

Deformation Invariant Image Matching

Haibin Ling

David W. Jacobs

Center for Automation Research, Computer Science Department
University of Maryland, College Park
{hbling, djacobs}@ umiacs.umd.edu

Abstract

We propose a novel framework to build descriptors of local intensity that are invariant to general deformations. In this framework, an image is embedded as a 2D surface in 3D space, with intensity weighted relative to distance in x - y . We show that as this weight increases, geodesic distances on the embedded surface are less affected by image deformations. In the limit, distances are deformation invariant. We use geodesic sampling to get neighborhood samples for interest points, then use a geodesic-intensity histogram (GIH) as a deformation invariant local descriptor. In addition to its invariance, the new descriptor automatically finds its support region. This means it can safely gather information from a large neighborhood to improve discriminability. Furthermore, we propose a matching method for this descriptor that is invariant to affine lighting changes. We have tested this new descriptor on interest point matching for two data sets, one with synthetic deformation and lighting change, another with real non-affine deformations. Our method shows promising matching results compared to several other approaches.

1. Introduction

We propose a novel framework for building image descriptions that are invariant to deformations. An intensity image is treated as a surface embedded in 3D space, with the third coordinate proportional to the intensity values with an *aspect weight* α and the first two coordinates proportional to x - y with weight $1-\alpha$. As α increases, the geodesic distance on the embedded surface becomes less sensitive to image deformations. In the limit when $\alpha \rightarrow 1$, the geodesic distance is exactly deformation invariant. Based on this idea, we use *geodesic sampling* to get sample points on the embedded surface, then build the *geodesic-intensity histogram* (GIH) as a local descriptor. GIH captures the spatial distribution of intensities on the embedded manifold. With




Figure 1. Two images to be matched. Note that the right bottom corner of the flag in (b) is folded.

$\alpha = 1$, it is exactly invariant to deformation. In matching experiments on data sets with both synthetic and real deformations, GIH demonstrates promising results in comparison to other approaches.

Our work builds on much recent work on invariant local descriptors. This work has found wide application in areas such as object recognition [5, 13], wide baseline matching [19, 23], and image retrieval [20]. However, this previous work focuses on invariance to specific transformation groups, such as affine transformations. Affine invariant matching is useful when viewpoint changes relative to a rigid object that has locally planar regions large enough to contain distinctive intensity variations. It is less appropriate for objects without planar parts. For example, with a white Lambertian object, all intensity variations are due to 3D shape variations that lead to non-affine deformations as viewpoint changes. We are also interested in matching images of non-rigid objects, such as a flag waving or an animal moving its body (see Figure 1).

We make two main contributions. 1) GIH is the first local image descriptor that is invariant to general deformations. 2) Embedding images as surfaces in 3D and varying their aspect weight provides a novel framework for dealing with image deformation. This also suggests methods for

finding deformation invariant feature points, and for building descriptors in which we trade off discriminability and deformation insensitivity with our choice of α . In addition, GIH does not require a specific interest region detector, since geodesic sampling automatically determines regions for gathering information.

The rest of the paper is organized as follows. Sec. 2 discusses related works. Sec. 3 discusses deformation invariant features. Sec. 3.1 provides intuitions about deformation invariance through a 1D example. Sec. 3.2 shows how the aspect weight α relates to deformations by studying its relation to curve lengths on embedded image surfaces. Sec. 3.3 talks about geodesic distances and their computation and Sec. 3.4 explains geodesic sampling. Sec. 3.5 introduces the proposed descriptor, the geodesic-intensity histogram. Sec. 3.6 discusses several practical issues in using GIH, including illumination change, the interest points, and the choice of α . Sec. 4 describes all the experiments and analyzes the results. Sec. 5 concludes.

2. Related Work

Mikolajczyk and Schmid [15] gave a review and performance evaluation of several local descriptors including steerable filters [4], moment invariants [24], complex filters [19, 1], scale invariant feature transform (SIFT) [14] and cross-correlation. Most relevant to us are descriptors that are invariant to image transformations. For example, Lindeberg [11] proposed extracting scale invariant regions via the extremum of the scale space. Lowe [14] proposed a fast and efficient way to compute the scale invariant feature transform (SIFT), which measures the gradient distribution in detected scale invariant regions. Mikolajczyk and Schmid [16] proposed an affine invariant interest point detector through combining a scale invariant detector and the second moment of Harris corners [6]. Other work on affine invariant features can be found in [8, 18]

Our method can be categorized with so-called distribution based descriptors, which use histograms to capture local image information. These include the spin image [10, 7], shape context [2], and PCA-SIFT [9]. Our method differs from all these in two ways. First, our method is invariant to all deformations. Second, our descriptor automatically detects its support region, i.e., it does not need a special deformation invariant detector.

