Published in final edited form as:

Comput Vis Image Underst. 2010 July ; 114(7): 827–834. doi:10.1016/j.cviu.2010.03.009.

Contour based object detection using part bundles

ChengEn Lu^{a,c,d,*}, Nagesh Adluru^b, Haibin Ling^a, Guangxi Zhu^{c,d}, and Longin Jan Latecki^a

^aDept. of Computer and Information Science, Temple University, 324 Wachman Hall, 1805 N Broad St., Philadelphia, PA 19122, USA

^bBiotechnology Center, University of Wisconsin-Madison, T129 Waisman Center, 1500 Highland Ave, Madison, WI 53705, USA

^cDept. of Electronics and Information Engineering, Huazhong University of Science and Technology, Wuhan 430074, China

^dDiv Commun and Intelligent Networks, Wuhan National Laboratory for Optoelectronics, Wuhan 430074, China

Abstract

In this paper we propose a novel framework for contour based object detection from cluttered environments. Given a contour model for a class of objects, it is first decomposed into fragments hierarchically. Then, we group these fragments into *part bundles*, where a part bundle can contain overlapping fragments. Given a new image with set of edge fragments we develop an efficient voting method using local shape similarity between part bundles and edge fragments that generates high quality candidate part configurations. We then use global shape similarity between the part configurations and the model contour to find optimal configuration. Furthermore, we show that appearance information can be used for improving detection for objects with distinctive texture when model contour does not sufficiently capture deformation of the objects.

Keywords

Part bundle; Shape context; Object detection

1. Background

The key role of contours and their shapes in object extraction and recognition in images is well established in computer vision and in visual perception. Extracting edges in digital images is relatively well-understood and there are robust detectors like [21,9]. However, it is often difficult to distinguish edge pixels belonging to meaningful object contours. The main challenge is due to the fact that most edge pixels represent background and irrelevant texture, while only a small subset of edge pixels corresponds to object contours. Further, the edge pixels do not simply form occluding contours but broken contour fragments due to noise and occlusion.

*Corresponding author at: Div Commun and Intelligent Networks, Wuhan National Laboratory for Optoelectronics, Wuhan, 430074, China.

Grouping edge pixels into contours using various saliency measures and cues has been studied for long and is still an active research field [37,30,28,15,27]. Once contours are identified they can be further grouped into objects by performing shape matching with model contours. For example, Shotton et al. [26] and Opelt et al. [24] use chamfer distance [2] to match fragments of contours learnt from training images to edge images. McNeil and Vijayakumar [23] represent parts learnt from semi-supervised training as point-sets and establish probabilistic point correspondences for the points in edge images. Ferrari et al. [6] use a network of nearly straight contour fragments and a sliding window search. Thayananthan et al. [29] modify shape context [1] to incorporate edge orientations and Viterbi optimization for better matching in clutter. Flezenswalb and Schwartz [5] present a shape-tree based elastic matching among two shapes and extend it to match a model and cluttered images by identifying contour parts (smooth curves) using a min-cover approach [4]. More recently, Zhu et al. [38] formulate the shape matching of contours (identified using [37]) in clutter as a set-set matching problem. They present an approximate solution to the hard combinatorial problem by using a voting scheme [33,19] and a relaxed context selection scheme using linear programming. They use shape context [1] as shape descriptor. Bai et al. [40] provide a skeleton-based model decomposition scheme to solve non-rigid object detection problem. Aside from the recent work mentioned above, there are many early studies that use geometric constraints for model-based object and shape matching [13,11,12].

In this paper we focus on object detection by grouping the edge fragments in an image according to a contour model. Like many previous studies [33,19,26,38,7], we follow a two-phase framework: (1) generating object hypotheses, and (2) picking the best one. Our main contributions are summarized below: For the first phase, we propose a novel *part bundle* model that allows us to use local shape similarity of contour parts to efficiently generate accurate hypotheses. Specifically, we decompose a given contour model into several part bundles, where each part bundle contains several overlapping contour fragments (see Fig. 1 for example decompositions). The model naturally handles hierarchical relations between fragments, as well as the AND/OR relations for co-existing and competing fragment pairs. By matching the contours composed of fragments in the part bundles to that of the edge fragments in a given image we can generate concise and accurate hypotheses for part-configurations. The part bundle model is closely related to the hierarchical Random Filed model [18] and the AND/OR graphs [14,39,3]. While our model avoids explicit graphtheoretic parsing it keeps all the necessary properties for inferring part-configurations.

