

Snakes, Shapes, and Gradient Vector Flow

Chenyang Xu and Jerry L. Prince

Image Analysis and Communications Laboratory
Department of Electrical and Computer Engineering
105 Barton Hall
The Johns Hopkins University
Baltimore, MD 21218

Tel: (410) 516-5192 Fax: (410) 516-5566 E-mail: prince@jhu.edu

ABSTRACT

Snakes, or active contours, are used extensively in computer vision and image processing applications, particularly to locate object boundaries. Problems associated with initialization and poor convergence to boundary concavities, however, have limited their utility. This paper presents a new external force for active contours, largely solving both problems. This external force, which we call *gradient vector flow* (GVF), is computed as a diffusion of the gradient vectors of a gray-level or binary edge map derived from the image. It differs fundamentally from traditional snake external forces in that it cannot be written as the negative gradient of a potential function, and the corresponding snake is formulated directly from a force balance condition rather than a variational formulation. Using several two-dimensional examples and one three-dimensional example, we show that GVF has a large capture range and is able to move snakes into boundary concavities.

IEEE Transactions on Image Processing, 359–369, March 1998.

Permission to publish this abstract separately is granted.

This research was supported by National Science Foundation Presidential Faculty Fellow Award MIP93-50336.

IP Editors' Information Classification Scheme (EDICS): IP 1.9

Corresponding Author: Jerry L. Prince

1 Introduction

Snakes [1], or *active contours*, are curves defined within an image domain which can move under the influence of internal forces coming from within the curve itself and external forces computed from the image data. The internal and external forces are defined so that the snake will conform to an object boundary or other desired features within an image. Snakes are widely used in many applications, including edge detection [1], shape modeling [2, 3], segmentation [4, 5], and motion tracking [4, 6].

There are two general types of active contour models in the literature today: *parametric active contours* [1] and *geometric active contours* [7, 8]. In this paper, we focus on parametric active contours, although we expect our results to have applications in geometric active contours as well. Parametric active contours synthesize parametric curves within an image domain and allow them to move toward desired features, usually edges. Typically, the curves are drawn toward the edges by *potential forces*, which are defined to be the negative gradient of a potential function. Additional forces, such as pressure forces [9], together with the potential forces comprise the *external forces*. There are also *internal forces* designed to hold the curve together (elasticity forces) and to keep it from bending too much (bending forces).

There are two key difficulties with parametric active contour algorithms. First, the initial contour must, in general, be close to the true boundary or else it will likely converge to the wrong result. Several methods have been proposed to address this problem including multiresolution methods [10], pressure forces [9], and distance potentials [11]. The basic idea is to increase the *capture range* of the external force fields and to guide the contour toward the desired boundary. The second problem is that active contours have difficulties progressing into *boundary concavities* [12, 13]. There is no satisfactory solution to this problem, although pressure forces [9], control points [12], domain-adaptivity [14], directional attractions [13], and the use of solenoidal fields [15] have been proposed. However, most of the methods proposed to address these problems solve only one problem while creating new difficulties. For example, multiresolution methods have addressed the issue of capture range, but specifying how the snake should move across different resolutions remains problematic. Another example is that of pressure forces, which can push an active contour into boundary concavities, but cannot be too strong or “weak” edges will be overwhelmed [16]. Pressure forces must also be initialized to push out or push in, a condition that mandates careful initialization.

In this paper, we present a new class of external forces for active contour models which addresses both problems listed above. These fields, which we call *gradient vector flow* (GVF) fields, are dense vector

fields derived from images by minimizing a certain energy functional in a variational framework. The minimization is achieved by solving a pair of decoupled linear partial differential equations which diffuses the gradient vectors of a gray-level or binary edge map computed from the image. We call the active contour that uses the GVF field as its external force a *GVF snake*. The GVF snake is distinguished from nearly all previous snake formulations in that its external forces cannot be written as the negative gradient of a potential function. Because of this, it cannot be formulated using the standard energy minimization framework; instead, it is specified directly from a force balance condition.

Particular advantages of the GVF snake over a traditional snake are its insensitivity to initialization and its ability to move into boundary concavities. As we show in this paper, its initializations can be inside, outside, or across the object's boundary. Unlike pressure forces, the GVF snake does not need prior knowledge about whether to shrink or expand towards the boundary. The GVF snake also has a large capture range, which means that, barring interference from other objects, it can be initialized far away from the boundary. This increased capture range is achieved through a diffusion process which does not blur the edges themselves, so multiresolution methods are not needed. The external force model that is closest in spirit to GVF is the distance potential forces of Cohen and Cohen [11]. Like GVF, these forces originate from an edge map of the image and can provide a large capture range. We show, however, that unlike GVF, distance potential forces cannot move a snake into boundary concavities. We believe that this is a property of all conservative forces which characterize nearly all snake external forces, and that exploring non-conservative external forces, such as GVF, is an important direction for future research in active contour models.

2 Background

2.1 Parametric Snake Model

A *traditional snake* is a curve $\mathbf{x}(s) = [x(s), y(s)]$, $s \in [0, 1]$, that moves through the spatial domain of an image to minimize the energy functional

$$E = \int_0^1 \frac{1}{2} (\alpha |\mathbf{x}'(s)|^2 + \beta |\mathbf{x}''(s)|^2) + E_{\text{ext}}(\mathbf{x}(s)) ds \quad (1)$$

where α and β are weighting parameters that control the snake's tension and rigidity, respectively, and $\mathbf{x}'(s)$ and $\mathbf{x}''(s)$ denote the first and second derivatives of $\mathbf{x}(s)$ with respect to s . The external energy function

E_{ext} is derived from the image so that it takes on its smaller values at the features of interest, such as boundaries. Given a gray-level image $I(x, y)$, viewed as a function of continuous position variables (x, y) , typical external energies designed to lead an active contour toward step edges are [1]:

$$E_{\text{ext}}^{(1)}(x, y) = -|\nabla I(x, y)|^2 \quad (2)$$

$$E_{\text{ext}}^{(2)}(x, y) = -|\nabla(G_\sigma(x, y) * I(x, y))|^2 \quad (3)$$

where $G_\sigma(x, y)$ is a two-dimensional Gaussian function with standard deviation σ and ∇ is the gradient operator. If the image is a line drawing (black on white), then appropriate external energies include [9]:

$$E_{\text{ext}}^{(3)}(x, y) = I(x, y) \quad (4)$$

$$E_{\text{ext}}^{(4)}(x, y) = G_\sigma(x, y) * I(x, y) \quad (5)$$

It is easy to see from these definitions that larger σ 's will cause the boundaries to become blurry. Such large σ 's are often necessary, however, in order to increase the capture range of the active contour.

