Technical Report UCLA CSD_TR040024

Robust Multi-Scale Shape Distance and Correspondence via Integral Invariants

Siddharth Manay, Byung-Woo Hong[*], Daniel Cremers, Anthony Yezzi[†], and Stefano Soatto

June 30, 2004

Abstract

We propose a framework for comparing shapes represented as closed planar contours that is by construction invariant with respect to Euclidean and similarity transformations, and insensitive to high-frequency perturbations as well as to high-energy localized deformations such as spikes and slivers. Our representation relies on a scale-space of integral invariants that is used as the basis for establishing local correspondence between contours being compared. We present an implementation of the optimal correspondence and the resulting geodesic distance.

1. Introduction

Our goal is to compare objects having different shape. This has obvious implication in shape classi£cation for object recognition, content-based image retrieval, medical diagnosis, etc. At this level of generality, this is a monumental task that admits no simple meaningful solution [17]. Therefore, before we proceed any further, we need to specify what we mean by "objects," how we describe their "shape," and concentrate our attention on particular ways in which they can "differ." Within the scope of this paper, by *objects* we mean closed planar contours[1] embedded in $\mathbb{R}^2$; their *shape* is the equivalence class of objects obtained under the action of a £nite-dimensional group, such as the Euclidean, similarity, af£ne or projective group [13]. Therefore, objects that are obtained from a closed planar contour by rotating, translating, and scaling it, for example, have the same (similarity) shape; all other objects have different shape. However, in comparing shapes, we want to be insensitive to certain accidents that can occur to an object: for instance, in Fig. 2, we want the "jagged," "smooth," "spiked," and "wedged" objects to be judged as having shapes that are similar to that of the original bowtie. We prefer not to use the word "noise" when referring to these accidents because, with the exception of the jaggedness, they are not obtained with standard additive, zero-mean, small variance perturbations. For the case of the spiked and wedged objects, for instance, the perturbation can be quite signi£cant in energy, and highly localized along the contour. Our goal is to de£ne a distance with respect to which the spiked bowtie in Fig. 2 (last row) is close to the bowtie in the same £gure.

The type of accidents we want to be resistant to can be lumped in three categories: "small deformations" that result in small set-symmetric differences between the interior of the curves being compared, "high-frequency noise" that affects a wide portion of the contour, and "localized changes" that signi£cantly affect the total arclength of the contour but are spatially localized, such as spikes or wedges. Many in the literature have addressed the comparison of objects under small deformations in a way that is invariant with respect to various groups (see Sect.1.1); fewer have addressed the sensitivity to high-frequency noise and yet fewer have addressed localized changes [22, 20]. In this paper, we plan to develop a framework that will allow addressing all of these accidents in one go. To this end, we plan to employ a recently proposed representation of shape in terms of so-called *integral invariants*, so that the distance between objects will by construction be invariant with

[*]University of Oxford, Oxford OX1 3BW, UK

[†]Georgia Institute of Technology, Atlanta GA 30332

[1]Many of our considerations can be extended to compact surfaces embedded in $\mathbb{R}^3$.

respect to the action of the chosen group; de£ning a *scale space* of such invariants allows us to address high-frequency noise in a principled way. Finally, establishing *point correspondence* among contours allows us to handle localized changes.

Given the wealth of existing work on invariance, scale space and correspondence, our work naturally relates with a large body of literature, as we describe in the next subsection. The reader should notice that we consider each object as one entity, and perform no analysis or decomposition, so there is no notion of hierarchy or compositionality in our representation, which is therefore intrinsically low-level.

1.1 Relation to prior work and contributions of this paper

There exists a vast literature on comparing shapes, represented as collection of points [2, 21, 29, 12], curves [14, 31, 19, 4, 8, 32], continuous curves reduced to various type of graph representations [33, 25, 26, 18, 15, 11]: we represent curves as continuous objects living in in£nite-dimensional spaces. Within this choice, many have addressed matching curves under various types of motion [2, 21, 29] and deformations [14, 31, 19, 4, 8, 7, 32, 6, 27], some involving a mapping from one curve to another that has some optimality property [5, 8, 14, 19, 31, 6, 27, 4, 32]. A representation of closed planar contours that is invariant with respect to the action of subgroups of plane-projective transformations that enjoyed great popularity during the past decade was based on differential invariants [22, 8, 14, 19, 31]. Those, however, exhibited dif£culties in handling high-frequency noise. Various approaches were presented to address this dif£culty, notably the introduction of a scale space of curvature [30] and more recently the introduction of invariant curves that are based not on differential operators, but on integral operators applied to the contour or the characteristic function of its interior [16, 20]. It is relatively straightforward to de£ne a scale space of such invariants, by choosing parametric families of kernels with appropriate scale parameters. We employ one such representation in our work, since it addresses the issue of invariance to £nite-dimensional group actions and high-frequency noise at the outset.