In the second phase, we compute global shape similarity of the part-configurations to the model contour using shape context [1] to pick the best part-configuration. Global shape matching is an expensive operation but we can afford to do this thanks to the concise and accurate hypotheses generated in the first phase using local shape matching of edge fragments to the part bundles. The illustration of our framework is shown in Fig. 2.

Another contribution we made for the second phase is to investigate if appearance can compensate for the lack of model deformability for objects with distinctive texture. Using the bag-of-features [36,35,32,10] model, we show that the detection performance can be significantly improved for objects with distinctive texture.

The rest of the paper describes the details of our framework followed by experimental evaluation. In Section 2 we describe the first phase of hypotheses generation. Section 3 describes how we select the best hypothesis in the second phase. In Section 4 we present experimental evaluation of our method, followed by conclusions and discussions in Section 5.

2. Hypotheses generation using part bundles

We aim to find the objects in the image using a shape matching scheme. The shape matching algorithm by Ferrari [8], called *kAS*, groups *k* pieces of adjacent segments (AS) together and treats the adjacent segments as an individual unit and performs shape matching between them. Zhu et al. [38] perform many-to-many matching between sets of model and image segments in their process of contour selection.

We propose a simple scheme that only allows one-to-one mapping between a model segment and an image segment. Our part bundle based decomposition of the model contour allows us to perform the one-to-one matching. We first explain how a given contour model is hierarchically decomposed into part bundles and then explain how the edge fragments are obtained from a given image. Using the part bundle model and edge fragments we generate good candidate part configurations of edge fragments in the image.

2.1. Part bundles model for shape decomposition

To model an object class, we first get its contour representation by tracing a sample image as in [8,38]. As stated in the above subsection, we try to apply the same criteria to the shape model as that in the image segments to get fairly straight contour fragments. Since the model contour is a complete contour, we only apply *rule 2* (described below) to break the model contour to several contours at high curvature contour points. Other methods like discrete curve evolution (DCE) [17] also represent a shape contour as linear segments connected by several salient contour points. Salient points are the contour points with high curvature. Fig. 3 shows an example which use our *rule 2* to decompose a shape model to several contour pieces. Thus, both the image fragments and model segments are processed with the same criteria, which increases the probability of their one-to-one matching.

Moreover, we obtain a more flexible model structure by manual interference based on above decomposition. We manually decompose contour segments at some salient points and the center of a contour segment to form a hierarchical structure. Let $F = \{f_i\}_{i=1}^m$ denote the set of such fragments. A part bundle B is a subset of F , such that at most one fragment in B is allowed to have correspondence to an edge fragment in E . In other words, all fragments in a bundle B compete with each other during shape-to-image matching. With this bundle decomposition, we now can write our contour model as

$$\mathcal{B} = \{B_k\}_{k=1}^{m'} \quad (1)$$

where m' is the number of part bundles in model $\mathcal{B}$, which is usually very small. For bundles we have the following constraints

$$B_k \subset F, \quad k=1, \dots, m',$$

$$B_i \cap B_j = \emptyset, \quad 1 \leq i < j \leq m'. \quad (2)$$

Sample bundle decompositions are shown in Fig. 1, we fix them as the models for ETHZ dataset.¹ It can be seen that a bundle can have fragments representing overlapping parts thus having redundancy and hierarchy. A cognitive motivation behind such decomposition scheme is that an object can be recognized even if some parts of it are missing, as can be observed in Fig. 4. There are several reasons why parts of objects can be missing in real images: missing edge information, occlusion, failures in contour grouping. The main constraint for the bundle design is to ensure shape flexibility, i.e., to accommodate for possible deformations and broken edge fragments in the edge image. Our experimental results presented in Section 4 show the superiority of our part bundle scheme over the control point scheme.

2.2. Extracting edge fragments from an image

Given is an input image I and its edge map extracted by the edge detection algorithm [22].

We extract a set of edge fragments, $E = \{e_i\}_{i=1}^n$, from the edge map in the following way. First, single pixel wide edge chains are obtained using the edge linking method of Peter Kovesi [16]. Then each branching chain from a junction point is treated as an edge segment. An example can be found in Fig. 5b, where different color represents different edge segments. The connected edges are partitioned to several branches at junction points.

