

Drums and Curve Descriptors

M. Zuliani, C. Kenney, S. Bhagavathy and B. S. Manjunath
Vision Research Lab, Dept. of Electrical and Computer Engineering
University of California, Santa Barbara
Santa Barbara, CA 93106
Email: {zuliani, kenney, sitaram, manj}@ece.ucsb.edu

Abstract

In this paper we present a new physically motivated curve descriptor based on the solution of Helmholtz's equation. The descriptor satisfies the six principles set by MPEG-7: it has a good retrieval accuracy, it is compact, it can be applied in general contexts, it has a reasonable computational complexity, it is robust and provides an hierarchical representation of the curve from coarse to fine. Moreover this descriptor generalizes straightforwardly to three dimensional surfaces. We tested the performance of the descriptor in the context of affine invariant curve matching using the *multiview curve dataset* (MCD), which consists of 40 curves (extracted from the MPEG-7 shape dataset) imaged under 14 different points of view. The results we obtained show that the proposed descriptor satisfies the MPEG-7 requirements and presents some advantages over some of the commonly used curve descriptors.

1 Introduction

The quest for efficient curve descriptors has been one of the leading themes in the computer vision community. In general, good curve descriptors should be invariant under an appropriate class of geometric transformations (like, for example, rotation-scaling-translation or affine), robust in the presence of noise, efficient to compute and easy to compare. Zhang et al. [13] classified the curve description approaches into two groups: *contour-based* and *region-based* methods. Each of these groups is further subdivided into two subgroups containing *global* or *structural* approaches. Some of the recently proposed descriptors fall in the contour-based category, as the *curvature scale space* (CSS) descriptor [8] (which has been standardized in the MPEG-7 framework) and the *shape context matrices* [1]. Some others belong to the class of region-based methods, like the descriptors based on *moments* (geometric [3], Zernike and Legendre [11]), on *region frequency representations* (Fourier descriptors [12]), on *medial axis transform* [9] and on *shock graphs* [10]. The goal of this paper is to develop a curve descriptor that satisfies the six principles set by MPEG-7 and a few other requirements, such as being rotation-scaling-translation (RST) invariant, having a clear physical interpretation and being easily extendable to the 3D case (for surface description). Note that we are interested in describing Jordan curves (i.e. curves that are closed and do not cross themselves).

The approaches mentioned before do not always fully satisfy the MPEG-7 requirements. The computation of the CSS descriptors is quite demanding, the algorithm con-

verges slowly if the curve is very complex (i.e. many points of the curve have a large curvature) and depends on some empirical parameters that need to be fine tuned. The comparison of CSS descriptors is not so simple and the method cannot be extended straightforwardly to the 3D case. Shape context matrices provide local curve descriptors that are not very compact (matrices) and their comparison is not very fast. Moreover an extension to 3D surfaces would increase both the size of descriptors and the complexity of the algorithm. Moment invariants of higher orders do not have a clear physical interpretation and the matching procedure requires a normalization process to compensate for the different dynamic range of the moments of different orders. Shock graphs are very suitable in contexts where the similarity between curves is defined in terms of structure, but are not the ideal solution if the notion of equivalence is defined within the class of some specific geometric transformation. Also in this case the computational complexity for extracting the descriptors and matching them is very high.

This paper is structured as follows. Section 2 introduces our proposed descriptor, it discusses its analytical properties and presents the numerical scheme used to compute the descriptor. Section 3 will describe a preprocessing step that aims at extracting the *shape* of a curve in order to obtain an affine invariant matching algorithm. Section 4 shows some experimental results to evaluate the performance of the descriptors and finally the conclusions are presented in Section 5.

