Shape Decomposition: Combining Minima Rule, Short-Cut Rule and Convexity

Anonymous CVPR submission

Paper ID 1204

Abstract

The aim of this work is to decompose shapes into parts which are consistent to human perception. We propose a novel shape decomposition method which utilizes the three perception rules suggested by psychology study: the Minima rule, the Short-cut rule and the convexity rule. Unlike the previous work, we focus on improving the convexity of the decomposed parts while minimizing the cut length as much as possible. The problem is formulated as a combinatorial optimization problem and solved by quadratic programming method. We test our approach on the MPEG-7 shape dataset, and the comparison results to previous work show that the proposed method can improve the part convexity while keeping the cuts short, and the decomposition is more consistent with human perception.

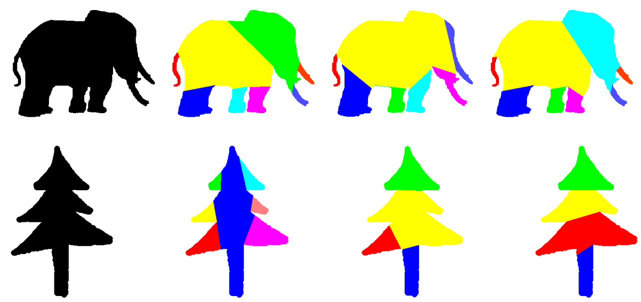

(a) (b) (c) (d)

Figure 1. Given the original shapes (a), comparison of our decomposition results (d) to the previous work [9] (b) and [3] (c).

1. Introduction

Part-based shape representation is a popular representation for objects in the community of computer vision [2]. Cognitive and psychological researches [1, 4, 15] have shown that there are several advantages of the part-based shape representation. First of all, once a whole shape is factorized into conditional independent components, it is computationally more flexible to account for shape deformation and articulation. Second, as occlusion is a common phenomenon, whereas decomposing an object into parts increases the chance of reliably detecting and recognizing the object if some of its "characteristic" parts are not occluded. Therefore, the part-based shape representation has received increasing attention.

In order to obtain a part-based shape model, the first question to answer is how to generate the shape parts. Most previous work on generating shape parts can be classified into two strategies: one is "bottom up" strategy which is grouping small shape elements into large shape parts [11], and the other is "top down" strategy which is shape decomposition. For the former strategy, people use bottom-up grouping method to learn parts as hierarchical shape vocabulary for object representation using shape fragments from a large number of shape instances, e.g. [11, 18]. This type of approaches consider the joint statistics between the object and curve fragments at different levels of hierarchies, whereas ignore shape perceptual properties such as convexity and cut length at all. As for the latter strategy, many people have studied the problem of shape decomposition which is partitioning one single shape into several parts under some generic constraints, such as the Minima Rule [4], the Short-cut Rule [16] and the Convexity [6, 17]. With different constraints, the same shape can have different partitions. The Minima Rule considers the curvature of the shape boundary curve or surface, and enforces that the shape is divided at places where the curvature is local minimum. It reflects a local constraint for shape decomposition. The Short-cut rule takes the cut length as a constraint and optimizes the decomposition by minimizing the total cut length. This is motivated by that human vision prefers to use the shortest possible cuts to parse shapes. Besides these two rules, convexity is also an important perceptual clue to determine visual parts [6]. These three rules share a common property which is "simplicity". Human perception can partition a shape into parts very easily and quickly. Thus it must make use of simple features and rules to make the quick decision. Based on these constraints, computer vision scientists have developed various optimization algorithms for solving the shape decomposition problem. For example, based on convexity, Gopalan *et al.* [3] proposed an algorithm to do approximate convex decomposition. Liu *et al.* [9] considered both convexity and the Short-Cut Rule in order to optimize the shape decomposition. Based on their work, Ren *et al.* [12] further encoded the Minimum Rule as well as the number of shape parts into the objective function during optimization and proposed a new optimization algorithm.

Besides these two classes, there are some previous work using symmetry [14] and Relatability [10] to find shape parts.

Although each method proposed its own objective function to optimize for shape decomposition, few of them can justify their goals quantitatively by experimental results. In other words, to our knowledge, there lacks of evaluation methods for shape decomposition in the literature. Most of these work just show the shape decomposition results by partitioned objects and leave the judgment to the readers. This fact makes the comparison of different methods difficult and therefore the claimed "improvement" is also subjective and vague.

