

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 15-12-2016		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 01Sep06 - 09May11	
4. TITLE AND SUBTITLE Final Report for Geometric Observers & Particle Filtering for Controlled Active Vision				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
				5d. PROJECT NUMBER 49414-NS.1	
6. AUTHOR(S) Allen R. Tannenbaum				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Georgia Tech Research Corporation 505 Tenth Street NW Atlanta, GA 30332-0420				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) US Army Research Office PO Box 12211 RTP, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Distribution is Unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Curvature driven flows based on the the minimization of certain geometric functionals have been used for a number of problems in visual control and computer vision in the past few years [103, 88, 109]. These flows themselves are very much motivated by ideas in optimal control; see [69]. We have summarized some of the key flows which we will need in the sequel.					
15. SUBJECT TERMS Final Report					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Allen R. Tannenbaum
a. REPORT U	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)

Final Report for Geometric Observers and Particle Filtering for Controlled Active Vision

by

Allen R. Tannenbaum
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia
(404) 894-7574

Table of Contents

Contents

1	Statement of Problem	3
2	Summary of Key Results	4
2.1	Brief Review of Geometric Flows in Vision and Image Processing	4
2.1.1	Algorithms and PDEs	4
2.1.2	Curve Evolution Using Level Sets	5
2.2	Active Vision and Visual Tracking	5
2.2.1	Snakes	6
2.2.2	Geometric Active Contours	6
2.2.3	Conformal (Geodesic) Active Contours	7
2.2.4	Conformal Area Minimizing Flows	8
2.3	Particle Filters	8
2.4	Optimal Transport	9
2.4.1	Registration and Optimal Transport	10
2.4.2	Variational Approach to Monge-Kantorovich Problem	10
2.4.3	Other Applications of Optimal Transport	12
3	Bibliography	12

1 Statement of Problem

The key objective of this project was the development of new methodologies for employing visual information in a feedback loop, the underlying problem of controlled active vision. Controlled active vision, and in particular visual tracking requires the integration of techniques from control theory, signal processing, and computer vision. For some time now the role of control theory in vision has been recognized. In particular, the branches of control that deal with system uncertainty, namely adaptive and robust, have been proposed as essential tools in coming to grips with the problems of both machine and biological vision.

Visual tracking provides a fundamental example of the need for controlled active vision. While tracking in the presence of a disturbance is a classical control problem, visual tracking raises new issues. First since cameras are part of the system, one must consider the nature of the disturbance from imaging sensors. The feedback signal may require some interpretation of the image, e.g., segmentation of a target from its background, or an inference about an occluder. In this project, we expressly emphasized active vision, because the result may be viewpoint dependent. In particular, calibration may influence the control law. Finally, as visual processing becomes more complex, the issue of processing time arises. Each of these problems must be answered before target detection, and visually-mediated control can be provided for advanced weapon systems.

There are a number of technical problems connected with using vision in a closed loop setting. In general, many of the mathematical formulations of the key issues of image understanding may be ill-posed. Many basic vision tasks such as segmentation and a robust theory of shape remain largely unsolved. For visual tracking, real-time processing becomes a major concern. Despite these formidable obstacles, progress is being made especially in integrating control ideas into the vision framework. Control is very powerful in treating uncertainty, and with the newer partial differential equation methodologies, there is a strong mathematical and systems-theoretic fit. In fact many of the recent curvature flows may be derived from principles of optimal control. Until very recently, much of the control work in vision was performed by researchers in computer vision usually with a computer science background. Now, however, researchers with a strong control/systems background are getting more actively involved given the opportunities for modern control ideas in vision. We believe that this welcome development will significantly help in alleviating many of the formidable problems remaining in using visual information in a feedback loop, and having more reliable robust systems in which the primary sensor is based largely on some imaging device.

The prevalence of biological vision in even very simple organisms, indicates its inherent utility for man-made devices. Cameras are rather simple, reliable passive sensing devices which are quite inexpensive per bit of data. Furthermore vision can offer information at a high rate with high resolution with a wide field of view and accuracy capturing multispectral information. Finally cameras can be used in a more active manner. Namely, one can include motorized lenses and mounted on mobile platforms which can actively explore the surrounds and suitably adapt their sensing capabilities. Computer vision has formulated several approaches for interpreting the signals, opening the possibility that control laws can be based on more abstracted descriptions. These problems all become manifest when one attempts to use a visual sensor in an uncertain environment, and to feed back in some manner the information. These issues represent some of the main challenges for our research in controlled active vision.

A major goal of our program is to ensure that *our research is directly driven by applications and that our research results have a direct impact on the military and industrial base*. Thus we need to ensure that all of the mathematically based results are amenable to robust and reliable computer implementations.

In this research program, we continued the use of curvature-driven partial differential equations to explore problems in visual control and controlled active vision, but we have also considered the sensor/estimation problem much more carefully. This has led to a new approach for tracking in which particle filtering was combined with level set ideas for the first time.

2 Summary of Key Results

2.1 Brief Review of Geometric Flows in Vision and Image Processing

Curvature driven flows based on the the minimization of certain geometric functionals have been used for a number of problems in visual control and computer vision in the past few years [103, 88, 109]. These flows themselves are very much motivated by ideas in *optimal control*; see [69]. We will summarize some of the key flows below which we will need in the sequel.

2.1.1 Algorithms and PDEs

We briefly review the major concepts involved in using partial differential equations (PDEs) for image processing and computer vision.

As explained in detail in [20], one can think of an image as a map $I : \mathfrak{D} \rightarrow \mathfrak{C}$, i.e., to any point $\mathbf{x}$ in the domain $\mathfrak{D}$, I associates a “color” $I(\mathbf{x})$ in a color space $\mathfrak{C}$. For ease of presentation we will mainly restrict ourselves to the case of a two-dimensional gray scale image which we can think of as a function from a domain $\mathfrak{D} = [0, 1] \times [0, 1] \subset \mathbf{R}^2$ to the unit interval $\mathfrak{C} = [0, 1]$.

The algorithms all involve solving the initial value problem for some PDE for a given amount of time. The solution to this PDE can be either the image itself at different stages of modification, or some other object (such as a closed curve delineating object boundaries) whose evolution is driven by the image.

For example, introducing an artificial time t , the image can be deformed according to

$$\frac{\partial I}{\partial t} = \mathcal{F}[I], \quad (1)$$

where $I(\mathbf{x}, t) : \mathfrak{D} \times [0, T] \rightarrow \mathfrak{C}$ is the evolving image, $\mathcal{F}$ is an operator which characterizes the given algorithm, and the initial condition is the input image I_0 . The processed image is the solution $I(\mathbf{x}, t)$ of the differential equation at time t . The operator $\mathcal{F}$ usually is a differential operator, although its dependence on I may also be nonlocal.

