Identifying Perceptually Salient Features on 2D Shapes

Lisa Larsson, Géraldine Morin, Antoine Begault, Raphaëlle Chaine, Jeannine
Abiva, Evelyne Hubert, Monica Hurdal, Mao Li, Beatriz Paniagua, Giang
Tran, et al.

HAL Id: hal-01071631

https://hal.inria.fr/hal-01071631

Submitted on 7 Oct 2014

Identifying Perceptually Salient Features on 2D Shapes

Lisa J. Larsson, Géraldine Morin, Antoine Begault, Raphaëlle Chaine, Jeannine Abiva, Evelyne Hubert, Monica Hurdal, Mao Li, Beatriz Paniagua, Giang Tran, and Marie-Paule Cani

Abstract Maintaining the local style and scale of 2D shape features during deformation, such as when elongating, compressing, or bending a shape, is essential for interactive shape editing. To achieve this, a necessary first step is to develop a robust classification method able to detect salient shape features, if possible in a hierarchical manner. Our aim is to overcome the limitations of existing techniques, which are not always able to detect what a user immediately identifies as a shape feature. Therefore, we first conduct a user study enabling us to learn how shape features are perceived. We then

Lisa J. Larsson
Department of Mathematics and Statistics, McGill University, 805 Sherbrooke Street West, Montreal, QC, Canada H3A 0B9 e-mail: lisa.powers@mail.mcgill.ca

Géraldine Morin
University of Toulouse - INPT - IRIT, IRIT-N7, 2 Rue Camichel, 31071 Toulouse, France e-mail: geraldine.morin@enseeiht.fr

Antoine Begault · Marie-Paule Cani
LJK, University Grenoble Alpes, CNRS and Inria, F-38000 Grenoble, France e-mail: antoine.begault@inria.fr; marie-paule.cani@inria.fr

Raphaëlle Chaine
Université de Lyon, Université Lyon 1, CNRS, LIRIS, UMR5205, F-69622 Lyon, France e-mail: raphaelle.chaine@liris.cnrs.fr

Jeannine Abiva
Department of Mathematics, University of Iowa, 14 MacLean Hall, Iowa City, IA 52242 e-mail: jeannine.abiva@gmail.com

Evelyne Hubert
INRIA Méditerranée, 2004 Route des Lucioles, Sophia Antipolis 06902, France e-mail: evelyne.hubert@inria.fr

Monica Hurdal · Mao Li
Department of Mathematics, Florida State University, 208 Love Building, 1017 Academic Way, Tallahassee, FL 32306 e-mail: mhurdal@math.fsu.edu; mli@math.fsu.edu

Beatriz Paniagua
Department of Psychiatry, University of North Carolina at Chapel Hill, 101 Manning Drive, Chapel Hill, NC 27514 e-mail: beatriz.paniagua@med.unc.edu

Giang Tran
Department of Mathematics, University of California, Los Angeles, 520 Portola Plaza, Los Angeles, CA 90095 e-mail: giangtran@math.ucla.edu

propose and compare several algorithms, all based on the medial axis transform or similar skeletal representations, to identify relevant shape features from this perceptual viewpoint. We discuss the results of each algorithm and compare them with those of the user study, leading to a practical solution for computing hierarchies of salient features on 2D shapes.

1 Introduction

Fig. 1 Real shapes with perceptually salient features, such as seeds on strawberries, ridges on leaves, and skin details on a chameleon.

Natural objects typically include a hierarchy of shape features, from fine-scale details to the main object. These features, such as those in the examples of Figure 1, are fundamental to our perception of the object. Similarly, digital models of shape incorporate a number of features of various scales. During interactive shape editing, artists should preferably not have to manually adjust or duplicate all these features when deforming the parent shape, e.g., by bending, twisting, stretching or shrinking it: this would be time-consuming, even for skilled artists. Therefore, research in interactive shape design has sought methods that enable adjustments on the parent shape while automatically preserving the style and aesthetics of small shape features.

This paper takes a necessary step towards this challenging goal: it proposes perceptually-based algorithms to detect shape features, in the case of 2D shapes. Once shape features are detected, they can be carefully manipulated during deformation processes. In this work, our first focus is defining what is meant by a shape feature. Therefore, we use the results of a perceptual user study to guide our definition of shape features. Secondly, we choose to use the medial axis transform to detect features on 2D shapes, since this is likely to ease subsequent generalization to 3D shapes, compared to other criteria, such as local curvature.

Throughout this paper, S will denote a shape in $\mathbb{R}^2$ that has a closed, piecewise smooth boundary. As a subset of $\mathbb{R}^2$, S is a compact set.

The organization of this paper is as follows: In Section 2, we present previous work. A user study conducted to understand human perception of shape

features is then motivated and described in Section 3; the results of this study are also presented. Then, two classes of algorithms for detecting shape features are proposed: a geometric algorithm is described in Section 4, followed by several methods enabling the detection of features at different resolutions in Section 5. Section 6 then presents numerical tests of the new algorithms on a variety of interesting shapes and compares this to an existing feature-detection algorithm. We conclude with a discussion of these results and directions for future work.

2 Previous Work

2.1 The medial axis transform

The medial axis can be defined equivalently in different ways, for instance, in terms of maximally inscribed balls or the shock set of the eikonal flow from the shape boundary [3, 7, 8, 16]. The following definition was presented in [3, 4].

Definition 1 (Medial Axis Transform). The Medial Axis Transform of $\mathcal{S}$ is given by the set of locations $\mathcal{M}$ internal to the object with more than one corresponding closest boundary point and their distance $\mathcal{R}$ from the boundary $\partial\mathcal{S}$.

Remark: When the shape $\mathcal{S}$ is non-convex, an "exterior medial axis" can also be defined. These are points in $\mathcal{S}^c$ that are equidistant to two or more closest points along $\partial\mathcal{S}$.

