SIGNAL MATCHING
THROUGH SCALE SPACE

Andrew Witkin
Demetri Terzopoulos
Michael Kass

Schlumberger Palo Alto Research
3340 Hillview Ave.
Palo Alto, CA 94304

ABSTRACT

Given a collection of similar signals that have been de-
formed with respect to each other, the general signal
matching problem is to recover the deformation. We for-
mulate the problem as the minimization of an energy mea-
sure that combines a smoothness term and a similarity
term. The minimization reduces to a dynamic system
governed by a set of coupled, first-order differential equa-
tions. The dynamic system finds an optimal solution at
a coarse scale and then tracks it continuously to a fine
scale. Among the major themes in recent work on vi-
sual signal matching have been the notions of matching
as constrained optimization, of variational surface recon-
struction, and of coarse-to-fine matching. Our solution
captures these in a precise, succinct, and unified form.
Results are presented for one-dimensional signals, a mo-
tion sequence, and a stereo pair.

I Introduction

Given a collection of similar signals that have been deformed
with respect to each other, the general signal matching problem
is to recover the deformation. Important matching problems in-
clude stereo vision, motion analysis, and a variety of registration
problems such as template matching for speech and vision.

We cast the problem as the minimization of an energy func-
tional $E(\mathbf{V})$ where $\mathbf{V}$ is the deformation. The energy functional
is the sum of two terms, one based on the correlation of the
deformed signals, and the other based on the smoothness of the
deformation.

In general, the energy functional $E(\mathbf{V})$ can be highly non-
convex, so that ordinary optimization methods become trapped
in local minima. Optimization by simulated annealing can be
attempted, but at severe computational expense. Instead, we
rely on *continuation methods* to solve the problem. By introduc-
ing a scale parameter σ, the minimization problem is embedded
within a larger space. A suitable minimum can be achieved
relatively easily for large σ because the signals and hence the
energy landscape are very smooth. The solution of the original
minimization problem is then obtained by continuously tracking
the minimum as σ tends to zero. This is analogous to a coarse-
to-fine tracking of extrema through scale-space in the sense of
[Witkin, 1983].

The entire procedure consists of solving the first-order dy-
namic system

$$\dot{\sigma} = -k_1 \exp(-k_2 |\nabla E(\mathbf{V}, \sigma)|), \quad \dot{\mathbf{V}} = -\nabla E(\mathbf{V}, \sigma),$$

where the dot denotes a time derivative, σ is the scale parameter
and k_1 and k_2 are constants. Given an initial crude estimate for
$\mathbf{V}$ at a coarse scale σ_0, the system minimizes E at σ_0 and follows
a trajectory of minima through finer scales, thereby increasing
the resolution of V. Any of a number of well known numerical

techniques can be used to solve for the trajectory. Through
a series of incremental deformations, correlations of deformed
signals are optimized and balanced against the smoothness of
the deformations while moving from coarse to fine scale. Thus,
the first-order system compactly unifies a number of important
yet seemingly disparate signal-matching notions.

In the remainder of this section, the relation of our technique
to previous work on matching is discussed. Then in section 2,
a framework for the minimization problem is introduced. In
section 3, the solution of the problem by continuation methods
and the resulting single differential equation are also developed.
Section 4 describes the specific similarity term employed and,
in section 5, the details of the smoothness term are discussed.
Finally, section 6 presents several examples of matching results
for one- and two-dimensional signals.

A. Background

An enormous amount of work has been done on signal match-
ing, giving precedent for several components of our approach.

Optimization of constrained deformations guided by correla-
tion or L_2 metrics can be found in prior work. In speech recog-
nition, the problem of time warping speech segments to match
input utterances with stored prototypes has been addressed in
this context. Dynamic programming has been used to com-
pute constrained warping functions (see [Rabiner and Schafer,
1978] and [Sankoff and Kruskal, 1983] Part II). This particular
optimization technique is readily applicable in matching situa-
tions involving sequentially ordered signals, such as speech, and
unilateral continuity constraints. However, its stringent require-
ments on the energy functional appear incompatible with the
unordered multi-dimensional signals and isotropic smoothness
constraints which are of primary concern to us.