Similarly, one can evolve a closed curve $\mathcal{C} \subset \mathfrak{D}$ representing the boundaries of some planar shape ($\mathcal{C}$ need not be connected and could have several components). In this case, the operator $\mathcal{F}$ specifies the normal velocity of the curve that it deforms. In many cases this normal velocity is a function of the curvature κ of $\mathcal{C}$, and of the image I evaluated on $\mathcal{C}$. A flow of the form

$$\frac{\partial \mathcal{C}}{\partial t} = \mathcal{F}(I, \kappa) \mathcal{N} \quad (2)$$

is obtained, where $\mathcal{N}$ is the inward unit normal to the curve $\mathcal{C}$.

Very often, the deformation is obtained as the steepest descent for some energy functional. For example, the energy

$$\mathcal{E}(I) = \frac{1}{2} \int \|\nabla I\|^2 dx dy \quad (3)$$

and its associated steepest descent, the heat equation,

$$\frac{\partial I}{\partial t} = \Delta I \quad (4)$$

correspond to the classical Gaussian smoothing.

The use of PDEs allows for the modelling of the crucial but poorly understood interactions between top-down and bottom-up vision. For example, in a variational framework, an energy $\mathcal{E}$ is defined globally while the corresponding operator $\mathcal{F}$ will influence the image locally. Algorithms defined in terms of PDEs treat images as continuous rather than discrete objects. This has a simplifying effect on the formalism, which becomes grid independent. On the other hand models based on nonlinear PDEs may be much harder to rigorously analyze and implement.

2.1.2 Curve Evolution Using Level Sets

Geometric active contours evolving according to an edge based and/or region based energy flow are very commonly used for image segmentation. In these methods, starting from an initial estimate, the curve deforms under the influence of various forces until it fits the object boundaries. The curve evolution equation is obtained by reducing an energy E_{image} as fast as possible, i.e., by doing a gradient descent on E_{image} . In general, E_{image} may depend on a combination of image based features and external constraints (smoothness, shape, etc.); see [82], [21] and the references therein. The level set methods of Osher and Sethian [87] offer a natural and numerically reliable implementation of such curve evolution equations. Level sets have the advantage of being parameter independent (i.e., they are implicit representation of the curve) and can handle topological changes in a very natural way.

We now briefly go over the level set representation of a given curve evolution equation. Let $\mathcal{C}(p, t) : [0, 1] \times [0, T) \rightarrow \mathbf{R}^2$ be a family of closed curves (i.e., $\mathcal{C}(0, t) = \mathcal{C}(1, t)$ for all $t \in [0, T)$), satisfying the following evolution equation:

$$\frac{\partial \mathcal{C}}{\partial t} = V \mathcal{N} \quad (5)$$

where, t is the time parameter. The basic idea of the level set approach is to embed the contour $\mathcal{C}(p, t)$ as the zero level set of a smooth and Lipschitz continuous function $\Phi : \mathbf{R}^2 \times [0, T) \rightarrow \mathbf{R}$. Assume that Φ is negative in the interior and positive in the exterior of the zero level set. We consider the zero level set, defined by

$$\{Z(t) \in \mathbf{R}^2 : \Phi(Z, t) = 0\}. \quad (6)$$

We have to find an evolution equation of Φ , such that the evolving curve $\mathcal{C}_t$ is given by the evolving zero level set $Z(t)$. By differentiating (6) with respect to t we obtain:

$$\nabla \Phi(Z, t) \cdot \frac{\partial Z}{\partial t} + \frac{\partial \Phi}{\partial t} = 0. \quad (7)$$

Note that for the zero level set, the following relation holds:

$$\frac{\nabla \Phi}{\|\nabla \Phi\|} = -\mathcal{N}. \quad (8)$$

In this equation, the left side uses terms of the surface Φ , while the right side is related to the curve $\mathcal{C}$. The combination of equations (5) to (8) gives

$$\frac{\partial \Phi}{\partial t} = V \|\nabla \Phi\|, \quad (9)$$

and the curve $\mathcal{C}$, evolving according to (5), is obtained by the zero level set of the function Φ , which evolves according to (9). This is a Hamilton-Jacobi equation which can be analyzed using viscosity theory [27]. Finally, given an initial curve, one must generate an initial level set function. A well known scheme [87] is to use a signed distance function.

2.2 Active Vision and Visual Tracking

The use of active contour methods was essential in our research program in controlled active vision. We now go over some of the key results pertaining to this methodology.

2.2.1 Snakes

The concept of *snakes* (also called *deformable* or *active contours*) was introduced by Witkin, Kass and Terzopoulos [65], and later developed by a number of researchers (see [26] and the references therein). They may be used for edge detection in the following manner. Given an image $I : \mathfrak{D} \subset \mathbf{R}^2 \rightarrow \mathfrak{C}$, one subjects an initial simple closed parameterized curve $\mathcal{C} : [0, 1] \rightarrow \mathfrak{D}$ to a steepest descent flow for an energy functional of the form:

$$\mathcal{E}(\mathcal{C}) = \int_0^1 \left\{ \frac{1}{2} w_1(p) \|\mathcal{C}_{pp}\|^2 + \frac{1}{2} w_2(p) \|\mathcal{C}_p\|^2 + W(\mathcal{C}(p)) \right\} dp. \quad (10)$$

(Here the p subscript indicates differentiation with respect to p .) The first two terms control the smoothness of the active contour $\mathcal{C}$. The contour interacts with the image through the potential function $W : \mathfrak{D} \rightarrow \mathbf{R}$ which is chosen to be small near the edges, and large everywhere else (it is a decreasing function of some edge detector). For example, one could take:

$$W(\mathbf{x}) = \frac{1}{1 + \|\nabla G_\sigma * I(\mathbf{x})\|^2}, \quad (11)$$

where G_σ denotes a Gaussian filter of standard deviation σ .

Minimizing $\mathcal{E}$ will therefore attract $\mathcal{C}$ toward the edges. The gradient flow is the fourth order nonlinear parabolic equation

$$\frac{\partial \mathcal{C}}{\partial t} = - (w_2(p) \mathcal{C}_{pp})_{pp} + (w_1(p) \mathcal{C}_p)_p + \nabla W(\mathcal{C}(p, t)). \quad (12)$$

This approach gives reasonable results; see [78] for a survey of snakes in medical image analysis. One limitation however is that the active contour or snake cannot change topology, i.e., it starts out being a topological circle and it will always remain a topological circle and will not be able to break up into two or more pieces, even if the image would contain two unconnected objects and this would give a more natural description of the edges. Special *ad hoc* procedures have been developed in order to handle splitting and merging [79].

2.2.2 Geometric Active Contours

Another disadvantage of the snake method is that it explicitly involves the parametrization of the active contour $\mathcal{C}$, while there is no obvious relation between the parametrization of the contour and the geometry of the objects to be captured. Geometric models have been developed in [18] to address this issue.

As in the snake framework, one deforms the active contour $\mathcal{C}$ by a velocity which is essentially defined by a curvature term, and a constant inflationary term weighted by a stopping function W . By formulating everything in terms of quantities which are invariant under reparametrization (such as the curvature and normal velocity of $\mathcal{C}$) one obtains an algorithm which does not depend on the parametrization of the contour. In particular, it can be implemented using level sets.

