A Shape Descriptor for Shapes with Boundary Noise and Texture

Wooi-Boon Goh and Kai-Yun Chan
School of Computer Engineering
Nanyang Technological University, Nanyang Avenue,
Singapore 639798
aswbgoh@ntu.edu.sg

Abstract

We present a part-based shape descriptor that incorporates both the description of the general shape form of each subpart and its geometric relationship with other connected parts. Associated with each descriptor is a saliency measure that weighs each part's visual significance. By incorporating this saliency measure into the shape matching process, we able to discriminate between shape forms as well as take boundary texture into consideration when computing shape similarity. This paper also describes a multi-resolution pyramidal framework for generating the required gradient vector field and vector field disparity map from which the shape descriptors, in the form of gradient vector field histograms, are derived. Experimental results involving silhouettes images are presented to demonstrate the various characteristics of the proposed shape descriptor, which includes its invariance to similarity transform and its ability to match composite shapes containing boundary noise and texture, limb articulation and occlusion.

1 Introduction

Specifying an accurate query in content-based image retrieval systems is a challenging task that typically involves the integration of multiple cues like shape, color, texture and the spatial relationship between such features. Shape is an important cue as it captures a prominent element of an object. However, the description of a shape is difficult to specify since humans may perceive two shapes to be similar despite their varying sizes, orientations and boundary characteristics. For example, two rabbit shapes are recognized to be rabbits despite one being a small smooth-coat and the other a large hairy-coat. The visual forms of these two rabbits are perceptually similar despite the differing boundary characteristics. The boundary variations can be viewed as undesirable boundary noise or in some cases, informative boundary texture that should be incorporated into the shape description.

There have been many approaches proposed to deal with boundary perturbations. Boundary-based methods such as [14] represent shapes by the locations of the maxima of its curvature scale space (CSS) image. Shapes are smoothed by selecting the

BMVC 2003 doi:10.5244/C.17.21

appropriate scale and then matched by shifting the CSS contours so that the major maxima of one image overlaps that of the other. In [11], shape boundaries are approximated using polygonal curves and are progressively simplified through discrete curve evolution based on a novel relevance measure. The weakness of the boundary-based approach is that it does not represent the interior of the shape [16] and is therefore very sensitive to the spatial reconfigurations of parts and local boundary perturbations.

Shape descriptors based on the medial-axis representation are equally susceptible to boundary noise despite incorporating both boundary and region information since many skeletonization algorithms produces false branches with even slight boundary perturbations. To overcome this problem, [13] proposed a multi-scale shape representation called the skeleton-space. It employs a scale-space-driven selection scheme, which is able to extract the medial axes representing the most salient shape features. Methods based on shock computation [15], [17] are also highly sensitive to boundary perturbations. As a result, [15] introduced an edit distance algorithm that is able to splice such noise-induced branches with a minimal edit cost by taking into account their corresponding change to the boundary model. This effectively transfers the de-noising and texture analysis process to the shape similarity computation stage. This is in contrast to the scale-space approaches [11], [13], [14], which perform shape regularization at the feature extraction stage, resulting in multiple shape descriptors at varying scales.

In this paper we present a part-based shape descriptor that incorporates both the description of the general shape form of each subpart as well as the local boundary perturbation (boundary texture). Associated with each shape part descriptor is a saliency measure that weighs its visual significance based on the proportion of the overall shape region it occupies. The saliency measure is incorporated into the shape matching process and this results in a shape descriptor that is relative robust to boundary noise but at the same time, it takes boundary texture into consideration when computing shape similarity. The proposed shape descriptor can be classified as a medial axis-based approach but unlike [12], [15], [17], we do not represent shapes by graphs or trees. Instead, the shape axes provide a means to partition the shape into subparts. Each part is then represented by two 1-D histograms that statistically describe its shape and its geometric relationship to other parts in its vicinity. Shape similarity is then computed by comparing these histograms.