Smoothness constraints have been popular in computational
vision. Consider the important problem of stereo matching. In
the past, dense disparity maps have been computed through a
two step process of local matching followed by smooth [Grimson,
1983], multiresolution [Terzopoulos, 1983], or piecewise contin-
uous [Terzopoulos, 1986] surface reconstruction from the sparse
disparities. The approach in the present paper unifies match-
ing and piecewise smooth reconstruction into a single iterative
optimization process.

Broit's [1981] work in registering a deformed image to a model
image resembles ours, in that matching is explicitly formulated
as a minimization problem involving a cost functional that com-
bines both a deformation constraint and a similarity measure.
His deformation model, which involves the strain energy of a
globally smooth elastic body, is more elaborate than the defor-
mation constraints inherent in the spring loaded subtemplate
matching technique of Fischler and Elschlager [1973] or the iter-
ative Gaussian smoothed deformation models proposed by Burr
[1981]. Our controlled-continuity deformation model provides us

with the additional capability to regulate the order of smoothness and to preserve discontinuities in the deformation.

Coarse-to-fine matching schemes have previously been treated as a multistage process in which a matching operation is performed at each successive level [Mori, et al., 1973; Hannah, 1974; Moravec, 1977; Marr and Poggio, 1979; Gennery, 1980]. We have extended this idea into a matching process which evolves continuously towards finer spatial scale. The idea of progressing continuously through scale space derives from Witkin [1983].

As our matching process computes the deformation iteratively, it is best to perform the similarity measurements by deforming the signals according to the current approximation of the deformation. This concern has also been addressed by the matching algorithms described in [Mori, et al., 1973; Burr, 1981; Broit, 1981; Quam, 1984].

II Framework

Consider a vector of n similar signals $\mathbf{f}^*(\mathbf{x}) = [f_1^*(\mathbf{x}), \ldots, f_n^*(\mathbf{x})]$ defined in d dimensional space $\mathbf{x} = [x_1, \ldots, x_d] \in \Re^d$, and a deformation mapping $\mathbf{V} : \Re^d \mapsto \Re^{nd}$, such that $\mathbf{V}(\mathbf{x}) = [\mathbf{v}_1(\mathbf{x}), \ldots, \mathbf{v}_n(\mathbf{x})]$, where each of the n disparity functions $\mathbf{v}_i :$ $\Re^d \mapsto \Re^d$ is a vector valued function $\mathbf{v}_i(\mathbf{x}) = [v_1(\mathbf{x}), \ldots, v_d(\mathbf{x})]^\top$. Given a set of deformed signals $\mathbf{f}$ such that $\mathbf{f}^*(\mathbf{x}) = \mathbf{f}(\mathbf{V}^*(\mathbf{x}))$, the matching problem is to recover the deformation $\mathbf{V}^*(\mathbf{x})$.

Suppose that the similarity between the signals $\mathbf{f}$ for a given deformation $\mathbf{V}$ is measured by a functional $Q(\mathbf{V}) : \Re^{nd} \mapsto \Re$ bounded from above by a value achieved by the best possible match. A reasonable objective is to find the deformation $\mathbf{U}$ which maximizes the quality of the match; i.e., to minimize $-Q(\mathbf{V})$ over possible deformations $\mathbf{V}$. Thus, $\mathbf{U}$ represents an optimal approximation to $\mathbf{V}^*$.

This minimization problem is clearly ill-posed in the absense of constraints on admissible deformations, since, e.g., degenerate or chaotic deformations can always be contrived that achieve the minimum value. Such constraints may be encoded by a second functional $S(\mathbf{V}) : \mathcal{H}^{nd} \mapsto \Re$, where $\mathcal{H}^{nd} \subset \Re^{nd}$ is the subset of admissible deformations.

Useful instances of similarity and constraint functionals will be formulated shortly. Their combination, however, leads to the following minimization problem: Find the deformation $\mathbf{U} \in \mathcal{H}^{nd}$ such that $E(\mathbf{U}) = \inf_{\mathbf{V} \in \mathcal{H}^{nd}} E(\mathbf{V})$, where the energy functional is given by

$$E = -(1 - \lambda)Q - \lambda S \qquad (1)$$

and where $\lambda \in (0, 1)$ is a weighting parameter.