The shape $\mathcal{S}$ can be reconstructed as the union of balls centered along $\mathcal{M}$ of radius $\mathcal{R}$. These balls are maximal and have at least two contact points on the surface. The exterior medial axis is not required for shape reconstruction.

It is often convenient to assume that $\mathcal{S}$ is simply connected; i.e., that it has no holes. This simplifies the topology of the medial axis, in particular guaranteeing no loops. In the remainder of this paper, we only focus on such simply connected shapes.

The medial axis representation has been widely studied in computational geometry. In practice, the medial axis can be approximated as a subset of the Voronoi diagram of points sampled on the boundary of the shape [1]. In [13], the authors edit the shock graph to systematically define the distance between two shapes for recognition.
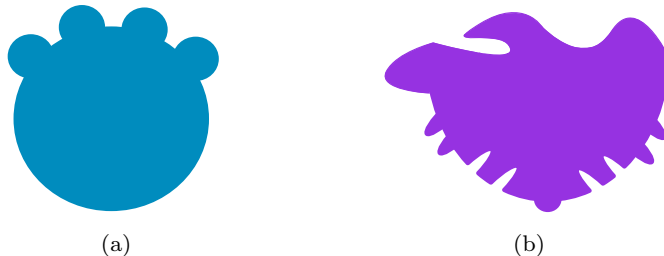
Fig. 2 Shapes with salient features. (a) The bumps protruding from the large sphere are salient. (b) A shape with many salient features.

2.2 Identifying salient features

In his seminal work [9], Michael Leyton characterizes features on 2D shapes using curvature of the shape boundary: a salient feature of a 2D shape is identified by the presence of two extrema of curvature of the same sign along this boundary. The two extrema should occur at the contact points of one medial ball of the medial axis (or of the external medial axis).

Figure 2 shows a number of salient shape features that follow this definition: the shape in (b) has many salient features; some are larger, such as what appears to be the head of a bird, and others are smaller in scale. In our user study, we probe whether, on average, parts like the bird's head are perceived to be a shape feature or part of the main shape. We do not restrict ourselves to any particular definition of salience here, but rather use the results of the perceptual study to guide our definition of shape feature.

Directly using skeletal representations to detect salient features has been proposed in previous work. For example, in [12], the authors use a modification of the medial axis transform, called the chord axis transform, and decompose shapes based on the extremal chord strength. In [18], the authors use both skeletal and boundary features and define a protrusion strength that is then used to decompose the shape. The originality of our approach compared to these methods is the fact that we build our work on a user study, enabling us to address the identification of perceptually salient shape features.

Furthermore, there are many algorithms designed for pruning the medial axis that can be applied to identifying shape features. These algorithms were developed because the medial axis transform is known to be very sensitive to noise on the shape boundary, resulting in spurious branches [1, 2, 5, 15, 17]. Much work has been devoted to the pruning of these non-informative branches. The resulting algorithms, aimed at detecting features in order to remove the less salient ones, should indeed be considered with respect to our goal. These approaches filter points on the medial axis depending on the geometric configuration of the contact points of the medial balls (e.g., the

radius of medial balls combined with the angle between a center and the contact points [2, 15], the radius of the ball circumscribing the contact points [5], or the area enclosed by the contact points [14, 17]). More recently the Scale Axis Transform (SAT) was developed, which uses a nonlinear scaling of the medial balls to produce a hierarchy of simplified medial axes [6, 11]. This method is recalled in more detail in Section 5, as we evaluated it with respect to the results of our user study, and then build on it to propose a solution that better matches our goals.

To adhere to the results of the user study, we propose to extend the SAT in order to take into account the length or thickness of medial branches. To compute the length, we rely on the extended distance function (EDF) [10], which measures the tubularity of a shape. A detailed description of EDF is given in Section 5; to compute thickness, WEDF, an original weighted version of EDF, is then proposed. Combining these branch characteristics with the SAT is one of the key points of our solution.

3 User Study on Shape Feature Perception

A user study was conducted to determine the criteria for the perceptual identification of features on a 2D shape. For that, given a collection of 2D shapes, we asked the users to identify any shape features they considered different from the main shape. This study was designed to gauge the consistency with which users identify shape features, and both motivates and validates the algorithms we construct. The user study consisted of a set of 44 shapes, that were printed on three sheets of paper. The results for shapes with semantic content—i.e., shapes that looked like plants, animals, or other naturally-occurring objects—were discarded to mitigate semantic bias. The order and the orientation of the shapes was random, to reduce the influence of similar shapes in adjacent positions.

Indeed, our aim in this study was to identify the geometric criteria that are important in identifying shape features. For the non-specific, simply connected shapes presented in the user study, we assumed that the partition of the shape into a main shape plus shape features was independent from the shape's orientation. Before conducting the user study, we formed four hypotheses about the properties of shape features in terms of the radius and orientation of the shape's medial axis:

(H1) *Any part of a shape delimited by two corresponding extrema of curvature of the same sign (i.e., occurring at the contact points of one medial ball), is a shape feature. This is Leyton's criteria [9].*

(H2) *A junction of the medial axis (local Y shape) is important but neither necessary nor sufficient to characterize a shape feature.*

(H3) *A change in thickness along the medial axis is more predictive of shape features than a change in direction. Yet, it is neither necessary nor sufficient to characterize a shape feature.*

(H4) *If the main shape is not of larger radius than the features, it is of larger length. Moreover, a part with smaller area is more likely to be a shape feature than a part with smaller length.*

Once the users obtained the set of printed shapes, we gave them the following instruction:

(Q) *We are trying to construct a hierarchy of a shape's features. Highlight what you think is NOT the main shape.*

The user study was given to 24 people, the sheets were scanned and aligned, and each pixel was averaged over the 24 forms. The average results in grayscale are presented in Figure 20. A summary of relevant results is given below:
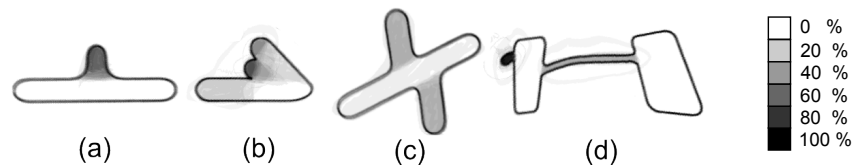

Fig. 3 Results of the user study for (H1). (From the left) The first two shapes validate (H1), as the majority of users choose shape features that are consistent with Leyton's definition of salience. The third shape disagrees with Leyton's definition and (H1), as only two of the four salient parts are identified as shape features. The fourth shape also disagrees with (H1), as the two larger areas are salient features by this definition.