2 Gradient Vector Field and Disparity Map

2.1 Generating the Gradient Vector Field

Gradient vector fields have been used for extracting medial axis such as the divergence-based skeleton of [9],[19] and the multi-scale approach of [8]. However, one unique requirement of our proposed shape descriptor is to ensure that local boundary perturbations will only produce short unconnected shape axes that remain close to the boundary. To achieve this, the gradient vector field in the interior of the shape must be derived mainly from the low spatial frequency components of the shape (i.e. the gross shape form) and the high spatial frequency components are emphasized on the boundary. To meet these requirements, we proposed using multi-resolution pyramids to generate the gradient vector field.

For a review of standard pyramid operations such as *REDUCE* and *EXPAND*, readers can refer to [5]. It is assumed that the input image is a binary silhouette image $I(x,y)$. Firstly, a Gaussian pyramid $G(l,x,y)$ of N+1 levels is created by iteratively applying the *REDUCE* operation N times on each consecutive output image, starting with $I(x,y)$. From the scalar Gaussian pyramid, we then derived the vectorial Gradient pyramid $\boldsymbol{H}(l,x,y)$, which consist of two pyramids $H^x(l,x,y)$ and $H^y(l,x,y)$ given by

$$H_l^x(x,y) = g_{\sigma_H}^x(x,y) * G_l(x,y) \text{ and } H_l^y(x,y) = g_{\sigma_H}^y(x,y) * G_l(x,y) \text{ for } 0 \le l \le N \qquad \textbf{(1)}$$

The convolution kernels $g_{\sigma_H}^x$ and $g_{\sigma_H}^y$ are first-order Gaussian derivatives in the x and y directions, respectively and are given by

$$g_{\sigma_H}^x(x,y) = -\frac{x}{\sigma_H^2}\exp(-\frac{x^2+y^2}{2\sigma_H^2}) \text{ and } g_{\sigma_H}^y(x,y) = -\frac{y}{\sigma_H^2}\exp(-\frac{x^2+y^2}{2\sigma_H^2}) \qquad \textbf{(2)}$$

Each level $\boldsymbol{V}_l$ of the Gradient Vector Field pyramid, for $0 \le l < N$, is constructed by iteratively applying *EXPAND* to a proportioned sum of the corresponding level $\boldsymbol{H}_l$ of the Gradient pyramid and *EXPAND* $[\boldsymbol{V}_{l+1}]$, starting with $l = N$-1. More formally,

$$\boldsymbol{V}_l = \alpha\, \boldsymbol{H}_l + (1-\alpha)\, EXPAND\,[\boldsymbol{V}_{l+1}] \qquad \text{for } 0 \le l < N \qquad \textbf{(3)}$$

where $\boldsymbol{V}_l = \boldsymbol{H}_l$ for $l = N$ and the parameter $\alpha \in [0,1]$ determines the smoothness of the gradient vector field within the object. A smaller α results in a smoother vector field in the interior of the shape (i.e. more low spatial frequency components) [10].

2.2 The Vector Field Disparity Map

In order to describe a complex shape by decomposing it into suitable parts, the shape axes must first be extracted from the gradient vector field. These shape axes can be located by detecting locations in the vector field where the local gradient vectors exhibit high directional disparity. Given such a disparity map, the shape axes are extracted by locating the local maxima in the disparity measure. Extending the idea in [2], the normalized Vector Field Disparity pyramid $D(l,x,y)$ is derived from $\boldsymbol{V}(l,x,y)$ where each level D_l is defined as

$$D_l(x,y) = \frac{g_{\sigma_D}(x,y) * \|\boldsymbol{V}_l(x,y)\| - \|g_{\sigma_D}(x,y) * \boldsymbol{V}_l(x,y)\|}{g_{\sigma_D}(x,y) * \|\boldsymbol{V}_l(x,y)\|} \quad \text{for } 0 \le l \le N \qquad \textbf{(4)}$$

The local disparity measure is computed within a weighted locality defined by the Gaussian kernel $g_{\sigma D}$ given by $\exp[-(x^2+y^2)/2\sigma_D^2]/\sqrt{2\pi}\sigma_D$. The normalized vector field disparity measure $D_l(x,y) \in [0,1]$ gives a value close to 1 in localities of high disparity such as at the center of a circle. In order to detect consistent shape axes over different scales of a shape (see Fig. 1b), a full resolution vector field disparity map $M(x,y)$ is obtained by iteratively applying *EXPAND* to the sum of D_l and *EXPAND* $[D_{l+1}]$, starting with at $l = N$-1 to $l = 0$, as shown in Fig. 1a. The final summed output is divided by N+1 to re-normalise the disparity map such that $M(x,y) \in [0,1]$.