Stabilization offers a general approach to a numerical solution through the construction of a discrete dynamic system whose fixed points include a discrete solution of the above optimization problem [Bakhvalov, 1977]. A simple dynamic system with this property is characterized by the differential equations

$$\dot{\mathbf{V}} + \nabla E = 0, \qquad (2)$$

where the dot denotes differentiation with respect to time t and ∇E denotes the gradient of E with respect to the free variables of the discrete deformation. Optimization occurs by dissipation of energy; energy cannot increase along the system's trajectory $\mathbf{V}(\mathbf{x}, t)$ in $\mathcal{H}^{nd}$, which follows the direction of the gradient of E. Although the trajectory terminates at a local minimum of E, there is no guarantee that the global minimum $\mathbf{U}$ will be attained by solving this initial value problem starting from an arbitrary initial condition $\mathbf{V}(\mathbf{x}, 0)$.

III Continuation over Scale

The key remaining difficulty is that for obvious choices of Q, such as linear correlation, E is likely to have many local minima, making the minimization problem highly non-convex and therefore extremely difficult to solve. There are two options: solving this hard problem directly (for example by simulated annealing) or simplifying the problem by choosing Q to be convex or nearly so. We pursue the second option because annealing is expensive.

A. Continuation Methods

Q may be smoothed by subjecting $\mathbf{f}$ to a smoothing filter of characteristic width σ. We observe empirically that the best solution for $\mathbf{v}$ as σ increases tends to be an increasingly smoothed version of the correct solution. These means that slightly deblurring $\mathbf{v}$ by reducing σ produces a slightly better solution close to the one just obtained. To the extent this is so, we can solve the problem using equation 2 by means of *continuation methods* [Dahlquist and Björck, 1974].

Continuation methods embed the problem to be solved,

$$\mathbf{g}(\mathbf{v}) = 0,$$

in a family of problems

$$\mathbf{g}(\mathbf{v}, s) = 0,$$

parameterized by s. Let $s_{i+1} = s_i + \Delta s$, $\mathbf{g}(\mathbf{v}, s_n)$ be the problem we wish to solve, presumably difficult, and $\mathbf{g}(\mathbf{v}, s_1)$ a readily solvable member of the family, and let

$$\mathbf{u}(s) = \mathbf{H}(\mathbf{g}, s, \mathbf{v}_0)$$

be the solution for $\mathbf{g}$ at s given $\mathbf{v}_0$ as an initial condition. Then $\mathbf{u}(s_n)$ is obtained from $\mathbf{u}(s_1)$ by the iteration

$$\mathbf{u}(s_{i+1}) = \mathbf{H}(\mathbf{g}, s_{i+1}, \mathbf{u}_i); \qquad i = 1, \ldots, n - 1$$

that is, each solution is used as an initial condition to obtain the next one.

For the current problem, the continuation parameter is σ, with $\Delta\sigma < 0$. We continue from an initial coarse scale σ_1 and an initial guess $\mathbf{V}_1$ by

$$\mathbf{V}_{i+1} = \mathbf{H}(\dot{\mathbf{V}} + \nabla E, \sigma_{i+1}, \mathbf{V}_i),$$

to a fine scale σ_n and a final answer $\mathbf{V}_n$. To visualize this method, imagine the energy landscape at each value of σ as a contoured surface in 3-space. The surfaces are stacked one above the other, so that the topmost surface is very smooth, while the lower ones become increasingly bumpy. Imaging a hole drilled at each local minimum on each surface. A ball bearing dropped onto the topmost surface will roll down to the bottom of the hill. At this point, it falls through to the next level, rolls down again, falls through again, and so on to the bottom.