The first hypothesis aims to verify the definition of salient shape feature given by Leyton (see Section 2). This hypothesis was both validated and refuted by the user study. Actually, Figure 3(a) and (b) confirm (H1), while (c) and (d) refute it.

In Figure 3 (a), nearly all users identify the nob on top as a shape feature, while the white area is the main shape. This feature is delimited by two extrema of curvature of the same sign. Figure 3 (b) also confirms (H1). Although some users thought of the lower and upper extremities as shape features, the majority considered those parts to be the main shape, while the only feature delimited by two extrema of curvature of the same sign was chosen to be a shape feature by the majority of the users. In Figure 3 (c), all four extremities should be considered shape features under (H1). However, this is not observed, as the two extremities that are aligned are considered by the majority of users to be the main shape. In Figure 3 (d), the users

both validate and refute (H1). The small nob on the left side of the shape is overwhelmingly characterized as a shape feature, this is in agreement with (H1). However, both the large white areas are delimited by two extrema of curvature of the same sign, yet they are considered the main shape. This is likely because of the larger size of these areas compared to the thin connecting strip. In summary, there are cases that both support and refute this hypothesis, so it cannot be used in isolation to identify shape features.
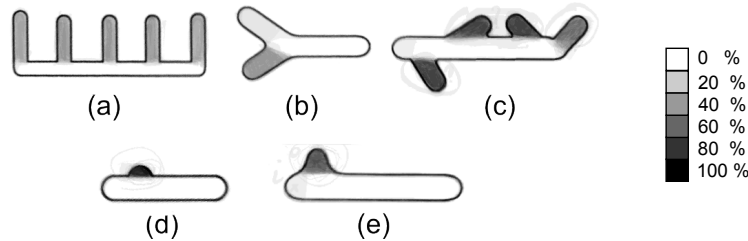

Fig. 4 Results of the user study for (H2). (From the top left) The first shape (a) confirms (H2) because of the first and last prong of the comb—there is no junction here, yet they are shape features. The shape in (b) also confirms (H2) because it was unclear which branch of the junction was the shape feature—many thought it was all the main shape. Shape (c) confirms (H2), as there is no junction for the right-most feature. In the bottom two figures, a junction is sufficient.

The second hypothesis is concerned with junctions in the medial axis. (H2) conjectures that a junction in the medial axis is important but neither sufficient nor necessary to characterize shape features. A junction refers to a meeting point of three edges of the medial axis. The results here were quite interesting. There were many cases (see Figure 4) where a junction from the primary medial axis branch was sufficient to characterize a shape feature (for example, (d) and (e)). However, for the Y shape (Figure 4 (b)) the users did not reach a consensus about which branch of the Y was the shape feature, and some users found there to be no shape features. It is difficult to draw broad conclusions from this as the Y shape has semantic information—it is shaped like a letter, which may influence the perception of shape features in this case. In the top row, Figure 4 (a) and (c) have repeated shape features. In the comb example (Figure 4 (a)) all the prongs are shape features, even though the first and last do not occur at a junction of the medial axis. Similarly, in Figure 4 (c) the right-most feature does not occur at a junction of the medial axis, so a junction is not necessary in this case. These results demonstrate that junctions in the medial axis are strong indicators of shape features, but not sufficient, and suggest that repeated features should also be taken into account. Hypothesis (H2) was therefore validated by the user study.

The third hypothesis considers whether it is more important to have a change in shape thickness or a change in the direction of the shape to classify

shape features. In Figure 5, there are three examples that support (H3).
The results were not unanimous: In Figure 5(a), some users highlighted the
bottom left part as a feature, but most decided that the main shape was the
V to the left, and the feature was the nob to the right. The left V has a
consistent thickness, but not a consistent orientation, while the thickness of
the right nob is significantly reduced. This supports (H3). Figure 5(b) also
depicts mixed opinions, though the majority of users selected the thin right
part of the shape to be a shape feature, which is the part with the smallest
thickness. The results were similar for Figure 5(c). These results validate
(H3). The example of the comb shape (Figure 4 (a)) illustrates that changes
in thickness are not necessary nor sufficient to determine a shape feature.
The thickness of the main part is similar to the thickness of the features in
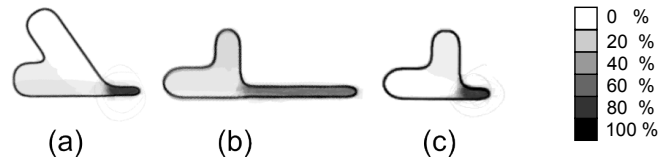this example.

Fig. 5 Results of the user study for (H3). All three figures confirm (H3), which says
that a change in thickness is more important than a change in direction in identifying
shape features.