In this paper, we focus on studying shape decomposition using generic rules. Each of the three rules captures certain aspects of part decomposition. The Minima rule reflects the local "salient" feature of the part junctions, the Short-cut rule indeed corresponds to the "compactness" of the desired part-based representation and convexity is easy to extract and identify for shape perception [5]. Considering these aspects, we propose a novel shape decomposition method by jointly considering the three generic rules to find an optimal shape decomposition. The way we formulate the problem is different from previous work, including [3,9,12]. Both [9] and [3] only applied one or two rules. Although Ren *et al.* [12] also considered all three rules, their focus is to minimize the number of parts which has overlap with the Short-Cut Rule and optimize the "visual naturalness" which is different from our goal. In addition, they did not provide a rigorous way to define and evaluate the "visual naturalness". Compared to previous work, we focus on improving the convexity of shape parts. Meanwhile, we also consider the Minima Rule and the Short-cut Rule. Specifically, we encode the contribution of each candidate cut to improving the part convexity into the objective function to optimize instead of requiring each part's convexity above a threshold as [9, 12] did.

Another contribution of the paper is that we propose a quantitative evaluation method to compare the decomposition results from different methods and therefore make the judgment clear. Specifically, we use the inner distance to measure the convexity and design a metric to measure the contribution of each candidate cut to improving the part convexity.

The rest of the paper is organized as follows. In Section 2, the shape decomposition problem is formulated as an optimization problem. Section 3 reviews the preliminary work which is related to our method. Section 4 introduces our approach and Section 5 shows the experimental results. Finally, Section 6 concludes the paper.

2. Problem Formulation

Given a planar shape S, a partition of S is defined as

$$S = \bigcup_i P_i, s.t., \forall P_i, P_i \in S; \forall i, j, P_i \bigcap P_j = \emptyset, \quad (1)$$

which means that S is composed of several parts $\{P_i\}$ and these parts do not overlap each other. On the other hand, a partition of S is associated with a set of cuts $\{C_j\} = \{\overline{p_{j1}p_{j2}}\}$ where each cut is a line segment $\overline{p_{j1}p_{j2}}$ and both points p_{j1} and p_{j2} lie on the boundary of S. These cuts are expected not to intersect each other. The boundary of each part P_i is composed of the boundary of S and a subset of $\{C_j\}$.

Based on the above notations, we can explain the three rules as follows:

- The Minima Rule [4] suggests that the cut points $\{p_{j1}, p_{j2}\}$ are located at the points where the curvature is local minimum.
- The Short-cut Rule [16] suggests to minimize the total length of the cuts, i.e, $\min \sum_j L(C_j)$ where $L(C_j)$ denotes the length of cut C_j which is usually computed as the Euclidean distance between points p_{j1} and p_{j2}.
- The Convexity Rule [6, 17] suggests to maximize the convexity of parts or minimize the concavity of parts. Based on this rule, Rosin [13] proposed a weighted average convexity to evaluate the decomposition quality as follows,

  $$Convexity(\{P_i\}) = \sum_i \frac{A_i}{A} Convexity(P_i). \quad (2)$$

  where A_i denotes the area of P_i and A is the area of shape S or $A = \sum_i A_i$.

Considering the three rules, the goal of shape decomposition can be formulated as

$$\min(\sum_j L(C_j) + \sum_i \frac{A_i}{A} Concavity(P_i)), \quad (3)$$

$s.t. \{p_{j1}, p_{j2}\}$ have local minimal curvatures and cuts do not intersect each other.

3. Preliminary Work

Given the above problem formulation, the remaining questions include : (i) how to measure $Concavity(P_i)$? (2)

how to encourage the cut points $\{p_{j1}, p_{j2}\}$ to have the local minimal curvatures? (3) how to optimize the objective function based on the measurements? We will introduce the preliminary work related to these questions in the following sections.

3.1. Convexity/Concavity Measurement

There are several choices for the measurement of convexity. One classical definition of convexity is the ratio of the area of the part to the area of its convex hull [13]. But this convexity measure is criticized due to its insensitivity to deep (but thin) protrusions of the boundary because it is area-based.

3.1.1 Inner Distance

Recently, the inner distance (ID) is becoming a popular measure for convexity [3, 8] because it is not sensitive the deep protrusions. Therefore, we define the convexity of a part P_i as the minimal ratio of the ED over ID of a pair of points within this part, i.e.,

$$Convexity(P_i) = min_{p,q \in P_i} ED(p,q)/ID(p,q). \quad (4)$$

If P_i is a convex, then given any point pair p and q, the inner distance is always equal to the Euclidean distance. Therefore, $Convexity(P_i) = 1$. On the other hand, if P_i is not convex, there must be a pair of points p and q such that their ED is less than ID. Therefore, $Convexity(P_i) < 1$. With this convexity measure, we can define the concavity measure as

$$Concavity(P_i) = 1 - Convexity(P_i). \quad (5)$$

3.1.2 Morse Function

Besides the above concavity measure, Liu *et al.* proposed a new measure [9] which is defined as follows:

$$concave(P_i) = \max_{p,q \in P_i} concave(p,q), \quad (6)$$
$$concave(p,q) = \min_{R \in \mathbf{R}(p,q)} \max_f \max_{t \in R} g_f(t), \quad (7)$$

where p and q are both points within part P_i, R denotes a path connecting p and q, $\mathbf{R}(p,q)$ denotes all paths connecting p and q, f denotes a Morse function corresponding to a projection function where the part P_i is projected, t denotes a point on the path R, and $g_f(t)$ represents a distance function between t and (p,q). Suppose $f(p) \geq f(q)$, it is defined as follows:

$$g_f(t) = \begin{cases} f(t) - f(p), & f(t) > f(p); \\ 0, & f(p) \geq f(t) \geq f(q); \\ f(q) - f(t), & f(t) < f(q). \end{cases} \quad (8)$$