B. A Scale Space Equation

This iteration solves a separate initial value problem at each step. A more attractive alternative is to collapse the continuation over σ into a single differential equation. Ideally, the solution should follow a curve $\mathbf{V}(\sigma)$ satisfying

$$|\nabla E(\mathbf{V}(\sigma))| \equiv 0;$$

i.e., a continuous curve of solutions over scale. A differential equation for this curve is

$$\mathbf{V}_\sigma = -(\nabla E)_\sigma (\nabla\nabla E)^{-1}. \qquad (3)$$

The solution to this equation tracks a given coarse-scale solution continuously to fine scale, in precise analogy with the coarse-to-fine tracking through scale space of [Witkin, 1983]. Unfortunately, it is impractical to solve this equation for arbitrary S and Q, since $\nabla\nabla E$ is high dimensional.

To construct an approximate equation, we introduce the quantity

$$N = -k_1 e^{-k_2|\nabla E|}, \tag{4}$$

so that $N = -k_1$ at a solution to equation 2, diminishing with distance from the solution at a rate determined by the space constant k_2. The equation

$$\dot{\sigma} = N, \quad \dot{\mathbf{V}} = -\nabla E \tag{5}$$

approximates the desired behavior. Far from a solution, where N is small, equation 5 approaches equation 2, changing $\mathbf{V}$ but not σ. Approaching a solution, σ begins to decrease. At a solution, $\dot{\mathbf{V}} = 0$ and $\dot{\sigma} = -k_1$. From an initial $\mathbf{V}(t_0), \sigma(t_0)$, the solution $\mathbf{V}(t), \sigma(t)$ moves through $\mathbf{V}$ at nearly constant scale until a minimum in E is approached, then it begins descending in scale staying close to a solution. [1]

Equation 5 finds a solution at the initial scale, then tracks it continuously to finer scales. To use equation 5, we choose a coarse scale $\sigma(t_0)$, a crude initial guess $\mathbf{V}(t_0)$, and a terminal fine scale σ_T. We then run the equation until $\sigma(t) = \sigma_T$, taking $\mathbf{V}(t)$ as the solution.

C. Ambiguous Solutions

From time to time, we expect to encounter instabilities in the solutions of equation 5, in the sense that a small perturbation of the data induces a large change in the solution curve's trajectory through scale space. These instabilities correspond to bifurcations of the solution curve, analogous to bifurcations that can be observed in Gaussian scale space. We have considered two approaches to dealing with them. First, by adding a suitable noise term to E, equation 5 becomes a hybrid of scale space continuation and simulated annealing. We believe that local ambiguities can be favorably resolved using low-amplitude noise, hence with little additional computational cost. A second approach is to regard these instabilities as genuine ambiguities whose resolution falls outside the scope of the method. In that case, a set of alternative solutions can be explored by the addition of externally controlled bias terms to E. These terms can reflect outside constraints of any kind, for example, those imposed by the operation of attentional proceseses.

In the following sections, we turn to specific choices for S and Q.

IV   Similarity Functional

In general the similarity measure Q should capture what is known about the specific matching problem. In many cases, the undeformed signals are sufficiently similar that a simple correlation measure suffices. In this section we formulate a generic choice for this class of problems. Note that, by assumption, it is the undeformed signals $\mathbf{f}(\mathbf{x})$ which are similar, so the quality of a potential solution $\mathbf{V}$ should be measured by the similarity of the signals $\mathbf{f}(\mathbf{V}(\mathbf{x}))$.

Consider the case of two signals f_i and f_j. A general family of similarity measures is obtained by integrating a local

[1] The solution to equation 5 oscillates around the exact solution (equation 3,) with frequency and amplitude controlled by k_1 and k_2. This oscillation can be damped by the addition of second order terms in t, but we have not found it necessary to do so in practice.

measure of similarity over position. If $K_{i,j}(\mathbf{x})$ is a local measure of the similarity of $f_i(\mathbf{v}_i(\mathbf{x}))$ and $f_j(\mathbf{v}_j(\mathbf{x}))$ around $\mathbf{x}$, then $Q_{i,j} = \int K_{i,j}(\mathbf{x})d\mathbf{x}$ is a global measure of similarity for f_i and f_j. By simply adding up pairwise similarities, a global measure of similarity can be constructed for n signals: $Q = \sum_{i \neq j} Q_{i,j}$.