Treating images as 2D surfaces embedded in 3D space is not new. Our work is particularly motivated by the Beltrami framework proposed in [22]. They treat images as 2D manifolds and operators on the manifold are applied directly on the surface for low level vision tasks such image enhancement. Our work focuses more on feature extraction. Also we are more interested in using large aspect weights to produce deformation invariance.

Geodesic distance has also been used in object recognition. For example, in [3] it is used to build bending invariant signatures for real surfaces. Our work is different in that we are using embedded surfaces that vary according to aspect weights. This achieves deformation invariance for 2D images, as opposed to bending invariance for 3D data.

This work can also be viewed as a general version of our previous work of using the inner-distance for shape matching [12]. If we treat a given shape as a binary image, the inner-distance between two boundary points is the same as the geodesic distance used in this paper for α near 1.

3. Deformation Invariant Features

In this section we first discuss deformation invariance within the framework of embedded image surfaces, then introduce geodesic sampling and the geodesic-intensity histogram, which is invariant to deformation. After that, some practical issues in using this descriptor are discussed.

3.1. Intuitions about Deformation Invariance

We consider deformation as homeomorphisms (continuous, one-to-one transformations) between two images. Intensities change position, but not their value, with deformations. To obtain distinctive descriptors we collect intensities from a neighborhood. Our problem is to do this in a way that is invariant to deformation.

To gain intuition, we consider a one dimensional image. Figure 2(a) shows two 1D images I_1, I_2 , with height denoting image intensity (dashed lines are the geodesics between marked points). They look quite different, but they are related by a deformation (composed of stretching and compressing in different places). Consider the images as 1D surfaces embedded in a 2D space, where intensity is scaled by α , the *aspect weight* and x is scaled by $1 - \alpha$. Using different α s produces the images in Figure 2 (b,c,d). We see that when α increases, the graphs of the two embedded images look more and more similar. It is natural to expect that they become exactly the same as $\alpha \rightarrow 1$. One way to explain this is that α controls the weight on the intensity compared to the weight on the image coordinate. A larger α means that we place more importance on the intensity, which is deformation invariant. So $\alpha = 1$ leads to a deformation invariant view of the image.

How does α work? Let p_1, q_1 be two points on I_1 , with their deformed counterparts p_2, q_2 on I_2 . Consider the geodesic distance g_1 between p_1 and q_1 on αI_1 and g_2 between p_2 and q_2 on αI_2 . Figure 2 (e) shows how g_1 and g_2 vary for different α , from which we see that g_1 and g_2 tend to become similar when α increases. This implies that the geodesic distance on the embedded surface tends to be deformation invariant when $\alpha \rightarrow 1$.

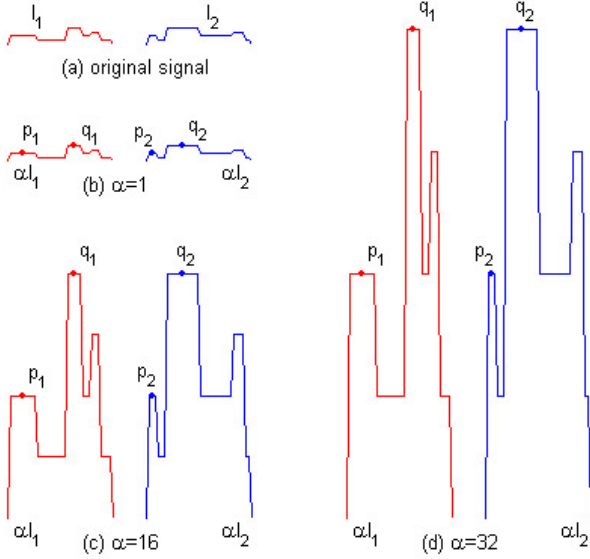


Figure 2. Deformation invariance for one dimensional images. Details in Section 3.1

This provide a solution for the problem of finding deformation invariant neighborhoods. For a given interest point, we can use the geodesic distance to find and sample within a neighborhood, and then build a descriptor based on the intensities of the sampled points. If the procedure is done with a very large α (in the limit approaching 1), then the descriptor is deformation invariant, which is exactly what we want.