The fourth hypothesis addresses the case where the main shape is not
as thick or as long as the shape features. In Figure 6, (a) and (b) show
cases where the main shape has the same thickness as the shape features.
This validates (H4), as in these two cases, the users identified shape features
that have the smallest length. Therefore, given a shape with nearly constant
thickness, it is necessary and sufficient to use length to successfully identify
shape features. In Figure 6 (c), a shape with non-constant thickness is shown.
In this case, length alone is not sufficient to characterize shape features, as
the thin feature is much longer than the main shape. For shapes with non-
constant thickness, a branch with a bigger area is more likely to be considered
the main shape.
In summary, these results show that:

- Two extrema of curvature of the same sign do not always identify shape
  features.
- Junctions in the medial axis are important for determining shape features,
  but are neither necessary nor sufficient.
- Repeated parts can be considered features even when there is no junction
  in the medial axis.
- Change in radius is more important than change in orientation to identify
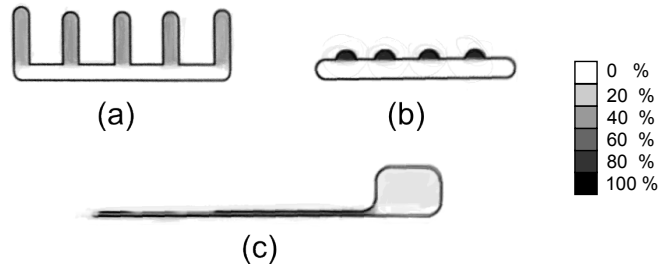 shape features.

Fig. 6 Results of the user study for (H4). All three figures confirm (H4), which says that the measure (in length or area) of the main shape is larger than that of the shape feature.

- The length of a part of a shape is necessary and sufficient to identify shape features when the thickness of the shape is nearly constant.
- A part of a shape with larger area is more likely to be considered the main shape than a lengthier part of smaller area.

These conclusions have driven the construction of a geometric feature-detection algorithm, based on the radius and orientation of the medial axis, which is presented next.

4 Perceptually-based Geometric Feature Detection

Our goal is to partition a 2D shape to separate the main part of the shape from the shape features. The geometric approach presented in this section directly builds on the results of the user study we just described: the input of the algorithm is a 2D shape and the output is the features, given in terms of their medial axis representation. The method involves first computing the medial axis, and then partitioning it using criteria that combine branching information with variations in radius of the medial balls.

4.1 Junctions in the Medial Axis

The most natural way to partition a 2D shape according to its medial axis is to consider points where the topology of the medial axis changes, that is, to categorize the branching parts as features. In Figure 7, the small branch of the medial axis corresponds to the raised bump on the main shape.

However, as confirmed by (H2) in the user-study, determining which branch corresponds to the feature and which branch corresponds to the main

Fig. 7 The medial axis (in white) for a shape (in blue). The junction in the medial axis indicates a branching, which is a criteria for identifying shape features.

shape requires extra information in addition to the detection of junctions: in Figure 7, there are three branches of the medial axis, and while two correspond to what is visually perceived as the main shape, one corresponds to the shape feature. Considering junctions of the medial axis in isolation is therefore not sufficient. In practice, we combine it with other information, namely changes of radius and of orientation among neighboring branches. This corresponds to the conditions for identifying features in Algorithm 1 below.

4.2 Changes in radius

As demonstrated in the user study (H2), considering junctions is important but neither sufficient nor necessary. Shape features may also arise along a single branch of the medial axis. Figure 8 depicts a simple example, namely a sphere with a small triangular outcropping, which can appear as a tail or a nose. The nose is a shape feature, however there is no branch in the medial axis that would lead to its detection. This suggests that along each branch, the radius must be considered, and the branch will be further subdivided if there is a significant change in thickness. A subdivision of the branch means that we split the branch into two branches and add a new junction between the two.
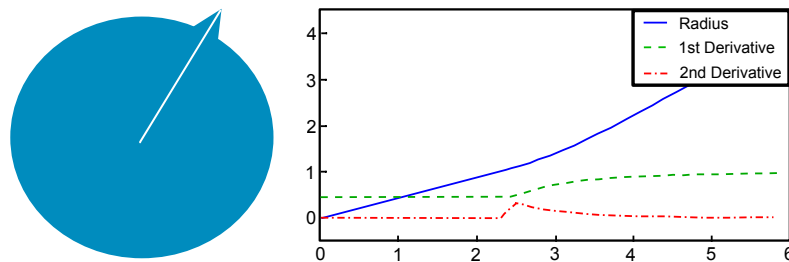

Fig. 8 Left: Although the medial axis has no junction, the nose is a shape feature. Right: The radius and its derivatives are shown. The calculation starts at the tip of the feature and proceeds towards the circle center.

One way to analyze the thickness is to consider the variation of the radius on the medial axis. Analyzing the behavior of the medial axis in Figure 8 one can observe that there is a sudden increase in the second derivative (see the curves in Figure 8, right). Thus, we propose to use the variation of the radius along the curve to further partition the medial axis. New junctions are added on simple branches when a second order variation of the radius is detected.

4.3 Resulting geometric algorithm

Once the medial axis has been partitioned by considering both junction and thickness information along branches, consistency of thickness and orientation at each junction are considered sequentially to compute the shape decomposition we are looking for. According to (H3), we first consider changes in thickness, and when there is no change in thickness, we identify features based on directions. Our method is summarized in Algorithm 1 below. Note that we implemented a discrete version of the radius-based partitioning method just discussed: the medial axis is sampled into a graph of nodes with associated radius. Adjacent nodes refer to neighboring points along this discretized graph. The second order derivative of the radius is approximated by the second discrete differences (denoted by Δ_2) on three successive values of the radius along a branch.

Results of this algorithm will be discussed in Section 6. This algorithm requires different parameter values for comparing changes of radius along a branch (Preprocessing), branch radii at junctions (Part I), and branch orientations (Part II). These parameters can be tuned to get results that consistently match the user study. Moreover the algorithm provides the relative importance of adjacent branches at each junction: this leads to different levels of resolution. Of course, branches of the same level are not necessarily topologically connected. Moreover, branch importance is computed based on local information at each junction. A drawback of this method could be its lack of computing global branch importance: branches of the same level may be very different in thickness, for example. Using this algorithm to detect hierarchies of shape features at different resolutions would require complex tuning, with different families of values for the three parameters. We therefore investigate an approach that uses a single resolution threshold, presented next.