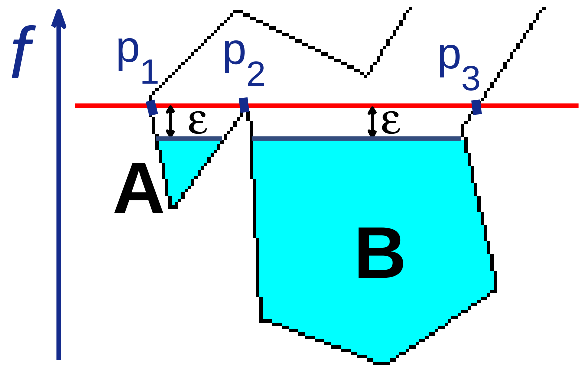

Figure 2. Illustration of a mutex pair of regions (A and B) and a candidate cut (red line). f is the Morse function.

By this measure, for each point pair (p,q), the concavity is defined by a path which can minimize the maximal perpendicular distance between the line passing (p,q) and the projected contour points between p and q w.r.t. all Morse functions. Although this definition is different from Eqn. 7, we can prove that the two definitions are inherently consistent for shapes without holes:

Theorem : Given a shape S without holes, for any point pair $(p,q) \in S$, the path R which corresponds to the inner distance between p and q is also the path which can minimize $\max_f \max_{t \in R} g_f(t)$ as defined in Eqn. 7.
The proof is provided in the supplemental material.

3.2. Mutex Pair and Candidate Cuts

The shape decomposition problem can be also viewed as a selection problem. Since there are infinite cuts inside a shape, the goal is to select a subset of cuts which can optimize an objective function. How to propose qualified candidate cuts is a challenging problem. Liu *et al.* [9] proposed a way to generate candidate cuts. The idea is to find pairs of components which cannot be kept together otherwise the concavity of the part containing both components will be high. Each pair of such components is defined as a "mutex pair" of regions. Specifically, it is a pair of regions A and B with

$$m(A,B) = \min_{p \in A, q \in B} concave(p,q) \quad (9)$$

above a threshold ϵ (See Figure 2). Given a fixed threshold ϵ, let MP denote all the mutex pairs to be separated and $|MP| = n_{mp}$. The motivation of generating candidate cuts is to separate these mutex pairs of regions. Figure 3 (a) shows an example of generated candidate cuts by this method. Due to the limited number of Morse functions being sampled (16 directions here), the "best" cut which can separate the left piece is missing.

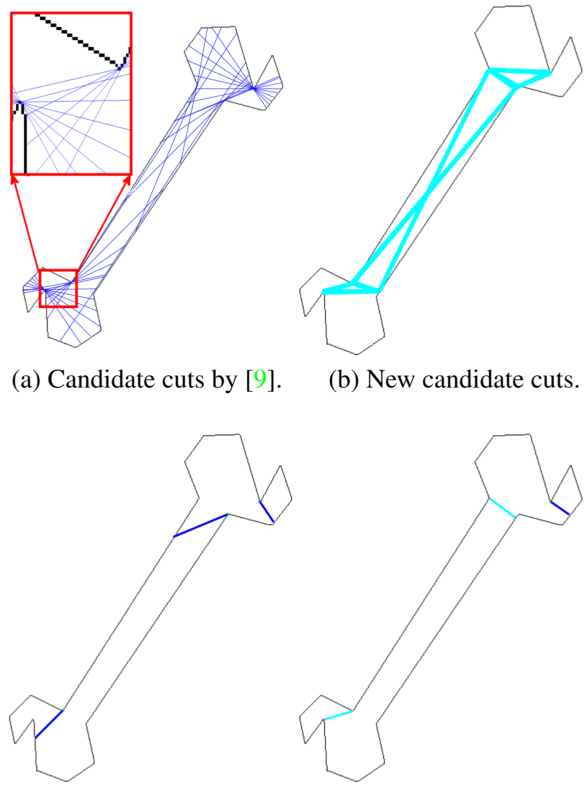

(a) Candidate cuts by [9]. (b) New candidate cuts.

(c) Shape decomposition by [9]. (d) Our result.

Figure 3. Comparisons of the candidate cuts by previous work and our new cuts as well as the resulted shape decompositions.