2 The Curve Descriptor

In 1966 the mathematician M. Kac published his famous paper entitled “Can One Hear the Shape of a Drum?” [5]. Kac was interested in understanding whether the knowledge of a drum’s modes of vibration was sufficient for univocally inferring its geometric structure. The problem posed by Kac can be related to the problem of identifying curve descriptors. In fact, if we imagine that the curve we want to label defines the contour of a drum, it is reasonable to think that the spectrum of such curve (in terms of modes of vibration) could be an appealing descriptor, given the fact that it can be easily made RST invariant and has a strong physical characterization. To this purpose the answer to Kac’s question becomes crucial, i.e. we would like to have the normal modes of vibration of a drum to identify univocally its geometry (so that we can establish a bijection between the space of the Jordan curves modulo a given transformation and the curve descriptors).

The problem posed by Kac remained unsolved until 1992 when the mathematicians C. S. Gordon, D. L. Webb and S. Wolpert proposed a pair of isospectral drums having the same area and perimeter but different contours. In other words “One Cannot Hear the Shape of a Drum” [4]. Even though for our purposes this fact is unfortunate, since it implies that there may exist curves that are not related by an RST transformation and nonetheless have the same spectrum (i.e. the same descriptor), the experiments presented in Section 4 will show how this problem has a limited impact in real life scenarios.

In the following subsections we will describe in detail the proposed curve descriptor and the numerical scheme used to compute it.

2.1 The Helmholtz’s Equation

Let Γ be a Jordan curve corresponding to the boundary of Ω , an open subset of \mathbb{R}^2 . The vibration of the membrane of a drum whose contour is defined by Γ is expressed by the

function $w(\mathbf{x}, t) : \bar{\Omega} \times \mathbb{R} \rightarrow \mathbb{R}$ which solves the wave equation:

$$\Delta w - \frac{1}{v^2} \frac{\partial w}{\partial t} = 0$$

where Δ denotes the Laplacian operator, t indicates time and v is a characteristic constant associated to the drum. This equation can be solved via separation of variables, assuming that w can be decomposed in a spatial part and in a temporal part according to $w(\mathbf{x}) = u(\mathbf{x})q(t)$. It can be shown that the spatial part solves the *Helmholtz equation*, an elliptic partial differential equation given by:

$$\Delta u - \lambda u = 0 \tag{1}$$

where λ is a suitable scalar. The corresponding boundary problem with Dirichlet conditions is:

$$\Delta u - \lambda u = 0 \tag{2a}$$

$$u(\mathbf{x}) = 0 \quad \text{for } \mathbf{x} \in \Gamma \tag{2b}$$

As explained in [2], the pairs $\{\lambda_k, u_k\}$ that solve the problem (2) are such that:

- All the eigenvalues are real and positive: $0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots$ with $\lambda_k \rightarrow \infty$ as $k \rightarrow \infty$
- The real eigenfunctions u_k can be chosen to form an orthonormal basis of $L^2(\Omega)$

Our idea is to use the first N_λ eigenvalues to build an RST-invariant descriptor for the curve Γ , as follows from the definition and lemma:

Definition 1. Let Γ be a Jordan curve and let $\lambda_1, \dots, \lambda_{N_\lambda}$ be the first N_λ eigenvalues that solve (2). The *Helmholtz curve descriptor (HCD)* associated to Γ is defined as:

$$\mathbf{F}(\Gamma) \stackrel{\text{def}}{=} \begin{bmatrix} \frac{\lambda_1}{\lambda_2} & \frac{\lambda_1}{\lambda_3} & \dots & \frac{\lambda_1}{\lambda_{N_\lambda}} \end{bmatrix}^T \in \mathbb{R}^{N_\lambda} \tag{3}$$

Lemma 1. Consider the two Jordan curves Γ_1 and Γ_2 related by an RST transformation:

$$\Gamma_2 = \{\mathbf{x}_2 \in \mathbb{R}^2 : \text{there exists } \mathbf{x}_1 \in \Gamma_1 \text{ such that } \mathbf{x}_2 = sR\mathbf{x}_1 + \mathbf{t}\}$$

where $s \in \mathbb{R}$ is the scaling factor, $R \in SO(2)$ is a rotation matrix and $\mathbf{t} \in \mathbb{R}^2$ is a translation vector. Then $\mathbf{F}(\Gamma_1) = \mathbf{F}(\Gamma_2)$.