5 Towards Multi-Resolution Feature Detection

In this section, we develop an algorithm that can identify a hierarchy of shape features of different scales. The starting point to achieve this goal is the Scale Axis Transform, which achieves multiscale feature detection using a nonlinear

---

Algorithm 1 Geometric Feature Detection Algorithm

Input: Given $\mathcal{S}$, and a three thresholds ε, T and α used for the second order variation of the radius of branches, the relative change of thickness, and the relative angle between branches at a junction, respectively.

Compute $\mathcal{M}$, the medial axis.

Preprocessing:
if there is a second order change of radius along the skeleton, that is $\Delta_2 R = R_{i+1} - 2R_i + R_{i-1} > \varepsilon$ **then**
   Split the branch at this point (add a junction).
end if

Main Loop:
for each junction x in the medial axis **do**

   Part I: Compare the radius between adjacent branches
   for each branch b_i entering the junction x **do**
      1. Get the intersection point x_i of b_i with the circle centered at x of radius $R(x)$.
      2. Compute the radius of the i^{th} branch $r_i = R(x_i)$.
   end for

   Find $r = \max_i r_i$ the maximum radius.
   for each branch b_i associated with the junction x **do**
     if $r - r_i > T$ **then**
       Label b_i a feature.
     else
       Keep b_i as the main shape.
     end if
   end for

   Part II: Compare the orientation between adjacent branches
   if All branches (b_i) are labeled as the main shape **then**
     for each branch b_i associated with the junction x **do**
       Compute the orientation θ_i of the branch b_i at x_i.
       if for all $j : |\theta_i - \theta_j| - \pi > \alpha$ **then**
         Label b_i a feature.
       else
         Keep b_i the main shape.
       end if
     end for
   end if
end for

---

scaling of the medial balls. We first study the standard SAT algorithm and identify why it does not match our goals. We then propose two extensions, based on different weightings of the scaling function that use length and area criteria, respectively. The input to the problem is again the 2D shape, and the output is a hierarchy of features based on a thresholding parameter.

5.1 Using the Scale Axis Transform

The *Scale Axis Transform* is a method for pruning spurious branches from the medial axis of 2D and 3D shapes [6, 11]. Given a shape, $\mathcal{S}$, the first step of the algorithm is to compute its medial axis $\mathcal{M}$. Next, each ball of the medial axis is scaled by a factor $s > 0$ (thus creating a non-linear scaling of $\mathcal{S}$), and the union of these scaled medial balls forms the s-scaled version of the shape $\mathcal{S}$. If $s > 1$, then the shape is dilated, if $s < 1$, then the shape is shrunk; $s = 1$ recovers the original shape. Under a dilation, $s > 1$, there will be medial balls that no longer contribute to the boundary of the s-scaled shape. These medial balls are removed. The inverse scaling $\frac{1}{s}$ is then applied to the new medial balls, recovering a medial axis representation of $\mathcal{S}$ where the parts of the medial axis corresponding to redundant medial balls are now pruned. This method is very effective for pruning parts of the medial axis that are artifacts of rasterized image boundaries and other boundary noise. Additionally, when $s > 1$ is progressively increased, this method yields a hierarchy of medial axes that correspond to progressively simpler versions of the input shape, $\mathcal{S}$.

In [6], the authors show that for any scaling $0 < s < 1$, the s-scaled shape is homotopy equivalent to the input shape $\mathcal{S}$. Moreover, for a class of scalings $s > 1$, the s-scaled shape is also shown to be homotopy equivalent. The s-scaled shape can be obtained as the appropriate level set of a multiplicatively-weighted distance function, where the weight is the distance from the surface to the closest point on the medial axis—the radius of the corresponding medial ball. These homotopy results are valid in any spatial dimension. In [11], the authors extend the algorithm for computing pruned medial surfaces in 3D.

Could the Scale Axis Transform be used for perceptual feature detection? Indeed, defining a hierarchy of shape features according to their persistence under the SAT non-linear scaling looks natural. However, since the SAT prunes features based on the relative size of their medial balls as compared to neighboring balls, a shape part will be identified as the *main shape* only of it is of larger radius than neighboring parts. For example, in Figure 9, the main part of the shape is of the same radius as the features, meaning the SAT will fail to identify the features correctly—the whole shape will be considered the main shape. Features that are attached to a support with the same radius will not be pruned using the Scale Axis Transform. The next

Fig. 9 Shapes with features (bumps) of the same radius as the main shape. The radius of the medial axis transform is nearly constant throughout the shape: the bumps will thus not be detected as salient features using the classical SAT. However, in such cases, the user study clearly indicates that the bump should still be identified as a feature, even without repetition.

two sections propose modifications of the SAT aimed at improving detection of shape features when the part perceived as the main part of the shape is relatively thin.

5.2 Length-weighted SAT

Results of the perceptual user study (H4) showed that the main shape may not be of larger radius than features, but that it should be associated with some larger measure, such as being longer. To be robust to such cases, our insight is to use the length of each medial branch to weight the SAT scaling. In order to characterize the length of a medial axis branch, we use the Extended Distance Function (EDF) proposed by Lui et. al [10].

Fig. 10 The medial axis of two shapes, the EDF along the medial axis (in color), and the EMA (red dot).

