Learning Basic Patterns from Repetitive Texture Surfaces Under Non-Rigid Deformations

Roman Filipovych and Eraldo Ribeiro

Computer Vision and Bio-Inspired Computing Laboratory
Department of Computer Sciences
Florida Institute of Technology
Melbourne, FL 32901, USA
{rfilipov,eribeiro}@fit.edu
http://www.cs.fit.edu/~eribeiro

Abstract. In this paper, we approach the problem of determining the basic components from repetitive textured surfaces undergoing free-form deformations. Traditional methods for texture modeling are usually based on measurements performed on fronto-parallel planar surfaces. Recently, affine invariant descriptors have been proposed as an effective way to extract local information from non-planar texture surfaces. However, affine transformations are unable to model local image distortions caused by changes in surface curvature. Here, we propose a method for selecting the most representative candidates for the basic texture elements of a texture field while preserving the descriptors' affine invariance requirement. Our contribution in this paper is twofold. First, we investigate the distribution of extracted affine invariant descriptors on a nonlinear manifold embedding. Secondly, we describe a learning procedure that allows us to group repetitive texture elements while removing candidates presenting high levels of curvature-induced distortion. We demonstrate the effectiveness of our method on a set of images obtained from man-made texture surfaces undergoing a range of non-rigid deformations.

Key words: texture learning, non-rigid motion, texture classification, dynamic texture

1 Introduction

In this paper, we investigate the problem of learning representative local texture components from repetitive textures. In particular, we focus on the specific case in which the learning stage is performed by observing a textured surface undergoing unknown free-form deformations. Figure 1 shows samples of such textured surfaces presenting a range of curvature deformations.

Perceiving and modeling the appearance of repetitive textures are important visual tasks with a number of applications including surface tracking [8, 14], texture classification [3, 12], and texture synthesis [16, 9]. However, obtaining accurate descriptions from non fronto-parallel texture fields is not a trivial

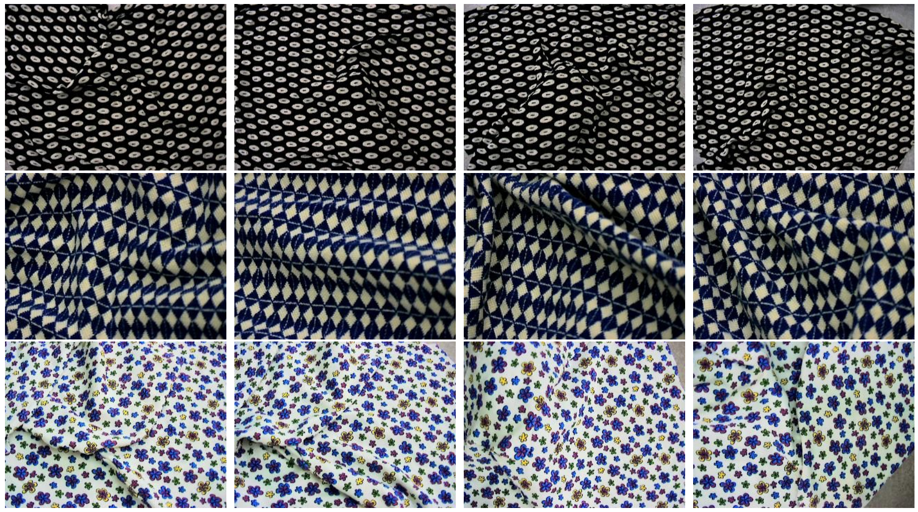

Fig. 1. Sample frames of sequences of surface textures under non-rigid deformation. Local planarity of texture primitive is often disturbed by abrupt changes in local surface curvature where the surface tends to curve into folds.

problem as the underlying texture appearance varies significantly with both the perspective geometry and the orientation of the observed surfaces [21, 9, 16, 12]. While this appearance variation represents a rich source of information for shape-from-texture methods [25, 19], it is also a main source of problems for standard texture learning methods. Indeed, changes in local curvature produce nonlinear warping of some image regions. Consequently, texture descriptors evaluated on these warped image regions are likely to produce unreliable measurements.