Proof. The lemma can be shown easily using the variational characterization of the eigenvalues. Let u_1, \dots, u_{k-1} be the first orthonormal eigenfunctions of (2a). Let's define for $h = 1, 2$ the subspaces:

$$S_k^h = \{w : w \in C^2(\Omega_h), w \neq 0, w(\mathbf{x})|_{\mathbf{x} \in \Gamma_h} = 0, \langle w | u_l \rangle = 0 \text{ for } 1 \leq l \leq k-1\}$$

These subspaces contain all the functions (continuous up to the second derivative) defined over Ω_h , not identically equal to zero in Ω_h , equal to zero on the boundary Γ_h and orthogonal to the first $k-1$ eigenmodes of (2a). The notation $\langle \cdot | \cdot \rangle$ denotes the inner product on

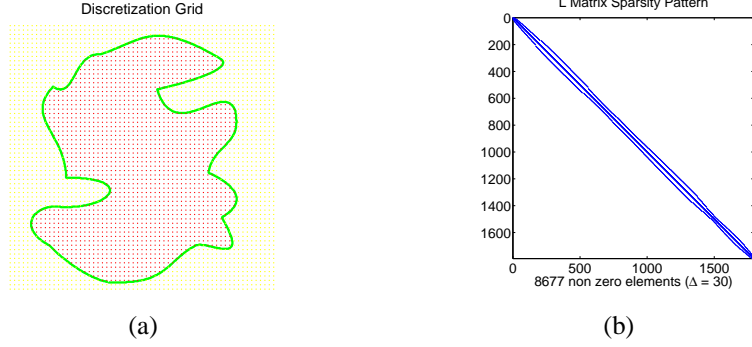


Figure 1: Image (a) shows an example of the discretization mesh used for the finite difference method. Figure (b) shows the structure of the matrix L .

$\Omega: \langle f|g \rangle = \int_{\Omega} f(\mathbf{x})g(\mathbf{x}) d\mathbf{x}$. It can be shown that the k^{th} eigenvalue can be obtained by minimizing the Rayleigh quotient:

$$\lambda_{h,k} = \min_{w \in S_k^h} \frac{\|\nabla w\|^2}{\|w\|^2}$$

Since $\mathbf{x}_2 = sR\mathbf{x}_1 + \mathbf{t}$ we can write:

$$\lambda_{2,k} = \min_{w \in S_k^2} \frac{\int_{\Omega_2} \|\nabla w(\mathbf{x}_2)\|_2^2 d\mathbf{x}_2}{\int_{\Omega_2} w(\mathbf{x}_2)^2 d\mathbf{x}_2} = \min_{w \in S_k^1} \frac{\int_{\Omega_1} s^2 \|\nabla w(\mathbf{x}_1)\|_2^2 s^2 d\mathbf{x}_2}{\int_{\Omega_1} w(\mathbf{x}_2)^2 s^2 d\mathbf{x}_2} = s^2 \lambda_{1,k}$$

It is clear that if we take the ratios that define the HCD components the scaling factor s^2 is cancelled out, and therefore we have that $\mathbf{F}_i(\Gamma_1) = \mathbf{F}_i(\Gamma_2)$ for any $1 \leq i \leq N_{\lambda} - 1$. \square

As mentioned before, it has been theoretically proven that there exist different curves that have the same spectrum. However this event is quite rare (where the notion of ‘‘rare’’ can be formalized more precisely) as the experiments presented in Section 4 will confirm. The similarity between the descriptors is defined in terms of their Euclidean distance:

$$d(\mathbf{F}(\Gamma_1), \mathbf{F}(\Gamma_2)) = \|\mathbf{F}(\Gamma_1) - \mathbf{F}(\Gamma_2)\|_2 \quad (4)$$

The choice of this distance function implicitly defines an hierarchical comparison of the curves. Because of the different dynamic ranges of the HCD components (it can be shown that the entries of $\mathbf{F}(\Gamma)$ decay according to $1/i^2$, where i is the component index) the coarse structure of the curve (captured by the first components of $\mathbf{F}(\Gamma)$) is naturally weighted more than the finer details (represented by the last components of $\mathbf{F}(\Gamma)$).