However, when an object undergoes modi£cations that entail a signi£cant change in arclength concentrated locally on the contour, integral invariants, which are bounded positive functions of arclength, cannot be meaningfully matched (see Fig. 2). Therefore, a natural palliative consists of re-parameterizing arclength via a "warping function." This idea has been used in various contexts and applied to various representations. Most notably, Sebastian et. al. has applied it to a differential invariant representation of the curve using curvature [22] with very convincing results for noise-free shapes. Our warping operation is very similar in principle to Sebastian et. al.'s "alignment curve." However, we perform the warping operation on the integral invariant representation, rather than on curvature. While to the novice this may not seem like a signi£cant difference at £rst, the outcome is different indeed, both at the theoretical and practical level. In fact, unlike curvature, the integral invariant is a *positive, bounded* function of arclength, as we will review in Sect. 2. Therefore, it can be interpreted as an (un-normalized) mass density, and the problem of transforming one mass density onto another is well-posed, theoretically grounded in the optimal transport framework of Monge-Kantorovic[1], and for it there are optimal algorithms that are proven to converge to the unique global minimum. Furthermore, the integral invariant is intrinsically less sensitive to noise than curvature, and therefore we can handle high-frequency noise effortlessly and perform comparison in a scale space framework that is naturally de£ned by the chosen representation. Our implementation of the optimal warp is based on a discretization of the original curves and dynamic programming. This is standard in the alignment/warping literature [20, 22, 8, 12, 28, 3].

Our contribution, therefore, draws on extensive work on integral invariants and alignment-based shape matching. By combining these two bodies of work, we have developed a framework for shape comparison that is by construction invariant with respect to the action of the Euclidean or similarity group, insensitive with respect to high-frequency distributed noise, and with respect to localized deformations such as spikes and slivers. The resulting approach has close ties with optimal transport theory, and we show encouraging experimental results using a discretized implementation of shape warping.

2 Integral invariants

We represent a shape as a closed planar curve $\gamma : \mathbb{S} \mapsto \mathbb{R}^2$. For a group G acting on $\mathbb{R}^2$, an *integral G-invariant* is an integral function of the curve that does not vary due to the action of any $g \in G$. Formally, for a kernal $h : \mathbb{R}^2 \times \mathbb{R}^2 \mapsto \mathbb{R}$,

$$I_\gamma(p) = \int_{\bar{\gamma}} h(p, x) d\mu(x) \tag{1}$$

is an integral invariant if

$$I_\gamma(p) = I_\gamma(gp) \, \forall g \in G. \tag{2}$$

These invariants share several desirable qualities of their more well-known cousins, differential invariants, including locality. However, unlike differential invariants, integral invariants are far more robust to noise. Speci£c examples include the

"observed transport" shape measure [20] and the local area invariant [16], which is described below. Both of these invariants are parameterized by r, the radius of a circular kernel. This allows the invariants to be extended to a discrete or continuous multi-scale functions, yielding a vector-valued invariant. However, the "observed transport" invariant is evaluated on the contour, while the local area invariant is evaluated on an area form. Further, domain of integration of the "observed transport" invariant is based on visibility, which is very sensitive to curve perturbations near the kernel center.

For this reason we favor the invariant suggested in [16]. Speci£cally, the invariant is based on a kernel $h(p, x) = \chi(B_r(p) \cap \bar{\gamma})(x)$, which represents the indicator function of the intersection of a small circle of radius r centered at the point p with the interior of the curve γ. For a radius r, the corresponding integral invariant

$$I^r_\gamma(p) \doteq \int_{B_r(p)\cap\bar{\gamma}} dx \tag{3}$$

can be thought of as a function from the interval $[0, 1]$ to the positive reals, bounded above by the area of the region bounded by the curve γ. Alternately, normalizing the functional by the area of B_r maps the interval $[0, 1]$ to $[0, 1]$. This is illustrated in Fig. 1 and examples are shown in Fig. 2, which demonstrates (in the £rst four rows) the effectiveness of this representation for comparing shapes under Euclidian transformations, in global additive noise, and under smoothing. However, Fig. 2 (last row) demonstrates that the local area invariant, like all other integral and differential invariants, is not robust to a highly localized perturbation such as a spike. In the next section we present a method to warp the parameterization of the curves so that such comparisons can be made robustly.

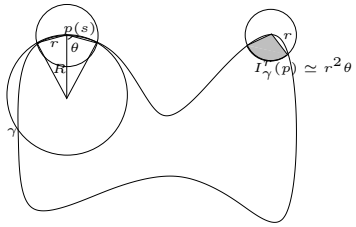

Figure 1: *Integral area invariant de£ned by eq. (3).*

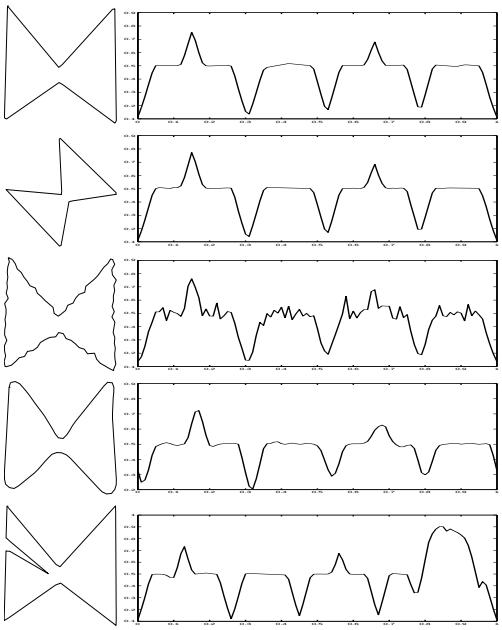

Figure 2: *Sample shapes and their integral invariants.* Kernal size is 15% of the curve's bounding box.. Note the added "hump" on the right side of the invariant due to the spike in the spiked bowtie (last row).