4. Our Approach

4.1. New Candidate Cuts

In Liu *et al.* [9]'s work, they only consider the Convexity Rule to propose a set of candidate cuts $\mathbf{C}_L$. However, the Minima Rule can also help to propose useful candidate cuts. We add a new set of candidate cuts $\mathbf{C}_M$ such that both cut points of each new cut have local minimum curvatures. Figure 3 (b) shows a set of new cuts generated by this rule. By combing these two sets of proposed candidate cuts $\mathbf{C}_L$ and $\mathbf{C}_M$, it can be seen that the set of candidate cuts $\mathbf{C}_p = \mathbf{C}_P \cup \mathbf{C}_M$ is more comprehensive and complete which will improve the final solution (See Figure 3 (c) and (d)).

4.2. Cut Income

If the length of a cut is thought as the cost we pay for choosing this cut, the contribution of a cut can make for reducing the concavity of parts can be viewed as the "income" of a cut. For each mutex pair mp and each cut C, let $I(mp, C)$ denotes the income of C for mp. For example, in Figure 4, without the red cut, the concavity of A and B is $f(p_s) - f(p_A)$ by definition of Eqn. 6. With the red cut, the

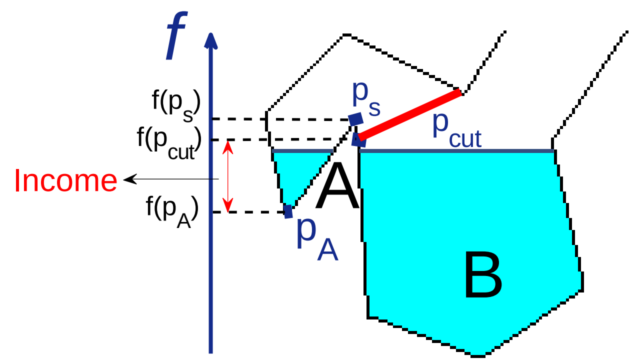

Figure 4. Illustration of the income of a cut (red line) for a mutex pair (A and B). Point p_s is the saddle point which corresponds to the mutex pair. p_{cut} is the cut point. p_A is the lowest point in part A w.r.t. the direction of Morese function f. The income of the red cut for mutex pair A and B is $f(p_{cut}) - f(p_A)$.

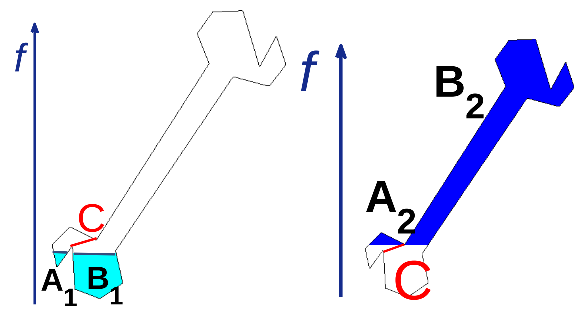

Figure 5. One cut can satisfy two mutex pairs of regions. The left figure shows mutex pair A_1 and B_1, the right figure shows mutex pair A_2 and B_2. Both can be satisfied by the red cut.

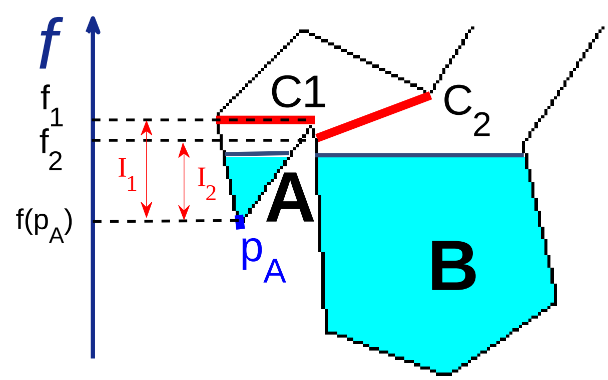

Figure 6. One mutex pair of regions A and B can be separated by two different cuts C_1 and C_2 with income I_1 and I_2 respectively.

concavity of the left part becomes $f(p_s) - f(p_{cut})$. So the reduction of the concavity is $f(p_{cut}) - f(p_A)$ which is the income of this cut for mutext pair A and B.

A cut can satisfy multiple mutex pairs (Figure 5) and a mutex pair can also be satisfied by multiple cuts(Figure 6). Let mp denote a mutex pair and $\mathbf{C}(mp)$ denote the set of candidate cuts which can satisfy mp. Let $\mathbf{M}(C)$ denote the set of mutex pairs which can be satisfied by C. So the real income of a candidate cut is determined by the mutex pairs it satisfies in the final shape decomposition. However, which mutex pairs it can really satisfy is unknown before the decomposition is finalized. If multiple cuts satisfying mp are chosen in the final solution, only the cut C^* which maximizes $I(mp, C)$ can get its income $I(mp, C^*)$ and other cuts have no income. For example, in Figure 6, both cuts C_1 and C_2 can satisfy mutex pair A and B, but the income of C_1 is larger. Thus if both cuts are chosen, only C_1 makes income for this mutex pair.