Our primary goal in this paper is to study the effects of surface curvature variation on local texture descriptors when the observed surface undergoes non-rigid deformations. In particular, we show how local texture appearance models can be estimated from deforming textured surfaces with high level of curvature distortion. The contribution of the study presented in this paper is twofold. First, we investigate the distribution of extracted affine invariant descriptors on a nonlinear manifold embedding. Here, we assume that the population of affine invariant descriptors lie on a lower dimensional manifold describing mainly variations in surface orientation and curvature. This lower dimensional manifold seems to describe the departure from local planarity of local affine invariant descriptors. Secondly, we describe a learning procedure that allows us to group repetitive texture elements while selecting the best set of candidates to represent the actual undistorted basic texture components.

The remainder of this paper is organized as follows. Section 2 provides a survey of the related literature. In Section 3, we discuss the effects of curvature distortion on affine invariant texture measurements. Section 4 describes the details of our texture primitive selection method. The preliminary results of our study are shown in Section 5. Finally, in Section 6, we present our conclusions and directions for future investigations.

2 Related literature

Modeling texture appearance is usually the initial step in the solution of many texture-related problems including recognition [3, 12], tracking [8, 14], and synthesis [16, 9]. Nevertheless, finding general representations for texture is a challenging problem. In fact, despite extensive research efforts by the computer vision community, there is no currently widely accepted method to model the complexity found in all available textures. State-of-the-art texture classification algorithms have successfully approached the texture representation problem by means of statistical descriptors. These descriptors can be obtained from measurements based on the response of convolution filters [13, 24, 26], image regions and pixel distributions [12], and frequency-domain measurements [7, 2].

Most texture modeling methods are based on measurements obtained from planar fronto-parallel texture fields [13, 24, 26, 7]. For example, Leung and Malik [13] introduced a filter bank-based descriptive model for textures that is capable of encoding the local appearance of both natural and synthetic textures. Since then, many extensions of this work have been proposed [24, 5]. Indeed, these methods achieve high classification rates due to their ability to learn representative statistical models of each texture. However, it is unclear how they would perform on non-rigid deforming surfaces.

There has been some recent research effort aimed at addressing the texture learning problem from non-rigid, non fronto-parallel texture images [21, 9, 16, 12]. For example, Chetverikov and Foldvari [4] use a frequency-domain affine-invariant representation for local texture regions. Bhalerao and Wilson [2] also use a frequency-based affine-invariant descriptor for texture segmentation. Mikolajczyk *et al.* [17] describe a local affine frame measurement for wide-baseline stereo followed by a comparison of affine region detectors. Zhu *et al.* [26] present a multilevel generative image model for learning texture primitives based on Gabor-like filter measurements. Hayes *et al.* [9] address the problem of learning non-rigid texture deformations on regular and quasi-regular lattices. They propose a feature matching algorithm to discover lattices of near-regular textures in images. Their final goal is to synthesize fields of texture from a set of examples.

The focus of this paper is on non-rigid deforming texture surfaces. Recent work by Lazebnik *et al.* [12] addresses the problem of learning texture models from images of non fronto-parallel texture surfaces. They propose a texture classification method based on learned texture components using affine-invariant descriptors. This approach is very effective under the assumptions of orthographic viewing and low-curvature surfaces. However, for surfaces with high levels of curvature deformation, the surface folds and bends will reduce the ability of affine-invariant descriptors to capture correct local texture representations. As a result, the learned appearance of basic texture components using this approach is likely to be less representative of the actual surface texture.

Depending on the curvature of the observed surfaces, the deformation of texture elements can present a significant degree of nonlinearity. Nonlinear manifold learning techniques such as Isomap [23], Local Linear Embedding (LLE) [20], and Laplacian Eigenmaps [1] are suitable candidates for the analysis of such deformations. For example, Souvenir and Pless [22] characterize deformations in magnetic resonance imaging. Nonlinear manifold learning is also a useful technique for the synthesis of dynamic textures. Liu *et al.* [15] approach the dynamic texture synthesis problem using nonlinear manifold learning and manifold traversing.