Although the shape axes produced by the described multi-resolution technique share many similar characteristics, they are not identical to the shocks of [17] nor the skeletons generated using medial axis transform [4]. As can be seen in Fig. 1c, minor

boundary perturbations such as noise and texture result in low saliency part axes that remain close to the boundary. As a result, boundary texture, though accounted for in our shape descriptor, does not contribute significantly to the shape similarity measure. This allows the matching of textured objects that share similar gross shape forms. Description of the procedures for shape axes extraction have been omitted due to space constraints but further details can be found in [10].

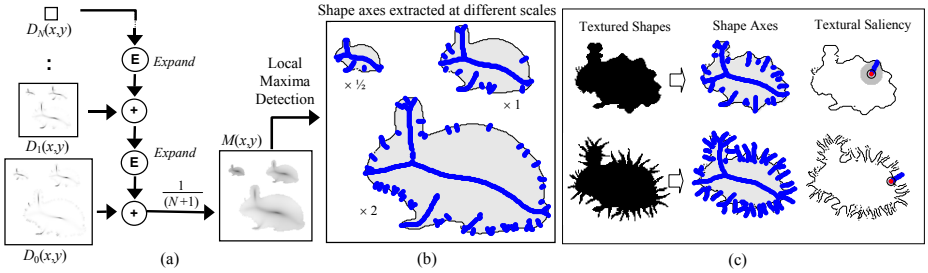

Fig. 1. (a) A multi-level integration technique for generating the vector field disparity map. (b) The shape axes extracted is relatively consistent with scales spanning 2 octaves. (c) Extracted shape axes for boundary texture remains close to the boundary, resulting in part axes that have lower saliency. Low spatial frequency boundary texture like bumps results in higher textural saliency then spiky boundary perturbations.

3 A Part-based Shape Descriptor

There is much evidence that human perception of shapes involves some form of part-based representation [9], [17]. A more local form of shape description allows for recognition that is more robust to occlusion, articulation of limbs and spatial rearrangement of parts. In our work, a shape is decomposed into subparts along segmented sections (*part axes*) of the extracted shape axes and each part is described by two 1-D histograms derived from the local gradient vector field (see Fig. 2a).

3.1 Part Description and Saliency

Using appropriate morphological operators and skeleton analysis algorithms, the shape axes are thinned and segmented at every intersections and junctions into Q continuous segments called *part axes*. With reference to Fig. 2b, let the k^{th} part axis $P_k = \{p_k^i\}_{i=1}^M$ be an ordered set of M discrete points starting at one end of a continuous segment and ending at the other. Let r_1 and r_M be the respective radial distances of points p_k^1 and p_k^M to their nearest edge points on the shape boundary. If the ordering of the points $\{p_k^i\}_{i=1}^M$ on P_k is such that $r_1 > r_M$, then the starting point p_k^1 is termed the *link node* of the part axis P_k.

Each part axis is associated with two normalized gradient vector field histograms. They are the *segment histogram* and the *link histogram*. The segment histogram S_k describes the general shape of the part, such as its length-to-width ratio, its convexity, its taper, etc. The link histogram L_k^1 contains information pertaining to the common space that the part axis k shares with other part axes within its vicinity. This relational information is very useful in differentiating similar looking protrusions that may be link

to the main shape body at different places and in varying configurations. Link histogram L^1_k is first constructed by determining the orientation associated with all C^1_k discrete image pixels within the circle of radius r_1, centered about the start point p^1_k. The histogram L^1_k with n bins representing the value range $[0,2\pi)$ cumulates the quantised orientation of all C^1_k gradient vectors. The histogram is made rotationally invariant by adding a value θ^1_k to all orientation values before cumulation. The angle θ^1_k is derived from the orientation of a straight line fitted along all part axis points $\{p^i_k\}_{i=1}^M \cap C^1_k$ that lie within the circle of radius r_1. The start point of the straight line is the end closest to p^1_k. The link histogram L^1_k is normalized with the value C^1_k. The segment histogram S_k is obtained by repeating the procedure described for extracting histogram L^1_k but this time summing all the computed histograms L^i_k for $1 \le i \le M$ over the entire length of the part axis P_k. More formally, for a part axis P_k, whose locality is defined by a region U_k, its segment histogram S_k and its associated normalisation value W_k are given by