A snake that minimizes E must satisfy the Euler equation

$$\alpha \mathbf{x}''(s) - \beta \mathbf{x}''''(s) - \nabla E_{\text{ext}} = 0 \quad (6)$$

This can be viewed as a force balance equation

$$\mathbf{F}_{\text{int}} + \mathbf{F}_{\text{ext}}^{(p)} = 0 \quad (7)$$

where $\mathbf{F}_{\text{int}} = \alpha \mathbf{x}''(s) - \beta \mathbf{x}''''(s)$ and $\mathbf{F}_{\text{ext}}^{(p)} = -\nabla E_{\text{ext}}$. The internal force \mathbf{F}_{int} discourages stretching and bending while the external potential force $\mathbf{F}_{\text{ext}}^{(p)}$ pulls the snake towards the desired image edges.

To find a solution to (6), the snake is made dynamic by treating \mathbf{x} as function of time t as well as s — i.e., $\mathbf{x}(s, t)$. Then, the partial derivative of \mathbf{x} with respect to t is then set equal to the left hand side of (6) as follows

$$\mathbf{x}_t(s, t) = \alpha \mathbf{x}''(s, t) - \beta \mathbf{x}''''(s, t) - \nabla E_{\text{ext}} \quad (8)$$

When the solution $\mathbf{x}(s, t)$ stabilizes, the term $\mathbf{x}_t(s, t)$ vanishes and we achieve a solution of (6). A numerical solution to (8) can be found by discretizing the equation and solving the discrete system iteratively (cf. [1]).

We note that most snake implementations use either a parameter which multiplies \mathbf{x}_t in order to control the temporal step-size, or a parameter to multiply ∇E_{ext} which permits separate control of the external force strength. In this paper, we normalize the external forces so that the maximum magnitude is equal to one, and use a unit temporal step-size for all the experiments.

2.2 Behavior of Traditional Snakes

An example of the behavior of a traditional snake is shown in Fig. 1. Fig. 1a shows a 64×64 -pixel line-drawing of a U-shaped object (shown in gray) having a boundary concavity at the top. It also shows a sequence of curves (in black) depicting the iterative progression of a traditional snake ($\alpha = 0.6, \beta = 0.0$) initialized outside the object but within the capture range of the potential force field. The potential force field $\mathbf{F}_{\text{ext}}^{(p)} = -\nabla E_{\text{ext}}^{(4)}$ where $\sigma = 1.0$ pixel is shown in Fig. 1b. We note that the final solution in Fig. 1a solves the Euler equations of the snake formulation, but remains split across the concave region.

The reason for the poor convergence of this snake is revealed in Fig. 1c, where a close-up of the external force field within the boundary concavity is shown. Although the external forces correctly point toward the object boundary, within the boundary concavity the forces point horizontally *in opposite directions*. Therefore, the active contour is pulled apart toward each of the “fingers” of the U-shape, but not made to progress downward into the concavity. There is no choice of α and β that will correct this problem.

Another key problem with traditional snake formulations, the problem of limited capture range, can be understood by examining Fig. 1b. In this figure, we see that the magnitude of the external forces die out quite rapidly away from the object boundary. Increasing σ in (5) will increase this range, but the boundary localization will become less accurate and distinct, ultimately obliterating the concavity itself when σ becomes too large.

Cohen and Cohen [11] proposed an external force model that significantly increases the capture range of a traditional snake. These external forces are the negative gradient of a potential function that is computed using a Euclidean (or chamfer) distance map. We refer to these forces as *distance potential forces* to distinguish them from the traditional potential forces defined in Section 2.1. Fig. 2 shows the performance of a snake using distance potential forces. Fig. 2a shows both the U-shaped object (in gray) and a sequence of contours (in black) depicting the progression of the snake from its initialization far from the object to its final configuration. The distance potential forces shown in Fig. 2b have vectors with large magnitudes far away from the object, explaining why the capture range is large for this external force model.

As shown in Fig. 2a, this snake also fails to converge to the boundary concavity. This can be explained by inspecting the magnified portion of the distance potential forces shown in Fig. 2c. We see that, like traditional potential forces, these forces also point horizontally in opposite directions, which pulls the snake apart but not downward into the boundary concavity. We note that Cohen and Cohen's modification to the basic distance potential forces, which applies a nonlinear transformation to the distance map [11], does

not change the direction of the forces, only their magnitudes. Therefore, the problem of convergence to boundary concavities is not solved by distance potential forces.

2.3 Generalized Force Balance Equations

The snake solutions shown in Figs. 1a and 2a both satisfy the Euler equations (6) for their respective energy model. Accordingly, the poor final configurations can be attributed to convergence to a local minimum of the objective function (1). Several researchers have sought solutions to this problem by formulating snakes directly from a force balance equation in which the standard external force $\mathbf{F}_{\text{ext}}^{(p)}$ is replaced by a more general external force $\mathbf{F}_{\text{ext}}^{(g)}$ as follows

$$\mathbf{F}_{\text{int}} + \mathbf{F}_{\text{ext}}^{(g)} = 0 \quad (9)$$

The choice of $\mathbf{F}_{\text{ext}}^{(g)}$ can have a profound impact on both the implementation and the behavior of a snake. Broadly speaking, the external forces $\mathbf{F}_{\text{ext}}^{(g)}$ can be divided into two classes: static and dynamic. Static forces are those that are computed from the image data, and do not change as the snake progresses. Standard snake potential forces are static external forces. Dynamic forces are those that change as the snake deforms.