3 Curvature-induced distortion of texture elements

In this section, we investigate the effect of surface curvature on affine-invariant descriptors. The aim of this section is twofold. First, we show how changes in local curvature can reduce the quality of affine-invariant descriptors. Secondly, we analyze the distribution of local affine-invariant descriptors using a nonlinear manifold learning technique.

We will represent the local texture appearance using the affine-invariant descriptor proposed in [12]. This descriptor is essentially a pixel gray-level intensity histogram calculated on a scale-invariant polar representation of an image subregion. It represents the radial frequency of normalized pixel intensities. The representation is normalized to minimize the effects of illumination. The polar representation of the pixel intensity maps rotations onto translations. The histogram representation that follows is translation invariant. This representation allows for full affine invariance when calculated at locations centered at image patches. Figure 2(a) illustrates the affine-invariant region extraction process, and is described in detail next.

Single texture patch distortion. We begin by analyzing the influence of surface curvature on a single affine-invariant descriptor calculated on a small texture subregion. The selected region contains a single planar texture element from the textured surface. Here, we artificially project the extracted region on a cylindrical surface observed under perspective viewing. We proceed by bending the surface to increase the local curvature while using the standard Euclidean distance to measure the error between the affine-invariant descriptor of a planar patch and its curved versions. The plot in Figure 2(b) shows the Euclidean distance measuring the departure from the original planar patch for increasing levels of curvature deformation. The plot shows that the information in the affine-invariant descriptor remains almost constant when the local curvature is low. However, the difference from the planar patch increases significantly for medium and high surface curvatures. This behavior is surely expected as, for curved surfaces, texture affine-invariance is only valid for small planar regions.

Distortion distribution of local affine-invariant descriptors. The previous analysis suggests that local texture affine invariance is not preserved if the

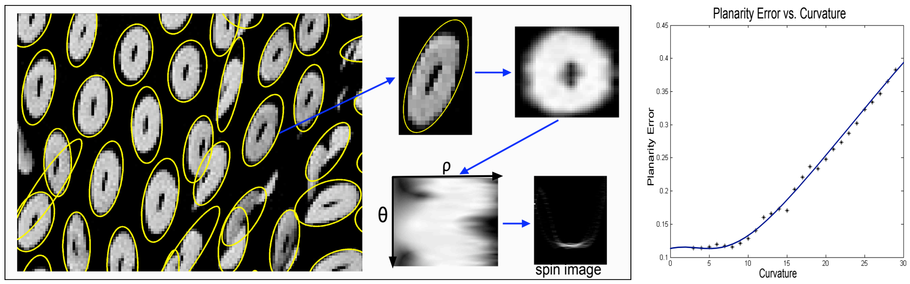

(a) (b)

Fig. 2. (a) affine-invariant region extraction. Elliptical regions centered at affine-invariant interest points are extracted from the image. The elliptical regions are normalized to a circular shape. A translation invariant radial histogram of pixel intensities is build to represent the extracted regions. (b) Deviation from the original texton with increase in curvature observed in the affine-invariant space.

texture region is deformed by the 3-D surface's local curvature. We now would like to turn our attention to the distribution of distortions of affine-invariant descriptors across a deforming surface. In this analysis, we commence by detecting a large set of interest points in a sequence of images of a non-rigid deforming surface. A sample image of the analyzed surface is shown in Figure 2(a). The set of interest point locations is obtained using the affine-invariant interest point detector proposed by Kadir and Brady [11]. This detector provides information about the affine scale of the neighborhood of each detected interest point. Once these image locations are at hand, we use the scale information provided by the detector to extract elliptical regions from the images. These regions are subsequently normalized to a standard rectangular size and represented by the rotation-invariant descriptor proposed by Lazebnik *et al.* [12] to remove the inherent rotational ambiguity. At this stage, we would like to point out the following: **(a)** we expect the distribution of local texture patches to form groups representing various basic texture components and their parts; **(b)** repetitive texture patterns can be compactly represented by statistical descriptors based on affine-invariant measurements. However, localization errors and incorrect scale axes are likely generate redundant groups in the distribution (i.e., non-centered similar patches will form separate groups); **(c)** the distribution of distorted texture elements can be assumed to lie on a low dimensional manifold modeling both patch location noise and curvature-induced distortions.