$$S_k = \sum_{i=1}^{M} L^i_k \quad \text{and} \quad W_k = \sum_{i=1}^{M} C^i_k \quad \text{and} \quad U_k = \bigcup_{i=1}^{M} C^i_k \tag{5}$$

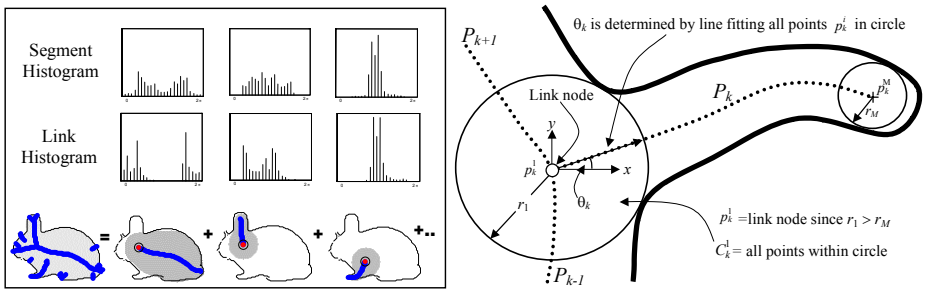

Fig. 2. (a) Decomposition of a rabbit shape into subparts based of its extracted shape axes. The gray-shaded area in each part shows the locality U_k, in which the segment and link histograms are computed. The size of this locality gives a good indication of the part's saliency. The small circle marking one end of each of the part axis is a link node. (b) Part axis P_k and relevant parameters for link node determination, histogram normalization and angular compensation for rotation invariance (see text for details).

To have the ability to simultaneously classify gross shape forms and take boundary noise and texture into consideration, some form of feature saliency must be incorporated into the shape similarity computation. Mismatch of a more salient part (e.g. rabbit's body) should carry a higher penalty then one that is small and perceptually insignificant, such as boundary noise or a small protrusion (e.g. rabbit's hind feet). In this context, a salient part is defined as one that covers a large region of the overall shape (see Fig. 2a). A useful measure that relates to this definition of saliency is the value W_k given in (5). Therefore, if a shape A has Q_A part axes, then the saliency $R_{A,k} \in (0, 1]$ of the part axes $P_{A,k}$ is given by

$$R_{A,k} = W_k \Big/ \sum_{i=1}^{Q_A} W_i \tag{6}$$

4 Matching Parts and Shapes

It is widely believed that in human perception, the triangular inequality does not hold in shape similarity comparison [3], especially in situations where a shape is a composite of several different shapes or when parts are occluded. Consequently, the distance between two shapes should not be evaluated based on just a metric distance. But from a computational viewpoint, a metric-based distance allow part-based descriptors to be organized into a database and retrieved through multi-dimensional index structures such as M-trees [6], for example. Similar to [3], we propose reconciling the two contrasting requirements with the use of two distinct measures called *part distance* and *shape distance*. The part distance is a metric distance used to measure similarity between two parts. On the other hand, the shape distance is a non-metric distance defined as the optimal combination of part distances between two shapes and it gives a global shape similarity measure that is more consistent with human perception.