We estimate the expected income of each candidate cut by stochastic analysis. Assume that the probability that each candidate cut being chosen is equal to $1/2$. For each mutex pair, we rank $\mathbf{C}(mp)$ by $I(mp, C)$ and get a sequence of candidate cuts as $C_1^{mp}, C_2^{mp}, \ldots, C_k^{mp}$ such that $I(mp, C_1^{mp}) \geq I(mp, C_2^{mp}) \geq \ldots \geq I(mp, C_k^{mp})$. Let $r(mp, C)$ be the ranking of the cut C w.r.t. mp. For example, $r(mp, C_1^{mp}) = 1$. If a subset of $\mathbf{C}(mp)$ is chosen in the final solution, only the cut with maximum $I(mp, C)$ counts, the others do not make any income for this mutex pair. The probability for cut C being counted for mutex pair mp depends on its ranking $r(mp, C)$. If $r(mp, C) = 1$, the probability for C being counted for mp is $1/2$. Because once C is chosen, it will be counted for mp no matter whether any other cut is chosen. The probability of chosing C is $1/2$. In general, it is easy to show the probability for C being counted for mp is

$$\begin{aligned} & Prob(C \text{ is counted for } mp) \\ = & \; Prob(C \text{ is the 1st chosen in ranked } \mathbf{C}(mp)) \\ = & \; \frac{1}{2^{r(mp,C)}} \end{aligned} \qquad (10)$$

Based on this observation, we can estimate the expected income of a cut C as follows:

$$\overline{I(C)} = \sum_{mp \in \mathbf{M}(C)} \frac{1}{2^{r(mp,C)}} I(mp, C). \qquad (11)$$

4.3. Optimization

Assume that there are n candidate cuts in the proposed candidate cut set $\mathbf{C}_p = \{C_j\}_{j=1}^n$, the final decomposition chooses a subset of $\mathbf{C}_p$, denoted by $\mathbf{C}^*$. With the expected income for each candidate cut, we can reformulate Eqn. 3 as follows:

$$\min_{\mathbf{C}^*} \{ \sum_{C_j \in \mathbf{C}^*} [L(C_j) - \overline{I(C_j)}] \} \qquad (12)$$

because minimization of the concavity of parts is equivalent to maximization of the reduction of concavity by the chosen cuts, i.e., the expected income of chosen cuts.

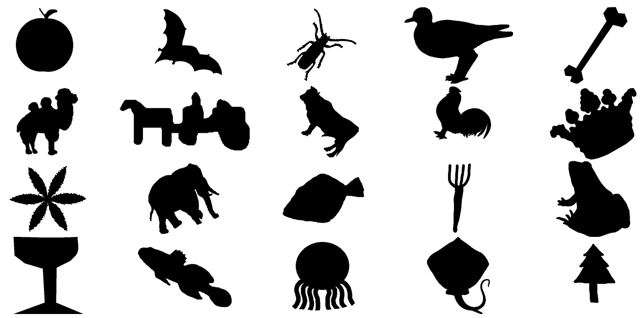

Figure 7. Example shapes of 20 categories from MPEG-7 shape dataset.

Design a binary vector $\mathbf{x}$ such that:

$$\mathbf{x}_j = 1 \iff C_j \in \mathbf{C}^*. \qquad (13)$$

Let vector $\mathbf{L}$ represent the cut length of $\mathbf{C}_p$ such that $\mathbf{L}_j = L(C_j)$. Let vector $\mathbf{I}$ represent the expected income of $\mathbf{C}_p$ such that $\mathbf{I}_j = \overline{I(C_j)}$. Design a penalty matrix $\mathbf{H}_{n \times n}$ such that if C_j and C_k intersects $\mathbf{H}(j,k) = +\infty$. Therefore, Eqn. 12 can be rewritten as

$$\begin{aligned} & \min_{\mathbf{x}} \mathbf{L}^T\mathbf{x} - a\mathbf{I}^T\mathbf{x} + \mathbf{x}^T\mathbf{H}\mathbf{x} \\ s.t. \quad & \mathbf{A}\mathbf{x} \geq 1, \mathbf{x} \in \{0,1\}^n, \end{aligned} \qquad (14)$$

where $\mathbf{A}$ is a matrix of size $n_{mp} \times n$. It denotes the relationship between the mutex pairs $MP = \{mp_k\}_{k=1}^{n_{mp}}$ and the candidate cuts $\mathbf{C}_p$. If a mutex pair mp_k can be satisfied by cut C_j, then $\mathbf{A}(k,j) = 1$, otherwise it is zero. a is parameter to adjust the impact of cuts' income.