We observe that if edge segments are long and straight, they are more likely to correspond to a contour part of the model. Based on this observation, we further group Peter Kovesi's edge linker result to generate fairly long and straight edge segments. At this step, no shape prior is taken into consideration for contour grouping. We apply following two rules.

Rule 1 (Linking): Connect two edge segments if the distance between two of their endpoints is less than D_T and the corresponding turn angle is less than K_T , and repeat the process until there are no more edge segments to link. In cases where there are more than two edge segments at a junction, edge segments with smaller turn angle should be connected first.

Rule 2 (Decomposing): Split a given edge segment at a point with turn angle larger than K_T , and continue doing this till there is no more edge segments to split.

In our work, we set $D_T = 10$ pixels in most cases, and K_T is measured by curvature with a certain scale. K_T is typically set to 0.5. In Fig. 5c, it can be seen that that short edge fragments are linked to form longer edge segments at junction points following *rule 1*, and in

¹We design the part bundles based on the hand drawn models provided in the ETHZ dataset, but they are not optimal. For example, we observe that in the giraffe category, the necks of many giraffes have different bending angles. If we add another giraffe model to this category with a different neck, the performance of our method on that object category improves significantly.

Fig. 5d, all edge segments are decomposed to fairly straight segments. We apply our edge rules to all the test images in the ETHZ dataset.

To summarize, we obtain a set of edge fragments $E = \{e_i\}_{i=1}^n$ by applying Peter Kovese's edge linker [16] and our two rules. A sample set of resulting edge fragments can be seen in Figs. 5d and 2a.

2.3. Efficient centroid voting with part bundles

Given a set of model fragments F and edge fragments E , the task now is to find a candidate subset of E such that the global shape composed by it matches to that of $\mathcal{B}$. Matching F and E is a hard combinatorial matching problem. In [38], *many-to-many* matching is utilized and a context selection scheme is introduced. They solve a relaxed version (that utilizes additivity of shape context representation) instead of the discrete version with linear programming.

In contrast by taking advantage of the exclusion property of part bundles and our edge linking rules, we can model the mapping from $\mathcal{B}$ to E as a *one-to-one mapping*. Denote such a mapping as $\pi : \{1, \dots, m'\} \rightarrow \{0_1, 0_2, \dots, 0_{m'}, 1, \dots, n\}$, such that $B_i \in \mathcal{B}$ is mapped to edge fragment $e_{\pi(i)} \in E$ or empty edge when $\pi(i) = 0_j^2$. It can be seen that each such

mapping can be considered as a part configuration obtained by the $\{e_{\pi(i)}\}_{i=1}^{m'}$ in the image and our task is to find an optimal one $\hat{\pi}$ that maximizes global shape similarity between the part-configuration and the model contour. Searching among all configurations is very

expensive since there are $\prod_{i=0}^{m'-1} (n+m' - i) = O((n+m')^{m'}) \approx O(n^{m'})$ possible configurations.

As mentioned in the previous section, our part bundle generates a very concise model decomposition, with $m' \approx 4$ for most objects. Taking advantage of this fact, we design an efficient voting scheme for searching in the configuration space. This allows us to explore the configurations explicitly using *any* high-level shape matching tool.

The basic idea is to use fragment correspondences between F and E to vote for the object centroid. Specifically, for an image fragment $e \in E$ and a model fragment $f \in F$, we first build the mapping between the sample points. Using sequences of tangent directions along e and f as their shape descriptors, we perform sequence matching with the Smith–Waterman algorithm [34]. Then, using the part-centroid relation from the model, we can project the expected centroid of the object to the image. The scale transformation is estimated by the length ratio of the image fragment e and the model fragment f . For each $f \in F$, we rank all fragments in E according to their similarities to f as computed above. Then we choose the best two matches for centroid voting. Consequently, the whole voting procedure generates $2m$ centroid votes (see Fig. 2c and d).