The part distance $d_p(P_{A,k}, P_{B,k})$ between two part axes $P_{A,k}$ and $P_{B,k}$ extracted from shapes A and B is given by the combined χ^2 distance between their respective n-bin link and segment histograms and is given by

$$d_p(P_{A,k}, P_{B,k}) = \frac{\beta}{2}\sum_{i=1}^{n}\frac{[L^1_{A,k}(i) - L^1_{B,k}(i)]^2}{L^1_{A,k}(i) + L^1_{B,k}(i)} + \frac{(1-\beta)}{2}\sum_{i=1}^{n}\frac{[S_{A,k}(i) - S_{B,k}(i)]^2}{S_{A,k}(i) + S_{B,k}(i)} \quad \textbf{(7)}$$

where $d_p(P_{A,k}, P_{B,k}) \in [0,1]$ and the parameter $\beta \in [0,1]$, determines the relative importance attached to the matching of the link and segment histograms. A value of $\beta = 0.5$ will give equal emphasis to both when computing the matching cost. Using (7), the part distance matrix PM_{AB} of size $Q_A \times Q_B$ is obtained by matching all part axes in shape A to those in shape B and we assume $Q_A \le Q_B$ (if not, shapes A and B are reversed). The matching costs in matrix PM_{AB} is then weighted by the part saliency of shape A as defined in (6), to yield a weighted part distance matrix WM_{AB} given by

$$WM_{AB}(i,j) = PM_{AB}(i,j)R_{A,i} \qquad \text{for } 1 \le i \le Q_A \text{ and } 1 \le j \le Q_B \quad \textbf{(8)}$$

The shape distance between two shapes A and B, where $Q_A \le Q_B$, is the optimal saliency weighted match between all parts in A and a subset of parts in B that results in the lowest total matching cost. Given the weighted part distance matrix WM_{AB}, we can determine the optimal match by minimizing the total cost of matching subjected to the constraint that a one-to-one match exist. Since $Q_A \le Q_B$, we can realize a one-to-one assignment by adding unit dummy costs to WM_{AB} to turn it into a square matrix. With a square WM_{AB} of size $Q_B \times Q_B$, we have essentially reduced this task to a matching problem for bipartite graphs and this can be solved in $O(Q_B{}^3)$ time by the Hungarian method [1], [7]. The resulting output from the Hungarian method is a permutation of $(P_{A,i}, P_{B,j})$ pairs that results in the lowest overall matching cost $d_h(A,B)$. The shape distance $d_s(A,B)$ between shapes A and B is then obtained by removing the added unit dummy costs and is given by $d_s(A,B) = d_h(A,B) - (Q_B - Q_A)$.

In order to handle reflection, a mirrored version of shape A must also be matched to shape B to see which results in a lower shape distance. To obtain a mirrored version of shape A, the link and segment histograms of all its part axes $P_{A,k}$ are reversed. This means $L^1_{A,k}(i)$ and $S_{A,k}(i)$ becomes $L^1_{A,k}(n-i+1)$ and $S_{A,k}(n-i+1)$ respectively, where n is the number of bins in the histograms. Choosing the lower of the two shape distances will ensure invariance to the mirror transform.

5 Experimental Results

5.1 Geometric Invariance

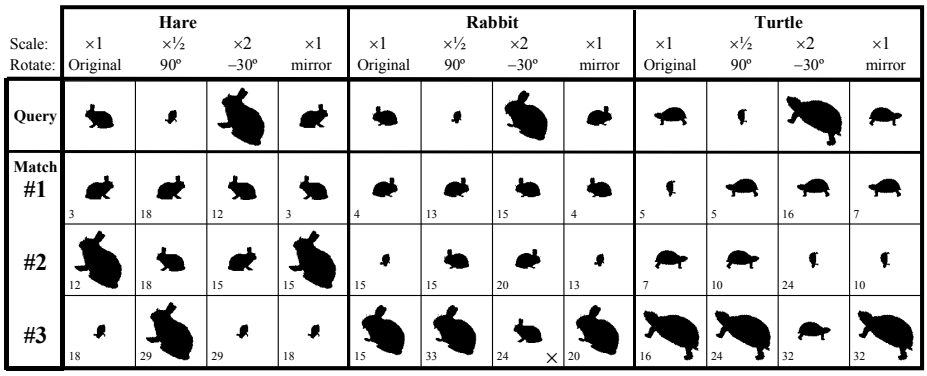

Fig. 3. Results demonstrating invariance to scaling, rotation and mirror reflection with shape distances (×1000) indicated at the bottom right corner of each image.