Several types of dynamic external forces have been invented to try to improve upon the standard snake potential forces. For example, the forces used in multiresolution snakes [10] and the pressure forces used in balloons [9] are dynamic external forces. The use of multiresolution schemes and pressure forces, however, adds complexity to a snake's implementation and unpredictability to its performance. For example, pressure forces must be initialized to either push out or push in, and may overwhelm weak boundaries if they act too strongly [16]. Conversely, they may not move into boundary concavities if they are pushing in the wrong direction or act too weakly.

In this paper, we present a new type of *static external force*, one that does not change with time or depend on the position of the snake itself. The underlying mathematical premise for this new force comes from the Helmholtz theorem (cf. [17]), which states that the most general static vector field can be decomposed into two components: an irrotational (curl-free) component and a solenoidal (divergence-free) component.¹ An external potential force generated from the variational formulation of a traditional snake must enter the force balance equation (6) as a static irrotational field, since it is the gradient of a potential function. Therefore, a more general static field $\mathbf{F}_{\text{ext}}^{(g)}$ can be obtained by allowing the possibility that it comprises *both* an irrotational component and a solenoidal component. Our previous paper [15] explored the idea of

¹Irrotational fields are sometimes called conservative fields; they can be represented as the gradient of a scalar potential function.

constructing a separate solenoidal field from an image, which was then added to a standard irrotational field. In the following section, we pursue a more natural approach in which the external force field is designed to have the desired properties of both a large capture range and the presence of forces that point into boundary concavities. The resulting formulation produces external force fields that can be expected to have both irrotational and solenoidal components.

3 Gradient Vector Flow Snake

Our overall approach is to use the force balance condition (7) as a starting point for designing a snake. We define below a new static external force field $\mathbf{F}_{\text{ext}}^{(g)} = \mathbf{v}(x, y)$, which we call the *gradient vector flow* (GVF) field. To obtain the corresponding dynamic snake equation, we replace the potential force $-\nabla E_{\text{ext}}$ in (8) with $\mathbf{v}(x, y)$, yielding

$$\mathbf{x}_t(s, t) = \alpha \mathbf{x}''(s, t) - \beta \mathbf{x}''''(s, t) + \mathbf{v} \quad (10)$$

We call the parametric curve solving the above dynamic equation a *GVF snake*. It is solved numerically by discretization and iteration, in identical fashion to the traditional snake.

Although the final configuration of a GVF snake will satisfy the force-balance equation (7), this equation does not, in general, represent the Euler equations of the energy minimization problem in (1). This is because $\mathbf{v}(x, y)$ will not, in general, be an irrotational field. The loss of this optimality property, however, is well-compensated by the significantly improved performance of the GVF snake.

3.1 Edge Map

We begin by defining an *edge map* $f(x, y)$ derived from the image $I(x, y)$ having the property that it is larger near the image edges.² We can use any gray-level or binary edge map defined in the image processing literature (cf. [18]); for example, we could use

$$f(x, y) = -E_{\text{ext}}^{(i)}(x, y) \quad (11)$$

where $i = 1, 2, 3$, or 4 . Three general properties of edge maps are important in the present context. First, the gradient of an edge map ∇f has vectors pointing toward the edges, which are normal to the edges at the edges. Second, these vectors generally have large magnitudes only in the immediate vicinity of the edges. Third, in homogeneous regions, where $I(x, y)$ is nearly constant, ∇f is nearly zero.

²Other features can be sought by redefining $f(x, y)$ to be larger at the desired features.

Now consider how these properties affect the behavior of a traditional snake when the gradient of an edge map is used as an external force. Because of the first property, a snake initialized close to the edge will converge to a stable configuration near the edge. This is a highly desirable property. Because of the second property, however, the capture range will be very small, in general. Because of the third property, homogeneous regions will have no external forces whatsoever. These last two properties are undesirable. Our approach is to keep the highly desirable property of the gradients near the edges, but to extend the gradient map farther away from the edges and into homogeneous regions using a computational diffusion process. As an important benefit, the inherent competition of the diffusion process will also create vectors that point into boundary concavities.

3.2 Gradient Vector Flow

We define the gradient vector flow field to be the vector field $\mathbf{v}(x, y) = (u(x, y), v(x, y))$ that minimizes the energy functional

$$\mathcal{E} = \iint \mu(u_x^2 + u_y^2 + v_x^2 + v_y^2) + |\nabla f|^2 |\mathbf{v} - \nabla f|^2 dx dy \quad (12)$$

This variational formulation follows a standard principle, that of making the result smooth when there is no data. In particular, we see that when $|\nabla f|$ is small, the energy is dominated by sum of the squares of the partial derivatives of the vector field, yielding a slowly-varying field. On the other hand, when $|\nabla f|$ is large, the second term dominates the integrand, and is minimized by setting $\mathbf{v} = \nabla f$. This produces the desired effect of keeping \mathbf{v} nearly equal to the gradient of the edge map when it is large, but forcing the field to be slowly-varying in homogeneous regions. The parameter μ is a regularization parameter governing the tradeoff between the first term and the second term in the integrand. This parameter should be set according to the amount of noise present in the image (more noise, increase μ).

We note that the smoothing term — the first term within the integrand of (12) — is the same term used by Horn and Schunck in their classical formulation of optical flow [19]. It has recently been shown that this term corresponds to an equal penalty on the divergence and curl of the vector field [20]. Therefore, the vector field resulting from this minimization can be expected to be neither entirely irrotational nor entirely solenoidal.