3 Shape distance and warping integral invariants

While invariants can be designed to be robust to group actions, and integral invariants are robust to noise, an invariant for highly localized, high energy perturbations of a shape, such as the addition of spikes, slivers, or other parts, or the "disappearance" of parts due to occlusion or change in con£guration (i.e. a hand with £ve £ngers extended vs. four £ngers extended) is much more problematic. A naive distance de£ned on an invariant would yield a relatively large distance between the two hands above (implying they are very different) while intuitively these shapes should have a small distance between them.

From the perspective of invariants, the large difference between shapes with local variation is due to the parameterization of the curves, as demonstrated in Fig. 2. If the spike could be re-parameterized to have a smaller arclength, then its segment of the invariant would contribute less to the distance between the two shapes. This warping of the parameterization of one shape, or more symmetrically both shapes, would allow the computation of an intuitively more satisfying distance while effectively inducing point correspondences between the two curves.

More formally, we ultimately wish to compute $D(\vec{I}_1, \vec{I}_2)$, a distance between two vector-valued integral invariants. We emphasize vector-valued invariants to ensure that this formulation generalizes to multi-scale analysis. Rather than compute a naive distance between two invariants,

$$D(\vec{I}_1, \vec{I}_2) = \arg\min_{s_0} \int_0^1 |\vec{I}_1(s) - \vec{I}_2(s)|^p \, ds, \tag{4}$$

we introduce added robustness by warping the parameterization of both curves via a function $\rho(\mu) = (h_1(\mu), h_2(\mu))$, $\rho : [0,1] \mapsto [0,1] \times [0,1]$, (called *alignment curve* in [22]). Each value of μ indexes a correspondence between $\gamma_1(h_1(\mu))$ and $\gamma_2(h_2(\mu))$. Now a distance between two invariants can be thought of as a cost functional that aims to ful£ll two criteria: it should put into correspondence points with similar integral invariant, and it should shrink or stretch the two contours as little as possible. Both of these goals can be represented by a distance of the form

$$D(\vec{I}_1, \vec{I}_2) = \arg\min_{h_1, h_2, s_0} \int |\vec{I}_1(h_1) - \vec{I}_2(h_2)|^p \, d\rho, \tag{5}$$

For simplicity, we consider the initial corresponding pair (i.e. s_0) to be known; we discuss how to £nd this pair in Sect. 3.2. The exponent p weights between the two goals:

$$\begin{array}{ll} p \to 0 & \text{minimize path length (stretching \& shrinking)} \\ p \to \infty & \text{similarity of invariants dominates} \end{array} \tag{6}$$

Expanding $d\rho$ in (5) as

$$d\rho = \sqrt{\left(\frac{dh_1}{d\mu}\right) + \left(\frac{dh_2}{d\mu}\right)} \, d\mu \tag{7}$$

(where μ is the arclength parameterization of ρ) one clearly sees that the norm in (5) imposes similarity of the invariants and the $d\rho$ term (speci£cally the portion under the square root) measures stretching or shrinking of the two contours.

In [22, 20] and others, the authors propose a cost functional which consists of two separate terms for measuring the stretching and the bending, respectively, with a parameter R – see equation (4) in [22]. The second component (weighted by R) also depends on the path length (i.e. shrinking and stretching effects) as well as similarity of invariants. Therefore, appropriate choices of R to impose only one of the two constraints do not exist. One could normalize the second term so that it would no longer depend on the length of ρ, thus measuring the mean dissimilarity of the invariants along the path. This separates stretching and similarity of invariants. However, since we are optimizing with respect to the path, this would be quite hard to implement.

3.1 Relationship to optimal transport

The classical problem of optimal transport put forth by Monge in the 18th century considers two measures μ and ν de£ned over $U \subset \Re^n$ and $V \subset \Re^n$ respectively of identical total mass

$$\int_U d\mu = \int_V d\nu$$

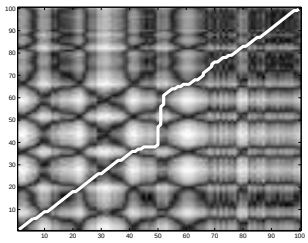

Figure 3: *Optimal path through the graph.* The path warps the parameterization of the hand (on the bottom of the graph) and the parameterization of the noisy four-£ngered hand (on the left side of the graph). Both shapes are shown in Fig. 4. The gray levels are the distance between points; lighter color indicates higher distance.

and seeks to £nd a mass preserving mapping $\psi : U \to V$, where $\nu(E) = \mu\left(\psi^{-1}(E)\right)$ for any measurable $E \subset V$. Of course there are many such mappings ϕ and the goal of the Monge problem is to £nd the one that minimizes the total transport cost

$$\int_U |\psi(x) - x| \, d\mu$$

which may be interpreted physically as the total work required to move the mass from its source distribution μ to its target distribution ν according to the transport scheme de£ned by ψ.

We may cast our problem into this framework due to the fact that that the integral invariants I we consider are positive and bounded and therefore may be interpreted as densities in $\mathbb{R}^1$ of £nite total mass. In other words given $I_1(s_1)$, an integral invariant computed on one curve (with arclength parameter s_1) and $I_2(s_2)$, an integral invariant computed on a second curve (with arclength parameter s_1) we may set $d\mu = I_1(s_1)ds_1$ and $d\nu = I_2(s_2)ds_2$ and $U = [0, L_1]$ and $V = [0, L_2]$ where L_1 and L_2 are the total arclengths of the £rst and second curves respectively. To obtain the mass balance condition required of this framework, we may simply employ a uniformly scaled version of the same integral invariant used for I_1 to calculate I_2 where the scale parameter is chosen to equalize their total integrals.