We now use Isomap [23] to obtain an embedded representation of the lower dimensional nature of the distortion distribution presented by the subregion descriptors. Isomap manifold embedding preserves "geodesic" distances between data points. Figure 3 illustrates the two-dimensional manifold learned using Isomap for the local affine-invariant patch distribution. The figure also shows the embedded image subregions back-projected onto the image domain. We ex-

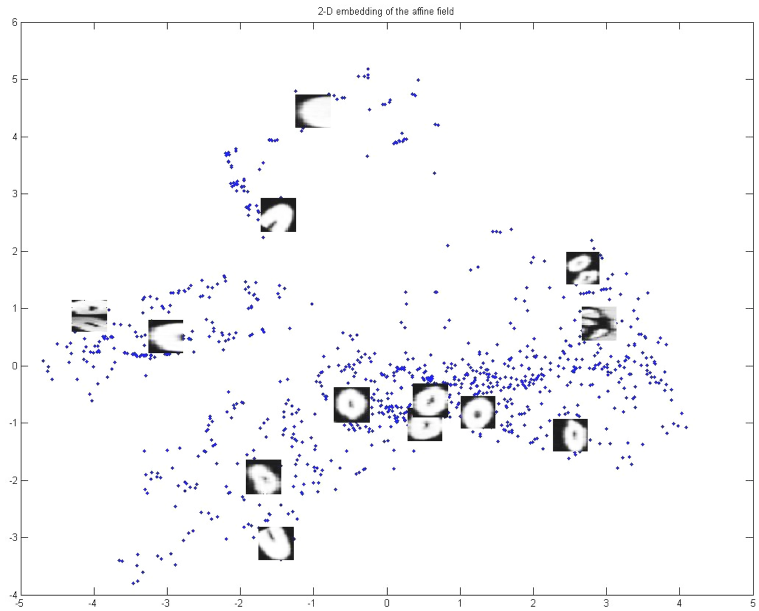

Fig. 3. 2-D Isomap embedding of the affine-invariant field. Undistorted texture components form the large cluster close to the center of the plot. Nonlinearly deformed elements also tend to form clusters. Border points usually represent noise.

tracted a large number of affine-invariant descriptors from the sequence of images shown in the first row of Figure 1. The figure shows a highly dense cluster near the center of the plot containing mostly locally planar patches. On the other hand, nonlinearly deformed patches tend to group themselves into clusters with respect to their deformation similarity. Finally, occluded or distorted elements form relatively sparse groups with large within-class variation.

4 Our Texture Component Learning Method

Our goal in this paper is to build an algorithm for learning compact representations of textures from images of surfaces undergoing non-rigid deformations. While this represents a significantly challenging capture setup, it also allows us to apply our method to more realistic imaging situations such as capturing the basic patterns from a piece of waving textured cloth, a moving animal's natural skin pattern, and clothes worn by a moving person. In this section, we use some of the insights obtained from the analysis presented in the previous section to derive a algorithm for selecting a set of representative components of a texture pattern. The main steps of our texture learning algorithm are given as follows.

Extraction of affine-invariant regions. The first step of our algorithm consists of extracting a large number of image subregions from a set of video frames of the observed surface. This step is subdivided into two main parts. First, we detect a large set of affine-invariant interest points on each image or video-frame. We use the Kadir and Brady's salient feature detector [11] as it provides information about the affine scale of image features. The salient feature detector outputs elliptical regions centered at each feature of interest. Subregions extracted using this detector can be normalized to a common scale invariant shape (e.g., a circle). The remaining rotation ambiguity can then be removed by representing the normalized subregions using the spin-image affine-invariant descriptor proposed in [12]. Accordingly, let $\mathbf{S} = (\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_N)$ be the set of affine-invariant descriptors obtained by this step, where N is the total number of subregions.