Using the *calculus of variations* [21], it can be shown that the GVF field can be found by solving the

following Euler equations

$$\mu \nabla^2 u - (u - f_x)(f_x^2 + f_y^2) = 0 \quad (13a)$$

$$\mu \nabla^2 v - (v - f_y)(f_x^2 + f_y^2) = 0 \quad (13b)$$

where ∇^2 is the Laplacian operator. These equations provide further intuition behind the GVF formulation. We note that in a homogeneous region (where $I(x, y)$ is constant), the second term in each equation is zero because the gradient of $f(x, y)$ is zero. Therefore, within such a region, u and v are each determined by Laplace's equation, and the resulting GVF field is interpolated from the region's boundary, reflecting a kind of competition among the boundary vectors. This explains why GVF yields vectors that point into boundary concavities.

3.3 Numerical Implementation

Equations (13a) and (13b) can be solved by treating u and v as functions of time and solving

$$u_t(x, y, t) = \mu \nabla^2 u(x, y, t) - (u(x, y, t) - f_x(x, y))(f_x(x, y)^2 + f_y(x, y)^2) \quad (14a)$$

$$v_t(x, y, t) = \mu \nabla^2 v(x, y, t) - (v(x, y, t) - f_y(x, y))(f_x(x, y)^2 + f_y(x, y)^2) \quad (14b)$$

The steady-state solution of these linear parabolic equations is the desired solution of the Euler equations (13a) and (13b). Note that these equations are decoupled, and therefore can be solved as separate scalar partial differential equations in u and v . The equations in (14) are known as *generalized diffusion equations*, and are known to arise in such diverse fields as heat conduction, reactor physics, and fluid flow [22]. Here, they have appeared from our description of desirable properties of snake external force fields as represented in the energy functional of (12).

For convenience, we rewrite Equation (14) as follows

$$u_t(x, y, t) = \mu \nabla^2 u(x, y, t) - b(x, y)u(x, y, t) + c^1(x, y) \quad (15a)$$

$$v_t(x, y, t) = \mu \nabla^2 v(x, y, t) - b(x, y)v(x, y, t) + c^2(x, y) \quad (15b)$$

where

$$b(x, y) = f_x(x, y)^2 + f_y(x, y)^2$$

$$c^1(x, y) = b(x, y)f_x(x, y)$$

$$c^2(x, y) = b(x, y)f_y(x, y)$$

Any digital image gradient operator (cf. [18]) can be used to calculate f_x and f_y . In the examples shown in this paper, we use simple central differences. The coefficients $b(x, y)$, $c^1(x, y)$, and $c^2(x, y)$, can then be computed and fixed for the entire iterative process.

To set up the iterative solution, let the indices i , j , and n correspond to x , y , and t , respectively, and let the spacing between pixels be Δx and Δy and the time step for each iteration be Δt . Then the required partial derivatives can be approximated as

$$\begin{aligned} u_t &= \frac{1}{\Delta t}(u_{i,j}^{n+1} - u_{i,j}^n) \\ v_t &= \frac{1}{\Delta t}(v_{i,j}^{n+1} - v_{i,j}^n) \\ \nabla^2 u &= \frac{1}{\Delta x \Delta y}(u_{i+1,j} + u_{i,j+1} + u_{i-1,j} + u_{i,j-1} - 4u_{i,j}) \\ \nabla^2 v &= \frac{1}{\Delta x \Delta y}(v_{i+1,j} + v_{i,j+1} + v_{i-1,j} + v_{i,j-1} - 4v_{i,j}) \end{aligned}$$

Substituting these approximations into (15) gives our iterative solution to GVF:

$$u_{i,j}^{n+1} = (1 - b_{i,j}\Delta t)u_{i,j}^n + r(u_{i+1,j}^n + u_{i,j+1}^n + u_{i-1,j}^n + u_{i,j-1}^n - 4u_{i,j}^n) + c_{i,j}^1\Delta t \quad (16a)$$

$$v_{i,j}^{n+1} = (1 - b_{i,j}\Delta t)v_{i,j}^n + r(v_{i+1,j}^n + v_{i,j+1}^n + v_{i-1,j}^n + v_{i,j-1}^n - 4v_{i,j}^n) + c_{i,j}^2\Delta t \quad (16b)$$

where

$$r = \frac{\mu\Delta t}{\Delta x\Delta y} \quad (17)$$

Convergence of the above iterative process is guaranteed by a standard result in the theory of numerical methods (cf. [23]). Provided that b , c^1 , and c^2 are bounded, (16) is stable whenever the Courant-Friedrichs-Lewy step-size restriction $r \leq 1/4$ is maintained. Since normally Δx , Δy , and μ are fixed, using the definition of r in (17) we find that the following restriction on the time-step Δt must be maintained in order to guarantee convergence of GVF:

$$\Delta t \leq \frac{\Delta x \Delta y}{4\mu} \quad (18)$$

The intuition behind this condition is revealing. First, convergence can be made to be faster on coarser images — i.e., when Δx and Δy are larger. Second, when μ is large and the GVF is expected to be a smoother field, the convergence rate will be slower (since Δt must be kept small).

Our 2-D GVF computations were implemented using MATLAB³ code. For an $N = 256 \times 256$ -pixel image on an SGI Indigo-2, typical computation times are 8 seconds for the traditional potential forces

³Mathworks, Natick MA

(written in C), 155 seconds for the distance potential forces (Euclidean distance map, written in MATLAB), and 420 seconds for the GVF forces (written in MATLAB, using \sqrt{N} iterations). The computation time of GVF can be substantially reduced by using optimized code in C or FORTRAN. For example, we have implemented 3-D GVF (see Section 5.2) in C, and computed GVF with 150 iterations on a $256 \times 256 \times 60$ -voxel image in 31 minutes. Accounting for the size difference and extra dimension, we conclude that written in C, GVF computation for a 2-D 256×256 -pixel image would take approximately 53 seconds. Algorithm optimization such as use of the multigrid method should yield further improvements.

4 GVF Fields and Snakes: Demonstrations

This section shows several examples of GVF field computations on simple objects and demonstrates several key properties of GVF snakes. We used $\alpha = 0.6$ and $\beta = 0.0$ for all snakes and $\mu = 0.2$ for GVF. The snakes were dynamically reparameterized to maintain contour point separation to within 0.5–1.5 pixels (cf. [24]). All edge maps used in GVF computations were normalized to the range $[0, 1]$.