Now the correspondence problem is formulated as £nding the optimal transport $\rho : U \to V$ of the two densities I_1 and I_2, and in this simple one-dimensional case, the optimal transport map ρ may be reformulated as a reparameterization of the £rst curve by rescaling by the ratio of the two curves' arclengths L_1/L_2. In a similar way, the inverse map ρ^{-1} may be rescaled by L_2/L_1 to obtain a reparameterization of the second curve. The distance between the two curves can then be measured by the cost of the optimal transport

$$\int_0^{L_1} |\rho(s_1) - s_1| I_1(s_1) \, ds_1$$

or equivalently

$$\int_0^{L_2} |\rho^{-1}(s_2) - s_2| I_2(s_2) \, ds_2.$$

Although there exist ef£cient numerical implementations for £nding the solution for the optimal transport map (see for instance [1]) that relies on successive re£nements of an initial guess, in the remainder of this paper we will employ a simple solution based on a discretization of the contours involved, which naturally results in a well-known dynamic programming approach. The implementation of the optimal transport framework to invariant warping is the subject of our continuing research.

3.2 Implementation

In Sect. 3 we presented a distance between invariants (and therefore shape) that depends on the choice of a warping function ρ. To complete the calculation to distance, and to establish a local correspondence between the curves, we must optimize distance with respect to ρ. This section brie¤y outlines the implementation of a well-known approach to globally optimize the correspondence for a discrete representation of the curves as ordered sets of points.

Our implementation is based on dynamic programming approaches similar to those employed by many in the shape, stereo, and registration (for medical imaging) communities [20, 22, 8, 12, 28, 3]. We £rst require that h_1 and h_2 be monotonic. A intuitive algorithm would be as follows. We £rst £nd an initial correspondence between a point on each curve (discussed

below). The "next" correspondence should be the choice of action that minimizes the distance; the actions are (1) locally contracting the £rst curve onto the second, (2) locally contracting the second curve onto the £rst, or (3) locally mapping the curves as one-to-one and onto. This sketch of the algorithm lends itself to a graph formulation, where each node of a periodic, directed graph is a correspondence between a point on each curve, and each edge links the current node to the possible "next" nodes. The edges are weighted by the distance between the invariants associated with the "next" node, c. f. eq. 5. For two curves with N points, the graph has N^2 nodes and $3N^2$ edges. Given the initial correspondence, the well known Dijkstra Algorithm will £nd the globally optimal shortest path back to the initial node. The resulting path is ρ and its distance is the distance between the shapes. An example of the result of this algorithm is shown in Fig. 3.

No fast algorithm exists to determine the best choice of the initial correspondence. Previous implementations (cited above) £x a point on the £rst curve ($h_1(0)$) and pair it with all possible choices of points on the second curve, calculating the path for each pair to determine the shortest. This exhaustive search can be avoided by observing that strong features, such as corners or convex/concave points, provide a heuristic way to propose point correspondences. These points are easily classi£ed in the invariant space. For instance, for the local area invariant, points with little or no curvature are in a ball around $I = .5$. We £nd a subset of points outside this ball, $\{s \,|\, |I_1(s) - .5| > T\}$ where T is some threshold (typically $T = .1$). These points, with their nearest neighbors on I_2, form a set of likely initial correspondences. In this way, we can £nd an initial correspondence and compute the warping function ρ and its associated distance for two curves with 100 points each in less than 1 second with MATLAB on a computer with an Intel 800 MHz processor.

4 Experiments

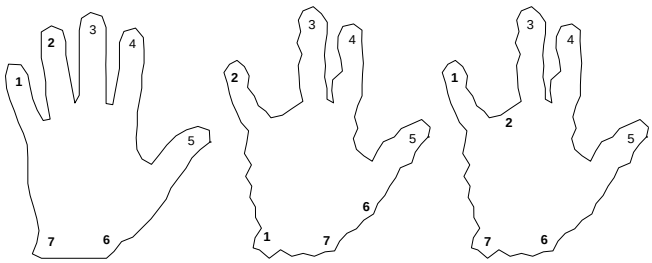

(a) Original (b) Curvature (c) Integral Inv.

Figure 4: *Correspondences on noisy hand with occluded £nger.* The numbers label corresponding segments of each curve. [LEFT] Original shape [CENTER], correspondences found via curvature, [RIGHT] correspondences found via integral invariant. Results are similar for all choices of p.

This section presents experiments that highlight the additional robustness to noise of integral-invariant-based shape distances and correspondences when compared to similarly implemented differential-invariant methods.

Fig. 4 demonstrates the difference in the correspondences induced by the integral-invariant-based distance discussed in Sect. 3 compared to a similar distance constructed on a differential invariant, namely curvature as in [22], on a partially occluded, noisy hand. The correspondence induced by curvature is not as robust to these local and global perturbations, and associates the ring £nger(a2) in the leftmost shape with the pinky(b2) in the middle shape; due to the ordering constraints in the optimization framework, this mismatch forces several other erroneous correspondences, such as the pinky(a1) with the heel(b1) of the hand. However, the correspondence induced by the integral invariant correctly associates all the features available in both shapes, forcing the association of the ring £nger(a2) with the knuckle(c2) the ring £nger should be connected to.