Manifold learning representation of affine descriptors. Our aim is to obtain a compact representation of the most significant repetitive patterns on the image. However, the nonlinear nature of the distortion present in the dataset does not allow for correct distance measurements in the original feature space. Additionally, spin-image descriptors carry a significant level of information redundancy. To accomplish a better description of the variation in the dataset, we assume that both the basic texture elements as well as their nonlinear deformations lie on a low-dimensional nonlinear manifold in which the two intrinsic dimensions of variability describe local surface orientation and curvature distortion. Based on this assumption, we perform Isomap [23] nonlinear manifold learning on the original distribution $\mathbf{S}$. Isomap allows for a reduction of the dimensionality of the data while preserving the manifold's intrinsic geometry, and is stable even for sparse sets of data points. The resulting reduced dimension dataset of extracted subregions is then given by $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N)$.

Learning representative texture components. The goal of this step is to determine the most representative classes of texture elements in $\mathbf{X}$. We begin by modeling the distribution of affine-invariant descriptors as a mixture of K Gaussian densities given by $p(\mathbf{x}|\Theta) = \sum_{i=1}^{K} \alpha_i p_i(\mathbf{x}|\theta_i)$, where $\mathbf{x}$ is an affine-invariant descriptor in the Isomap manifold space, α_i represent the mixing weights such that $\sum_{i=1}^{K} \alpha_i = 1$, Θ represents the collection of parameters $(\alpha_1, \ldots, \alpha_K, \theta_1, \ldots, \theta_K)$, and p_i is a multivariate Gaussian density function parameterized by $\theta_i = (\mu_i, \sum_i)$, such that:

$$p(\mathbf{x}|j) = \frac{1}{(2\pi)^{\frac{d}{2}} |\Sigma_j|^{1/2}} \exp\{-\frac{1}{2} (\mathbf{x} - \mu_j)^{\mathsf{T}} \Sigma_j^{-1} (\mathbf{x} - \mu_j)\} \tag{1}$$

where μ_j and Σ_j are the mean vector and covariance matrix of the texture element j, respectively. Each mixture component represents a set of texture descriptors of similar appearance on the image. The mixture of Gaussians model parameters can be estimated by the Expectation-Maximization (EM) algorithm [6].

A straightforward consequence of the mixture of Gaussians modeling is that texture descriptors closer to the mean vector in each class will present less curvature-induced distortion when compared to descriptors that are further away from the class mean. In order to obtain sharper representations of the learned texture components, we select a single descriptor from each cluster to represent a basic texture component in the image. In other words, a set of texture components is selected as:

$$\boldsymbol{\tau}_j = \arg\max_{\mathbf{x}_i} p(\mathbf{x}_i|j) \qquad j = 1, \ldots, K \tag{2}$$

A set of basic components is obtained by this process and is used to create a dictionary representation $\mathbf{d} = \{\boldsymbol{\tau}_1, \ldots, \boldsymbol{\tau}_K\}$. However, the nonlinear nature of the surface distortion will compromise the representativeness of some of the learned mixture components. As a result, the learned clusters might not represent actual texture components but a geometrically warped version of them. Next, we propose a way to remove these non-representative elements from our dataset of learned texture primitives.

Ranking of learned texture components in the dictionary. Our main goal in this step is to distinguish between distributions of affine transformed basic texture elements and their nonlinearly deformed counterparts. The nature of the nonlinear deformations is mostly anisotropic (i.e., directional appearance distortion). Consequently, we expect to have a relatively small number of data points belonging to clusters of non-affine distorted elements. Thus, distributions with low prior probability will most likely correspond to regions that were distorted by nonlinear transformations. Elements falling within such distributions can be safely discarded and therefore removed from the dictionary.