3.2. Curve Lengths on Embedded Surfaces

In this section we show that as α increases, the lengths of curves change less and less when an object deforms. Let $I_1(x, y)$ be an image defined as $I : R^2 \rightarrow [0, 1]$. Let $I_2(u, v)$ be a deformation of I_1 . Because deformation is a homeomorphism, and so it is invertible, we can write $u = u(x, y), v = v(x, y), x = x(u, v), y = y(u, v)$, and $I_2(u, v) = I_1(x(u, v), y(u, v))$.

Denote the embedding of an image $I(x, y)$ with aspect weight α as $\sigma(I; \alpha) = (x' = (1 - \alpha)x, y' = (1 - \alpha)y, z' = \alpha I(x, y))$. Denote σ_1, σ_2 as the embeddings of I_1, I_2 respectively

$$\begin{aligned} \sigma_1 &= (x' = (1 - \alpha)x, y' = (1 - \alpha)y, z' = \alpha I_1(x, y)) \\ \sigma_2 &= (u' = (1 - \alpha)u, v' = (1 - \alpha)v, w' = \alpha I_2(u, v)) \end{aligned}$$

Let γ_1 be a regular curve on $\sigma_1, t \in [a, b]$, and γ_2 the deformed version of this curve on σ_2

$$\gamma_1(t) = (x'(t), y'(t), z'(t))$$

$$\begin{aligned} &= ((1 - \alpha)x(t), (1 - \alpha)y(t), \alpha I(x(t), y(t))) \\ \gamma_2(t) &= (u'(t), v'(t), w'(t)) \\ &= ((1 - \alpha)u(t), (1 - \alpha)v(t), \alpha I(u(t), v(t))) \end{aligned}$$

Where

$$w'(t) = \alpha I_2(u(t), v(t)) = \alpha I(t) = \alpha I_1(x(t), y(t)) = z(t)$$

because the intensity is invariant to deformation,

Now we can study the length of γ_1, γ_2 , denoted as l_1, l_2 respectively. We have

$$\begin{aligned} l_1 &= \int_a^b \sqrt{x_t'^2 + y_t'^2 + z_t'^2} dt \\ &= \int_a^b \sqrt{(1 - \alpha)^2 x_t'^2 + (1 - \alpha)^2 y_t'^2 + \alpha^2 I_t'^2} dt \quad (1) \\ l_2 &= \int_a^b \sqrt{u_t'^2 + v_t'^2 + w_t'^2} dt \\ &= \int_a^b \sqrt{(1 - \alpha)^2 u_t'^2 + (1 - \alpha)^2 v_t'^2 + \alpha^2 I_t'^2} dt \quad (2) \end{aligned}$$

Where the subscripts denote partial derivatives, e.g., $x_t \doteq dx/dt, u_t \doteq \partial u/\partial t$, etc....

From (1) and (2) it is clear that for a large α , the curve length is dominated by the intensity changes along the curve. In the limit when $\alpha \rightarrow 1$, l_1, l_2 converge to the same value. Also, the length of curves with constant intensities tend to be trivial compared to lengths of curves with non-constant intensities.

In the rest of the paper, when talking about deformation invariance, we implicitly assume that $\alpha \rightarrow 1$.

3.3. Geodesic Distance and Level Sets

It follows from the last subsection that the *geodesic distance*, which is the distance of the shortest path between two points on the embedded surfaces, is deformation invariant. Given an interest point $p_0 = (x_0, y_0)$, the geodesic distances from it to all other points on the embedded surface $\sigma(I; \alpha)$ can be computed using the level set framework [21]. Points with identical geodesic distances from p_0 are treated as level curves. For images defined on discrete grids, the fast marching algorithm ([21]) provides an efficient method of computing these curves.

Figure 3 shows two example results of the geodesic distances computed for real images. It shows that when α is small (in (c),(d)), the geodesic distances are almost like Euclidean distances in the image plane. With a large α (in (e),(f)), the geodesic distance captures the geometry of image intensities and automatically adapts to deformation.