The above formulation considers the cut length, the expected income of cuts, the intersection of cuts and the mutex pairs to be separated. If we relax $\mathbf{x}$ to be linear, $\mathbf{x}_i \in [0,1]$, this problem becomes a standard quadratic programming problem as:

$$\begin{aligned} & \min_{\mathbf{x}} \mathbf{x}^T\mathbf{H}\mathbf{x} + (\mathbf{L}^T - a\mathbf{I}^T)\mathbf{x} \\ s.t. \quad & \mathbf{A}\mathbf{x} \geq 1, \mathbf{x} \in [0,1]^n. \end{aligned} \qquad (15)$$

By solving Eqn. 15, we can get a soft assignment of $\mathbf{x}$ and then iteratively choose a cut set $\mathbf{C}^*$ based on $\mathbf{x}$ as [9].

5. Experimental Results

To evaluate the proposed method on 2D shape decomposition, we choose 20 categories from the MPEG-7 shape dataset [7] as Figure 7 shows. For each category, we choose 20 shapes. In total, there are 20×20 shapes for evaluation.

We propose two measures to evaluate the performance of the decomposition result: cut length and convexity which correspond to the first two items in the designed objective function Eqn. 14. During the implementation, we choose

16 Morse functions and set the threshold ϵ for generating mutex pairs as $0.05\mathfrak{D}$ where $\mathfrak{D}$ is the maximum distance between points of the shape. These settings are same for both our method and [9]. The parameter a in our method is set as 0.1.

5.1. Shape Decomposition

In order to test the performance of our method, we implement two shape decomposition methods of [9] and [3], then compare them to our method. Figure 8 shows the comparison result in terms of the cut length and convexity. For each shape, we first obtain the three different composition results and then evaluate the cut length and the convexity by the definition Eqn. 2 in Section 3. For each category, we calculate the mean and standard variance of the cut length across 20 instances for each method respectively. The result is shown in Figure 8 (a). Similarly, the convexity is computed and shown in Figure 8 (b). It can be seen that our cut length(red bar) is shorter than the other two methods (blue and green) in most cases while our convexity is about same as those of other two methods. This proves that our optimization method can achieve same convexity with shorter cuts compared to previous work.

5.2. Human Perception

In order to verify whether the decomposition results are consistent with human perception, we conduct an experiment which asks people to decompose the shape examples. For each of the 20 categories from the MPEG-7 shape dataset, we select 3 shapes for testing. And for each selected shape, there are 3 people who decompose it manually. The cuts provided by the participants are taken as the set of "ground truth" cuts, called $\mathbf{C}^+$. To measure the discrepancy between our decomposition result and the "ground truth" cuts, we define the distance between two cuts $C_1 = \overline{p_1p_2}$ and $C_2 = \overline{p_1p_2}$ as follows:

$$D_1(C_1, C_2) = \min\{ED(p_1, p_3) + ED(p_2, p_4), ED(p_1, p_4) + ED(p_2, p_3)\}. \quad (16)$$

For each cut C_i from $\mathbf{C}^*$, its distance to the ground truth is defined as

$$D_2(C_i, \mathbf{C}^+) = \min_{C \in \mathbf{C}^+} D_1(C_i, C). \quad (17)$$

and the distance between the cut sets $\mathbf{C}^*$ and $\mathbf{C}^+$ is defined as

$$D(\mathbf{C}^*, \mathbf{C}^+) = \frac{1}{|\mathbf{C}^*|} \sum_{C_i \in \mathbf{C}^*} D_2(C_i, \mathbf{C}^+). \quad (18)$$

This is the average distance between the proposed cut and the nearest cut in the ground truth and called as "cut distance".

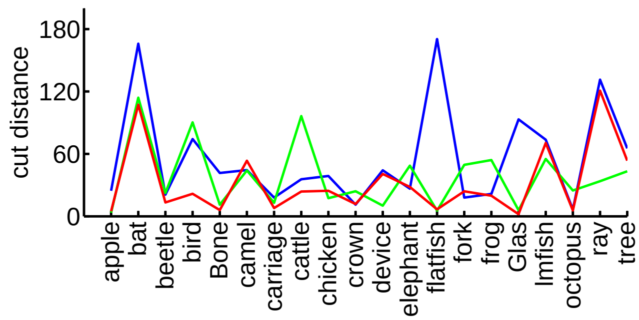

Figure 9. Distances between the human cuts and the shape decomposition results by the proposed method (red), [9](green) and [3](blue) respectively for the MPEG-7 shape dataset.

Figure 9 shows the computed "cut distance" of our result as well as [9] and [3]. It can be seen that our result is closer to the ground truth compared to these two previous work for most categories. This shows that combining the three generic rules is useful to learn perceptually meaningful shape parts. From Figure 9, we can see that our method performs best for "bird" and worst for "ray" (as shown in the Figure 10).