More specifically, the model of [18] is given by

$$V = W(\mathbf{x})(\kappa + c), \quad (13)$$

where both the velocity V and the curvature κ are measured using the inward normal $\mathcal{N}$ for $\mathcal{C}$. Here, as previously, W is small at edges and large everywhere else, and c is a constant, called the *inflationary parameter*. When c is positive, it helps push the contour through concavities, and can speed up the segmentation process. When it is negative, it allows expanding “bubbles,” i.e., contours which expand rather than contract to the desired boundaries. We should note that there is no canonical choice for the constant c , which has to be determined experimentally.

In practice, $\mathcal{C}$ is deformed using the Osher-Sethian level set method described in Section 2.1.2. Geometric active contours have the advantage that they allow for topological changes (splitting and merging) of the active contour $\mathcal{C}$. The main problem with this model is that the desired edges are not steady states for the flow (13). The effect of the factor $W(\mathbf{x})$ is merely to slow the evolving contour $\mathcal{C}_t$ down as it approaches an edge, but it is not the case that the $\mathcal{C}_t$ will eventually converge to anything like the sought-for edge as $t \rightarrow \infty$. Some kind of artificial intervention is required to stop the evolution when $\mathcal{C}_t$ is close to an edge.

2.2.3 Conformal (Geodesic) Active Contours

In [66, 19], the authors propose a novel technique that is both geometric and variational. In this approach one defines a Riemannian metric g^W on $\mathfrak{D}$ from a given image $I : \mathfrak{D} \rightarrow \mathbf{R}$, by conformally changing the standard Euclidean metric to,

$$g^W = W(\mathbf{x})^2 \|\mathrm{d}\mathbf{x}\|^2. \quad (14)$$

Here W is defined as above in (11). The length of a curve in this metric is

$$\mathcal{L}^W(\mathcal{C}) = \int_{\mathcal{C}} W(\mathcal{C}(s)) \, \mathrm{d}s. \quad (15)$$

(Here $\mathrm{d}s$ denotes arc-length.) Curves which minimize this length will prefer to be in regions where W is small, which is exactly where one would expect to find the edges. So, to find edges, one should minimize the W -weighted length of a closed curve $\mathcal{C}$, rather than some “energy” of $\mathcal{C}$ (which depends on a parametrization of the curve).

To minimize $\mathcal{L}^W(\mathcal{C})$, one computes a gradient flow in the L^2 sense. Since the first variation of this length functional is given by

$$\frac{\mathrm{d}\mathcal{L}^W(\mathcal{C})}{\mathrm{d}t} = - \int_{\mathcal{C}} V \{ W\kappa - \mathcal{N} \cdot \nabla W \} \, \mathrm{d}s,$$

where V is the normal velocity measured in the Euclidean metric, and $\mathcal{N}$ is the Euclidean unit normal, the corresponding L^2 gradient flow is

$$V_{\text{conf}} = W\kappa - \mathcal{N} \cdot \nabla W. \quad (16)$$

Note that this is not quite the curve shortening flow in the sense of [48, 49] on $\mathbf{R}^2$ given the Riemannian manifold structure defined by the conformally Euclidean metric g^W . Indeed, a simple computation shows that in that case one would have

$$V = W^{-2}(W\kappa - \mathcal{N} \cdot \nabla W). \quad (17)$$

Thus the term “geodesic active contour” used in [19] is a bit of a misnomer, and so we prefer the term “conformal active contour” as in [66]. However, following standard practice in the vision community, we will use both terms interchangeably in this report.

Contemplation of the conformal active contours leads to another interpretation of the concept “edge.” Using the landscape metaphor one can describe the graph of W as a plateau (where $\|\nabla I\|$ is small) in which a canyon has been carved (where $\|\nabla I\|$ is large). The edge is to be found at the bottom of the canyon. Now if W is a Morse function, then one expects the “bottom of the canyon” to consist of local minima of W alternated by saddle points. The saddle points are connected to the minima by their unstable manifolds for the gradient flow of W (the ODE $\mathbf{x}' = -\nabla W(\mathbf{x})$.) Together these unstable manifolds form one or more closed curves which one may regard as the edges which are to be found.

Comparing (16) to the evolution of the geometric active contour (13) we see that we have the new term $-\mathcal{N} \cdot \nabla W$, the normal component of $-\nabla W$. If the contour $\mathcal{C}_t$ were to evolve only by $V = -\mathcal{N} \cdot \nabla W$, then it would simply be deformed by the gradient flow of W . If W is a Morse function, then one can use the λ -lemma from dynamical systems [89, 81] to show that for a generic choice of initial contour the $\mathcal{C}_t$ will converge to the union of unstable manifolds “at the bottom of the canyon,” possibly with multiplicity more than one. The curvature term in (13) counteracts this possible doubling up and guarantees that $\mathcal{C}_t$ will converge smoothly to some curve which is a smoothed out version of the heteroclinic chain.

2.2.4 Conformal Area Minimizing Flows

Typically, in order to get expanding bubbles, an inflationary term is added in the model (16) as in (13). Many times segmentations are more easily performed by seeding the image with bubbles rather than contracting snakes. The conformal active contours will not allow this since very small curves will simply shrink to points under the flow (16). To get a curve evolution which will force small bubbles to expand and converge toward the edges, it is convenient to subtract a weighted area term from the length functional $\mathcal{L}^W$, namely

$$\mathcal{A}^W(\mathcal{C}) = \int_{R_{\mathcal{C}}} W(\mathbf{x}) d\mathbf{x}$$

where $d\mathbf{x}$ is 2D Lebesgue measure, and $R_{\mathcal{C}}$ is the region enclosed by the contour $\mathcal{C}$.

The first variation of this weighted area is [106, 92]):

$$\frac{d}{dt} \mathcal{A}^W(\mathcal{C}_t) = - \int_{\mathcal{C}_t} W(\mathcal{C}(s)) V ds \quad (18)$$

where, as before, V is the normal velocity of $\mathcal{C}_t$ measured with the inward normal.

The functional which one now tries to minimize is

$$\mathcal{E}^W(\mathcal{C}) = \mathcal{L}^W(\mathcal{C}) + c \mathcal{A}^W(\mathcal{C}), \quad (19)$$

where $c \in \mathbf{R}$ is a constant called the *inflationary parameter*.

To obtain steepest descent for $\mathcal{E}^W$ one sets

$$V_{\text{act}} = V_{\text{conf}} + cW = (\kappa + c)W(\mathbf{x}) - \mathcal{N} \cdot \nabla W. \quad (20)$$

For $c = 1$ this is a conformal length/area minimizing flow (see [106]). As in the model of [18] the inflationary parameter c may be chosen as positive (snake or inward moving flow) or negative (bubble or outward moving flow).