2.2 Numerical Scheme

The second order finite difference scheme we used to solve (2) provided a good compromise between accuracy and computational complexity. The step size of the $N \times N$ discretization mesh is calculated according to:

$$h = \frac{\max_{\mathbf{x} \in \Gamma} \|\mathbf{x} - \mathbf{m}(\Omega)\|}{\Delta}$$

where $\mathbf{m}(\Omega)$ is the center of gravity of the region Ω and Δ is a parameter that defines the mesh resolution. The spatial derivatives are approximated by the second order central difference formula:

$$\frac{\partial^2 u}{\partial x^2}(\mathbf{x}) \approx \frac{u(x+h, y) - 2u(x, y) + u(x-h, y)}{h^2}$$

$$\frac{\partial^2 u}{\partial y^2}(\mathbf{x}) \approx \frac{u(x, y+h) - 2u(x, y) + u(x, y-h)}{h^2}$$

Using these assumptions the problem can be represented in a linear algebra fashion according to:

$$L\mathbf{u} - \lambda\mathbf{u} = 0$$

where the linear operator L is given by the sparse symmetric matrix:

$$L = h^{-2} \begin{bmatrix} A & I_N & 0 & \dots & 0 \\ I_N & A & I_N & \dots & 0 \\ 0 & I_N & A & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & A \end{bmatrix} \in \mathbb{R}^{N^2 \times N^2} \quad A = \begin{bmatrix} -4 & 1 & 0 & \dots & 0 \\ 1 & -4 & 1 & \dots & 0 \\ 0 & 1 & -4 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & -4 \end{bmatrix} \in \mathbb{R}^{N \times N}$$

Note that I_N is the $N \times N$ identity matrix and the vector $\mathbf{u} \in \mathbb{R}^{N^2}$ is such that:

$$\mathbf{u}_{Np+q} = u(x_p, y_q)$$

where x_p and y_q are the mesh point coordinates. The size of the problem can be reduced by removing the entries of vector \mathbf{u} that correspond to points outside of the domain Ω (and consequently the corresponding row/columns in the matrix L). Figure 1 shows an example of the discretization mesh and of the matrix L for a certain curve Γ . In our implementation the eigenvalues/eigenvectors are computed using the Fortran library ARPACK [6] (accessed through Matlab) that takes advantage of the sparse and symmetric structure of L .

3 Achieving Affine Invariance

The descriptors we have introduced in Section 2 are RST-invariant. However very often it is necessary to match curves in an affine invariant fashion. As an example consider planar curves imaged from two different viewpoints using a distant¹ camera. In this case the perspective distortion can be approximated by an affine transformation. We will briefly summarize a procedure that allows to map a curve in a normalized coordinate system where affine-related curves become congruent modulo a geometric rotation. For a more thorough discussion containing also the proof of Theorem 1 refer to [14].

Let's first introduce the following quantities:

- Let $V(\Omega) \stackrel{\text{def}}{=} \int_{\Omega} d\mathbf{x}^2$ be the area of Ω , where $d\mathbf{x}^2$ is the infinitesimal area element.
- Let $\mathbf{m}(\Omega) \stackrel{\text{def}}{=} \frac{1}{V(\Omega)} \int_{\Omega} \mathbf{x} d\mathbf{x}^2$ be the centroid of Ω .

¹Where distant has to be intended with respect to the camera focal length.