4.1 Convergence to Boundary Concavity

In our first experiment, we computed the GVF field for the same U-shaped object used in Figs. 1 and 2. The results are shown in Fig. 3. Comparing the GVF field, shown in Fig. 3b, to the traditional potential force field of Fig. 1b, reveals several key differences. First, like the distance potential force field (Fig. 2b), the GVF field has a much larger capture range than traditional potential forces. A second observation, which can be seen in the closeup of Fig. 3c, is that the GVF vectors within the boundary concavity at the top of the U-shape have a downward component. This stands in stark contrast to both the traditional potential forces of Fig. 1c and the distance potential forces of Fig. 2c. Finally, it can be seen from Fig. 3b that the GVF field behaves in an analogous fashion when viewed from the inside of the object. In particular, the GVF vectors are pointing upward into the “fingers” of the U shape, which represent concavities from this perspective.

Fig. 3a shows the initialization, progression, and final configuration of a GVF snake. The initialization is the same as that of Fig. 2a, and the snake parameters are the same as those in Figs. 1 and 2. Clearly, the GVF snake has a broad capture range and superior convergence properties. The final snake configuration closely approximates the true boundary, arriving at a sub-pixel interpolation through bilinear interpolation of the GVF force field.

4.2 Streamlines

Streamlines are the paths over which free particles move when placed in an external force field. By looking at their streamlines, we can examine the capture ranges and motion inducing properties for various snake external forces. Fig. 4 shows the streamlines of points arranged on a 32×32 grid for the traditional potential forces, distance potential forces, and GVF forces used in the simulations of Figs. 1, 2, and 3.

Several properties can be observed from these figures. First, the capture ranges of the GVF force field and the distance potential force field are clearly much larger than that of the traditional potential force field. In fact, both distance potential forces and GVF forces will attract a snake that is initialized on the image border. Second, it is clear that GVF is the only force providing both a downward force within the boundary concavity at the top of the U-shape and an upward force within the “fingers” of the U-shape. In contrast, both traditional snake forces and distance potential forces provide only sideways forces in these regions. Third, the distance potential forces appear to have boundary points that act as regional points of attraction. In contrast, the GVF forces attract points uniformly toward the boundary.

4.3 Snake Initialization and Convergence

In this section we present several examples that compare different snake models with the GVF snake, showing various effects related to initialization, boundary concavities, and subjective contours. The object under study is the line drawing drawn in gray in both Figs. 5 and 6. This figure may depict, for example, the boundary of a room having two doors at the top and bottom and two alcoves at the left and right. The open doors at the top and bottom represent subjective contours that we desire to connect using the snake (cf. [1]).

The snake results shown in Figs. 5b–d all used the initialization shown in Fig. 5a. We first note that for this initialization, the traditional potential forces were too weak to overpower the snake's internal forces, and the snake shrank to a point at the center of the figure (result not shown). To try to fix this problem, a balloon model with outward pressure forces just strong enough to cause the snake to expand into the boundary concavities was implemented; this result is shown in Fig. 5b. Clearly, the pressure forces also caused the balloon to bulge outward through the openings at the top and bottom, and therefore the subjective contours are not reconstructed well.

The snake result obtained using the distance potential force model is shown in Fig. 5c. Clearly, the capture range is now adequate and the subjective boundaries at the top and bottom are reconstructed well.

But this snake fails to find the boundary concavities at the left and right, for the same reason that it could not proceed into the top of the U-shaped object of the previous sections. The GVF snake result, shown in Fig. 5d, is clearly the best result. It has reconstructed both the subjective boundaries and the boundary concavities quite well. The slight rounding of corners, which can also be seen in Figs. 5b and 5c, is a fundamental characteristic of snakes caused by the regularization coefficients α and β .⁴

The snake results shown in Figs. 6b–d all used the initialization shown in Fig. 6a, which is deliberately placed across the boundary. In this case, the balloon model cannot be sensibly applied because it is not clear whether to apply inward or outward pressure forces. Instead, the result of a snake with traditional potential forces is shown in Fig. 6b. This snake stops at a very undesirable configuration because its only points of contact with the boundary are normal to it and the remainder of the snake is outside the capture range of the other parts of the boundary. The snake resulting from distance potential forces is shown in Fig. 6c. This result shows that although the distance potential force snake possesses an insensitivity to initialization, it is incapable of progressing into boundary concavities. The GVF snake result, shown in Fig. 6d, is again the best result. The GVF snake appears to have both an insensitivity to initialization and an ability to progress into boundary concavities.

5 Gray-level Images and Higher Dimensions

In this section, we describe and demonstrate how GVF can be used in gray-level imagery and in higher dimensions.

5.1 Gray-level Images

The underlying formulation of GVF is valid for gray-level images as well as binary images. To compute GVF for gray-level images, the edge-map function $f(x, y)$ must first be calculated. Two possibilities are $f^{(1)}(x, y) = |\nabla I(x, y)|$ or $f^{(2)}(x, y) = |\nabla(G_\sigma(x, y) * I(x, y))|$, where the latter is more robust in the presence of noise. Other more complicated noise-removal techniques such as median filtering, morphological filtering, and anisotropic diffusion could also be used to improve the underlying edge map. Given an edge-map function and an approximation to its gradient, GVF is computed in the usual way using Equation (16).

Fig. 7a shows a gray-level image of the U-shaped object corrupted by additive white Gaussian noise; the

⁴The effect is only caused by α in this example since $\beta = 0$.

signal-to-noise ratio is 6 dB. Fig. 7b shows an edge-map computed using $f(x, y) = f^{(2)}(x, y)$ with $\sigma = 1.5$ pixels, and Fig. 7c shows the computed GVF field. It is evident that the stronger edge-map gradients are retained while the weaker gradients are smoothed out, exactly as would be predicted by the GVF energy formulation of (12). Superposed on the original image, Fig. 7d shows a sequence of GVF snakes (plotted in a shade of gray) and the GVF snake result (plotted in white). The result shows an excellent convergence to the boundary, despite the initialization from far away, the image noise, and the boundary concavity.