6. Conclusion

We proposed a method to solve the shape decomposition problem for learning perceptually meaningful parts. By jointly considering three generic rules: the Minima rule, the Short-cut rule and convexity, we formulate the shape decomposition problem as an optimization problem and design a new metric "cut income" to measure the contribution of candidate cuts for improving the convexity of decomposed cuts. By using this metric, the original problem is solved as a quadratic programming problem. The experimental results show that our approach is promising to keep a good tradeoff between cut length and convexity.

References

[1] I. Biederman. Recognition-by-components - a theory of human image understanding. *Psychological Review*, 94(2):115–147, 1987. 1

[2] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(9):1627–1645, 2010. 1

[3] R. Gopalan, P. Turaga, and R. Chellappa. Articulation-invariant representation of non-planar shapes. In *Proceedings of the 11th European Conference on Computer Vision: Part III*, pages 286–299, 1927030, 2010. Springer-Verlag. 1, 2, 3, 6, 7

[4] D. Hoffman and W. Richards. Parts of recognition. *Cognition*, 18:65–96, 1984. 1, 2

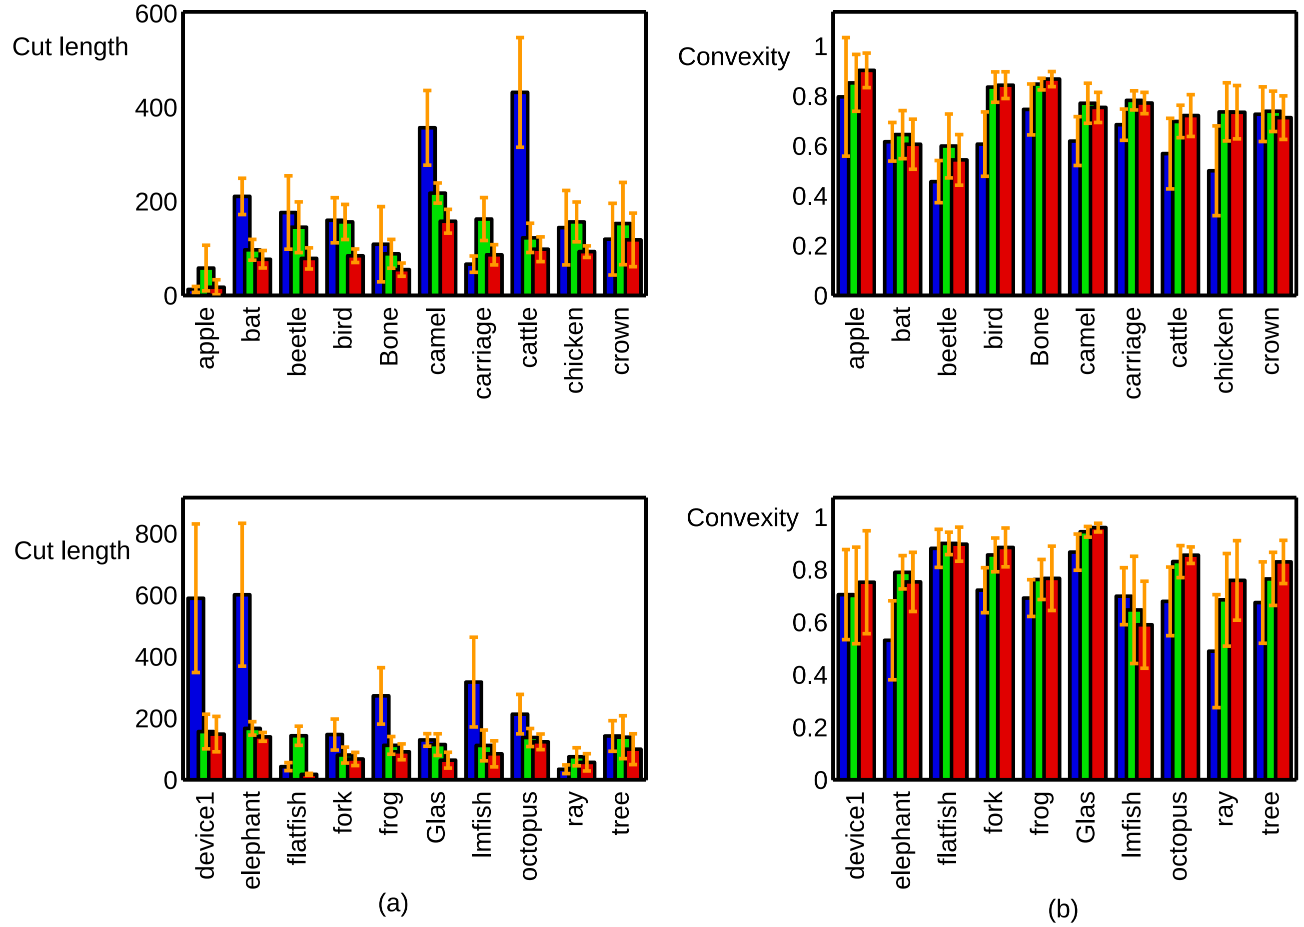

Figure 8. (a) compares the cut length of the shape decomposition results from [3](blue), [9](green) and our method(red). The yellow lines denote the standard variance. (b) compares the convexity of the decomposed parts from [3](blue), [9](green) and our method(red)