In practice, for expanding flows (negative c , weighted area maximizing flow), one expands the bubble just using the inflationary part

$$V = cW$$

until the active contour is sufficiently large, and then “turns on” the conformal part V_{conf} which brings the contour to its final position. Again as in [18], the curvature part of V_{act} also acts to regularize the flow. Finally, there is a detailed mathematical analysis of (20) in [66] as well as extensions to three dimensional space in which case the curvature κ is replaced by the mean curvature H in equation (20).

2.3 Particle Filters

We have proposed the use of particle filters in conjunction with geometric active contours.

Let $x_t \in \mathbf{R}^n$ be a state vector evolving according to the following equation:

$$x_{t+1} = f_t(x_t, u_t)$$

where u_t is i.i.d. random noise with known pdf. At discrete times, observations $Y_t \in \mathbf{R}^p$ become available. These measurements are related to the state vector via the observation equation:

$$Y_t = h_t(x_t, v_t)$$

where v_t is measurement noise which is known *a priori*. It is assumed that the initial state distribution denoted by $\pi_0(dx)$, the state transition kernel by $K_t(x_t, dx_{t+1})$ and the observation likelihood given the state, by $g_t(Y_t|x_t)$ are known. The particle filter (PF) [47, 36] is a sequential Monte Carlo method which produces at each time t , a cloud of n particles, $\{x_t^{(i)}\}_{i=1}^n$, whose empirical measure closely “follows” $\pi_t(dx_t|Y_{0:t})$, the posterior distribution of the state given past observations (denoted by $\pi_{t|t}(dx)$).

The PF was first introduced in [47] as the Bayesian Bootstrap filter and its first application to tracking in computer vision was the CONDENSATION algorithm [59]. The particle filter [36] recursively approximates the posterior distribution of the state at any time t given the past observations, by Monte Carlo sampling. It works for any linear or non-linear, Gaussian or non-Gaussian dynamical system for which π_0 , $K_t(x_t, dx_{t+1})$ is known and can be sampled from and $g_t(y_t|x_t)$ is known.

The PF starts with sampling n times from the initial state distribution $\pi_0(dx)$ to approximate it by $\pi_0^n(dx) = \frac{1}{n} \sum_{i=1}^n \delta_{x_0^{(i)}}(dx)$ and then *implements the Bayes’ recursion* at each time step. Now, the distribution of x_{t-1} given observations up to time $t-1$ can be approximated by $\pi_{t-1|t-1}^n(dx) = \frac{1}{n} \sum_{i=1}^n \delta_{x_{t-1}^{(i)}}(dx)$. The **prediction step** samples the new state $\bar{x}_t^{(i)}$ from the distribution $K_{t-1}(x_{t-1}^{(i)}, \cdot)$. The empirical distribution of this new cloud of particles, $\pi_{t-1|t-1}^n(dx) = \frac{1}{n} \sum_{i=1}^n \delta_{\bar{x}_t^{(i)}}(dx)$ is an approximation to the conditional probability distribution of x_t given observations up to time $t-1$ (*prediction distribution*).

In the **update step**, each particle is weighted in proportion to the likelihood of the observation at t , Y_t , i.e.

$$w_t^{(i)} = \frac{g_t(Y_t|\bar{x}_t^{(i)})}{\sum_{i=1}^n g_t(Y_t|\bar{x}_t^{(i)})}$$

$\bar{\pi}_{t|t}^n(dx) = \frac{1}{n} \sum_{i=1}^n w_t^{(i)} \delta_{\bar{x}_t^{(i)}}(dx)$ is then an estimate of $\pi_{t|t}$ (*filtering distribution*). One resamples n times with replacement from $\bar{\pi}_{t|t}^n(dx)$ to obtain the empirical estimate $\pi_{t|t}^n(dx) = \frac{1}{n} \sum_{i=1}^n \delta_{x_t^{(i)}}(dx)$. Note that both $\bar{\pi}_{t|t}^n$ and $\pi_{t|t}^n$ approximate $\pi_{t|t}$ but the resampling step is used because it increases the sampling efficiency by eliminating samples with very low weights.

2.4 Optimal Transport

The mass transport problem was first formulated by Gaspar Monge in 1781, and concerned finding the optimal way, in the sense of minimal transportation cost, of moving a pile of soil from one site to another. This problem was given a modern formulation in the work of Kantorovich [64], and so is now known as the *Monge–Kantorovich problem*. The problem of optimal transport has appeared in econometrics, fluid dynamics, automatic control, transportation, statistical physics, shape optimization, expert systems, and meteorology [94].

Very importantly optimal transport naturally fits into certain problems in computer vision [41]. In particular, for the general visual tracking problem, a robust and reliable object and shape recognition system is of major importance. A key way to carry this out is via *template matching*, which is the matching of some object to another within a given catalogue of objects. Typically,

the match will not be exact and hence some criterion is necessary to measure the “goodness of fit.” For a description of various matching procedures, see [53] and the references therein. The matching criterion can also be considered a *shape metric* for measuring the similarity between two objects. In fact, in his doctoral thesis [41], Fry uses optimal transport ideas to construct precisely such a shape metric for planar shapes whose boundaries can be described as closed contours.

2.4.1 Registration and Optimal Transport

We have shown how ideas from optimal transport may be used to formulate a new approach to image interpolation and in particular, new approaches to the computation of optical flow, image morphing, and registration.

Image registration is a process of aligning images so that corresponding features can be easily related. The images may come from different modalities or from the same modality at different times or both. The collection of papers in [111] gives an excellent overview of the field.

Registration proceeds in several steps. First, each image or data set to be matched should be individually calibrated, corrected for imaging distortions and artifacts, and cleared of noise. Next, a measure of similarity between the data sets must be established, so that one can quantify how close an image is from another after transformations are applied. Such a measure may include the similarity between pixel intensity values, as well as the proximity of predefined image features such as implanted fiducials, anatomical landmarks, surface contours, and ridge lines. Next, the transformation that maximizes the similarity between the transformed images is found. Often this transformation is given as the solution of an optimization problem where the transformations to be considered are constrained to be of a predetermined class $\mathcal{U}$. In the case of *rigid registration*, $\mathcal{U}$ is the set of Euclidean transformations. Many deformations are not of the class (e.g., the swelling of tissues in the body or the elastic part of the motion of a jelly fish). Therefore a more realistic and more challenging problem is *elastic registration* where $\mathcal{U}$ is the set of smooth diffeomorphisms. For example, in the medical imaging context, due to anatomical variability, elastic deformation maps are also useful when comparing images from different patients. Finally, once an optimal transformation is obtained, it is used to fuse the image data sets.

2.4.2 Variational Approach to Monge-Kantorovich Problem

There have been a number of algorithms considered for computing an optimal transport map. For example, methods have been proposed based on linear programming [94], and on Lagrangian mechanics closely related to ideas from the study of fluid dynamics [13]. An interesting geometric method has been formulated by Cullen and Purser [29]. In our case for image tracking, an effective, robust method may be based on gradient descent and the concept of “polar factorization”; see [15, 42, 76] for details about polar factorizations.