Another demonstration of GVF applied to gray-scale imagery is shown in Fig. 8. Fig. 8a shows a magnetic resonance image (short-axis section) of the left ventricle of a human heart, and Fig. 8b shows an edge map computed using $f(x, y) = f^{(2)}(x, y)$ with $\sigma = 2.5$. Fig. 8c shows the computed GVF, and Fig. 8d shows a sequence of GVF snakes (plotted in a shade of gray) and the GVF snake result (plotted in white), both overlaid on the original image. Clearly, many details on the endocardial border are captured by the GVF snake result, including the papillary muscles (the bumps that protrude into the cavity).

5.2 Higher Dimensions

GVF can be easily generalized to higher dimensions. Let $f(\mathbf{x}) : \mathbf{R}^n \rightarrow \mathbf{R}$ be an edge map defined in \mathbf{R}^n . The GVF field in \mathbf{R}^n is defined as the vector field $\mathbf{v}(\mathbf{x}) : \mathbf{R}^n \rightarrow \mathbf{R}^n$ which minimizes the energy functional

$$\mathcal{E} = \int_{\mathbf{R}^n} \mu |\nabla \mathbf{v}|^2 + |\nabla f|^2 |\mathbf{v} - \nabla f|^2 d\mathbf{x} \quad (19)$$

where the gradient operator ∇ is applied to each component of \mathbf{v} separately. Using the calculus of variations, we find that the GVF field must satisfy the Euler equation

$$\mu \nabla^2 \mathbf{v} - (\mathbf{v} - \nabla f) |\nabla f|^2 = \mathbf{0} \quad (20)$$

where ∇^2 is also applied to each component of the vector field \mathbf{v} separately.

A solution to these Euler equations can be found by introducing a time variable t and finding the steady-state solution of the following linear parabolic partial differential equation

$$\mathbf{v}_t = \mu \nabla^2 \mathbf{v} - (\mathbf{v} - \nabla f) |\nabla f|^2 \quad (21)$$

where \mathbf{v}_t denotes the partial derivative of \mathbf{v} with respect to t . Equation (21) comprises n decoupled scalar linear second order parabolic partial differential equations in each element of \mathbf{v} . Therefore, in principle, it can be solved in parallel. In analogous fashion to the 2-D case, finite differences can be used to approximate the required derivatives and each scalar equation can be solved iteratively.

A preliminary experiment using GVF in three dimensions was carried out using the object shown in Fig. 9a, which was created on a 64^3 grid, and rendered using an isosurface algorithm. The 3-D GVF field was computed using a numerical approximation to (21) and $\mu = 0.15$. This GVF result on the two planes shown in Fig. 9b, is shown projected onto these planes in Figs. 9c and d. The same characteristics observed in 2-D GVF field are apparent here as well.

Next, a deformable surface (cf. [3]) using 3-D GVF was initialized as the sphere shown in Fig. 9e, which is neither entirely inside nor entirely outside the object. Intermediate results after 10 and 40 iterations of the deformable surface algorithm are shown in Figs. 9f and g. The final result after 100 iterations is shown in Fig. 9h. The resulting surface is smoother than the isosurface rendering because of the internal forces in the deformable surface model.

6 Summary and Conclusion

We have introduced a new external force model for active contours and deformable surfaces, which we called the gradient vector flow (GVF) field. The field is calculated as a diffusion of the gradient vectors of a gray-level or binary edge map. We have shown that it allows for flexible initialization of the snake or deformable surface and encourages convergence to boundary concavities.

Further investigations into the nature and uses of GVF are warranted. In particular, a complete characterization of the capture range of the GVF field would help in snake initialization procedures. It would also help to more fully understand the GVF parameter μ , perhaps finding a way to choose it optimally for a particular image, and to understand the interplay between μ and the snake parameters α and β . Finally, the GVF framework might be useful in defining new connections between parametric and geometric snakes, and might form the basis for a new geometric snake.

Acknowledgments

The authors would like to thank Dzung Pham, Sandeep Gupta, and Prof. Joel Spruck for their discussions concerning this work, and the reviewers for providing several key suggestions for improving the paper.

References

- [1] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int. J. Computer Vision*,

- 1(4):321–331, 1987.
- [2] D. Terzopoulos and K. Fleischer. Deformable models. *The Visual Computer*, 4:306–331, 1988.
- [3] T. McInerney and D. Terzopoulos. A dynamic finite element surface model for segmentation and tracking in multidimensional medical images with application to cardiac 4D image analysis. *Computerized Medical Imaging and Graphics*, 19(1):69–83, 1995.
- [4] F. Leymarie and M. D. Levine. Tracking deformable objects in the plane using an active contour model. *IEEE Trans. on Pattern Anal. Machine Intell.*, 15(6):617–634, 1993.
- [5] R. Durikovic, K. Kaneda, and H. Yamashita. Dynamic contour: a texture approach and contour operations. *The Visual Computer*, 11:277–289, 1995.
- [6] D. Terzopoulos and R. Szeliski. Tracking with Kalman snakes. In A. Blake and A. Yuille, editors, *Active Vision*, Artificial Intelligence, pages 3–20. The MIT Press, Cambridge, Massachusetts, 1992.
- [7] V. Caselles, F. Catte, T. Coll, and F. Dibos. A geometric model for active contours. *Numerische Mathematik*, 66:1–31, 1993.
- [8] R. Malladi, J. A. Sethian, and B. C. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Trans. on Pattern Anal. Machine Intell.*, 17(2):158–175, 1995.
- [9] L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, March 1991.
- [10] B. Leroy, I. Herlin, and L. D. Cohen. Multi-resolution algorithms for active contour models. In *12th International Conference on Analysis and Optimization of Systems*, pages 58–65, 1996.
- [11] L. D. Cohen and I. Cohen. Finite-element methods for active contour models and balloons for 2-D and 3-D images. *IEEE Trans. on Pattern Anal. Machine Intell.*, 15(11):1131–1147, November 1993.
- [12] C. Davatzikos and J. L. Prince. An active contour model for mapping the cortex. *IEEE Trans. on Medical Imaging*, 14(1):65–80, March 1995.
- [13] A. J. Abrantes and J. S. Marques. A class of constrained clustering algorithms for object boundary extraction. *IEEE Trans. on Image Processing*, 5(11):1507–1521, November 1996.