One interesting issue is that since real images are defined on discrete grids, the fast marching method we use implicitly assumes that the surface is piecewise constant (constant




Figure 3. Geodesic distances computed via fast marching. The marked point in (a) corresponds to the marked point in (b) after deformation. (c),(e) shows the geodesic distances of all pixels in (a) from the marked point, with different α 's. Darker intensities mean large distances. (d),(f) shows the same thing for the marked point in (b). Note that image structures of (a) and (c) are captured in the distance map in (e) and (f).

within the region of each pixel). The image can also be interpolated as a smooth surface, in which case the arguments above still hold.

3.4. Deformation Invariant Sampling

Geodesic level curves provide us a way to find deformation invariant regions surrounding interest points. These regions can be used as support regions for extracting deformation invariant descriptors. To derive invariant descriptors, we must also sample these regions using geodesic distances, to find deformation invariant sample points. In the following Δ is used to denote the sampling interval.

Geodesic sampling for 2D images is done in two steps. First, the level curves are extracted at intervals of Δ . Sec-




Figure 4. 2D geodesic sampling. $\alpha = 0.98$. A large Δ is used for better illustration. The interest point is the marked star in the center of the sampling. The sampled points are marked on the geodesic level curves.

ond, points are sampled from each level curve at intervals of Δ . Figure 4 gives examples of 2D geodesic sampling. Note that the sampling along uniform intensity regions is sparser than along regions with large intensity variation. Intuitively, this implies that deformations (such as stretching) will not change the number of sample points, although it may change their locations. We sample densely, so that changes in the location of sample points do not have much effect on the resulting histogram.

3.5. The Geodesic-Intensity Histogram

Now we introduce the *geodesic-intensity histogram* (GIH), which is a deformation invariant descriptor extracted from geodesic sampling. It captures the joint distribution of the geodesic distance and the intensity of the sample points. Since both the geodesic distance and the intensity are deformation invariant, so is the GIH. It is based on spin images, which produce a related descriptor using Euclidean distance ([7, 10]).

Given an interest point p , together with a sample point set P_p obtained via geodesic sampling, the GIH H_p at p is a normalized two dimensional histogram obtained through the following steps:

1. Divide the 2D intensity-geodesic distance space into $K \times M$ bins. Here K is the number of intensity intervals, and M the number of geodesic distance intervals. The geodesic intervals can be segmented either linearly or at log scale.
2. Insert all points in P_p into H_p : $\forall 1 \leq k \leq K, \forall 1 \leq m \leq M$, $H_p(k, m) = \#\{q \in P_p : (I(q), g(q)) \in B(k, m)\}$
Here $I(q)$ is the intensity at q , $g(q)$ is the geodesic distance at q (from p), and $B(k, m)$ is the bin corresponding to the k th intensity interval and the m th geodesic interval.
3. Normalize each column of H_p (representing the same geodesic distance). Then normalize the whole H_p .

Figure 5 displays examples of the geodesic-intensity histograms of two points with deformation. The two histograms are quite similar, although the deformation between the two images is quite large.




Figure 5. Geodesic-intensity histograms, $\alpha = 0.98$, $K = 10$, $M = 5$. (a), (b) for points in Figure 3 (a),(b) respectively.

Given two geodesic-intensity histogram H_p, H_q , the similarity between them is measured using the χ^2 distance:

$$\chi^2(p, q) \equiv \frac{1}{2} \sum_{k=1}^K \sum_{m=1}^M \frac{[H_p(k, m) - H_q(k, m)]^2}{H_p(k, m) + H_q(k, m)} \quad (3)$$

3.6. Practical Issues

Dealing with illumination change. We use an affine model for lighting change ([15]), i.e., $aI(x, y) + b$ for the illumination change of the pixel at (x, y) . There are two steps to make GIH insensitive to lighting change. 1) GIH is made invariant to lighting in the same way as [15]. That is, when building the histogram, the intensity is normalized by subtracting the mean and then divided by the standard deviation, where the mean and deviation is estimated on the sampled point set H_p . 2) Compensate for the effect of lighting change on the geodesic sampling. For large α , the intensity change dominates the geodesic distance (1,2). So the change of geodesic distance is approximately linear with rate a under the lighting model, which is equivalent to

changing α to $a\alpha$. So when we compare two interest points, we compare several GIH's that use different α 's, and pick the match with minimal χ^2 distance (3).