The correspondences in the optimal configuration $(\hat{\pi}(1), \dots, \hat{\pi}(m))$ should consistently estimate the centroid of the object. Assuming there exists a centroid whose K -NN centroids

² 0_j denotes a dummy fragment that B_j maps to for handling missing parts.

correspond to the optimal fragment-bundle correspondences, we just need to examine part-configurations obtained from K -NN centroids around all $2m$ centroids (see Fig. 2e). A group

of K -NN centroids gives us $\left(\frac{K}{m'}\right)^{m'}$ part-configurations. So we have to examine a total of $2m\left(\frac{K}{m'}\right)^{m'}$ hypotheses or part configurations. We denote the set of these hypotheses as $\mathcal{H}$. The number of configurations to be explored is exponential in the number of part bundles

but usually $m' \approx 4$ for most of the models. Also since $m \ll n$, $2m\left(\frac{K}{m'}\right)^{m'} \ll O(n^{m'})$ as long as $K \ll nm'$. Theoretically, a larger K implies higher probability to get the optimal configuration, thus, the larger the value of K , the better the performance is. However, the search time increases significantly with the increasing of the value of K , and we observe that the performance improves only slightly when K is larger than 20. This is explained by the fact that we have a concise model structure and the improved edge linking method allows one-to-one mapping. We chose $K = 20$ empirically to strike a balance between performance and time complexity.

Note that the proposed searching approach uses a brute force scheme, which provides a more robust configuration estimation than traditional solutions that usually rely on mode seeking or relax the matching problem. Though more reliable, brute force solutions are usually forbidden due to their high computational complexity. The key factor that makes our solution computationally trackable is the part bundle model and the one-to-one mapping scheme. In particular, we benefit from the fact that the number of part bundles m' is small. This fact is a general property that is rooted in human visual perception, since humans decompose a given shape only in a small number of parts of visual form [25].

3. Hypotheses evaluation

3.1. Utilizing local and global shape information

Our part configurations are groups of image edge fragments and our models are based on contour parts. Hence it is natural to use shape constraints for evaluating the part configurations. Instead of just using geometric relationship between contour parts, we use global shape similarity to judge the quality of candidates.

We perform a coarse-to-fine shape matching to evaluate the hypotheses (part configurations). We first eliminate hypotheses $h_i \in \mathcal{H}$ whose cumulative part-to-fragment similarities measured using [34] is below a certain similarity threshold. Then for the remaining hypotheses we use Shape Context [1] which is a well known global shape descriptor and has demonstrated excellent performance for shape analysis tasks. We use it to estimate the similarity between a model and a hypothesis. The shape distance between them is defined by the their shape context distance [1] $d_{sc}(h_i, \mathcal{B})$. We then pick the best candidate according to d_{sc} :

$$\operatorname{argmin}_{h_i} d_{sc}(h_i, \mathcal{B}) \quad (3)$$

3.2. Combining shape and appearance

In addition to shape, appearance information often provides distinctive cues for object detection and recognition. Appearance can be helpful in cases of objects with distinctive texture and where the object model does not capture the variance in the shape. In this subsection, we describe our investigation along this line.

We use the popular bag-of-features [36,32] framework for appearance analysis. First, local patch descriptors are generated from training images using SIFT [20], which are used to build a dictionary $\mathcal{D}=\{W_i\}_{i=1}^M$ using k -means. Second, all local features in a window are mapped to visual words in $\mathcal{D}$, and the window is then represented by its word frequency vector. Third, we collect all such word frequency vectors to form a training set, and use support vector machine (SVM) [31] to learn a classifier.

For object class, Giraffes, the neck in the model does not capture the variance in the neck poses of the test images. Further, Giraffes have distinctive texture. For this class we evaluate the hypotheses by combining appearance and shape. Again we first eliminate part-configurations whose cumulative part-fragment similarities is below a certain threshold. Then for each of the remaining hypothesis h , we use the texture inside the smallest bounding box enclosing the fragments to get its visual word frequency representation. Then we measure its relevance to a class using the confidence of the learned SVM model, denoted as $d_{app}(h)$. Combining shape context and appearance similarity, we have the following weighted evaluation criterion for picking the best hypothesis:

$$\operatorname{argmin}_{h_i} d_{SC}(h_i, \mathcal{B}) + \lambda d_{app}(h_i) \quad (4)$$

where λ balances the weight between shape and appearance.