In Fig. 5 the results of matching and retrieving noisy shapes (shown on the left side) from a database (shown across the top) are shown. We especially highlight several pairs where representation by differential invariants leads to mismatches. These include the noisy bowties in the £rst and second rows, which, via curvature-based-distance, are as close to the hands (rightmost columns) as they are to the correct bowties (leftmost columns). For the same bowtie shapes, distances calculated with the integral invariant more consistently favor bowties (and bowtie-like £sh!) over hands and £sh. Examination of the data shows several such cases where integral invariants are more robust than curvature on noisy shapes.

Shape interpolation, commonly known as *morphing*, has been recognized as a powerful tool especially for visual effects including computer graphics and computer animation [10, 9, 24, 23]. The general problem is to £nd natural intermediate

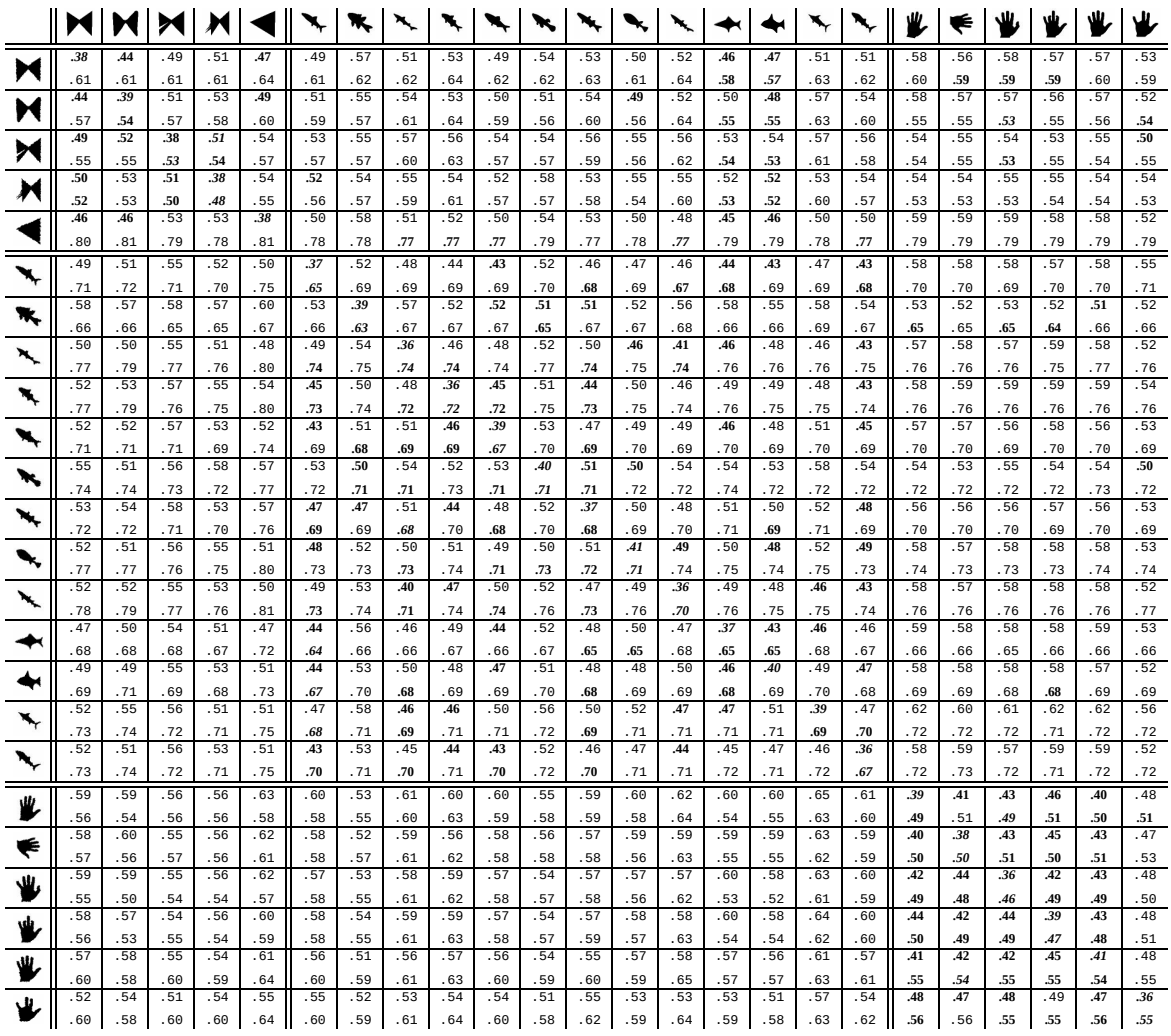

	.38	***.44***	.49	.51	***.47***	.49	.57	.51	.53	.49	.54	.53	.50	.52	***.46***	***.47***	.51	.51	.58	.56	.58	.57	.57	.53
	.61	.61	.61	.61	.64	.61	.62	.62	.64	.62	.62	.63	.61	.64	***.58***	***.57***	.63	.62	.60	***.59***	***.59***	***.59***	.60	.59
	.44	***.39***	.51	.53	***.49***	.51	.55	.54	.53	.50	.51	.54	***.49***	.52	.50	***.48***	.57	.54	.58	.57	.57	.56	.57	.52
	.57	***.54***	.57	.58	.60	.59	.57	.61	.64	.59	.56	.60	.56	.64	***.55***	***.55***	.63	.60	.55	.55	***.53***	.55	.56	***.54***
	.49	***.52***	***.38***	***.51***	.54	.53	.55	.57	.56	.54	.54	.56	.55	.56	.53	.54	.57	.56	.54	.55	.54	.53	.55	***.50***
	.55	.55	***.53***	***.54***	.57	.57	.57	.60	.63	.57	.57	.59	.56	.62	***.54***	***.53***	.61	.58	.54	.55	***.53***	.55	.54	.55
	.50	.53	***.51***	***.38***	.54	***.52***	.54	.55	.54	.52	.58	.53	.55	.55	.52	***.52***	.53	.54	.54	.54	.55	.55	.54	.54
	.52	.53	***.50***	***.48***	.55	.56	.57	.59	.61	.57	.57	.58	.54	.60	***.53***	***.52***	.60	.57	.53	.53	.53	.54	.54	.53
	.46	***.46***	.53	.53	***.38***	.50	.58	.51	.52	.50	.54	.53	.50	.48	***.45***	***.46***	.50	.50	.59	.59	.59	.58	.58	.52
	.80	.81	.79	.78	.81	.78	.78	***.77***	***.77***	***.77***	.79	.77	.78	***.77***	.79	.79	.78	***.77***	.79	.79	.79	.79	.79	.79
	.49	.51	.55	.52	.50	***.37***	.52	.48	.44	***.43***	.52	.46	.47	.46	***.44***	***.43***	.47	***.43***	.58	.58	.58	.57	.58	.55
	.71	.72	.71	.70	.75	***.65***	.69	.69	.69	.69	.70	***.68***	.69	***.67***	***.68***	.69	.69	***.68***	.70	.70	.69	.70	.70	.71
	.58	.57	.58	.57	.60	.53	***.39***	.57	.52	***.52***	***.51***	***.51***	.52	.56	.58	.55	.58	.54	.53	.52	.53	.52	***.51***	.52
	.66	.66	.65	.65	.67	.66	***.63***	.67	.67	.67	***.65***	.67	.67	.68	.66	.66	.69	.67	***.65***	.65	***.65***	***.64***	.66	.66
	.50	.50	.55	.51	.48	.49	.54	***.36***	.46	.48	.52	.50	***.46***	***.41***	***.46***	.48	.46	***.43***	.57	.58	.57	.59	.58	.52
	.77	.79	.77	.76	.80	***.74***	.75	***.74***	***.74***	.74	.77	***.74***	.75	***.74***	.76	.76	.76	.75	.76	.76	.76	.75	.77	.76
	.52	.53	.57	.55	.54	***.45***	.50	.48	***.36***	***.45***	.51	***.44***	.50	.46	.49	.49	.48	***.43***	.58	.59	.59	.59	.59	.54
	.77	.79	.76	.75	.80	***.73***	.74	***.72***	***.72***	***.72***	.75	***.73***	.75	.74	.76	.75	.75	.74	.76	.76	.76	.76	.76	.76
	.52	.52	.57	.53	.52	***.43***	.51	.51	***.46***	***.39***	.53	.47	.49	.49	***.46***	.48	.51	***.45***	.57	.57	.56	.58	.56	.53
	.71	.71	.71	.69	.74	.69	***.68***	***.69***	***.69***	***.67***	.70	***.69***	.70	.69	.70	.69	.70	.69	.70	.70	.69	.70	.70	.69