Fig. 3 shows the closest three matches for each of the twelve 257×257 pixel-sized images comprising of three basic shapes (i.e. hare, rabbit and turtle) in various orientations and scales (up to two octaves) when each was matched to every other. Except for the image marked (×), all shapes were correctly matches to their respective shape type. Errors in the gradient vector field histograms are more significant when the shape is small due to the reduced number of sampling points making up the histograms. It is not surprising that within the best 3 matches, the largest-scaled rabbit was not matched to the smallest-scaled rabbit but to a scale (×1) hare because the shape distance between the hare and rabbit is small. Note: For all experimental results presented, the following parameter settings were adotped: σ_H =1.0, σ_D = 1.5, α = 0.8, β = 0.5, n = 24,

5.2 Robustness to Boundary Noise and Texture

Fig. 4 shows the shape axes can be extracted relatively consistently from the multi-resolution gradient vector field despite high levels of boundary noise. Skeletons extracted from the popular Euclidean distance transform [9], [19] on the other hand, give rise to many spurious medial branches. In graph-based shape descriptors like those of [12],[15],[17], complicated medial branch pruning and edit distance algorithms must be employed to perform shape matching.

Fig. 5 shows the three closest matches for twelve 129×129 pixel-sized images comprising of four basic shapes (i.e. hare, rabbit, pear and candle) with three different boundary textures (i.e. smooth, wavy and spiky) when each was matched to every other. Only 3 mismatches between closely similar shapes (i.e. hare and rabbit) were observed. As seen in Fig. 2c, part axes of lower spatial frequency boundary textures such as large wavy bumps have higher saliency values and would therefore contribute more significantly to the shape matching score. As a result, the shape distance between a wavy-textured rabbit and a wavy-textured hare would be much lower than say a wavy-rabbit and a spiky-rabbit.

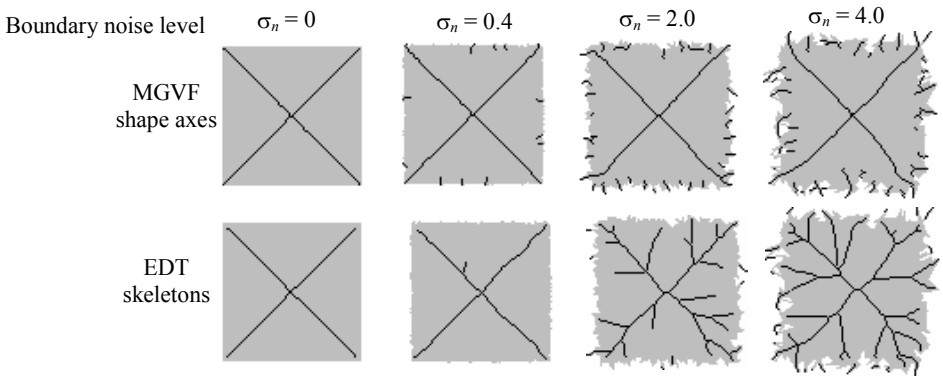

Fig. 4. Shape axes and skeletons extracted from the multi-resolution gradient vector field (MGVF) and Euclidean distance transform (EDT) when the boundary perturbed in the *x* and *y* direction by zero-mean Gaussian noise with varying standard deviation σ_n.

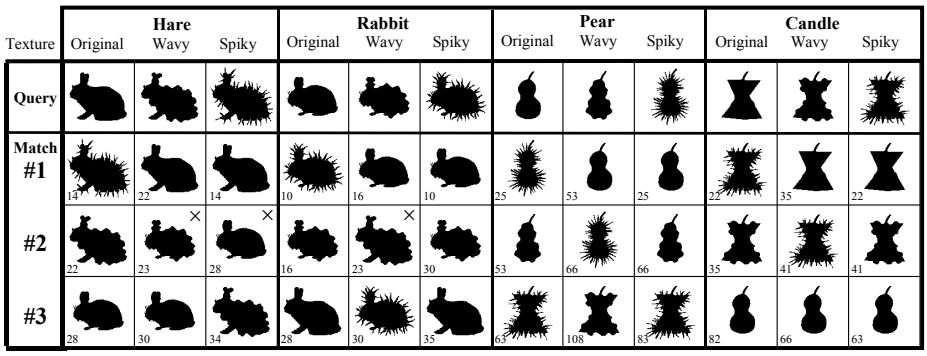

Fig. 5. Matching results obtained when comparing shapes with boundary texture.

