent of $A_{\phi_{\text{aff}}} := \int \int \phi_{\text{aff}} dx dy$. Note that the inflationary force $\mathcal{C}_t = \phi \vec{\mathcal{N}}$ in the Euclidean case of active contours is obtained from the (Euclidean) gradient descent of $A := \int \int \phi dx dy$, where ϕ is a regular edge detector. Formal results regarding existence and uniqueness of solutions to (38) can be derived following [1, 11, 12, 13, 36, 37].

Figure 8 illustrates simulations of these active contour models (the implementation is as in [12, 13, 36, 37, 46, 47, 48], based on the level-sets formulation [59]). The original curve is a square surrounding both objects. This square becomes non-convex and then splits in order to detect both objects. As mentioned before, the affine invariant active contours just presented have similar behavior than the ones in [12, 37], with the additional property of being affine invariant. Additional examples on gradient-based geometric active contours may be found in these references.

6 Concluding remarks

The problems of affine invariant detection and image denoising were 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, 65, 66]. 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 [53, 54], this function is the simplest possible with




Figure 8: Examples of the affine invariant active contours. The original image is presented on the left and the one with the corresponding objects detected by the affine active contours on the right. The original curve is a square surrounding both objects. Although the image contains high noise due to JPEG compression, both objects are detected. (This is a color figure.)

this characteristic.

We then presented two models for affine invariant active contours, extending the results presented in [12, 13, 36, 37] 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. Based on the same theory of affine invariant gradient descent and curve metrics, we proved that the affine geometric heat flow is minimizing area, in an affine invariant form, as fast as possible, and so is a steepest-descent flow. The affine edge maps were used to extend the image flows in [1, 67, 70], obtaining a complete affine invariant flow for image denoising and simplification.

The schemes here presented should be used as first steps in (“planar”) shape recognition systems. This will help to reduce non-intrinsic noise frequently added to the algorithms due to the use of non-invariant edge and object detection schemes. The affine-invariant recognition of contours bounding objects can then proceed by using the affine signature curves and their affine-invariant numerical approximations developed in [9, 10]. The affine

invariant image denoising (and simplification) scheme can be used as pre-processing for the algorithm recently introduced in [79].

We conclude by noting that the 3D extension of the affine active contours work is clear, one can use a modified volume functional in this case. Moreover, we plan to compare the Euclidean and affine methods on some more realistic medical imagery and to use the affine invariant object detection and image denoising schemes for recognition tasks in future publications. Study of the behavior of zero-crossings [78] associated with the affine invariant gradient is the subject of current research as well.

Acknowledgments

Images in this paper are from the public domain image archives at Caltech and courtesy of Prof. Andrew Zisserman from Oxford University. Dr. Ronny Kimmel from UC Berkeley provided us software support for the computation of the affine curvature.

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**, pp. 845-866, 1992.
- [3] S. Angenent, “Parabolic equations for curves on surfaces, Part II. Intersections, blow-up, and generalized solutions,” *Annals of Mathematics* **133**, pp. 171-215, 1991.
- [4] S. Angenent, G. Sapiro, and A. Tannenbaum, “On the affine heat flow for non-convex curves,” pre-print, 1995.
- [5] C. Ballester, V. Caselles, and M. Gonzalez, “Affine invariant segmentation by variational method,” *Technical Report, U. of Illes Balears*, 1994.
- [6] A. Blake and A. Yuille, *Active Vision* MIT Press, Cambridge, 1992.