	.55	.51	.56	.58	.57	.53	***.50***	.54	.52	.53	***.40***	***.51***	***.50***	.54	.54	.53	.58	.54	.54	.53	.55	.54	.54	***.50***
	.74	.74	.73	.72	.77	.72	***.71***	***.71***	.73	***.71***	***.71***	***.71***	.72	.72	.74	.72	.72	.72	.72	.72	.72	.72	.73	.72
	.53	.54	.58	.53	.57	***.47***	***.47***	.51	***.44***	.48	.52	***.37***	.50	.48	.51	.50	.52	***.48***	.56	.56	.56	.57	.56	.53
	.72	.72	.71	.70	.76	***.69***	.69	***.68***	.70	***.68***	.70	***.68***	.69	.70	.71	***.69***	.71	.69	.70	.70	.70	.69	.70	.69
	.52	.51	.56	.55	.51	***.48***	.52	.50	.51	.49	.50	.51	***.41***	***.49***	.50	***.48***	.52	***.49***	.58	.57	.58	.58	.58	.53
	.77	.77	.76	.75	.80	.73	.73	***.73***	.74	***.71***	***.73***	***.72***	***.71***	.74	.75	.74	.75	.73	.74	.73	.73	.73	.74	.74
	.52	.52	.55	.53	.50	.49	.53	***.40***	***.47***	.50	.52	.47	.49	***.36***	.49	.48	***.46***	***.43***	.58	.57	.58	.58	.58	.52
	.78	.79	.77	.76	.81	***.73***	.74	***.71***	.74	***.74***	.76	***.73***	.76	***.70***	.76	.75	.75	.74	.76	.76	.76	.76	.76	.77
	.47	.50	.54	.51	.47	***.44***	.56	.46	.49	***.44***	.52	.48	.50	.47	***.37***	***.43***	***.46***	.46	.59	.58	.58	.58	.59	.53
	.68	.68	.68	.67	.72	***.64***	.66	.66	.67	.66	.67	***.65***	***.65***	.68	***.65***	***.65***	.68	.67	.66	.66	.65	.66	.66	.66
	.49	.49	.55	.53	.51	***.44***	.53	.50	.48	***.47***	.51	.48	.48	.50	***.46***	***.40***	.49	***.47***	.58	.58	.58	.58	.57	.52
	.69	.71	.69	.68	.73	***.67***	.70	***.68***	.69	.69	.70	***.68***	.69	.69	***.68***	.69	.70	.68	.69	.69	.68	***.68***	.69	.69
	.52	.55	.56	.51	.51	.47	.58	***.46***	***.46***	.50	.56	.50	.52	***.47***	***.47***	.51	***.39***	.47	.62	.60	.61	.62	.62	.56
	.73	.74	.72	.71	.75	***.68***	.71	***.69***	.71	.71	.72	***.69***	.71	.71	.71	.71	***.69***	***.70***	.72	.72	.72	.71	.72	.72
	.52	.51	.56	.53	.51	***.43***	.53	.45	***.44***	***.43***	.52	.46	.47	***.44***	.45	.47	.46	***.36***	.58	.59	.57	.59	.59	.52
	.73	.74	.72	.71	.75	***.70***	.71	***.70***	.71	***.70***	.72	***.70***	.71	.71	.72	.71	.72	***.67***	.72	.73	.72	.71	.72	.72
	.59	.59	.56	.56	.63	.60	.53	.61	.60	.60	.55	.59	.60	.62	.60	.60	.65	.61	***.39***	***.41***	***.43***	***.46***	***.40***	.48
	.56	.54	.56	.56	.58	.58	.55	.60	.63	.59	.58	.59	.58	.64	.54	.55	.63	.60	***.49***	.51	***.49***	***.51***	***.50***	***.51***
	.58	.60	.55	.56	.62	.58	.52	.59	.56	.58	.56	.57	.59	.59	.59	.59	.63	.59	***.40***	***.38***	***.43***	***.45***	***.43***	.47
	.57	.56	.57	.56	.61	.58	.57	.61	.62	.58	.58	.58	.56	.63	.55	.55	.62	.59	***.50***	***.50***	***.51***	***.50***	***.51***	.53
	.59	.59	.55	.56	.62	.57	.53	.58	.59	.57	.54	.57	.57	.57	.60	.58	.63	.60	***.42***	***.44***	***.36***	***.42***	***.43***	.48
	.55	.50	.54	.54	.57	.58	.55	.61	.62	.58	.57	.58	.56	.62	.53	.52	.61	.59	***.49***	***.48***	***.46***	***.49***	***.49***	.50
	.58	.57	.54	.56	.60	.58	.54	.59	.59	.57	.54	.57	.58	.58	.60	.58	.64	.60	***.44***	***.42***	***.44***	***.39***	***.43***	.48
	.56	.53	.55	.54	.59	.58	.55	.61	.63	.58	.57	.59	.57	.63	.54	.54	.62	.60	***.50***	***.49***	***.49***	***.47***	***.48***	.51
	.57	.58	.55	.54	.61	.56	.51	.56	.57	.56	.54	.55	.57	.58	.57	.56	.61	.57	***.41***	***.42***	***.42***	***.45***	***.41***	.48
	.60	.58	.60	.59	.64	.60	.59	.61	.63	.60	.59	.60	.59	.65	.57	.57	.63	.61	***.55***	***.54***	***.55***	***.55***	***.54***	.55
	.52	.54	.51	.54	.55	.55	.52	.53	.54	.54	.51	.55	.53	.53	.53	.51	.57	.54	***.48***	***.47***	***.48***	.49	***.47***	***.36***
	.60	.58	.60	.60	.64	.60	.59	.61	.64	.60	.58	.62	.59	.64	.59	.58	.63	.62	***.56***	.56	***.55***	***.55***	***.56***	***.55***