5.3 Shape-based Image Retrieval

In this experiment, an image database of 35 silhouettes images forming 7 different classes was queried using 7 test images specially designed to demonstrate various features of the proposed shape descriptors. (*HS*,*EP*,*CP*,*WS*) demonstrates the ability to match similar shape classes despite boundary noise. (*CR*,*CD*) show the ability to handle some amount of occlusion. (*HS*,*GD*) show the descriptor's ability to handle limb articulation because the gradient vector field histograms are locally re-oriented along the part axes. (*GD*) demonstrates the ability to do partial matching using a non-metric shape distance measure. By relaxing triangular inequality [3], it is possible to retrieve multiple simple shapes that make up a composite shape. Notice the dude image retrieved in match #4 is the closest matching dude in the database after accounting for mirror-reflection, occlusion and articulation of its limbs.

Fig. 6c shows two image retrieval results using the shape context descriptor [1]. The shape context performs well with when the boundary perturbation is small and isolated, like with (*CP*) but performs poorly with significant boundary texture, like with (*EP*). The problem of assigning boundary points under such circumstances is non-trivial unless a significant amount of boundary smoothing is applied with the potential risk of merging salient parts that are in close spatial proximity.

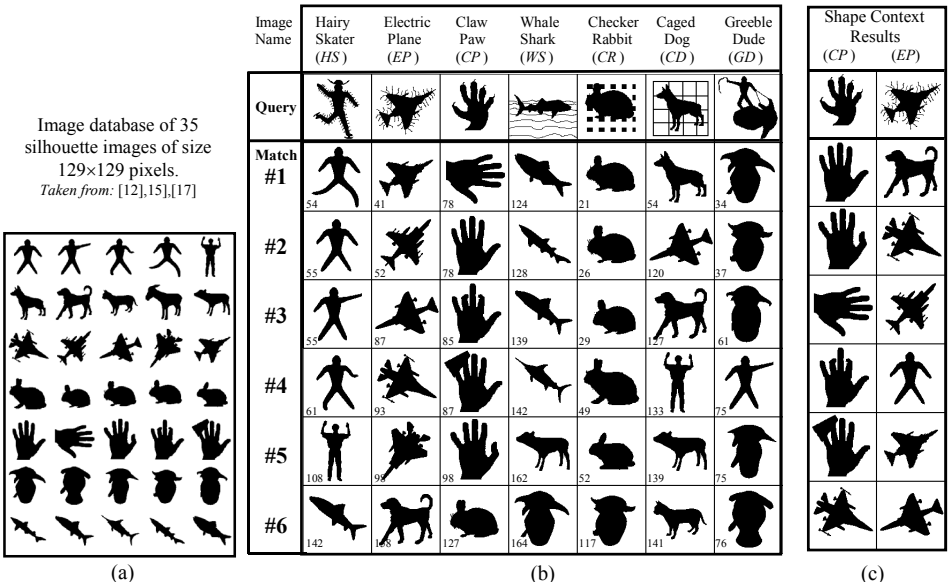

Fig. 6. (a) The image database. (b) Six best ordered retrievals for each of 7 queries using the MGVF shape descriptor. (c) Six best ordered retrievals for each of 2 queries using the shape context descriptor of Belongie et al. [1], with 5 bins for $log(r)$ over the range of 0.125α to 2α and 12 equally-spaced radial bins. The matching cost function is given by $\Sigma_i C_{i,\pi(i)}$ and is computed over 100 boundary points. The results were obtained with the Matlab™ code of Serge Belongie at *http://www.cs.berkeley.edu/~sjb/shape/*.