5.2.1 The Extended Distance Function

In [10], the authors propose to compute at each point of the medial axis a quantity EDF (the Extended Distance Function), which measures the distance to the closest extremity of the shape corresponding to a degree-1 node

on the longest path of the medial axis containing a given medial point. The set of medial balls corresponding to this longest path is called a *tube*, and EDF is the distance to the closest extremity of the *tube*. This quantity is defined for each point in the medial axis, and may be infinite if the medial axis contains loops (this happens if there are holes in $\mathcal{S}$). The authors also introduce a notion of the center of a shape based on its medial axis representation and EDF. They consider all paths along the medial axis, starting from one degree-1 node, and reaching to another degree-1 node of the medial axis. The midpoint of the path that is of maximal length is called the EMA (Extended Medial Axis). If the medial axis is acyclic (no loops), then the EMA is finite, else it is infinite. Figure 10 shows an example of EDF and EMA along the medial axis.

5.2.2 Feature Detection

Let $\mathcal{S}$ be simply connected. We propose a scaling factor based on the extended distance function. We partition the medial axis into branches. A branch is defined as a 1D subset of the medial axis on which EDF is continuous (such as straight parts in Figure 10). On a branch b, we consider the factor

$$\widetilde{\ell}(b) = \sup_{x \in b} \left\{ \mathrm{EDF}(x) \right\}.$$

We propose to tune the scaling factor of the SAT algorithm using the branch length factor $\ell(b) \in (0, 1]$:

$$\ell(b) = \frac{\widetilde{\ell}(b)}{\max_{b'} \widetilde{\ell}(b')}.$$

This leads to the following algorithm:

Algorithm 2 The SAT-EDF Feature Detection Algorithm

Input: Given $\mathcal{S}$, and some scaling threshold $s > 1$.
1. Compute $\mathcal{M}$, the medial axis.
2. For each branch b of the medial axis, compute $\ell(b)$, the EDF-based length.
3. For each point x_i along $\mathcal{M}$, scale the medial ball of radius $R(x_i)$ by a factor $s\,\ell(b)$, where b is the branch that contains x_i.
4. If the medial ball of x_i is contained in the union of other medial balls, label this medial point as belonging to a shape feature.

5.3 Area-weighted SAT

The length is important in determining shape features in cases where the medial balls have nearly constant radii. Similarly, the area can also be important, as stated in (H4). Parts of shapes can be very long, yet have small medial radius. In these cases, the results of the user study suggest that the parts of the shape that have a smaller area are more likely to be a shape feature. For these reasons, we propose a weighted version of EDF that is a proxy for the area of a shape and can be easily computed using the medial representation. This weighted EDF is then used to adjust the scaling of the SAT algorithm to incorporate the area corresponding to a medial branch.

5.3.1 The Weighted Extended Distance Function

We propose an extension of EDF, weighted EDF (WEDF), for simply connected 2D shapes that takes into account not only the length of the shape but also its thickness. We can similarly define the WEMA which generalizes the EMA. WEDF is a weighted variant of EDF, and corresponds to a weighted integral of the medial radii along branches of the medial axis. For each point along the medial axis, there is a path such that the area of the union of the medial balls along this path is maximal. One can compute the area of the part of this *tube* that lies to the left and to the right of the chosen medial point. WEDF is the smaller of these two areas. Recall that EDF was similarly defined, in terms of medial length not shape area (see Section 5.2.1).

  We first give some auxiliary definitions that will aid in defining WEDF. These quantities are illustrated in Figure 11.

i). For any point x along the medial axis that has exactly two contact points we define two angles $\theta_1(x)$ and $\theta_2(x)$ as follows: Form a line between the two contact points and x. The angle between these lines and the tangent of the medial axis at x yield $\theta_1(x)$ and $\theta_2(x)$. Either the tangent or negative tangent can be used to calculate these angles; this leads to different values of $\theta_1(x)$ and $\theta_2(x)$, however, the same value of $\sin(\theta_1(x))$ and $\sin(\theta_2(x))$ will be obtained. The set of medial points x with more than two corresponding boundary points has measure zero, so it suffices to consider these two angles.

ii). Let f be a connected curve within the medial axis, with endpoints a and c. At these endpoints, $A_f(a)$ and $A_f(c)$ refer to the area of the circle sector of the endpoint. This circle sector is defined by all boundary points of $\mathcal{S}$ that are a distance $R(x)$ from the endpoint a or c, respectively. It can be computed as follows: Compute the one-sided limits $\theta_1(a) = \lim_{x \to a} \theta_1(x)$ and $\theta_2(a) = \lim_{x \to a} \theta_2(x)$ in radians, where the tangent is calculated to point towards the extremity of the shape. Then $A_f(a) = \frac{\theta_1(a) + \theta_2(a)}{2} R(a)^2$.

See Figure 11 for an illustration of θ_1, θ_2 and A_f.

Definition 2. Let f be any connected curve within the medial axis, and let a and c be the end points of f. We define the **weight** of f by:

$$W(f) = \int_f R(x) \left(\sin \theta_1(x) + \sin \theta_2(x)\right) ds(x) + A_f(a) + A_f(c),$$
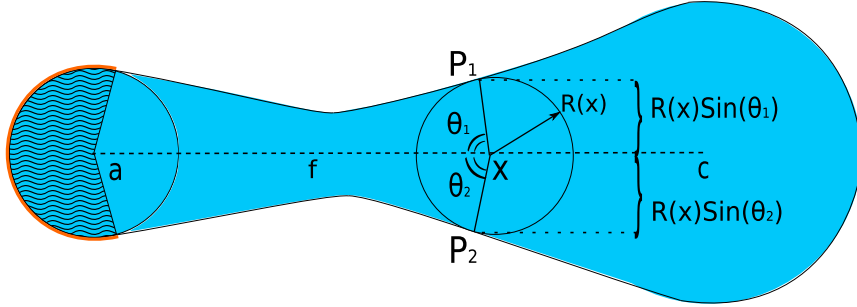
where s is the arclength parametrization of f.

Fig. 11 Computation of WEDF. The dashed area on the left corresponds to the quantity $A_f(a)$ of the medial endpoint a. The quantities relevant to compute WEDF, $R(x) \sin(\theta_1)$ and $R(x) \sin(\theta_2)$, are illustrated.