The remaining distributions may represent two cases. The first case corresponds to classes of affine transformed elements that are representative of the texture. The second case represents classes consisting of nonlinearly distorted regions. Our experiments have shown that the distribution of nonlinearly distorted elements tend to have high within-class variation. Based on this observation, we rank the remaining dictionary elements based on the decreasing order of the within-class variation of their classes (i.e., the determinant of the covariance matrix for the class, $|\Sigma_j|$). Finally, the top-ranked elements are selected as the ones that represent classes of locally planar regions:

$$\mathbf{d} = \{\boldsymbol{\tau}_j\} \quad \text{such that} \quad |\Sigma_j| \geq |\Sigma_{j+1}| \tag{3}$$

5 Experimental results

Our experiments were divided into two main parts. First, we evaluated our algorithm on video sequences of a number of deforming texture surfaces. The surfaces used in our experiments consisted of patterned fabrics bought from a local shop. To produce the deformations, we have waved and deformed the fabrics manually while recording the video sequences. Three of these patterns are shown in Figure 4. Secondly, we show qualitative results of extracted basic texture elements using a standard K-Means learning approach similar to the one proposed in [12].

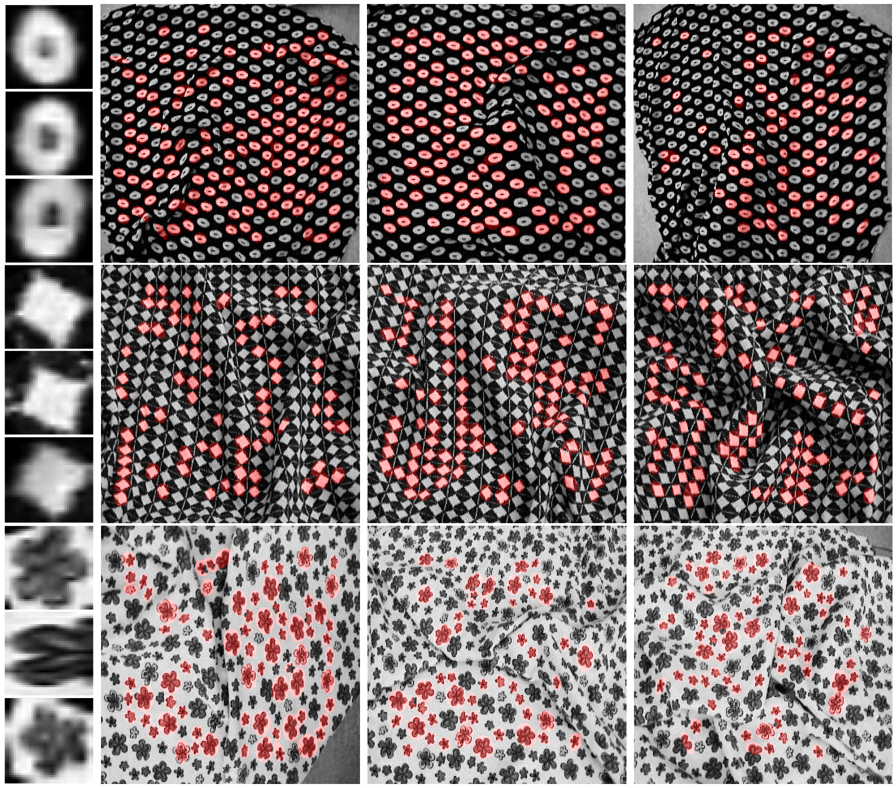
(a) (b)

Fig. 4. Learning locally planar regions. Three distinct textures. Column (a) shows the three top-ranked learned texture elements (top to down). Column (b) shows the corresponding learned elements (in red color) mapped on a sample video frame.

We commenced by extracting a large set of subregions from a sample of frames of the video sequence. Our current method does not use any temporal information and a sparse set of frames was usually sufficient for our algorithm to work. We extracted approximately 2,000 local affine-invariant descriptors from the set of images of the fabric surfaces. The feature extraction stage was followed by a ten-dimensional Isomap embedding of the corresponding affine-invariant descriptors. In our experiments, the EM learning step was performed using diagonal covariance matrices. After the EM learning step was performed, the algorithm automatically selected the classes with the highest prior probabilities such that the sum of priors formed 60% of the total population. For every such a class, the basic texture element was selected using Equation 2. Finally, the ranking stage was performed to remove noisy components and rank the most representative ones. Here, the 75% top-ranked elements were selected by the program. These