6 Conclusions

We have presented a novel histogram-based shape descriptor that is derived from the gradient vector field of a shape. This part-based shape descriptor encapsulates the general shape form of an object, its boundary texture and their respective saliency measure. Experimental results show that it is able to discriminate shapes with various boundary textures and is also invariant to scaling, rotation and mirror reflection. Relatively robust shape-based image retrieval was obtained on a small image database of 35 silhouette images using query images that demonstrated the shape descriptor's ability to handle boundary noise, limb articulation, occlusion and partial matching. We also proposed a multi-resolution pyramidal technique for extracting the shape descriptor.

References

[1] Belongie, S., Malik, J., Puzicha, J.*, Shape Matching and Object Recognition using Shape Contexts*. IEEE Trans. On PAMI, Vol. 24, No. 24 (2002) 509-522

[2] Ben-Arie, J., Wang, W., *Shape Description and Invariant Recognition Employing Connectionist Approach*. Intl. J. of Patt. Recog. and AI, Vol. 16, No.1 (2002) 69-83

[3] Berretti, S., Bimbo, A.D., Pala, P., *Retrieval by Shape Similarity with Perceptual Distance and Effective Indexing*. IEEE Trans. On Multimedia, Vol. 2, No. 4 (2000) 225-239

[4] Blum, H., *A Transformation for Extracting New Descriptors of Shape*. In: Wathen-Dunn, W. (ed.): Models for the Perception of Speech and Visual Form, MIT Press, Cambridge, MA (1967) 363-380

[5] Burt, P.J., *The Pyramid as a Structure for Efficient Computation*. In: Rosenfeld, A. (ed.): Multiresolution Image Processing & Analysis. Springer-Verlag, Berlin Heidelberg New York (1984) 6-35

[6] Ciaccia, P., Patella, M., F, Zezula, P., *M-tree: An Efficient Access Method for similarity Search in Metric Spaces*. In Proc. Of Intl. Conf. on Very Large Databases (1997) 522-525

[7] Clark, J., Holton, D.A., *A First Look at Graph Theory*. World Scientific Pub. S'pore (1991)

[8] Cross, A.D.J., Hancock, E.R., *Scale Space Vector Fields for Symmetry Detection*. Image and Vision Computing, Vol. 17 (1999) 337-345

[9] Dimitrov, P., Phillips, C., Siddiqi, K., *Robust and Efficient Skeletal Graphs*. In Proc. CVPR, Vol. 1 (2000) 417-423

[10] Goh, W.B. Chan, K.Y., *Shape Description using Gradient Vector Field Histogram*. In Proc. Scale-Space 2003, (2003) 713-728

[11] Latecki, L.J., Lakamper, R., *Convexity Rule for Shape Decomposition based on Discrete Contour Evolution*. CVIU, Vol. 73, No.3 (1999) 441-454

[12] Liu, T., Geiger, D., *Approximate Tree Matching and Shape Similarity*. In Proc. International Conference on Computer Vision (1999) 456-462

[13] Ogniewicz, R.L., *Skeleton-Space: A Multiscale Shape Description combining Region and Boundary Information*. In Proc. CVPR, Seattle, WA. (1994) 746-751

[14] Mokhtarian, F., *Silhouette-based Isolated Object Recognition through Curvature Scale Space*. IEEE Trans. On PAMI, Vol. 17, No. 5 (1996) 539-544

[15] Sebastian, T.B., Klein, P.N., Kimia, B.B., *Recognition of Shapes by Editing Shock Graphs*. In Proc. International Conference on Computer Vision, (2001) 755-762

[16] Sebastian, T.B., Kimia, B.B., *Curves vs Skeletons in Object Recognition*. In Proc. International Conference on Image Processing, Vol. 3 (2001) 22-25

[17] Sharvit, D., Chan, J., Tak, H., Kimia, B., *Symmetry-based Indexing of Image Databases*. J. Visual Comm. and Image Rep., Vol. 9, No. 4 (1998) 366-380

[18] Siddiqi, K., Kimia, B.B., *Parts of Visual Form: Computational Aspects*. IEEE Trans. on Pattern Analysis & Machine Intelligence, Vol. 17, No. 3 (1995) 239-251

[19] Siddiqi, K.S., Bouix, S., Tannenbaum, A., Zucker, S.W., *Hamilton-Jacobi Skeletons*, IJVC, Vol. 48, No. 3 (2002) 215-231