- [14] C. Davatzikos and J. L. Prince. Convexity analysis of active contour models. In *Proc. Conf. on Info. Sci. and Sys.*, pages 581–587, 1994.
- [15] J. L. Prince and C. Xu. A new external force model for snakes. In *1996 Image and Multidimensional Signal Processing Workshop*, pages 30–31, 1996.
- [16] H. Tek and B. B. Kimia. Image segmentation by reaction-diffusion bubbles. In *Proc. Fifth Int. Conf. on Computer Vision*, pages 156–162, 1995.
- [17] P. M. Morse and H. Feshbach. *Methods of Theoretical Physics*. McGraw-Hill Book Company, New York, 1953.
- [18] A. K. Jain. *Fundamentals of Digital Image Processing*. Prentice Hall, Engelwood Cliffs, NJ, 1989.
- [19] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [20] S. N. Gupta and J. L. Prince. Stochastic models for DIV-CURL optical flow methods. *IEEE Signal Processing Letters*, 3(2):32–35, 1996.
- [21] R. Courant and D. Hilbert. *Methods of Mathematical Physics*, volume 1. Interscience, New York, 1953.
- [22] A. H. Charles and T. A. Porsching. *Numerical Analysis of Partial Differential Equations*. Prentice Hall, Engelwood Cliffs, NJ, 1990.
- [23] W. F. Ames. *Numerical Methods for Partial Differential Equations*. Academic Press, Boston, 3rd edition, 1992.
- [24] S. Lobregt and M. A. Viergever. A discrete dynamic contour model. *IEEE Trans. on Medical Imaging*, 14(1):12–24, March 1995.

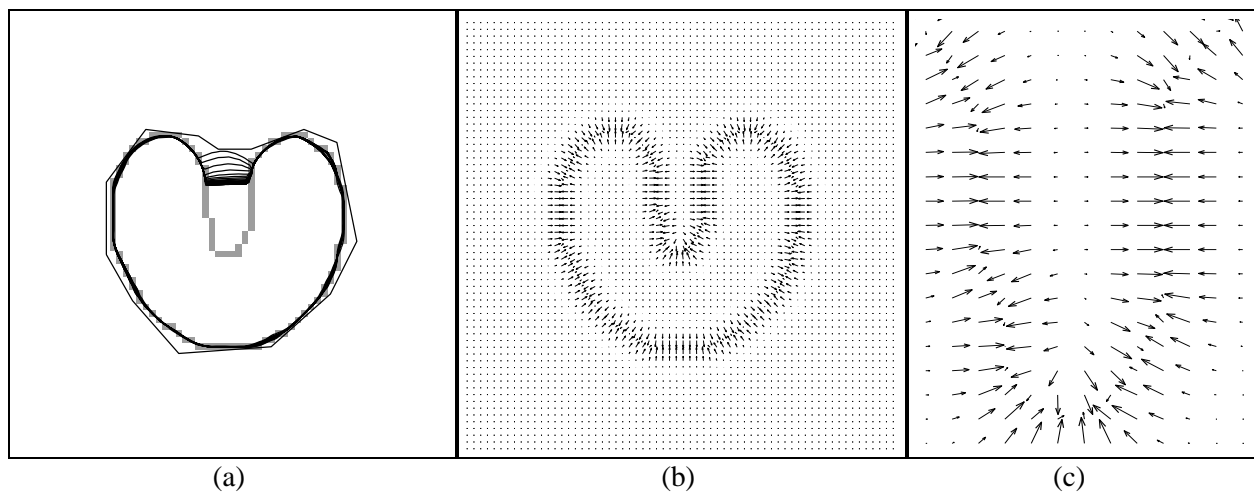


Figure 1: (a) The convergence of a snake using (b) traditional potential forces, (c) shown close-up within the boundary concavity.

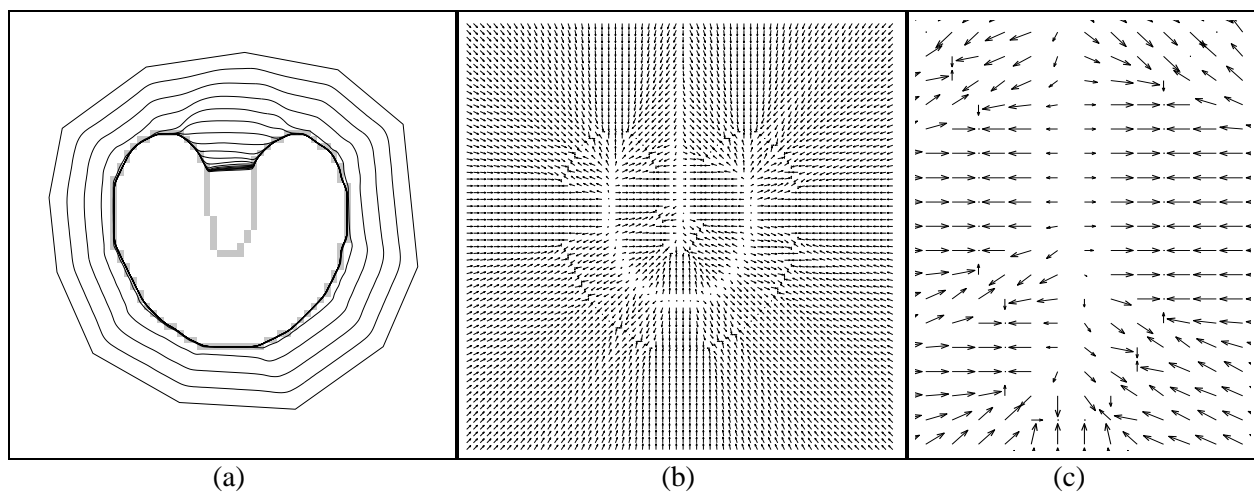


Figure 2: (a) The convergence of a snake using (b) distance potential forces, (c) shown close-up within the boundary concavity.