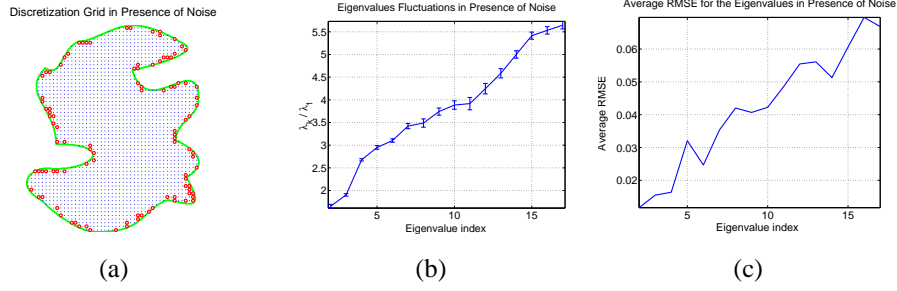


Figure 2: Image (a) shows the mesh points that discretize the domain $\Omega(\cdot)$ and the points modified by the morphological noise operator (\circ). Plot (b) shows the fluctuation of the eigenvalues in presence of morphological noise for a given curve. Plot (c) displays the RMSE of the eigenvalues for a set of 10 shapes averaged over 20 realizations of the morphological noise.

- Let $\Sigma(\Omega) \stackrel{\text{def}}{=} \frac{1}{V(\Omega)} \int_{\Omega} [\mathbf{x} - \mathbf{m}(\Omega)][\mathbf{x} - \mathbf{m}(\Omega)]^T d\mathbf{x}^2$ be the covariance of Ω .

We now have all the ingredients to define the shape of a Jordan curve:

Definition 2. Let Γ be a Jordan curve. The shape of Γ is defined as:

$$S(\Gamma) \stackrel{\text{def}}{=} \left\{ \mathbf{s} \in \mathbb{R}^2 : \mathbf{s} = \Sigma(\Omega)^{-\frac{1}{2}} [\mathbf{x} - \mathbf{m}(\Omega)] \text{ for } \mathbf{x} \in \Gamma \right\} \quad (5)$$

This definition is important because it allows us to relate affine-transformed curves, as stated in the following theorem:

Theorem 1. Let Γ_1 and Γ_2 be two Jordan curves related by an affine transformation:

$$\Gamma_2 = \left\{ \mathbf{x}_2 \in \mathbb{R}^2 : \exists \mathbf{x}_1 \in \Gamma_1 \text{ such that } \mathbf{x}_2 = A\mathbf{x}_1 + \mathbf{b} \right\}$$

where $A \in \mathbb{R}^{2 \times 2}$ is a non-singular matrix and $\mathbf{b} \in \mathbb{R}^2$. Then the shapes of Γ_1 and Γ_2 are geometrically congruent via a 2-dimensional rotation.

Therefore if we want to use the descriptors introduced in Section 2 in the context of affine-invariant matching we just have to extract the shape of a curve Γ and calculate the RST-invariant descriptors of $S(\Gamma)$.

4 Experimental Results

4.1 Numerical Stability in Presence of Noise

In this subsection we are interested in studying the numerical stability of the eigenvalues of the Laplacian in presence of noise. The descriptors we proposed are based on the first N_λ smallest eigenvalues which capture the slowest modes of vibration of the domain Ω and therefore are the less sensitive in presence of small boundary variations. On the other hand the smallest eigenvalues of a matrix are those more affected by the finite precision mathematical operations.