Figure 5: *Noisy shape recognition from a database of 24 shapes.* The upper number in each cell is the distance computed via the local-area integral invariant; the lower number is the distance computed via curvature invariant. The bold, italics number in each row represents the best match for the noisy shape at the left of that row; the four remaining bold numbers represent the next four best matches See text for more details.

shapes from one shape to the other. In general shape interpolation is an ill-posed problem in a sense that there are numerous ways to interpolate between two shapes and little has been discussed on notions of optimality in a quantitative way. In our framework, an optimal correspondence is obtained between two shapes and a linear interpolation is applied to primitive shape properties, which are relative distance and angle from the center point based on the correspondence. One example of interpolation is demonstrated as shown in Figure 6.

5 Conclusions

In this conference paper we have addressed two open questions in the computation of shape distance. Multi-scale *integral* group-invariant descriptions of shapes are employed for robustness to high-frequency noise and small deformations. Then shape warping, the application of reparameterization to induce correspondences, is used to account for large localized changes such as occlusions and con£guration changes. We embed both of these concepts in a formulation of a shape distance, and outline how distance and optimal correspondence are jointly computed via fast dynamic programming algorithms. To conclude, this model's robustness is demonstrated for applications such as shape matching on real shape data. Our ongoing research into these models includes implementing the optimization in an optimal-transport framework.