Figure 3: (a) The convergence of a snake using (b) GVF external forces, (c) shown close-up within the boundary concavity.




Figure 4: Streamlines originating from an array of 32×32 particles in (a) a traditional potential force field, (b) a distance potential force field, and (c) a GVF force field.




Figure 5: (a) An initial curve and snake results from (b) a balloon with an outward pressure, (c) a distance potential force snake, and (d) a GVF snake.




Figure 6: (a) An initial curve and snake results from (b) a traditional snake, (c) a distance potential force snake, and (d) a GVF snake.




Figure 7: (a) A noisy 64×64 -pixel image of a U-shaped object; (b) the edge map $|\nabla(G_\sigma * I)|^2$ with $\sigma = 1.5$; (c) the GVF external force field; and (d) convergence of the GVF snake.




Figure 8: (a) A 160×160 -pixel magnetic resonance image of the left ventricle of a human heart; (b) the edge map $|\nabla(G_\sigma * I)|^2$ with $\sigma = 2.5$; (c) the GVF field (shown subsampled by a factor of two); and (d) convergence of the GVF snake.

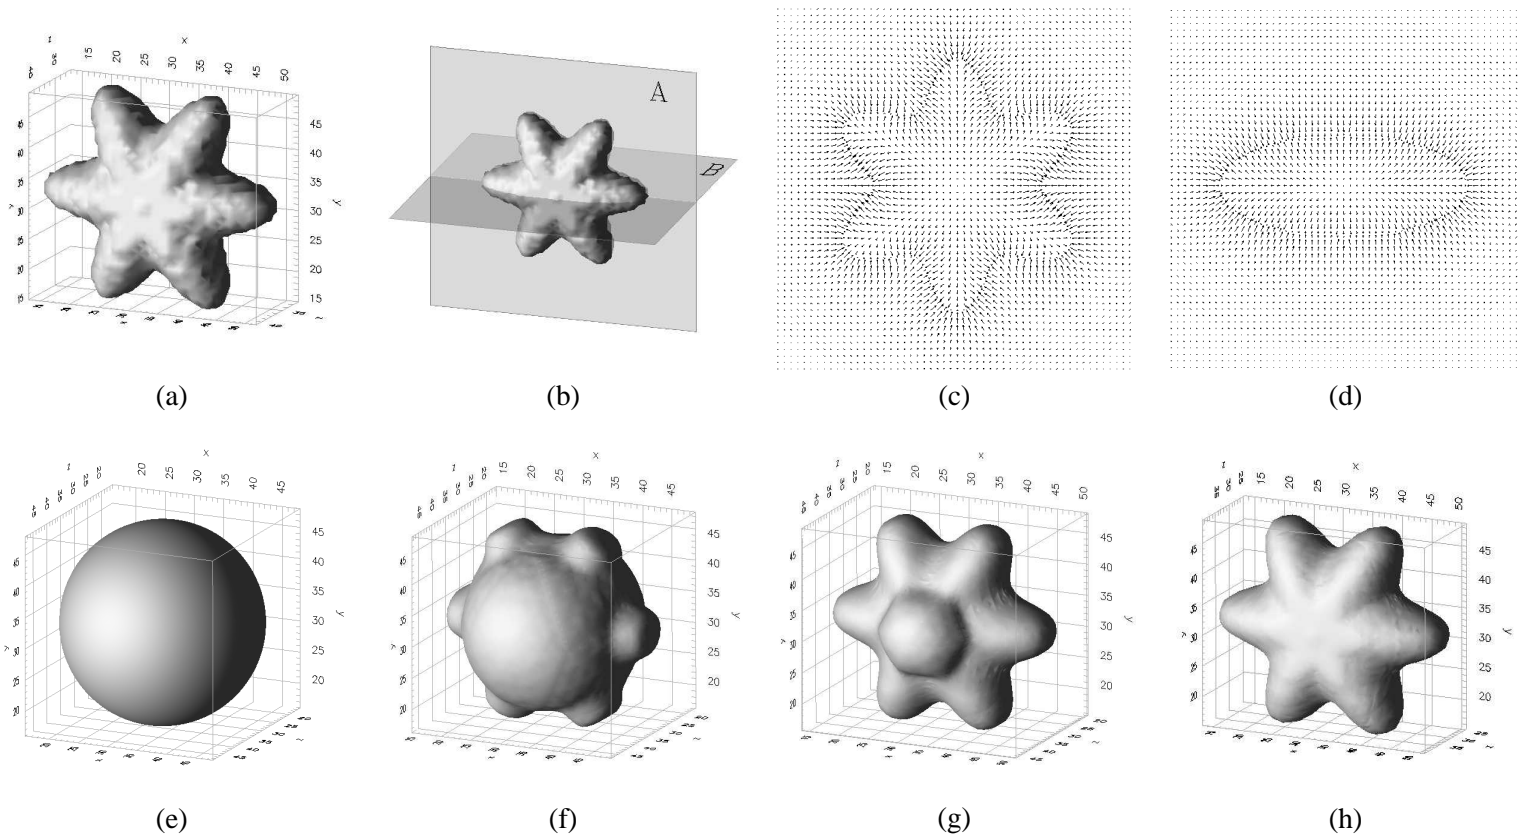


Figure 9: (a) Isosurface of a 3-D object defined on a 64^3 grid; (b) positions of planes A and B on which the 3-D GVF vectors are depicted in (c) and (d), respectively; (e) the initial configuration of a deformable surface using GVF and its positions after (f) 10, (g) 40, and (h) 100 iterations.