Interest points. GIH does not require a special interest point detector since it automatically locates the support region. However, there are problems with using some feature points. First, deformation invariance makes points within a constant region indistinguishable. Second, for real images the intensity on edges or corners may vary due to sampling. We have found that *extreme points*, where images have local intensity extremum, are less affected by the above factors (they are locally unique in the continuous cases). The extreme point can be viewed as a deformation invariant version of the DoG point proposed by Lowe [14], which is scale invariant. In Sec. 4.3 we tested the performance of GIH using several different interest point operators.

Choosing α . In the following experiments we will use a very large α (0.98) because we want to deal with large deformations. However, in domains involving only small deformations, a relatively smaller α might be a better choice. Smaller α 's can lead to descriptors that are somewhat insensitive to deformations, but that provide more information since they do not treat images related by large deformations as identical. It is obvious that GIH with $\alpha = 0$ becomes equivalent to spin images [10, 7].

4. Experiments

In this section we will describe our experiments using the GIH for interest point matching. Experiments were conducted on two groups of image pairs. One contains synthetic deformation as well as illumination change, the other contains real non-affine deformations. We have two experiments. The first one compares the GIH's matching ability to several other approaches. The second experiment studies the performance of GIH using several different kinds of interest points including the proposed extreme points.

4.1. Experimental Setup

Data set We evaluate the proposed method using two groups of images. The first group contains eight image pairs with synthetic deformation and illumination change (see Figure 9, the original images are from the Berkeley segmentation dataset¹). The deformation is created by mapping the original images to non-flat surfaces and viewing them from different viewpoints. The lighting change is generated through an affine model (intensities limited to $[0..1]$). The second group contains three pairs of images with real deformations (see Figure 10).

Interest point We use Harris-affine points [16] for the matching experiments. The interest point is detected using

¹<http://www.cs.berkeley.edu/projects/vision/grouping/segbench/>

the online code provided by Mikolajczyk [17]. One reason for this choice is its affine invariance. This makes the other descriptors invariant to affine transformation, although it is not necessary for our descriptor. The other reason is that [17] also provides executable codes for several state-of-art descriptors that we can use for comparison. For each image, we pick the 200 points extracted by the detector with the largest cornerness.

Evaluation criterion For each pair of images together with their interest points, we first obtained the ground truth matching (automatically for synthetic images, manually for real images). Then, for efficiency we removed those points in image 1 with no correct matches. After that, every interest point in image 1 is compared with all interest points in image 2 using the descriptors to be compared. An interest point p_1 in image 1 is treated as a correct match of another point p_2 in image 2 if the deformation of p_1 is within a three pixel distance of p_2 . The detection rate among the top N matches is used to study the performance. The detection rate is defined in a way similar to [15]:

$$r = \frac{\#correct\ matches}{\#possible\ matches} = \frac{\#correct\ matches}{\#points\ in\ image\ 1} \quad (4)$$

4.2. Matching Experiment

In this experiment we will study the performance of GIH in comparison with several other methods. The experiments are conducted on both the synthetic and real deformation data sets. All of them use the Harris-Affine interest point.

Mikolajczyk and Schmid [15, 17] provided convenient online code for several state-of-the-art local descriptors. The descriptors are normalized to enable direct comparison using sum of square differences (SSD). Benefitting from their code, we compare the geodesic-intensity histogram with steerable filters [4], SIFT [14], moments invariants [24], complex filters [19] and spin images [10].

The main difference between the evaluation here and that in [15] lies in that [15] focused more on the evaluation of region-like descriptors. For example, some of their experiments use interest regions instead of interest points. Furthermore, their matching criterion between two features is also related to their support regions. Also note that the Harris-Affine point is chosen because it provides affine invariant support regions to the descriptors we will compare to, although it is not necessary for GIH (see Sec. 4.3).

We tested two versions of the geodesic-intensity histogram. Version one uses $\alpha = 0.98, K = 13, M = 8$. This tests the ability of the GIH. The other version is a degenerate version where $\alpha = 0, K = 10, M = 5$. This demonstrates that GIH becomes like spin images for $\alpha = 0$.