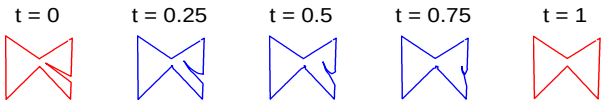

Figure 6: *Morphing based on correspondences found via integral-invariants.*

References

[1] S. Angentent, S. Haker, and A. Tannenbaum. Minimizing ¤ows for the Monge–Kantorovich problem. *SIAM J. on Math. Anal.*, 35(1):61–97, 2003.

[2] N. Ayache and O. Faugeras. HYPER: A new approach for the recognition and positioning of two-dimensional objects. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 8(1):44–54, Jan. 1986.

[3] M. Bakircioglu, U. Grenander, N. Khaneja, and M. I. Miller. Curve matching on brain surfaces using frenet distances. *Human Brain Mapping*, 6:329–333, 1998.

[4] R. Basri, L. Costa, D. Geiger, and D. Jacobs. Determining the similarity of deformable shapes. *Vision Research*, 38:2365–2385, 1998.

[5] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 24(24), Apr 2002.

[6] I. Cohen, N. Ayache, and P. Sulger. Tracking points on deformable objects using curvature information. In *Proc. European Conf. on Comp. Vision*, pages 458–466, 1992.

[7] A. DelBimbo and P. Pala. Visual image retrieval by elastic matching of user sketches. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 19(2):121–132, Feb. 1997.

[8] Y. Gdalyahu and D. Weinshall. Flexible syntactic matching of curves and its application to automatic hierarchical classication of silhouettes. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 21(12):1312–1328, 1999.

[9] A. Kaul and J. Rossignac. Solid-interpolating deformations: Construction and animation of pips. In *Proc. Eurographics*, pages 493–505, 1991.

[10] J. Kent, W. Carlson, and R. Parent. Shape transformation for polyhedral objects. In *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 47–54, 1992.

[11] P. N. Klein, S. Tirthapura, D. Sharvit, and B. Kimia. A tree-edit-distance algorithm for comparing simple, closed shapes. In *Symposium on Discrete Algorithms*, pages 696–704, 2000.

[12] L. J. Latecki and R. Lakämper. Shape similarity measure based on correspondence of visual parts. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 22(10):1185–1190, 2000.

[13] H. Le and D. G. Kendall. The riemannian structure of euclidean shape spaces: a novel environment for statistics. *The Annals of Statistics*, 21(3):1225–1271, 1993.

[14] H. Liu and M. Srinath. Partial shape classi£cation using contour matching in distance transforms. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 12(2):1072–1079, Nov. 1990.

[15] T. Liu and D. Geiger. Approximate tree matching and shape similarity. In *ICCV*, pages 456–462, 1999.

[16] S. Manay, B. Hong, A. Yezzi, and S. Soatto. Integral invariant signatures. Submitted to ECCV '03, Oct. 2003.

[17] D. Mumford. Mathematical theories of shape: do they model perception? In *In Geometric methods in computer vision, volume 1570*, pages 2–10, 1991.

[18] M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 21(11):1105–1120, 1999.

[19] A. Pikaz and I. Dinstein. Matching of partially occluded planar curves. *Patt. Rec.*, 28(2):199–209, Feb. 1995.

[20] A. Pitiot, H. Delingette, A. Toga, and P. Thompson. Learning Object Correspondences with the Observed Transport Shape Measure. In *Information Processing in Medical Imaging IPMI'03*, 2003.

[21] J. Schwartz and M. Sharir. Identi£cation of partially obscured objects in two and three dimensions by matching noisy characteristic curves. *Int. J. Rob. Res.*, 6(2):29–44, 1987.

[22] Thomas B. Sebastian, Philip N. Klein, and Benjamin B. Kimia. Alignment-based recognition of shape outlines. *Lecture Notes in Computer Science*, 2059:606–??, 2001.

[23] T. Sederberg, P. Gao, G. Wang, and H. Mu. 2D shape blending: An intrinsic solution to the vertex path problem. In *Computer Graphics (SIGGRAPH '93 Proceedings)*, volume 27, pages 15–18, 1993.

[24] T. Sederberg and E. Greenwood. A physically based approach to 2D shape blending. In *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 25–34, 1992.

[25] D. Sharvit, J. Chan, H. Tek, and B. Kimia. Symmetry-based indexing of image databases. In *IEEE Workshop on Content-based Access of Image and Video Libraries*, pages 56–62, 1998.

[26] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. In *ICCV*, pages 222–229, 1998.

[27] H. Tagare, D. O'Shea, and A. Rangarajan. A geometric correspondence for shape-based non-rigid correspondence. In *Intl. Conf. on Comp. Vision*, pages 434–439, 1995.

[28] C. Tomasi and R. Manduchi. Stereo without search. In *Proc. European Conf. on Comp. Vision*, pages 452–465, 1996.

[29] S. Umeyama. Parameterized point pattern matching and its application to recognition of object families. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 15(2):136–144, Feb. 1993.

[30] A. P. Witkin. Scale-space £ltering. *Int. Joint. Conf. Arti£cial Intelligence*, pages 1019–1021, 1983.

[31] H. Wolfson. On curve matching. *IEEE Trans. on Patt. Anal. and Mach. Intell.*, 12(5):483–489, May 1990.

[32] L. Younes. Optimal matching between shapes via elastic deformations. *Image and Vision Computing*, 17:381–389, 1999.

[33] S. Zhu and A. Yuille. Forms: A ¤exible object recognition and modeling system. *Int. J. Comp. Vision*, 20(3):187–212, 1996.