A number of possibilities exist for the local similarity measure $K_{i,j}(\mathbf{x})$. Normalized cross-correlation produces good results for several matching problems that we have examined. If $W_\gamma(\mathbf{x})$ is a window function where γ denotes the width parameter, $\mu_i(\mathbf{x}) = \int f_i(\mathbf{v}_i(\mathbf{y} - \mathbf{x}))W_\gamma(\mathbf{y})d\mathbf{y}$, and $\nu_i(\mathbf{x}) = \int [f_i(\mathbf{v}_i(\mathbf{x} - \mathbf{y})) - \mu_i(\mathbf{x} - \mathbf{y})]^2 W(\mathbf{y}) d\mathbf{y}$, then the normalized cross-correlation can be written

$$\begin{aligned} K_{i,j}(\mathbf{x}) &= [\nu_i(\mathbf{x})\nu_j(\mathbf{x})]^{-1/2} \int \Big\{ [f_i(\mathbf{v}_i(\mathbf{x} - \mathbf{y})) - \mu_i(\mathbf{x} - \mathbf{y})] \\ &\quad \times [f_j(\mathbf{v}_j(\mathbf{x} - \mathbf{y})) - \mu_j(\mathbf{x} - \mathbf{y})]W(\mathbf{y}) \Big\} d\mathbf{y}. \end{aligned}$$

The resulting functional $Q(\mathbf{V})$ generally has many local minima. In order to apply the continuation method, we compute Q for signals $\mathbf{f}$ which have been smoothed by Gaussians of standard deviation σ. The resulting functional $Q_\sigma(\mathbf{V})$ can then be made as smooth as is desired.

The correlation window size W_γ should be large enough to provide an accurate local estimate of the mean and variance of the signals, but small enough that non-stationarities in the signals do not become a problem. A convenient way to set γ to a reasonable value is to make it a fixed multiple of the average autocorrelation widths of the smoothed signals. Then γ can be regarded as a function of σ.

Note that $Q_\sigma(\mathbf{V})$ must be recomputed at each iteration, with the signals resampled to reflect the current choice of $\mathbf{V}$. If the deformation is very small, the distortion induced by failing to resample can be ignored, but the value of such resampling in stereo matching, for example, is well established [Mori, et al., 1973; Quam, 1984]. The Gaussian signal smoothing should also take place on the resampled functions $f(\mathbf{v}(\mathbf{x}))$.

V   Smoothness Functionals

The functional $S(\mathbf{V})$ places certain restrictions on admissible deformations in order to render the minimization problem better behaved. Perhaps the simplest possible restriction, and one that has been used often in the past, is to limit possible disparities between signals to prespecified ranges. A deformation can then be assigned within disparity bounds on a point by point basis according to maximal similarity criteria. Although simple, such limited searches are unfortunately error prone, since they are based on purely local information.

This problem can be resolved by imposing global constraints on the deformation that are more restrictive yet remain generic. Such constraints may be based on a priori expectations about deformations; for example, that they are coherent in some sense. In particular, admissible deformations may be characterized according to the controlled-continuity constraints defined in [Terzopoulos, 1986]. These constraints, which are based on generalized splines, restrict the admissible space of deformations to a class of piecewise continuous functions. Not only is the deformation's order of continuity controllable, but discontinuities of various orders (e.g., jump, slope, and curvature discontinuities) are permitted to form, subject to an energy penalty.

A general controlled-continuity constraint is imposed on the deformation by the functional

$$S(\mathbf{v}) = -\sum_{m=0}^{p} \int_{\Re^d} w_m(\mathbf{x}) \sum_{j_1 + \cdots + j_d = m} \frac{m!}{j_1! \ldots j_d!} \left| \frac{\partial^m \mathbf{v}(\mathbf{x})}{\partial x_1^{j_1} \ldots \partial x_d^{j_d}} \right|^2 d\mathbf{x}.$$

Figure 2 shows a more challenging example in which 4 signals are matched simultaneously. The signals are intensity profiles from a complex natural image. On the left the four signals are shown superimposed at several points in the matching process. As before, the original signals appear at the top and the final result is at the bottom. Note that coarse-scale features are aligned first in the matching process while fine scale features are matched later. The four corresponding deformation functions $v_i(x)$ are shown to the right.