A Receiver Operating Characteristics (ROC) based criterion is used which is similar to the one in [15]. Instead

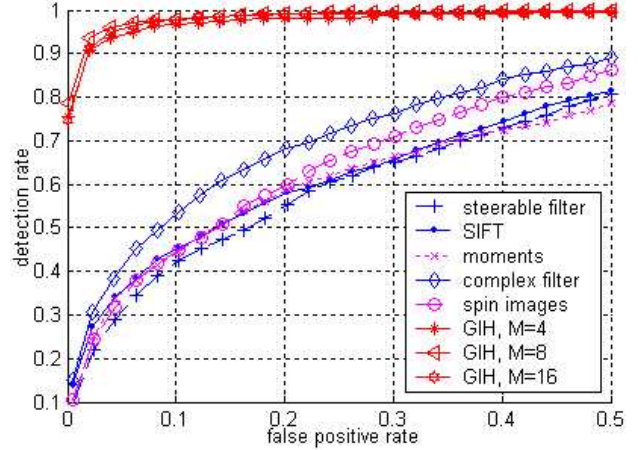


Figure 6. Experiment results on the synthetic deformation data set (see Figure 9).




Figure 7. Experiment results on the real deformation data set (see Figure 10).

of using the false positive rate, we study the detection rates among the top N matches, as N varies.

Figure 6 displays the ROC curves for the experiment on the synthetic deformation data set and Figure 7 for the real deformation data set. From the ROC curves we can see that GIH performs better than other methods in both data sets regardless of illumination changes. Note that for $\alpha = 0$, the performance drops a lot, with the performance similar to spin images (with no affine invariant support region).

4.3. Interest Points

This experiment is to test the performance of GIH using several kinds of interest points. In addition to extreme points and DoG [14] points, we also tested on Harris corners

[6] and Harris-Affine points [16]. We use the code provided at [17] except for the extreme points. The experiment is conducted on the synthetic deformation data set. For each image, 200 points are picked with the largest detector responses (cornerness, for example). For extreme points, the response is computed through Laplace-of-Gaussian filtering. The same parameters for GIH are used for all kinds of interest points, $\alpha = 0.98$, $K = 13$, $M = 8$.

Since different interest point detectors may generate different number of correct correspondences, the ROC curves is plotted as the detection rate versus the false positive rate instead of N as in the previous experiment. The false positive rate is defined as (similar to [15])

$$r^{false} = \frac{\#false\ matches}{(\#points\ in\ image\ 1)(\#points\ in\ image\ 2)}$$

Figure 8 shows the ROC curves. From the figure we can see that GIH works better than the others for small false positive rate less than 0.03 (this roughly corresponds to the top 6 matches). For large false positive rates, DoG performs the best. The Harris corner works the worst with GIH, which is consistent with our previous discussion.

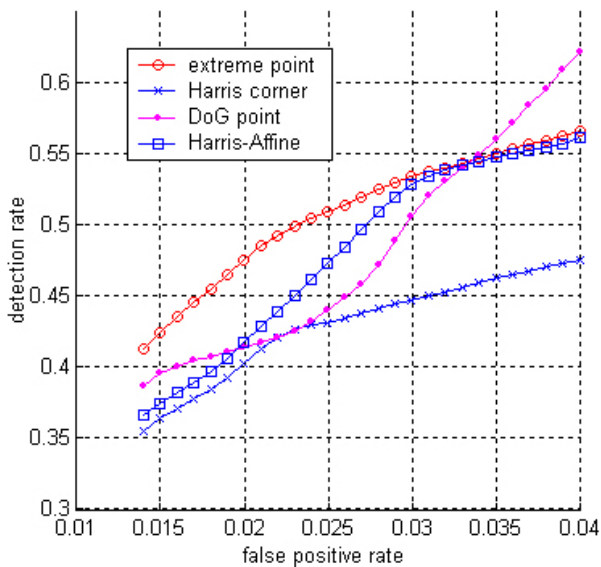


Figure 8. GIH using different kinds of interest points. The false positive rate 0.04 roughly corresponds to $N = 8$ as in Figure 7 and 6.

5. Conclusions

In this paper we proposed a novel deformation invariant feature, the geodesic-intensity histogram, for intensity images. Images are treated as 2D surfaces embedded in 3D

spaces. We then showed that the geodesic distance along the surface is invariant to deformation when the embedding aspect weight $\alpha \rightarrow 1$. The geodesic-intensity histogram is a 2D histogram measuring the geodesic distances and the intensities surrounding an interest point. With the geodesic sampled neighborhood points and an $\alpha \rightarrow 1$, the proposed histogram becomes deformation invariant. After that, we discussed practical issues including how to deal with illumination change and the option of choosing α to balance deformation invariance and discriminativity. The proposed descriptor is tested on data sets with both synthetic and real deformations. In all the experiments the new descriptor performs excellently in comparison with several other methods.