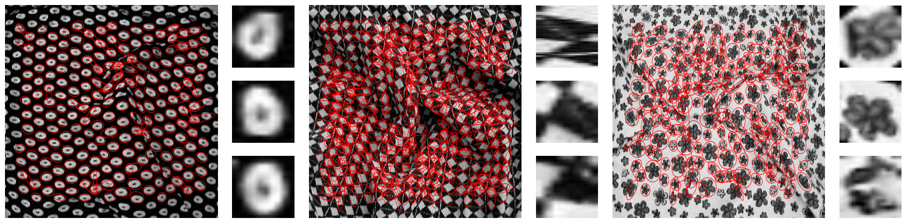

Fig. 5. Texture components learned using standard K-Means method. For every texture The affine-invariant space was clustered using k-means clustering algorithm. The learned texture elements are shown on the right-hand side of each sample frame.

thresholds were chosen experimentally. We selected a single patch descriptor whenever the calculations produced no resulting components. Our results were consistent for different values of K for the EM algorithm. It is common that among the top-ranked texture primitives there will be several exemplars of the same element. Such a situation may occur due to imprecisions in the region detection process as well as due to illumination changes. Figure 4 shows the results for three types of patterned fabrics used in our experiments. Figure 4(a) shows (from top to down) the ranked sequence of learned texture components obtained by our method. Figure 4(b) shows three frame images with the learned locally planar texture components superimposed on the fabric's surface.

At this point, we would like to turn the reader's attention to the third row in Figure 4. In the figure, the learned texture components represent a texture with flower-shaped elements as well as the first three top-ranked elements discovered by our method. While texture elements with large area are more likely to be distorted, some of their parts still may be locally planar. In the case of this texture, a part of a larger element received the second highest "planarity" rank. This suggests that our method is able to discover locally planar regions even in the presence of strong deformations by selecting those distributions that are most resistant to distortions.

In the second part of our experiments, we provide a qualitative comparison between the results obtained by our algorithm and typical results obtained by clustering the affine-invariant feature space using the K-Means clustering algorithm. Figure 5 shows a sample of results from this experiment. Basically, we use the same feature extraction steps as in the previous experiment. However, the nonlinear manifold learning step was not performed, and the EM algorithm clustering was substituted by the K-Means. Additionally, there was no texture component selection procedure performed on the clustered results. The figure shows a sample image-frame of each texture with some of the extracted affine regions indicated by red ellipses. The learned texture components produced by this procedure are shown on the righ-hand side of each texture sample. In this case, the resulting learned texture components present significant levels of deformation when compared to the ones obtained by our method.

These results show that our method is able to distinguish between locally planar texture elements and their nonlinearly distorted versions. For illustration purposes we selected only the first three top-ranked elements for every texture. Increasing the size of the dictionary of representative elements leads to classifying more regions as locally planar. However, lower ranked texture elements usually come from distributions with large covariance and are likely to contain noisy data. The investigation of this trade-off is important and is left for future work.

6 Conclusions and future work

In this paper, we proposed a method for learning the basic texture primitives of patterned surfaces distorted by non-rigid motion. The proposed algorithm uses nonlinear manifold learning to capture the intrinsic dimensionality of curvature distortion on non-rigid deforming texture surfaces. A selection procedure for determining the most representative local texture components was presented. A qualitative comparison between our texture primitive learning method and a standard K-Means learning procedure was also presented. Our experiments demonstrated the effectiveness of our method on a set of images obtained from patterned fabric surfaces undergoing a range of non-rigid deformations.

Interesting future directions include the further investigation of the effects of curvature on the local texture measurements as well as the introduction of both spatio-temporal information and inherent texture repetitiveness into the nonlinear manifold learning stage [10, 18]. Extending the method to include color information is also a possibility. We plan to apply our ideas to the classification of deforming texture surfaces. Studies aimed at developing these ideas are in hand and will be reported in due course.

Acknowledgments This research was supported by U.S. Office of Naval Research under contract: N00014-05-1-0764.