C. Motion Sequence

Figure 3 shows two frames from a motion sequence showing M. Kass moving against a stationary background. The frames are separated in time by about 1.5 seconds. Results of the matching process are shown as follows: the original image has been mapped onto a surface that encodes estimated speed as elevation. The raised area shows the region in which the algorithm detects motion.

D. Stereo Matching

Figure 4 contains a stereogram showing a potato partly occluding a pear. The matching results are rendered as two shaded surfaces with depth computed from the disparity. An image coordinate grid is mapped onto the first surface and the left image is mapped onto the second. The reconstructed surfaces are rendered from an oblique viewpoint showing the computed surface discontinuities. Only those portions of the scene visible in the original stereogram are shown.

VII   Conclusion

The main contribution of this paper is two-fold. First, we introduced the notion of tracking the solution to the matching problem continuously over scale. Second, we developed a single system of first-order differential equations which characterize this process. The system is governed by an energy functional which balances similarity of the signals against smoothness of the deformation. The effectiveness of this approach has been demonstrated for both one- and two-dimensional signals.

References
Bakhvalov, N.S., "Numerical Methods," Mir, Moscow (1977).

Broit, C., "Optimal Registration of Deformed Images," Ph.D. Thesis, Computer and Information Science Dept., University of Pennsylvania, Philadelphia, PA (1981).

Burr, D.J., "A dynamic model for image registration," Computer Graphics and Image Processing, **15** (1981) 102–112.

Dahlquist, G., and Björck, A., "Numerical Methods," N. Anderson (trans.), Prentice-Hall, Englewood Cliffs, NJ (1974).

Fischler, M.A., and Elschlager, R.A., "The representation and matching of pictorial structures," IEEE Trans. Computers, **C-22** (1973) 67–92.

Gennery, D.B., "Modeling the Environment of an Exploring Vehicle by means of Stereo Vision," Ph.D. thesis, Stanford Artificial Intelligence Laboratory, also Artificial Intelligence Laboratory Memo 339 (1980).

Figure 3: Two frames from a motion sequences about 1.5 seconds apart. The original image has been texture mapped onto a surface that encodes speed as elevation. The raised area is moving while the background remains stationary.

Figure 1: Matching two one-dimensional signals. The signals are measurements of the resistivity of a geological structure as a function of depth at two different locations. From top to bottom, the signals first appear in their original form, then partially deformed at intermediate stages of the matching process, finally showing the end result. Above the signals are shown the deformation function **V** and the correlation gradient ∇Q.

The positive integer p indicates the highest order generalized spline that occurs in the functional, and this determines the maximum order of continuity (C^{p-1}) of the admissible deformations. The nonnegative continuity control functions $\mathbf{w}(\mathbf{x}) = [w_0(\mathbf{x}), \ldots, w_p(\mathbf{x})]$ determine the placement of discontinuities. A discontinuity of order $q < p$ is permitted to occur at $\mathbf{x}_0$ by forcing $w_i(\mathbf{x}_0) = 0$ for $i \geq q$ (see [Terzopoulos, 1986] for details).

The $p = 2$ order controlled-continuity constraint is employed in our implementations to date. If, for convenience, a "rigidity" function $\rho(\mathbf{x})$ and a "tension" function $[1 - \tau(\mathbf{x})]$ are introduced such that $w_0(\mathbf{x}) = 0$, $w_1(\mathbf{x}) = \rho(\mathbf{x})[1 - \tau(\mathbf{x})]$, and $w_2(\mathbf{x}) = \rho(\mathbf{x})\tau(\mathbf{x})$, then it is natural to view the functionals as characterizing "generalized piecewise continuous splines under tension." In particular, for the case of n signals in 1 dimension, the functional reduces to

$$S(\mathbf{v}) = -\int_{\Re} \rho(x)\Big\{[1 - \tau(x)]|\mathbf{v}_x|^2 + \tau(x)|\mathbf{v}_{xx}|^2\Big\}\, dx,$$

while for the case of 2 signals in 2 dimensions, it becomes

$$S(\mathbf{v}) = -\int\int_{\Re^2} \rho(x,y)\Big\{[1 - \tau(x,y)](|\mathbf{v}_x|^2 + |\mathbf{v}|_y^2) + \tau(x,y)(|\mathbf{v}_{xx}|^2 + 2|\mathbf{v}_{xy}|^2 + |\mathbf{v}_{yy}|^2)\Big\}\, dx\, dy,$$