In order to motivate our approach for computing an optimal transport map developed in our Army Research Office sponsored program, we will consider the specific problem of elastic registration as in [52, 9] in which the similarity between two images is measured by their L^2 Kantorovich–Wasserstein distance. Finding “the best map” from one image to another then leads to an optimal transport problem.

In the approach of [52] one thinks of a gray scale image as a Borel measure¹ μ on some open domain $\mathfrak{D} \subset \mathbf{R}^d$, where, for any $E \subset \mathfrak{D}$, the “amount of white” in the image contained in E is given by $\mu(E)$. Given two images $(\mathfrak{D}_0, \mu_0)$ and $(\mathfrak{D}_1, \mu_1)$ one considers all maps $u : \mathfrak{D}_0 \rightarrow \mathfrak{D}_1$ which preserve the measures, i.e. which satisfy²

$$u_{\#}(\mu_0) = \mu_1, \tag{21}$$

¹This can be physically motivated. For example, in some forms of MRI the image intensity is the actual proton density.

²If X and Y are sets with σ -algebras $\mathcal{M}$ and $\mathcal{N}$, and if $f : X \rightarrow Y$ is a measurable map, then we write $f_{\#}\mu$ for the pushforward of any measure μ on $(X, \mathcal{M})$, i.e., for any measurable $E \subset Y$ we define $f_{\#}\mu(E) = \mu(f^{-1}(E))$.

and one is required to find the map (if it exists) which minimizes the Monge-Kantorovich transport functional (see the exact definition below).

More precisely, assuming that the cost of moving a mass m from $\mathbf{x} \in \mathbf{R}^d$ to $\mathbf{y} \in \mathbf{R}^d$ is $m \cdot \Phi(\mathbf{x}, \mathbf{y})$, where $\Phi : \mathbf{R}^d \times \mathbf{R}^d \rightarrow \mathbf{R}^+$ is called the *cost function*, Kantorovich [64] defined the cost of moving the measure μ_0 to the measure μ_1 by the map $u : \mathfrak{D}_0 \rightarrow \mathfrak{D}_1$ to be

$$M(u) = \int_{\mathfrak{D}_0} \Phi(\mathbf{x}, u(\mathbf{x})) d\mu_0(\mathbf{x}). \quad (22)$$

The infimum of $M(u)$ taken over all measure preserving $u : (\mathfrak{D}_0, \mu_0) \rightarrow (\mathfrak{D}_1, \mu_1)$ is called the *Kantorovich-Wasserstein distance* between the measures μ_0 and μ_1 . The minimization of $M(u)$ constitutes the mathematical formulation of the Monge-Kantorovich optimal transport problem.

If the measures μ_i and Lebesgue measure $d\mathcal{L}$ are absolutely continuous with respect to each other, so that we can write $d\mu_i = m_i(\mathbf{x})d\mathcal{L}$ for certain strictly positive densities $m_i \in L^1(\mathfrak{D}_i, d\mathcal{L})$, and if the map u is a diffeomorphism from $\mathfrak{D}_0$ to $\mathfrak{D}_1$, then the measure preservation property (21) is equivalent with *mass preservation*:

$$m_0(\mathbf{x}) = \det(Du(\mathbf{x})) \cdot m_1(u(\mathbf{x})), \quad (\text{for almost all } \mathbf{x} \in \mathfrak{D}_0). \quad (23)$$

The Jacobian equation (23) implies that if a small region in $\mathfrak{D}_0$ is mapped to a larger region in $\mathfrak{D}_1$, there must be a corresponding decrease in density in order to comply with mass preservation. In other words, expanding an image darkens it.

The L^2 Monge-Kantorovich problem corresponds to the cost function $\Phi(\mathbf{x}, \mathbf{y}) = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2$. A fundamental theoretical result for the L^2 case [15, 43, 70] is that there is a unique optimal mass preserving u , and that this u is characterized as the gradient of a convex function p , i.e., $u = \nabla p$. General results about existence and uniqueness may be found in [4] and the references therein.

To reduce the Monge-Kantorovich cost $M(u)$ of a map $u^0 : \mathfrak{D}_0 \rightarrow \mathfrak{D}_1$, in [52] we consider a rearrangement of the points in the domain of the map in the following sense: the initial map u^0 is replaced by a family of maps u^t for which one has $u^t \circ s^t = u^0$ for some family of diffeomorphisms $s^t : \mathfrak{D}_0 \rightarrow \mathfrak{D}_0$. If the initial map u^0 sends the measure μ_0 to μ_1 , and if the diffeomorphisms s^t preserve the measure μ_0 , then the maps $u^t = u^0 \circ (s^t)^{-1}$ will also send μ_0 to μ_1 .

Any sufficiently smooth family of diffeomorphisms $s^t : \mathfrak{D}_0 \rightarrow \mathfrak{D}_0$ is determined by its velocity field, defined by $\partial_t s^t = v^t \circ s^t$.

If u^t is any family of maps, then, assuming $u^t_{\#}\mu_0 = \mu_1$ for all t , one has

$$\frac{d}{dt}M(u^t) = \int_{\mathfrak{D}_0} \langle \Phi_{\mathbf{x}}(\mathbf{x}, u^t(\mathbf{x})), v^t(\mathbf{x}) \rangle d\mu_0(\mathbf{x}). \quad (24)$$

The steepest descent is achieved by a family $s^t \in \text{Diff}_{\mu_0}^1(\mathfrak{D}_0)$, whose velocity is given by

$$v^t = -\frac{1}{m_0(\mathbf{x})} \mathcal{P}(\Phi_{\mathbf{x}}(\mathbf{x}, u^t(\mathbf{x}))). \quad (25)$$

Here $\mathcal{P}$ is the Helmholtz projection, which extracts the divergence-free part of vector fields on $\mathfrak{D}_0$, i.e., for any vector field w on $\mathfrak{D}$ one has $w = \mathcal{P}[w] + \nabla p$.

From $u^0 = u^t \circ s^t$ one gets the transport equation

$$\frac{\partial u^t}{\partial t} + v^t \cdot \nabla u^t = 0 \quad (26)$$

where the velocity field is given by (25). In [9], it is shown that the initial value problem (25), (26) has short time existence for $C^{1,\alpha}$ initial data u^0 , and a theory of global weak solutions in the style of Kantorovich is developed.

For image registration, it is natural to take $\Phi(\mathbf{x}, \mathbf{y}) = \frac{1}{2}\|\mathbf{x} - \mathbf{y}\|^2$ and $\mathfrak{D}_0 = \mathfrak{D}_1$ to be a rectangle [52]. Extensive numerical computations show that the solution to the unregularized flow converges to a limiting map for a large choice of measures and initial maps. Indeed, in this case, we can write the minimizing flow in the following “nonlocal” form:

$$\frac{\partial u^t}{\partial t} = -\frac{1}{m_0} \{u^t + \nabla(-\Delta)^{-1} \operatorname{div}(u^t)\} \cdot \nabla u^t. \quad (27)$$

The technique has been mathematically justified in [9] to which we refer the reader for all of the relevant details.