4. Experimental results

We present results on the ETHZ shape classes ([8]). This dataset has 5 different object categories with 255 images in total. Many objects are surrounded by extensive background clutter and have interior contours. All categories have significant intra-class variations, scale changes, and illumination changes. The advantage of our method is to be able to identify good part-configurations among clutter. The ability to handle deformations depends on the capacity of global shape matching module or in other words the hypothesis evaluation method.

4.1. Experiments using shape information

Similar to the experimental setup in [38], we use only the single hand-drawn models provided for each class. For each model we introduce breakpoints at high curvature contour points. The part bundles created from the contour models can be seen in Fig. 1. In addition to longer contour segments, we need to select shorter ones, since contour parts may be missing in edge images.

Since we compare our results to the state-of-the-art [38], we use Precision vs. Recall (P/R) curves for quantitative evaluation. To compare with the results of [7] that are evaluated by detection rate (DR) vs. false positive per image (FPPI), we translate their results into P/R values as in [38]. Fig. 6 shows P/R curves for the three methods: Contour Selection method of [38] in black, that of [7] in green³, and our method in red. Our approach performs better than method in [7] on three categories, outperforms the Contour Selection method [38] on three categories, and outperforms both methods on two categories (“applelogos” and “swans”).

We also compare the precision to [38,7] at the same recall values in Table 1. The precision/recall pairs of [38,7] are quoted from [38]. We use the same criterion for correct detection as in [7]. Sample successful detections, failure cases and false-positives are shown in Fig. 7.

On a PC with 2.3 GHz Pentium (D) CPU and 3 GB RAM, the average processing time of our method in MATLAB implementation is around 25 s per image.

4.2. Experiments using shape and appearance information

We now present results of using appearance information on the class of giraffes that have distinctive texture because of their skin. We pick the first 16 images from the class to obtain the code book $\mathcal{D}$ of 100 visual words. We only selected the SIFT features inside the ground truth bounding boxes. The detection is based on a nearest neighbor retrieval according to the distance d_{mix} with $\lambda = 8$.

Fig. 8 shows the P/R curves using only shape and combining shape and appearance. Fig. 9 shows the improvement of using appearance information on two sample images. As can be seen using appearance information can significantly improve the detection performance.

5. Conclusions

In this paper we study the problem of contour based object detection. A new contour model, part bundles, is proposed that decomposes contour fragments into subsets. We show that this decomposition efficiently groups competing fragments and therefore enables brute force search hypothesis voting, which generates high quality and concise candidate detection sets. With these candidate sets, we use global shape similarity for false positive pruning and achieved excellent performance on ETHZ dataset. In addition, we show that appearance information, which can be easily integrated into our approach, can help to achieve further improvement for objects with distinctive texture when contour model and shape matching do not sufficiently capture the deformability of objects. Integrating such additional cues is straightforward in our framework.

Acknowledgments

This work was supported in part by the NSF Grants IIS-0812118, BCS-0924164, IIS-0916624, AFOSR Grant FA9550-09-1-0207, National Natural Science Foundation of China 60903096, 60496315, 60802009 and National 863 plans projects 2008AA01Z204, 2009AA01Z205. N. Adluru is supported by Computation and Informatics in Biology and Medicine (CIBM) program and Morgridge Institute for Research at the UW – Madison.

³For interpretation of color in Figs. 1–3 and 5–9, the reader is referred to the web version of this article.