- [7] W. Blaschke, *Vorlesungen über Differentialgeometrie II*, Verlag Von Julius Springer, Berlin, 1923.
- [8] F. L. Bookstein, "Fitting conic sections to scattered data," *Comp. Graph. Image Process.* **9**, pp. 56-71, 1979.
- [9] E. Calabi, P. J. Olver, and A. Tannenbaum, "Affine geometry, curve flows, and invariant numerical approximations," to appear in *Advances in Mathematics*.
- [10] E. Calabi, P. J. Olver, C. Shakiban, A. Tannenbaum, and S. Haker, "Differential and numerical invariant signature curves applied to object recognition," to appear in *Int. Journal of Computer Vision*.
- [11] V. Caselles, F. Catte, T. Coll, F. Dibos, "A geometric model for active contours," *Numerische Mathematik* **66**, pp. 1-31, 1993.
- [12] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic active contours," *International Journal of Computer Vision* **22:1**, pp. 61-79, 1997. Also in *Proc. ICCV*, Cambridge, MA, June 1995.
- [13] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert, "Minimal surfaces: A three dimensional segmentation approach," *IEEE-PAMI*, **19:4**, pp. 394-398, April 1997.
- [14] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert, "Three dimensional object modeling via minimal surfaces," *Proc. ECCV*, Cambridge, UK, April 1996.
- [15] Y. G. Chen, Y. Giga, and S. Goto, "Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations," *J. Differential Geometry* **33**, pp. 749-786, 1991.
- [16] D. Chopp, "Computing minimal surfaces via level set curvature flows," *LBL TR-University of Berkeley*, 1991.
- [17] L. D. Cohen, "On active contour models and balloons," *CVGIP: Image Understanding* **53**, pp. 211-218, 1991.

- [18] I. Cohen, L. D. Cohen, and N. Ayache, "Using deformable surfaces to segment 3D images and infer differential structure," *CVGIP: Image Understanding* **56**, pp. 242-263, 1992.
- [19] T. Cohignac, C. Lopez, and J. M. Morel, "Integral and local affine invariant parametrizations and applications to shape recognition," *Proc. 12th IEEE Int. Conf. Pattern Recognition*, Jerusalem, 1994.
- [20] M. G. Crandall, H. Ishii, and P. L. Lions, "User's guide to viscosity solutions of second order partial linear differential equations," *Bulletin of the American Math. Society* **27**, pp. 1-67, 1992.
- [21] B. A. Dubrovin, A. T. Fomenko, and S. P. Novikov, *Modern Geometry - Methods and Applications I*, Springer-Verlag, New York, 1984.
- [22] C. L. Epstein and M. Gage, "The curve shortening flow," in *Wave Motion: Theory, Modeling, and Computation*, A. Chorin and A. Majda, Editors, Springer-Verlag, New York, 1987.
- [23] L. C. Evans and J. Spruck, "Motion of level sets by mean curvature, I," *J. Differential Geometry* **33**, pp. 635-681, 1991.
- [24] 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.
- [25] O. Faugeras, "Cartan's moving frame method and its application on the geometry and evolution of curves in the Euclidean, affine, and projective planes," *INRIA TR 2053*, September 1993.
- [26] 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.
- [27] D. Forsyth, J. L. Mundy, A. Zisserman, and Brown, "Projectively invariant representations using implicit curves," *Image Vis. Comput.* **8**, pp. 130-136, 1990.

- [28] D. Forsyth, J. L. Mundy, A. Zisserman, C. Coelho, A. Heller, and C. Rothwell, “Invariant description of object representation and pose,” *IEEE PAMI* **13** pp. 971-991, 1991.
- [29] P. Fua and Y. G. Leclerc, “Model driven edge detection,” *Machine Vision and Applications*, **3**, pp. 45-56, 1990.
- [30] M. Gage and R. S. Hamilton, “The heat equation shrinking convex plane curves,” *J. Differential Geometry* **23**, pp. 69-96, 1986.
- [31] M. Grayson, “The heat equation shrinks embedded plane curves to round points,” *J. Differential Geometry* **26**, pp. 285-314, 1987.
- [32] M. Grayson, “Shortening embedded curves,” *Annals of Mathematics* **129**, pp. 285-314, 1989.
- [33] H. W. Guggenheimer, *Differential Geometry*, McGraw-Hill Book Company, New York, 1963.
- [34] M. E. Gurtin, *Thermomechanics of Evolving Phase Boundaries in the Plane*, Oxford Univ. Press, New York, 1993.
- [35] M. Kass, A. Witkin, and D. Terzopoulos, “Snakes: Active contour models,” *International Journal of Computer Vision* **1**, pp. 321-331, 1988.
- [36] S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. Yezzi, “Gradient flows and geometric active contour models,” *Proc. ICCV*, Cambridge, MA, June 1995.
- [37] S. Kichenassamy, A. Kumar, P. J. Olver, A. Tannenbaum, and A. Yezzi, “Conformal curvature flows: From phase transitions to active vision,” *Archive for Rational Mechanics and Analysis* **134**, pp. 275-301, 1996.
- [38] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, “Toward a computational theory of shape: An overview,” *Lecture Notes in Computer Science* **427**, pp. 402-407, Springer-Verlag, New York, 1990.