This new metric integrates the radius along the path f, and leads naturally to definitions of the weighted analogs of EDF and EMA.

Definition 3. Let f be a simple path with endpoint a and c.

i). We first define the Weighted Extended Distance Function at a point x relative to a simple path f as:

$$\mathrm{WEDF}_f(x) = \min_{y=a,c} \left(\int_x^y R(t) \left(\sin \theta_1(t) + \sin \theta_2(t)\right) ds(t) + A_f(y) \right).$$

ii). We call the Weighted Extended Distance Function, or **WEDF** the following quantity:

$$\mathrm{WEDF}(x) = \sup_{f \ni x} \left\{ \mathrm{WEDF}_f(x) \right\}.$$

In this paper we do not consider shapes with holes, or equivalently, we do not consider medial axes with loops. The definition of WEDF above was made under this assumption of no loops. The measure EDF is defined in a similar way, but allows EDF to become infinite if there are loops in the medial axis [10]. We do not pursue this extension because computationally,

we require that WEDF be finite—the area of the shape is finite even when there are holes. Intuitively, the WEDF measures area, whereas EDF measures the length along the medial axis.

To define the Weighted Extended Medial Axis (WEMA), consider the path $\widetilde{f}$ in the medial axis of maximal weight $\widetilde{f} = \text{argmax}\, W(f)$. The **WEMA** is the point x such that $\text{WEDF}_{\widetilde{f}}(x) = W(\widetilde{f})/2$. For an example of the computation of WEDF and WEMA and its comparison with EDF and EMA, see Figure 12.

(a) EDF   (b) WEDF
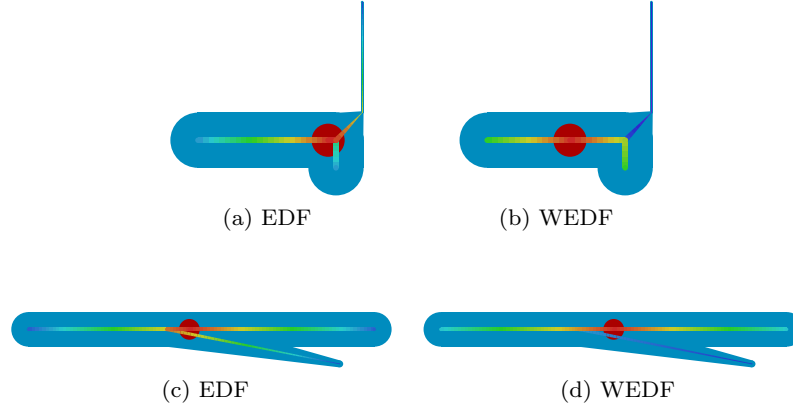
(c) EDF   (d) WEDF

Fig. 12 Comparison of EDF (a) and WEDF (b) for a shape with a thin feature and for a shape with a slowly-changing radius (c)-(d).

5.3.2 Construction of WEDF

To compute WEDF, we first discretize the medial axis and represent it as a weighted graph. Then, starting from degree-1 nodes, we compute WEDF and work inwards along the medial axis, updating WEDF for each new discretized medial point. At each junction, we choose the maximum previous WEDF value to proceed. This algorithm is analogous to the scheme for computing EDF [10]. The discrete algorithm is summarized below:

1. Initialize each degree-1 node x_i with $\text{WEDF}(x_i) = A_f(x_i)$ where f is a path bearing the node x_i. Other nodes are initialized with infinite WEDF.
2. For each degree-1 node x_i, if its neighbor x_j is not degree-1, update

$$\text{WEDF}(x_j) = \min\left(\text{WEDF}(x_j), \text{WEDF}(x_i) + |x_i - x_j| * \frac{g(x_i) + g(x_j)}{2}\right),$$

where $g(x) = R(x) \left(\sin \theta_1(x) + \sin \theta_2(x) \right)$.

3. At any junction, wait until there is only one adjacent medial point with infinite WEDF. When this is true, update the junction point using the largest of the finite adjacent WEDF values (and the formula in step 2).

4. Terminate the algorithm when all discrete medial points have been visited (equivalently, when all WEDF values are finite).

5.3.3 Feature detection

Assume that $\mathcal{S}$ is simply connected. We propose a scaling factor based on WEDF that is analogous to the EDF-based scaling. We partition the medial axis into branches. A branch is defined as a 1D subset of the medial axis on which WEDF is continuous. On a branch b, we consider the factor

$$\widetilde{\omega}(b) = \sup_{x \in b} \left\{ \text{WEDF}(x) \right\}.$$

We propose to tune the scaling factor of the SAT algorithm using the branch area proxy $\omega(b) \in (0, 1]$:

$$\omega(b) = \frac{\widetilde{\omega}(b)}{\max_{b'} \widetilde{\omega}(b')}.$$

This leads to the generalization of the Scale Axis Transform for feature detection described in Algorithm 3. For a comparison of the SAT, SAT-EDF and SAT-WEDF algorithms, see Figure 13.

Algorithm 3 The SAT-WEDF Feature Detection Algorithm

Input: Given $\mathcal{S}$, and some scaling threshold $s > 1$.

1. Compute $\mathcal{M}$, the medial axis.

2. For each branch b of the medial axis, compute $\omega(b)$, the WEDF-based area.

3. For each point x_i along $\mathcal{M}$, scale the medial ball of radius $R(x_i)$ by a factor $s \, \omega(b)$, where b is the branch that contains x_i.

4. If the medial ball of x_i is contained in the union of other medial balls, label that a medial point corresponding to a shape feature.