References

1. Belongie S, Malik J, Puzicha J. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2002; 24:522–705.
2. Borgefors G. Hierarchical chamfer matching: a parametric edge matching algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1988; 10(6):849–865.
3. Chen, Y., Zhu, L., Lin, C., Yuille, A., Zhang, H. Rapid inference on a novel and/or graph for object detection, segmentation and parsing. In: Platt, J.Koller, D.Singer, Y., Roweis, S., editors. *Advances in Neural Information Processing Systems*. Vol. 20. Cambridge, MA: MIT Press; 2008. p. 289-296.
4. Felzenszwalb, P., McAllester, D. A min-cover approach for finding salient curves. *CVPRW: Proceedings of the Conference on Computer Vision and Pattern Recognition Workshop*, IEEE Computer Society; 2006; Washington, DC, USA.
5. Felzenszwalb PF, Schwartz J. Hierarchical matching of deformable shapes. *CVPR*. 2007
6. Ferrari V, Fevrier L, Jurie F, Schmid C. Groups of adjacent contour segments for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2008; 30(1):36–51. [PubMed: 18000323]
7. Ferrari V, Jurie F, Schmid C. Accurate object detection with deformable shape models learnt from images. *CVPR*. 2007:1–8.
8. Ferrari V, Tuytelaars T, Gool LJV. Object detection by contour segment networks. *ECCV*. 2006:14–28.
9. Galun M, Basri R, Brandt A. Multiscale edge detection and fiber enhancement using differences of oriented means. *ICCV '07: Proceedings of the Eleventh IEEE International Conference on Computer Vision*, IEEE Computer Society. 2007
10. Grauman, K., Darrell, T. The pyramid match kernel: discriminative classification with sets of image features; *IEEE Int'l Conference on Computer Vision II*; 2005. p. 1458-1465.
11. Grimson WEL. The combinatorics of object recognition in cluttered environments using constrained search. *Artificial Intelligence*. 1990; 44(1–2):121–165.
12. Grimson, WEL. *Object Recognition by Computer: The Role of Geometric Constraints*. Cambridge, MA, USA: MIT Press; 1990.
13. Grimson WEL, Lozano-Pérez T. Localizing overlapping parts by searching the interpretation tree. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1987; 9(4):469–482. [PubMed: 21869405]
14. Han, F., Zhu, S-C. Bottom-up/top-down image parsing by attribute graph grammar. *Proceedings of the Tenth IEEE International Conference on Computer Vision*, IEEE Computer Society; 2005; Washington, DC, USA. p. 1778-1785.
15. Hoiem, D., Stein, A., Efros, AA., Hebert, M. Recovering occlusion boundaries from a single image; *International Conference on Computer Vision (ICCV)*; 2007 Oct.
16. Kovesi, PD. MATLAB and Octave functions for computer vision and image processing, School of Computer Science & Software Engineering. The University of Western Australia; 2008. <<http://www.csse.uwa.edu.au/~pk/research/matlabfns/>>
17. Latecki LJ, Lakamper R. Convexity rule for shape decomposition based on discrete contour evolution. *CVIU*. 1999; 73:441–454.
18. Latecki, LJ., Lu, C., Sobel, M., Bai, X. Multiscale random fields with application to contour grouping; *Neural Information Processing Systems Conference (NIPS)*, Vancouver; 2008 Dec.
19. Leibe, B., Seemann, E., Schiele, B. Pedestrian detection in crowded scenes. *Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, IEEE Computer Society; 2005; Washington, DC, USA. p. 878-885.
20. Lowe DG. Object recognition from local scale-invariant features. *ICCV*. 1999
21. Martin D, Fowlkes C, Tal D, Malik J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. *Proceedings of the 8th Int'l Conference on Computer Vision*. 2001 Jul.2:416–423.

22. Martin D, Fowlkes C, Malik J. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 2001; 26:530–549.
23. McNeill G, Vijayakumar S. Part-based probabilistic point matching using equivalence constraints. *Neural Information Processing Systems*. 2006:969–976.
24. Opelt, A., Pinz, A., Zisserman, A. A boundary-fragment-model for object detection; *Proceedings of the European Conference on Computer Vision*; 2006.
25. Siddiqi K, Kimia BB. Parts of visual form: computational aspects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 1995; 17:239–251.
26. Shotton, J., Blake, A., Cipolla, R. Contour-based learning for object detection. *ICCV '05: Proceedings of the Tenth IEEE International Conference on Computer Vision (ICCV'05)*, vol. 1, IEEE Computer Society; 2005; Washington, DC, USA. p. 503-510.
27. Stein, A., Hoiem, D., Hebert, M. Learning to find object boundaries using motion cues; *IEEE International Conference on Computer Vision (ICCV)*; 2007 Oct.
28. Tamrakar, A., Kimia, BB. No grouping left behind: from edges to curve fragments; *ICCV '07: Proceedings of the Eleventh IEEE International Conference on Computer Vision*, IEEE Computer Society; 2007.
29. Thayananthan A, Stenger B, Torr PHS, Cipolla R. Shape context and chamfer matching in cluttered scenes. *CVPR*. 2003:127–133.
30. Trinh N, Kimia BB. A symmetry-based generative model for shape. *ICCV '07: Proceedings of the Eleventh IEEE International Conference on Computer Vision*, IEEE Computer Society. 2007
31. Vapnik, V. *The Nature of Statistical Learning Theory*. New York: Springer; 1995.
32. Varma M, Zisserman A. A statistical approach to texture classification from single images. *IJCV*. 2005; 62(1–2):61–81.
33. Wang L, Shi J, Song G, Shen I. Object detection combining recognition and segmentation. *ACCV07*. 2007:189–199.
34. Waterman ST. Identification of common molecular subsequences. *Journal of Molecular Biology*. 1981; 147:195–197. [PubMed: 7265238]
35. Willamowski, J., Arregui, D., Csurka, G., Dance, C., Fan, L. Categorizing nine visual classes using local appearance descriptors; *ICPR Workshop Learning for Adaptable Visual Systems*; 2004.
36. Zhang J, Marszalek M, Lazebnik S, Schmid C. Local features and kernels for classification of texture and object categories: a comprehensive study. *International Journal of Computer Vision*. 2007; 73(2):213–238.
37. Zhu, Q., Song, G., Shi, J. Untangling cycles for contour grouping; *ICCV '07: Proceedings of the Eleventh IEEE International Conference on Computer Vision*, IEEE Computer Society; 2007.
38. Zhu Q, Wang L, Wu Y, Shi J. Contour context selection for object detection: a set-to-set contour matching approach. *ECCV08*. 2008
39. Zhu S-C, Mumford D. *A Stochastic Grammar of Images*. Now Publishers. 2007
40. Bai X, Wang X, Latecki LJ, Liu W, Tu Z. Active Skeleton for Non-rigid Object Detection *ICCV*. 2009