- [39] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, "Shapes, shocks, and deformations, I," *International Journal of Computer Vision* **15**, pp. 189-224, 1995.
- [40] R. Kimmel, "Invariant framework for differential affine signatures," *Proc. ICPR '96*.
- [41] R. Kimmel, A. Amir, A. M. Bruckstein, "Finding shortest paths on surfaces using level sets propagation," *IEEE-PAMI*, 17(1):635-640, June, 1995.
- [42] R. Kimmel, N. Kiryati, A. M. Bruckstein, "Distance maps and weighted distance transforms," *Journal of Mathematical Imaging and Vision*, Special Issue on Topology and Geometry in Computer Vision, to appear.
- [43] T. Lindeberg, *Scale-Space Theory in Computer Vision*, Kluwer, 1994.
- [44] T. Lindeberg and J. Garding, "Shape-adapted smoothing in estimation of 3D depth cues from affine distortions of local 2D structures," *Proc. ECCV*, Stockholm, Sweden, May 1994.
- [45] R. Malladi and J. A. Sethian, "A unified approach to noise removal, image enhancement, and shape recovery," *IEEE Trans. Image Processing*, to appear.
- [46] R. Malladi, J. A. Sethian and B. C. Vemuri, "Evolutionary fronts for topology independent shape modeling and recovery," *Proc. of the 3rd ECCV*, Stockholm, Sweden, pp. 3-13, 1994.
- [47] R. Malladi, J. A. Sethian and B. C. Vemuri, "Shape modeling with front propagation: A level set approach," *IEEE Trans. on PAMI* **17**, pp. 158-175, 1995.
- [48] R. Malladi, J. A. Sethian and B. C. Vemuri, "A fast level set based algorithm for topology independent shape modeling," *Journal of Mathematical Imaging and Vision*, special issue on Topology and Geometry, Ed. A. Rosenfeld and Y. Kong, to appear.
- [49] T. McInerney and D. Terzopoulos, "Topologically adaptable snakes," *Proc. ICCV*, Cambridge, MA, June 1995.

- [50] D. Mumford and J. Shah, “Optimal approximations by piecewise smooth functions and variational problems,” *Comm. Pure and App. Math.* **42**, 1989.
- [51] J. L. Mundy and A. Zisserman (Eds.), *Geometric Invariance in Computer Vision*, MIT Press, 1992.
- [52] W. J. Niessen, B. M. ter Haar Romeny, L. M. J. Florack, and A. H. Salden, “Nonlinear diffusion of scalar images using well-posed differential operators,” Technical Report, Utrecht University, The Netherlands, October 1993.
- [53] P. J. Olver, *Applications of Lie Groups to Differential Equations*, second ed., Springer-Verlag, New York, 1993.
- [54] P. J. Olver, *Equivalence, Invariants, and Symmetry*, Cambridge University Press, Cambridge-UK, 1995.
- [55] P. J. Olver, G. Sapiro, and A. Tannenbaum, “Differential invariant signatures and flows in computer vision: A symmetry group approach,” in [62].
- [56] 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.
- [57] P. J. Olver, G. Sapiro, and A. Tannenbaum, “Affine invariant detection: Edges, active contours, and segments,” *Proc. Computer Vision Pattern Recognition*, San Francisco, June 1996.
- [58] P. J. Olver, G. Sapiro, and A. Tannenbaum, “Invariant geometric evolutions of surfaces and volumetric smoothing,” *SIAM J. of Appl. Math.*, February 1997.
- [59] 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.