Figure 2: Simultaneous matching of four signals. The signals are intensity profiles from a complex natural image. On the left the four signals are shown superimposed at several points in the matching process. As before, the original signals appear at the top and the final result is at the bottom. Note that coarse-scale features are aligned first in the matching process while fine scale features are matched at the end. The four corresponding deformation functions $v_i(x)$ are shown to the right.

where $x_1 = x$ and $x_2 = y$.

VI   Results

A.   Implementation Notes

Discretization of the continuous variational form of the matching problem can be carried out using standard methods. Although finite element methods offer the greatest flexibility, for simplicity we employ standard multidimensional finite difference formulas for uniform meshes to approximate the spatial derivatives in $S(\mathbf{V})$. These approximations yield local computations analogous to those in [Terzopoulos, 1983].

Equation 5 is a standard first-order initial value problem, for which solution methods abound. We have employed numerical methods of varying sophistication, each giving satisfactory results. In order of sophistication, these include Euler's method, a fourth-order Runge-Kutta method, and Adams-Moulton predictor-corrector methods. The latter offer the advantage that the step size can be automatically adapted, making them particularly robust [Dahlquist and Björck, 1974].

B.   One-dimensional Signals

The method is applicable to matching n signals, each of which is d-dimensional. Figure 1 shows the simplest case, that of matching two one-dimensional signals. The signals are measurements of the resistivity of a geological structure as a function of depth at two different locations. From top to bottom, the signals first appear in their original form, then partially deformed at intermediate stages of the matching process, finally showing the end result. The deformation function **V** and the correlation gradient ∇Q are shown above the signals.

Figure 4: A stereogram showing a potato partly occluding a pear (the images are reversed for free fusing). The matching results are rendered as two shaded surfaces with depth computed from the disparity. An image coordinate grid is mapped onto the first surface and the left image is mapped onto the second. The reconstructed surfaces are rendered from an oblique viewpoint showing the computed surface discontinuities. Only those portions of the scene visible in the original stereogram are shown.

Grimson, W.E.L., "An implementation of a computational theory of visual surface interpolation," *Computer Vision, Graphics, and Image Processing,* **22** (1983) 39–69.

Hannah, M.J., "Computer matching of areas in stereo images," Stanford Artificial Intelligence Laboratory memo, AIM-239, July (1974).

Marr, D., and Poggio, T. "A theory of human stereo vision," *Proc. Roy. Soc. Lond.*B, **204** (1979) 301–328.

Moravec, H.P., "Towards automatic visual obstacle avoidance," *Fifth International Joint Conference on Artificial Intelligence,* Cambridge, Massachusetts (1977) 584.

Mori, K., Kidodi, M., and Asada, H., "An iterative prediction and correction method for automatic stereo comparison," *Computer Graphics and Image Processing,* **2** (1973) 393–401.

Quam, L.H., "Hierarchical warp stereo," Proc. Image Understanding Workshop, New Orleans, LA, October, 1984, 149–155.

Rabiner, L.R., and Schafer, R.W., *Digital Processing of Speech Signals,* Prentice-Hall, NJ (1978).

Sankoff, D., and Kruskal, J.B., (eds.), *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison,* Addison-Wesley, Reading, MA, (1983).

Terzopoulos, D., "Multilevel computational processes for visual surface reconstruction," *Computer Vision, Graphics, and Image Processing,* **24** (1983) 52–96.

Terzopoulos, D., "Regularization of inverse visual problems involving discontinuities," *IEEE Trans. Pattern Analysis and Machine Intelligence,* **PAMI-8** (1986).

Witkin, A., "Scale Space Filtering," Proceedings of International Joint Conference on Artificial Intelligence, Karlsruhe (1983) 1019-1021.