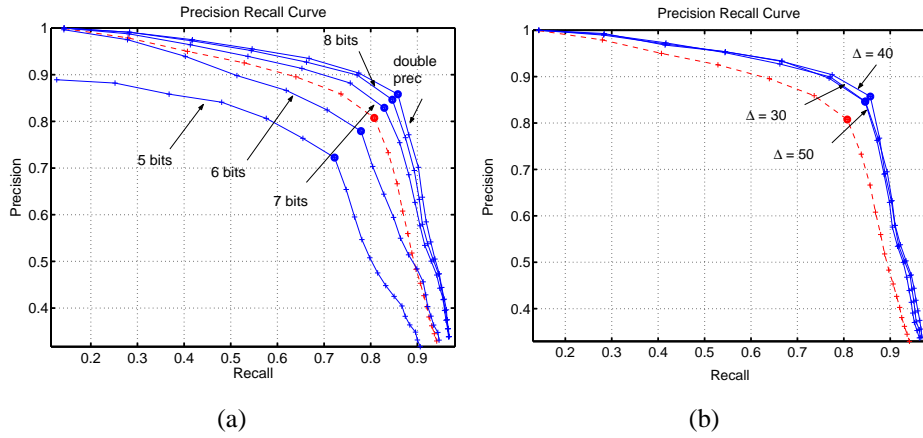


Figure 3: Precision recall curve calculated over the MCD for the Helmholtz curve descriptors for different numbers of bits used to represent the descriptors (a) and for different mesh steps (b). In both examples $N_\lambda = 17$. The dashed line shows the performance of CSS curve descriptors calculated using MPEG-7 XM software.

To evaluate the effect of small morphological changes in the curve we choose a set of 10 normalized curves and applied to each curve a noise morphological operator that modifies with probability 0.8 the membership of the mesh points in a 8-pixel neighborhood (i.e. the fact that a point (x_p, y_q) belongs to Ω or not). The average RMSE of the eigenvalues (with respect to the noise free curve) has been calculated. The results for $\Delta = 30$ (which are shown in Figure 2) demonstrate that the eigenvalues with larger indices are those more affected by the morphological noise. This can be explained observing that these eigenvalues are related to faster modes of vibration. However the effect of the morphological noise does not produce drastic variations of the value of the eigenvalues and consequently of the components of the descriptors.

4.2 Precision Recall Performance

The MCD comprises 40 curve categories, each corresponding to a curve drawn from an MPEG-7 curve class. Each category in the dataset contains 14 curve samples that correspond to different perspective distortions of the original curve. To construct the MCD we printed on white paper 40 curves taken from the MPEG-7 dataset. For each curve 7 samples were taken from different view angles using a digital camera. The contours were extracted from the iso-intensity level set decomposition of the images [7]. By adding random rotations and reflections to these samples, the number of samples in each category is doubled to 14. Some of the curves comprising the MCD are shown in Figure 5.

The performance of the descriptors is evaluated using the precision-recall curve calculated over the MCD. Each curve Γ is used in turn as the query. Let $A(\Gamma, T)$ denote the set of T retrievals (based on the smallest distances (4) from Γ in the descriptor space) and $R(\Gamma)$ the set of 14 images in the dataset relevant to Γ . The *precision* is defined by $P(\Gamma, T) \stackrel{\text{def}}{=} \frac{|A(\Gamma, T) \cap R(\Gamma)|}{T}$, and the *recall* by $C(\Gamma, T) \stackrel{\text{def}}{=} \frac{|A(\Gamma, T) \cap R(\Gamma)|}{14}$, where $|\cdot|$ denotes cardinality. The precision recall curve is plotted by averaging precision and recall over all Γ ,

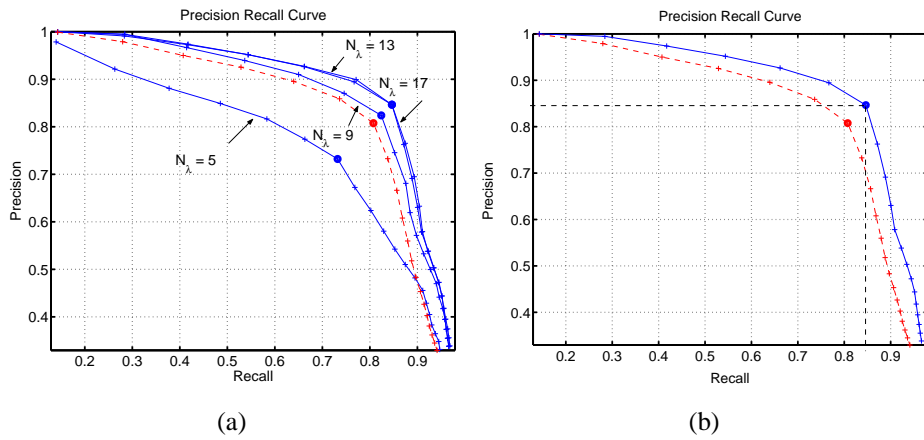


Figure 4: Precision recall curves for different HCD lengths (a). Figure (b) shows the performance of the 12 bytes long descriptor. The dashed line shows the performance of CSS curve descriptors calculated using MPEG-7 XM software.