Acknowledgements: We would like to thank Krystian Mikolajczyk and Cordelia Schmid for the feature extraction code used in the comparison. The original images in the synthetic deformation experiment are from the Berkeley segmentation dataset. This work is supported by NSF (ITR- 03258670325867). Haibin Ling is also supported by the Horvitz Assistantship.

References

- [1] A. Baumberg. “Reliable feature matching across widely separated views”, *CVPR*, I:774-781, 2000.
- [2] S. Belongie, J. Malik and J. Puzicha. “Shape Matching and Object Recognition Using Shape Context,” *IEEE Trans. on PAMI*, 24(24):509-522, 2002.
- [3] A. Elad(Elbaz) and R. Kimmel. “On Bending Invariant Signatures for Surfaces”, *IEEE Trans. on PAMI*, 25(10):1285-1295, 2003.
- [4] W. Freeman and E. Adelson. “The design and use of steerable filters”, *IEEE Trans. on PAMI*, 13(9): 891-906, 1991.
- [5] R. Fergus, P. Perona and A. Zisserman. “Object Class Recognition by Unsupervised Scale-Invariant Learning”, *CVPR*, II:264-271, 2003.
- [6] C. Harris and M. Stephens, “A combined corner and edge detector”, *Alvey Vision Conference*, 147-151, 1988.
- [7] A. Johnson, M. Hebert, “Using spin images for efficient object recognition in cluttered 3D scenes”. *IEEE Trans. on PAMI*, 21(5):433-449, 1999.
- [8] T. Kadir, A. Zisserman, and M. Brady, “An affine invariant salient region detector”, *ECCV*, I:228-241, 2004.
- [9] Y. Ke and R. Sukthankar. “PCA-SIFT: a more distinctive representation for local image descriptors”, *CVPR*, II:506-513, 2004.

- [10] S. Lazebnik, C. Schmid, and J. Ponce, "A sparse texture representation using affine-invariant regions," *CVPR*, 2:319-324, 2003.
- [11] T. Lindeberg, "Feature Detection with Automatic Scale Selection," *IJCV*, 30(2), pp. 79-116, 1998.
- [12] H. Ling and D. W. Jacobs, "Using the Inner-Distance for Classification of Articulated Shapes", *CVPR*, II:719-726, 2005.
- [13] D. Lowe, "Object recognition from local scale-invariant features," *ICCV*, 1150-1157, 1999.
- [14] D. Lowe, "Distinctive Image Features from Scale-Invariant Keypoints," *IJCV*, 60(2), pp. 91-110, 2004.
- [15] K. Mikolajczyk and C. Schmid, "A Performance Evaluation of Local Descriptors," *CVPR*, pp. 257-264, 2003.
- [16] K. Mikolajczyk and C. Schmid, "Scale & Affine Invariant Interest Point Detectors," *IJCV*, 60(1), pp. 63-86, 2004.
- [17] K. Mikolajczyk, "Affine Covariant Features", <http://www.robots.ox.ac.uk/vgg/research/affine/>, Visual Geometry Group, University of Oxford, 2004.
- [18] M. Petrou and A. Kadyrov, "Affine Invariant Features from the Trace Transform," *IEEE Trans. on PAMI*, 26(1):30-44, 2004.
- [19] F. Schaffalitzky and A. Zisserman. "Multi-view matching for unordered image sets", *ECCV*, pp 414-431, 2002.
- [20] C. Schmid and R. Mohr, "Local Grayvalue Invariants for Image Retrieval", *IEEE Trans. on PAMI*, 19(5):530-535, 1997.
- [21] J. A. Sethian. "A Fast Marching Level Set Method for Monotonically Advancing Fronts", *Proc. Nat. Acad. Sci.*, 93(4), 1996.
- [22] N. Sochen, R. Kimmel and R. Malladi, "A General Framework for Low Level Vision", *IEEE Trans. on Image Processing*, 7(3):310-318, 1998.
- [23] T. Tuytelaars and L. Van Gool. "Matching widely separated views based on affine invariant regions", *IJCV*, 59(1):61-85, 2004.
- [24] L. Van Gool, T. Moons, and D. Ungureanu. "Affine / photometric invariants for planar intensity patterns", *ECCV*, pp 642-651, 1996.




Figure 9. Four of the eight pairs of images with synthetic deformation and illumination change.




Figure 10. Images with real non-affine deformation.