2.4.3 Other Applications of Optimal Transport

In our research, we focused on the uses of ideas from optimal transport for problems in controlled active vision and visual tracking. However, given the potential power of these ideas in systems and control, we would like to list some other applications of Monge-Kantorovich:

1. Lyapunov theory is essential in studying nonlinear system stability and controller synthesis. In some very interesting work, Rantzer [95] has formulated a dual to Lyapunov’s second theorem. The idea is that the Lyapunov function is regarded as the “cost to go” in an optimal transport problem and is dual (in the sense of linear programming) to the density function typically studied in Monge-Kantorovich theory. These ideas give a powerful new tool in studying nonlinear system analysis.
2. Shape optimization is another area of use for optimal transport [39]. For example, given two densities and an insulating medium into which we place a fixed amount of conducting material one can consider the problem of the optimal placement of the conducting material to minimize the heating induced by the flow. This can be put into the Monge-Kantorovich optimal transport framework. Similar remarks apply to problems in compression molding, where one considers an incompressible plastic material being pressed between two plates in which one wants to track the air-plastic interface.
3. One of the most beautiful applications of optimal transport is in meteorology, in particular, semigeostrophic models. These are concerned with large scale stratified flows and front formation [29]. The idea is that meteorologists want to model how fronts arise in large-scale weather patterns. Tracking such fronts is a key goal, and semigeostrophic equations seem to give a reasonable mathematical model for the creation of such fronts. This leads naturally to optimal mass transport equations.

3 Bibliography

References

- [1] L. Alvarez, P. L. Lions, and J. M. Morel, “Image selective smoothing and edge detection by nonlinear diffusion,” *SIAM J. Numer. Anal.* **29** (1992), pp. 845-866.
- [2] L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel, “Axioms and fundamental equations of image processing,” *Arch. Rational Mechanics* **123** (1993), pp. 200-257.
- [3] L. Ambrosio and M. Soner, “Level set approach to mean curvature flow in arbitrary codimension,” *Journal of Differential Geometry* **43** (1996), pp. 693-737.
- [4] L. Ambrosio, “Lecture notes on optimal transport theory,” *Euro Summer School, Mathematical Aspects of Evolving Interfaces*, CIME Series of Springer Lecture Notes, Springer, 2000.

- [5] S. Angenent, "Parabolic equations for curves on surfaces, Part I. Curves with p -integrable curvature," *Annals of Mathematics* **132** (1990), pp. 451-483.
- [6] S. Angenent, "Parabolic equations for curves on surfaces, Part II. Intersections, blow-up, and generalized solutions," *Annals of Mathematics* **133** (1991), pp. 171-215.
- [7] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis, "On area preserving maps of minimal distortion," in *System Theory: Modeling, Analysis, and Control*, edited by T. Djaferis and I. Schick, Kluwer, Holland, 1999, pages 275-287.
- [8] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis, "Laplace-Beltrami operator and brain surface flattening," *IEEE Trans. on Medical Imaging* **18** (1999), pp. 700-711.
- [9] S. Angenent, S. Haker, and A. Tannenbaum, "Minimizing flows for the Monge-Kantorovich problem," *SIAM J. Math. Analysis* **35** (2003) pp. 61-97.
- [10] S. Angenent, G. Sapiro, and A. Tannenbaum, "On the affine heat equation for non-convex curves," *Journal of the American Mathematical Society* **11** (1998), pp. 601-634.
- [11] R. Bajcsy and S. Kovacic, "Multiresolution elastic matching," *Computer Vision, Graphics and Image Processing* **46** (1989), pp. 1-21.
- [12] J. Barron, D. Fleet, and S. Beauchemin, "Performance of optical flow techniques," *Int. Journal of Computer Vision* **12** (1994), pp. 43-77.
- [13] J.-D. Benamou and Y. Brenier, "A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem," *Numerische Mathematik* **84** (2000), pp. 375-393.
- [14] A. Blake and M. Isard, *Active Contours*, Springer-Verlag, New York, 1998.
- [15] Y. Brenier, "Polar factorization and monotone rearrangement of vector-valued functions," *Com. Pure Appl. Math.* **64** (1991), pp. 375-417.
- [16] H. Blum, "Biological shape and visual science," *J. Theor. Biology* **38** (1973), pp. 205-287.
- [17] R. Brockett and A. Blake, "Estimating the shape of a moving contour," *Proceedings of IEEE CDC*, 1994.
- [18] V. Caselles, F. Catte, T. Coll, and F. Dibos, "A geometric model for active contours in image processing," *Numerische Mathematik* **66** (1993), pp. 1-31.
- [19] V. Caselles, R. Kimmel, and G. Sapiro, "Geodesic snakes," *Int. J. Computer Vision* **22** (1997), pp. 61-79.
- [20] V. Caselles, J. Morel, G. Sapiro, and A. Tannenbaum, "Introduction To The Special Issue On Partial Differential Equations And Geometry-driven Diffusion In Image Processing And Analysis," *IEEE Trans. on Image Processing* **7** (1998), pp. 269-273.
- [21] T. Chan and L. Vese, "Active Contours without edges," *IEEE Trans. on Image Processing* **10** (2001), pp. 266-277.
- [22] G. Charpiat, O. Faugeras, and R. Keriven, "Approximations of shape metrics and application to shape warping and empirical shape statistics," *Foundations of Computational Mathematics* (2004), pp. OF1-OF58.
- [23] Y.-G. Chen, Y. Giga and S. Goto, "Uniqueness and existence of viscosity solutions of generalized mean curvature flow equations," *J. Diff. Geo.* **33**, pp. 749-786, 1991.
- [24] G.E. Christensen, R.D. Babbitt, M.I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. on Image Processing* **5** (1996), pp. 1435-1447.
- [25] I. Cohen and I. Herlin, "Curves matching using geodesic paths," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 741-746, 1998.
- [26] L. D. Cohen, "On active contour models and balloons," *CVGIP: Image Understanding* **53** (1991), pp. 211-218.