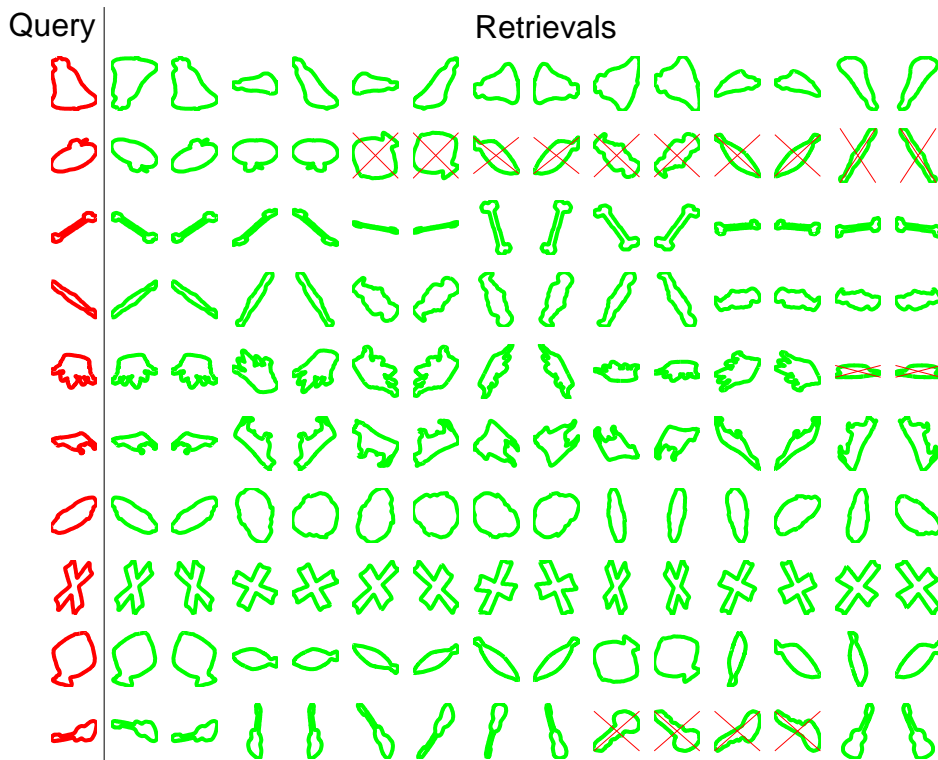


Figure 5: Retrieval results obtained using 10 random queries from the MCD. The curve scale has been normalized for visualization purposes. The experiment was carried out using the following parameters: $\Delta = 30$, $N_\lambda = 13$, 8 bits. The crosses mark retrievals that do not belong to the correct category.

for different values of T (in all the plots the point corresponding to $T = 14$ is represented by a thick circle).

Figure 3(a) shows the precision recall curves for different numbers of bits used to represent the components of the descriptors. In this experiment $N_\lambda = 17$. Using 8 bits the size of the HCD is 16 bytes and the performance is very close to the case where a double precision representation is used. From Figure 3(b) we deduce that the mesh resolution does not affect considerably the descriptor performance.

The precision recall curves plotted in Figure 4(a) show how most of the information about the curve structure can be condensed in a 12-dimensional curve descriptor. Note that not always the strong perspective distortions that characterize the MCD can be approximated by an affine transformation. Therefore the experiments carried out using this dataset measure also the robustness of the descriptors in presence of non modelled curve transformations. This is confirmed by the results shown in Figure 4(b), where using a 12 dimensional curve descriptors quantized using 8 bits ($\Delta = 30$) we have that for $T = 14$ the precision is higher than 85% and the recall higher than 86%. This means that in more than 85% of the cases a curve in the top 14 retrievals is from the same class of the query.