- [60] E. J. Pauwels, P. Fiddelaers, and L. J. Van Gool, "Shape-extraction for curves using geometry-driven diffusion and functional optimization," *Proc. ICCV*, Cambridge, MA, June 1995.
- [61] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Machine Intell.* **12**, pp. 629-639, 1990.
- [62] B. Romeny, Editor, *Geometry Driven Diffusion in Computer Vision*, Kluwer, 1994.
- [63] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D* **60**, pp. 259-268, 1992.
- [64] G. Sapiro, R. Kimmel, D. Shaked, B. B. Kimia, and A. M. Bruckstein, "Implementing continuous-scale morphology via curve evolution," *Pattern Recog.* **26:9**, pp. 1363-1372, 1993.
- [65] G. Sapiro and A. Tannenbaum, "On affine plane curve evolution," *Journal of Functional Analysis* **119:1**, pp. 79-120, 1994.
- [66] G. Sapiro and A. Tannenbaum, "Affine invariant scale-space," *International Journal of Computer Vision* **11:1**, pp. 25-44, 1993.
- [67] G. Sapiro and A. Tannenbaum, "Image smoothing based on an affine invariant flow," *Proceedings of Conference on Information Sciences and Systems*, Johns Hopkins University, March 1993.
- [68] G. Sapiro and A. Tannenbaum, "On invariant curve evolution and image analysis," *Indiana University Mathematics Journal* **42:3**, 1993.
- [69] G. Sapiro and A. Tannenbaum, "Area and length preserving geometric invariant scale-spaces," *IEEE Trans. PAMI* **17:1**, pp. 67-72, 1995.
- [70] G. Sapiro, A. Tannenbaum, Y. L. You, and M. Kaveh, "Experiments on geometric image enhancement," *First IEEE-International Conference on Image Processing*, Austin-Texas, November 1994.

- [71] J. Shah, "Recovery of shapes by evolution of zero-crossings," *Technical Report, Math. Dept. Northeastern Univ*, Boston MA, 1995.
- [72] H. M. Soner, "Motion of a set by the curvature of its boundary," *J. of Diff. Equations* **101**, pp. 313-372, 1993.
- [73] R. Szeliski, D. Tonnesen, and D. Terzopoulos, "Modeling surfaces of arbitrary topology with dynamic particles," *Proc. CVPR*, pp. 82-87, 1993.
- [74] J. A. Sethian, "Curvature and the evolution of fronts," *Commun. Math. Phys.* **101**, pp. 487-499, 1985.
- [75] J. A. Sethian, "A review of recent numerical algorithms for hypersurfaces moving with curvature dependent flows," *J. Differential Geometry* **31**, pp. 131-161, 1989.
- [76] H. Tek and B. B. Kimia, "Image segmentation by reaction-diffusion bubbles," *Proc. ICCV*, Cambridge, MA, June 1995.
- [77] D. Terzopoulos, A. Witkin, and M. Kass, "Constraints on deformable models: Recovering 3D shape and nonrigid motions," *Artificial Intelligence* **36**, pp. 91-123, 1988.
- [78] V. Torre and T. Poggio, "On edge detection," *IEEE Trans. PAMI* **8**, pp. 147-163, 1986.
- [79] L. Van Gool, T. Moons, and D. Ungureanu, "Affine/photometric invariants for planar intensity patterns," *Proc. ECCV*, pp. 642-651, Cambridge, UK, April 1996.
- [80] I. Weiss, "Geometric invariants and object recognition," *International Journal of Computer Vision*, pp. 207-231, 1993.
- [81] R. T. Whitaker, "Volumetric deformable models: Active blobs," *ECRC TR* **94-25**, 1994.
- [82] R. T. Whitaker, "Algorithms for implicit deformable models," *Proc. ICCV'95*, Cambridge, MA, June 1995.

- [83] A. Yezzi, S. Kichenassamy, A. Kumar, P. Olver, and A. Tannenbaum, "Geometric snakes for edge detection and segmentation of medical imagery," *IEEE Trans. Medical Imaging* **16**, pp. 199-210, 1997.
- [84] A. Yezzi, S. Kichenassamy, P. Olver, and A. Tannenbaum, "A gradient surface approach to 3D segmentation," *Proceedings of 49th IS&T*, 1996.
- [85] S. C. Zhu, T. S. Lee, and A. L. Yuille, "Region competition: Unifying snakes, region growing, energy/Bayes/MDL for multi-band image segmentation," *Proc. ICCV*, Cambridge, MA, June 1995.