- [27] M. G. Crandall, H. Ishii, and P. L. Lions, "User's guide to viscosity solutions of second order partial differential equations," *Bulletin of the American Mathematical Society* **27** (1992), pp. 1-67.
- [28] D. Cremers and S. Soatto, "A pseudo-distance for shape priors in level set segmentation," *Proceedings of the IEEE International Workshop on Variational, Geometric and Level Set Methods in Computer Vision*, pp. 169-176, 2003.
- [29] M. Cullen and R. Purser, "An extended Lagrangian theory of semigeostrophic frontogenesis," *J. Atmos. Sci.* **41** (1984), pp. 1477-1497.
- [30] B. Dacorogna and J. Moser, "On a partial differential equation involving the Jacobian determinant," *Ann. Inst. H. Poincaré Anal. Non Linéaire*, **7** (1990), pp. 1-26.
- [31] C.A. Davatzikos, J.L. Prince, and R.N. Bryan, "Image registration based on boundary mapping," *IEEE Trans. Med. Imaging* **15** (1996), pp. 112-115.
- [32] P. DelMoral, "A uniform convergence theorem for the numerical solving of the nonlinear filtering problem," *Journal of Applied Probability* **35** (1998), pp. 873-884.
- [33] P. Del Moral and L. Miclo, "Branching and interacting particle systems approximations of Feynman-Kac formulae with applications to non-linear filtering," *Séminaire de Probabilités XXXIV, Lecture Notes in Math.* **1729**, Springer (2000), pp. 1-145.
- [34] M. P. Do Carmo, *Differential Geometry of Curves and Surfaces*, Prentice-Hall, Inc., New Jersey, 1976.
- [35] A. Doucet, "On sequential Monte Carlo sampling methods for Bayesian filtering," Technical Report CUED/F-INFENG/TR. 310, Cambridge University Department of Engineering, 1998; url = "citeseer.ist.psu.edu/article/doucet00sequential.html".
- [36] A. Doucet, N. deFreitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*, Springer-Verlag, New York, 2001.
- [37] J. C. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory*, McMillan, New York, 1991.
- [38] A. Duci, A. J. Yezzi, S. K. Mitter, S. Soatto, "Region matching with missing parts," *Proceedings of the European Conference on Computer Vision*, Lecture Notes in Computer Science **2352** (2002), pp. 48-64.
- [39] L. C. Evans, "Partial differential equations and Monge-Kantorovich," in *Current Developments in Mathematics*, edited by S. T. Yau, 1997.
- [40] L. C. Evans and J. Spruck, "Motion of level sets by mean curvature, I," *J. Diff. Geo.* **33**, pp. 635-681, 1991.
- [41] D. Fry, *Shape Recognition Using Metrics on the Space of Shapes*, Ph.D. Thesis, Harvard University, 1993.
- [42] W. Gangbo, "An elementary proof of the polar factorization of vector-valued functions," *Arch. Rational Mechanics Anal.* **128** (1994), pp. 381-399.
- [43] W. Gangbo and R. McCann, "The geometry of optimal transportation," *Acta Math.* **177** (1996), pp. 113-161.
- [44] W. Gangbo and R. McCann, "Shape recognition via Wasserstein distance," Technical Report, School of Mathematics, Georgia Institute of Technology, 1999.
- [45] M. Gage and R. S. Hamilton, "The heat equation shrinking convex plane curves," *J. Differential Geometry* **23** (1986), pp. 69-96.
- [46] J.C. Gee, C. Barillot, L. Le Briquer, D.R. Haynor, R. Bajcsy, "Matching structural images of the human brain using statistical and geometrical features," *Proc. 3rd International Conference Visualization in Biomedical Computing (VBC'94)* (1994), pp. 191-204.

- [47] N. J. Gordon, D. J. Salmond, and A. F. Smith, "Novel Approach to Nonlinear/NonGaussian Bayesian State Estimation," *IEE Proceedings-F (Radar and Signal Processing)* **140**(2) (1993), pp. 107-113.
- [48] M. Grayson, "The heat equation shrinks embedded plane curves to round points," *J. Differential Geometry* **26** (1987), pp. 285-314.
- [49] M. Grayson, "Shortening embedded curves," *Annals of Mathematics* **129** (1989), pp. 71-111.
- [50] U. Grenander, *General Pattern Theory*, Oxford University Press, 1993.
- [51] H. W. Guggenheimer, *Differential Geometry*, McGraw-Hill Book Company, New York, 1963.
- [52] S. Haker, L. Zhu, A. Tannenbaum, and S. Angenent, "Optimal mass transport for registration and warping" **60** (2004), pp. 225-240.
- [53] R. Haralick and L. Shapiro, *Computer and Robot Vision*, Addison-Wesley, New York, 1992.
- [54] A. Hassanien and M. Nakajima, "Image morphing of facial images transformation based on Navier elastic body splines," *Computer Animation Proceedings* **8** (1988), pp. 119-125.
- [55] G. Hermosillo, C. Chef d'Hotel, and O. Faugeras, "A variational approach to multi-modal image matching," INRIA Report Number 4117, 2001.
- [56] E. C. Hildreth, "Computations underlying the measurement of visual motion," *Artificial Intelligence* **23** (1984), pp. 309-354.
- [57] B. K. P. Horn, *Robot Vision*, MIT Press, Cambridge, Mass., 1986.
- [58] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, 23:185-203, 1981.
- [59] M. Isard and A. Blake, "CONDENSATION – conditional density propagation for visual tracking," *International Journal of Computer Vision* **29** (1998), pp. 5-28.
- [60] J. Jackson, A. Yezzi, and S. Soatto, "Tracking deformable moving objects under severe occlusions," *IEEE Conference on Decision and Control*, 2004.
- [61] O. Juan, R. Keriven, and G. Postelnicu, "Stochastic mean curvature motion in computer vision: stochastic active contours," INRIA Report, France, 2004.
- [62] T. Kaijser, "Computing the Kantorovich distance for images," *Journal of Mathematical Imaging and Vision* **9** (1998), pp. 173-191.
- [63] K. Kanatani, "Analysis of 3-D rotation fitting," *IEEE Trans. Pattern Analysis and Machine Intelligence* **16** (1994), pp. 543-549.
- [64] L. V. Kantorovich, "On a problem of Monge," *Uspekhi Mat. Nauk.* **3** (1948), pp. 225-226.
- [65] M. Kass, A. Witkin, and D. Terzopoulos, "Snakes: active contour models," *International Journal of Computer Vision* (1988), pp. 321-331.
- [66] S. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum, and A. Yezzi, "Conformal curvature flows: from phase transitions to active vision," *Archive for Rational Mechanics and Analysis* **134** (1996), pp. 275-301.
- [67] 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, 1991.
- [68] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, "Shapes, shocks, and deformations, I," *Int. J. Computer Vision* **15** (1995), pp. 189-224.
- [69] B. B. Kimia, A. Tannenbaum, and S. W. Zucker, "Optimal control methods in computer vision and image processing," in *Geometry Driven Diffusion in Computer Vision*, edited by Bart ter Haar Romeny, Kluwer, 1994.