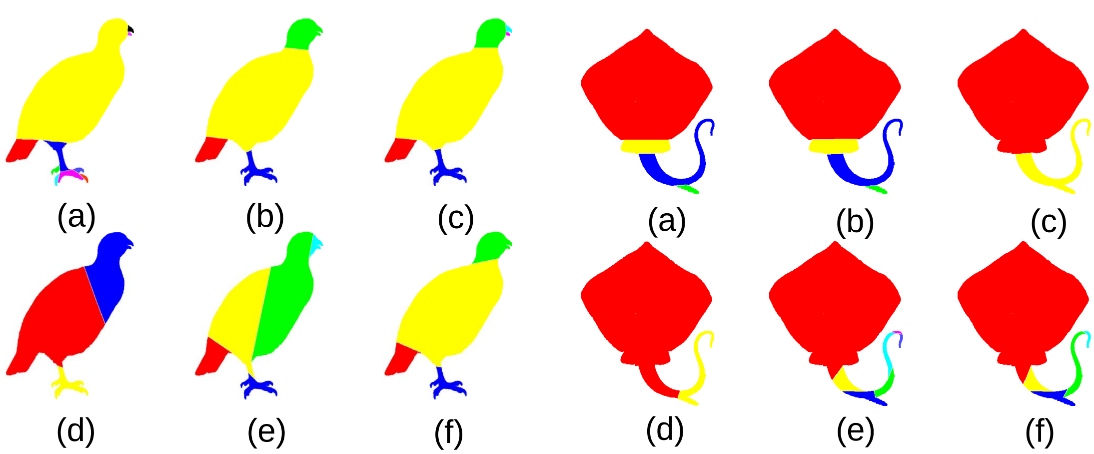

Figure 10. The first row (a)-(c) shows the examples of human decomposition results, (d) is the result of [3], (e) is the result of [9] and (f) shows our result.

[5] G. Kanisza and W. Gerbino. *Convexity and symmetry in figureground organization*. Vision and Artifact, M. Henle, ed.

New York Springer, 1976. 2

[6] L. J. Latecki and R. Lakamper. Convexity rule for shape decomposition based on discrete contour evolution. *Computer Vision and Image Understanding*, 73(3):441–454, 1999. 1, 2

[7] L. J. Latecki, R. Lakamper, and U. Eckhardt. Shape descriptors for non-rigid shapes with a single closed contour. In *Computer Vision and Pattern Recognition*, volume 5, pages 424–429, 2000. 5

[8] H. B. Ling and D. W. Jacobs. Shape classification using the inner-distance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(2):286–299, 2007. 3

[9] H. Liu, W. Liu, and L. L. J. Convex shape decomposition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 97–104, 2010. 1, 2, 3, 4, 5, 6, 7

[10] X. Mi and D. DeCarlo. Separating parts from 2d shapes using relatability. In *IEEE 11th International Conference on Computer Vision*, 2007. 2

[11] M. Pilu and R. B. Fisher. Model-driven grouping and recognition of generic object parts from single images. *Robotics and Autonomous Systems*, 21(1):107–122, 1997. 1

[12] Z. Ren, J. Yuan, C. Li, and W. Liu. Minimum near-convex decomposition for robust shape representation. In *International Conference on Computer Vision*, 2011. 2

[13] P. L. Rosin. Shape partitioning by convexity. *IEEE Transactions on Systems Man and Cybernetics Part a-Systems and Humans*, 30(2):202–210, 2000. 2, 3

[14] K. Siddiqi and B. B. Kimia. Parts of visual form - computational aspects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(3):239–251, 1995. 2

[15] K. Siddiqi, K. J. Tresness, and B. B. Kimia. Parts of visual form: Psychophysical aspects. *Perception*, 25(4):399–424, 1996. 1

[16] M. Singh, G. D. Seyranian, and D. D. Hoffman. Parsing silhouettes: The short-cut rule. *Perception and Psychophysics*, 61(4):636–660, 1999. 1, 2

[17] L. L. Walker and J. Malik. Can convexity explain how humans segment objects into parts? *Journal of Vision*, 3(9), 2003. 1, 2

[18] L. Zhu, Y. Chen, A. Torralba, W. Freeman, and A. Yuille. Part and appearance sharing: Recursive compositional models for multi-view multi-object detection. In *Computer Vision and Pattern Recognition*, pages 1919–1926, 2010. 1