(a) a bottle model

(b) a giraffe model

(c) an apple logo model

(d) a mug model

(e) a swan model

Fig. 1.

The part bundles of the hand-drawn contour models. Different parts and the part bundles are shown in different colors. (a) a bottle model (b) a giraffe model (c) an apple logo model (d) a mug model (e) a swan model.

Fig. 2.

Illustration of our framework: (a) The edge fragments are generated using edge linking software [16] (see Section 2.2). (b) Contour based model is decomposed into part bundles (see Section 2.1). (c) Contour fragments similarity, the color varies from blue to red indicates similarity from weak to strong. (d) Using contour fragment similarity we perform efficient centroid voting (see Section 2.3). (e) and (f) The part configurations generated using K -NN centroids around each centroid are evaluated using shape and appearance information and the best part-configuration is picked (see Sections 2.3 and 3), the black dots are the centroids estimated by mapping a model fragment to a image fragment.

Fig. 3.
A shape model is decomposed to several pieces at salient contour points using *rule 2* of our edge linking criteria.

Fig. 4.
Parts of the objects are missed both due to missing edge information and due to broken edge links. Objects can still be recognized using shape even if some parts are missing.

Fig. 5.

Decomposing and linking of edge segments in a real image with our two rules described in the text.

Fig. 6. Precision/Recall curves of the Contour Selection method by Zhu et al. [38] (black) and the method by Ferrari et al. [7] (green) on ETHZ shape classes. Our method (red) performs significantly better on two categories viz. Applelogos and Swans.

Fig. 7. Sample detection results. The edge map is overlaid on the image in white. The detected fragments are shown in black. The corresponding model parts are shown in top left corners. The bottom most row (enclosed in red) shows some failure cases and false positives.

Fig. 8.

The Precision/Recall curves using shape based evaluation (d_{sc}) and combining shape and appearance (d_{mix}) on the class of giraffes. We also show the results of [38] for comparison. The combination of shape and appearance significantly improves the results.

Fig. 9.

Two examples to show the improvement obtained by combining appearance information.
 Top row: Results of using only shape. Bottom row: Results of combining appearance with shape.

Table 1

Precision/recall pairs of the method of [7] and the Contour Selection method of [38] are quoted from [38]. Our precision is higher than [7] on four classes (except for “mugs”) and higher than [38] on three classes (except for “bottles” and “giraffes”).

Class (Recall) (%)	[7] (%)	[38] (%)	Our method (%)
Applelogos (86.4)	32.6	49.3	56.0
Bottles (92.7)	33.3	65.4	48.0
Mugs (83.4)	40.9	25.7	27.8
Giraffes (70.3)	43.9	69.3	50.0
Swans (93.9)	23.3	31.3	59.0