References

1. M. Belkin and P. Niyogi. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Neural Information and Processing Systems (NIPS)*, pages 585–591, 2001.
2. A. Bhalerao and R. Wilson. Affine invariant image segmentation. In *British Machine Vision Conference (BMVC)*, 2004.
3. M. Carcassoni, E. Ribeiro, and E. R. Hancock. Texture recognition through modal analysis of spectral peak patterns. In *IEEE International Conference on Pattern Recognition (ICPR)*, volume 1, page 10243, 2002.
4. D. Chetverikov and Z. Foldvari. Affine-invariant texture classification. In *IEEE International Conference on Pattern Recognition (ICPR)*, page 3901, Washington, DC, USA, 2000. IEEE Computer Society.
5. O. Cula and K. J. Dana. 3D texture recognition using bidirectional feature histograms. *International Journal of Computer Vision*, 59:33–60, 2004.

6. A. Dempster, N.M.Laird, and D.B.Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal Royal Stat. Soc., Series B*, 39(1):1–38, 1977.
7. M. N. Do and M. Vetterli. Wavelet-based texture retrieval using generalized Gaussian density and Kullback-Leibler distance. *IEEE Transactions on Image Processing*, 11(2):146–158, 2002.
8. I. Guskov, S. Klibanov, and B. Bryant. Trackable surfaces. In *ACM SIGGRAPH/Eurographics symposium on Computer animation*, pages 251–257, 2003.
9. J. Hays, M. Leordeanu, A. A. Efros, and Y. Liu. Discovering texture regularity as a higher-order correspondence problem. In *European Conference on Computer Vision (ECCV)*, volume 2, pages 522–535, 2006.
10. O. C. Jenkins and M. J. Matari. A spatio-temporal extension to isomap nonlinear dimension reduction. In *Int'l Conference on Machine Learning*, page 56, 2004.
11. T. Kadir, A. Zisserman, and M. Brady. An affine invariant salient region detector. In *European Conference on Computer Vision (ECCV)*, pages Vol I: 228–241, 2004.
12. S. Lazebnik, C. Schmid, and J. Ponce. A sparse texture representation using local affine regions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1265–1278, 2005.
13. T. Leung and J. Malik. Recognising surfaces using three-dimensional textons. In *IEEE International Conference on Computer Vision*, pages 1010–1017, 1999.
14. W.-C. Lin and Y. Liu. Tracking dynamic near-regular textures under occlusion and rapid movements. In *European Conference on Computer Vision*, May 2006.
15. C. B. Liu, R. S. Lin, N. Ahuja, and M. H. Yang. Dynamic textures synthesis as nonlinear manifold learning and traversing. In *British Machine Vision Conference (BMVC)*, pages 859–868, 2006.
16. Y. Liu, W.-C. Lin, and J. H. Hays. Near regular texture analysis and manipulation. *ACM Transactions on Graphics (SIGGRAPH 2004)*, 23(3):368 – 376, August 2004.
17. K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. V. Gool. A comparison of affine region detectors. *International Journal of Computer Vision*, 65(1-2):43–72, 2005.
18. A. Rahimi, B. Recht, and T. Darrell. Learning appearance manifolds from video. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 868–875, 2005.
19. E. Ribeiro and E. Hancock. Shape from periodic texture using the eigenvectors of local affine distortion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(12):1459–1465, December 2001.
20. S. T. Roweis and L. K. Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290:2323–2326, December 2000.
21. F. Schaffalitzky and A. Zisserman. Geometric grouping of repeated elements within images. In *Shape, Contour and Grouping in Computer Vision*, pages 165–181, 1999.
22. R. Souvenir and R. Pless. Isomap and nonparametric models of image deformation. In *IEEE Workshop on Motion and Video Computing - Vol. 2*, pages 195–200, 2005.
23. J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, Dec. 2000.
24. M. Varma and A. Zisserman. A statistical approach to texture classification from single images. *International Journal of Computer Vision*, 62:61–81, 2004.
25. R. White and D. A. Forsyth. Retexturing single views using texture and shading. In *European Conference on Computer Vision (ECCV)*, pages 70–81, 2006.
26. S. C. Zhu, C. en Guo, Y. Wang, and Z. Xu. What are textons? *International Journal of Computer Vision*, 62(1-2):121–143, 2005.