Figure 5 shows the retrieval results obtained using 10 random queries from the MCD. It is interesting to note that for the category *apple* (second row) the “wrong” curves that are retrieved share a similar shape and this is true also for the mismatched curves in the category “guitar” (last row), where the four mismatches belong to the category “key”.

5 Conclusions

In this paper we proposed a new curve descriptor that satisfies the six principles set by MPEG-7. The experimental results on the MCD show that:

- The HCD has a good retrieval accuracy and if applied to the shape of a curve it can deal with affine transformed curves.
- The HCD is very compact (satisfactory performances can be obtained with descriptors that are only 12 bytes long; the average size of the CSS descriptor is 14 bytes per curve) and has a strong physical characterization.
- The HCD can be applied in general contexts, since the algorithm parameters do not need to be fine tuned (the experiments show a low performance fluctuation for relevant changes in the mesh step).
- The computational complexity is limited, both for the descriptor calculation and (especially) for the descriptor comparison.
- The descriptors can tolerate non-modelled geometric curve distortions (in fact the shape extraction procedure does not always compensate for the strong perspective distortion of the curves contained in the MCD).
- The Euclidean distance function and the decreasing dynamic range of the descriptor coefficients naturally provide an intrinsically hierarchical comparison of the curves.

Moreover the proposed descriptor as well as the shape extraction procedure generalize straightforwardly to 3D surfaces. Future work involves the extension of this descriptor to 3D surfaces and a more quantitative evaluation of the problem of isospectral shapes.

Acknowledgments. The authors would like to thank prof. Shiv Chandrasekaran in the ECE Department of UCSB for the many discussion. This project was partly supported by the grant ONR #N00014-04-1-0121 from the Office of Naval Research.

References

- [1] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, April 2002.
- [2] L. C. Evans. *Partial Differential Equations (Graduate Studies in Mathematics, 19)*. American Mathematical Society, 1998.
- [3] J. Flusser and T. Suk. Pattern recognition by affine moment invariants. *Pattern Recognition*, 26(1):167–174, 1993.
- [4] C. Gordon, D. L. Webb, and S. Wolpert. One cannot hear the shape of a drum. *Bulletin of the American Mathematical Society*, 27(1):134–138, July 1992.
- [5] M. Kac. Can one hear the shape of a drum? *American Mathematical Monthly*, 73(2):1–23, 1966.
- [6] R.B. Lehoucq, D.C. Sorensen, and C. Yang. *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*. SIAM Publications, 1998.
- [7] J. L. Lisani. *Shape Based Automatic Images Comparison*. PhD thesis, Université Paris IX-Dauphine, July 2001.
- [8] F. Mokhtarian and M. Bober. *Curvature Scale Space Representation: Theory, Applications, and MPEG-7 Standardization*. Kluwer Academic Publishers, 2003.
- [9] B. S. Morse. *Computation of Object Cores from Grey-level Images*. PhD thesis, University of North Carolina at Chapel Hill, 1994.
- [10] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. *International Journal of Computer Vision*, 35(1):13–32, November 1999.
- [11] M.R. Teague. Image analysis via the general theory of moments. *Journal of the Optical Society of America*, 70(8):920–930, 1979.
- [12] D. S. Zhang and G. Lu. Generic fourier descriptors for shape-based image retrieval. In *Proceedings of IEEE International Conference on Multimedia and Expo*, volume 1, pages 425–428, Lausanne, Switzerland, August 2002.
- [13] D. S. Zhang and G. Lu. Review of shape representation and description techniques. *Pattern Recognition*, 37(1):1–19, 2004.
- [14] M. Zuliani, C. Kenney, and B. S. Manjunath. A definition of shape for jordan manifolds. Technical Report TR-VRL-03-2004, Vision Research Lab, Department of Electrical and Computer Engineering, University of California, Santa Barbara, 2004.