- [70] M. Knott and C. Smith, "On the optimal mapping of distributions," *J. Optim. Theory* **43** (1984), pp. 39-49.
- [71] J. J. Koenderink, "The structure of images," *Biological Cybernetics* **50** (1984), pp. 363-370.
- [72] Y. Lauzier, K. Siddiqi, A. Tannenbaum, and S. Zucker, "Area and length minimizing flows for segmentation," *IEEE Trans. Image Processing* **7** (1998), pp. 433-444.
- [73] S.-Y. Lee, K.-Y. Chwa, J. Hahn, and S.Y. Shin, "Image morphing using deformation techniques," *J. Visualization and Computer Animation* **7** (1996), pp. 3-23.
- [74] A. Leow, M.-C. Chiang, H. Protas, P. Thompson, L. Vese, and H. Huang, "Linear and non-linear geometric object matching with implicit representation," *Proceedings of the IEEE International Conference on Pattern Recognition*, pages 710-713, 2004.
- [75] M. Leventon, W. E. L. Grimson, and O. Faugeras, "Statistical shape influence in geodesic active contours," *Proceedings of IEEE CVPR* (2000), pp. 1316-1324.
- [76] R. McCann, "Polar factorization of maps on Riemannian manifolds," *Geom. Funct. Anal.*, **11** (2001) pp. 589-608.
- [77] R. McCann, "A convexity principle for interacting gases," *Adv. Math.*, **128** (1997), pp. 153-179.
- [78] T. McInerney and D. Terzopoulos, "Deformable models in medical image analysis: a survey," *Medical Image Analysis* **1** (1996), pp. 91-108.,
- [79] T. McInerney and D. Terzopoulos, "Topologically adaptable snakes," *Int. Conf. on Computer Vision*, pp. 840-845, 1995.
- [80] M. Miller, G. Christensen, Y. Amit, U. Grenander, "Mathematical textbook of deformable neuroanatomies," *Proc. National Academy of Science* **90**, pp. 11944-11948.
- [81] J. Milnor, *Morse theory*, Princeton University Press, 1963.
- [82] D. Mumford and J. Shah, "Optimal Approximation by piecewise smooth functions and associated variational problems," *Commun. Pure Applied Mathematics* **42** (1989), pp. 577-685.
- [83] M. Niethammer and A. Tannenbaum, "Dynamic geodesic snakes for visual tracking," *Proceedings CVPR* **1** (2004), pp. 660-667.
- [84] M. Niethammer, A. Tannenbaum, and S. Angenent, "Dynamic active contours in controlled active vision," submitted to the *IEEE Transactions on Automatic Control*, 2005.
- [85] M. Niethammer, A. Tannenbaum, and P. Vela, "Geometric observers for dynamically evolving curves," submitted to *IEEE CDC*, 2005.
- [86] P. Olver, G. Sapiro, and A. Tannenbaum, "Invariant geometric evolutions of surfaces and volumetric smoothing," *SIAM J. Applied Math.* **57** (1997), pp. 176-194.
- [87] S. J. Osher and J. A. Sethian, "Fronts propagation with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics* **79** (1988), pp. 12-49.
- [88] S. Osher and R. Fedkiw, *Level Set Methods and Dynamic Implicit Surfaces*, Springer-Verlag, 2003.
- [89] J. Palis and W. de Melo, Welington, *Geometric Theory of Dynamical Systems*, Springer-Verlag, New York, 1982.
- [90] N. Paragios and R. Deriche, "Geodesic active regions: a new framework to deal with frame partition problems in computer vision," *Journal of Visual Communication and Image Representation* **13** (2002), pp. 249-268.
- [91] P. Perona and J. Malik, "Scale-space and edge detection using anisotropic diffusion," *IEEE Trans. Pattern Anal. Machine Intell.* **12** (1990), pp. 629-639.

- [92] E. Pichon, A. Tannenbaum, and R. Kikinis, "Statistically based flow for image segmentation," *Medical Imaging Analysis* **8** (2004), pp. 267-274.
- [93] J.-P. Pons, G. Hermosillo, R. Keriven, O. Faugeras, "How To deal with point correspondences and tangential velocities in the level set framework," *Proceedings of the IEEE International Conference on Computer Vision*, pp. 894-899, 2003.
- [94] S. Rachev and L. Rüschendorf, *Mass Transportation Problems*, Volumes I and II, Probability and Its Applications, Springer, New York, 1998.
- [95] A. Rantzer, "A dual to Lyapunov's stability theorem," *Systems and Control Letters*, 2001.
- [96] Y. Rathi, N. Vaswani, A. Tannenbaum, and A. Yezzi, "Tracking moving and deforming shapes using a particle filter," *Proceedings of CVPR*, 2005.
- [97] B. ter Haar Romeny (editor), *Geometry-Driven Diffusion in Computer Vision*, Kluwer, Holland, 1994.
- [98] G. Sapiro, *Geometric Partial Differential Equations and Image Analysis* Cambridge University Press, 2001.
- [99] G. Sapiro and A. Tannenbaum, "On affine plane curve evolution," *Journal of Functional Analysis* **119** (1994), pp. 79-120.
- [100] G. Sapiro and A. Tannenbaum, "Area and length preserving geometric invariant scale-spaces," *IEEE Pattern Analysis and Machine Intelligence* **17** (1995), pp. 67-72.
- [101] G. Sapiro and A. Tannenbaum, "Invariant curve evolution and image analysis," *Indiana University J. of Mathematics* **42** (1993), pp. 985-1009.
- [102] G. Sapiro and A. Tannenbaum, "Invariant curve evolution and image analysis," *Indiana University J. of Mathematics* **42** (1993), pp. 985-1009.
- [103] J. A. Sethian, *Level Set Methods and Fast Marching Methods*, Cambridge University Press, 1999.
- [104] J. A. Sethian, "Curvature and the evolution of fronts," *Commun. Math. Phys.* **101** (1985), pp. 487-499.
- [105] J. Shah, "A common framework for curve evolution, segmentation, and anisotropic diffusion," *Proceedings of IEEE CVPR'96*, 1996, pp. 136-142.
- [106] K. Siddiqi, Y. Lauziere, A. Tannenbaum, and S. Zucker, "Area and length minimizing flows for shape segmentation," *IEEE TMI* **7** (1998), pp. 433-443.
- [107] G. Strang, *Introduction to Applied Mathematics*, Wellesley-Cambridge Press, Wellesley, Mass., 1986.
- [108] D. Terzopoulos and R. Szeliski, "Tracking with Kalman snakes," in *Active Vision* edited by A. Blake and A. Yuille, MIT Press, pages 3-20, 1992.
- [109] A. Tannenbaum, "Three snippets of curve evolution theory in computer vision," *Mathematical and Computer Modelling Journal* **24** (1996), pp. 103-119.
- [110] A. Tannenbaum and A. Yezzi, "Visual tracking, active vision, and gradient flows," in *The Confluence of Vision and Control*, edited by G. Hager and D. Kriegman, *Lecture Notes in Control and Information Sciences* **237**, Springer-Verlag, New York, 1998.
- [111] A. Toga, *Brain Warping*, Academic Press, New York, 1999.
- [112] N. Vaswani, *Change Detection in Stochastic Shape Dynamical Models with Applications in Activity Modeling and Abnormality Detection*, Ph.D. Thesis, ECE Dept, University of Maryland at College Park, August 2004.
- [113] N. Vaswani, Y. Rathi, A. Yezzi, and A. Tannenbaum, "Particle filtering for continuous closed curves," submitted to *IEEE CDC*, 2005.