---

6 Algorithm Results and Comparisons

In this section, we compare the four algorithms we have discussed for detecting shape features—the geometric algorithm (cf. Algorithm 1), the standard Scale Axis Transform, the SAT-EDF algorithm (cf. Algorithm 2), and the SAT-WEDF algorithm (cf. Algorithm 3). Parts of the medial axis correspond-
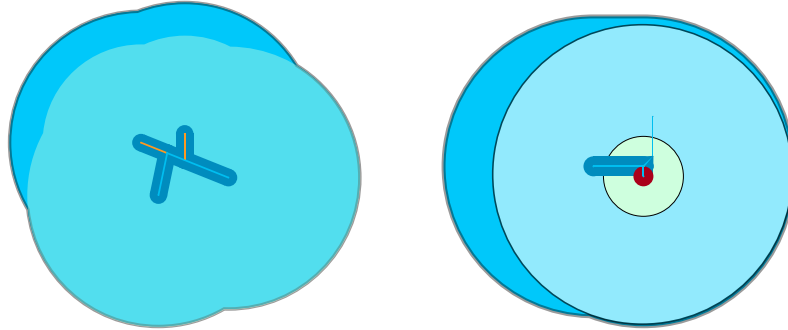
Fig. 13 These figures show dilations for the SAT (darker blue in the back) and SAT-EDF algorithms (light blue). On the left figure, since the radius is similar on all branches, the SAT will not detect any features. For SAT-EDF, the dilation is decreasing, and therefore the yellow parts of the medial axis are eventually considered features of the main shape (the blue part of the medial axis). On the right, the SAT dilation is indicated in the back as well, and the other regions correspond only to the dilation of the red point, for SAT-EDF (the color is the same blue as on the left figure) and SAT-WEDF (give by a lighter color), respectively. For SAT-EDF, we can see that the shape is less dilated on the parts that are not labeled as the main axis by EDF, and is rapidly labeled as a feature of the main shape by the dilation process. This is not the case for the SAT-WEDF region corresponding to the red point—it will never be eaten up since it belongs to the biggest branch in terms of area. This figure explains the results of Figure 19.

ing to the identified features are highlighted in yellow, whereas the part of the axis corresponding to the main shape are in blue. The geometric algorithm was based on local medial information—radius and orientation of the medial branches. The SAT algorithm was the medial pruning method that was the inspiration for the algorithms proposed [6]. As the Scale Axis Transform is not able to detect shape features when the medial radii are nearly constant, SAT-EDF corrects for this by using branch length in the SAT scaling. Similarly, to incorporate both length and thickness, the SAT-WEDF algorithm uses a proxy for branch area in the SAT scaling. The shapes on which we test the algorithms are similar to the shapes in the user study and have only the main shape and one level of features to simplify the SAT-based results. Below, we compare the results of these algorithms on a series of test shapes, and highlight the strengths and weaknesses of each approach.

Note that for some of our tests (e.g., in Figure 10), we used a simplified skeletal representation instead of the exact medial axis. We did not notice any change in the quality of results.

In Figure 14, the results of all four algorithms are shown on a shape with a thin shape feature. In this case, all algorithms are consistent, and correctly identify the thin branch as the shape feature.

In Figure 15, the results of the algorithms are shown on a different shape. In this case, the geometric algorithm, SAT-EDF and SAT-WEDF agree, and
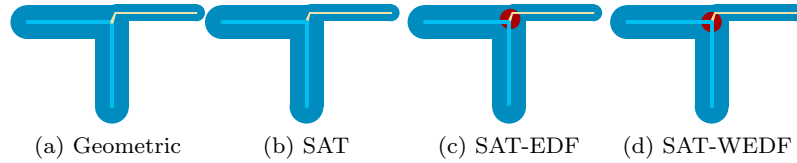
(a) Geometric   (b) SAT   (c) SAT-EDF   (d) SAT-WEDF

Fig. 14 Comparison of Algorithm Results. (a) The Geometric Algorithm, (b) The SAT, (c) EDF-Adjusted SAT with EMA in red, (d) WEDF-Adjusted SAT with WEMA in red.

correctly identify the shorter branch as a shape feature. However, the SAT algorithm cannot identify the smaller branch because the shape has near-constant thickness. This shows a case where the three algorithms proposed here outperform the SAT.
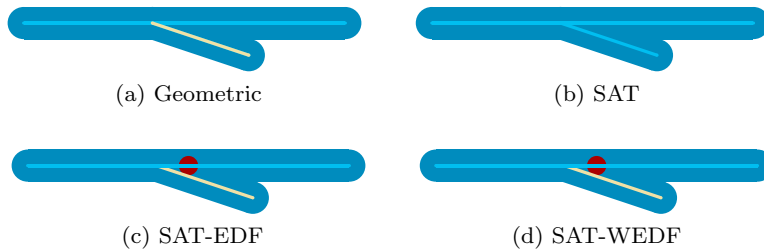

(a) Geometric   (b) SAT

(c) SAT-EDF   (d) SAT-WEDF

Fig. 15 Comparison of Algorithm Results. (a) The Geometric Algorithm, (b) The SAT, (c) EDF-Adjusted SAT with EMA in red, (d) WEDF-Adjusted SAT with WEMA in red.

Figure 16 depicts a case where the geometric algorithm outperforms the SAT-based algorithms. The geometric algorithm chooses the shape features which are most consistent with the user study (see Figure 3). The SAT algorithm does not detect shape features due to the nearly constant medial radius of the shape. The SAT-EDF and SAT-WEDF algorithms produce the same results (due to the near-constant medial radius), choosing the longest path as the main shape and the smaller branches as features, neglecting orientation.

On Figure 17 however, the local character of the geometric algorithm leads to a non-intuitive choice for the feature: the aligned segments are chosen as the main shape, despite their small size compared to a large feature. The SAT-based algorithms are able to identify the larger left part of the shape as part of the main shape, because they employ a global strategy to identify shape features. However, here all SAT-based algorithms use the same scaling factor: whereas SAT-WEDF is already able to identify the right part as the
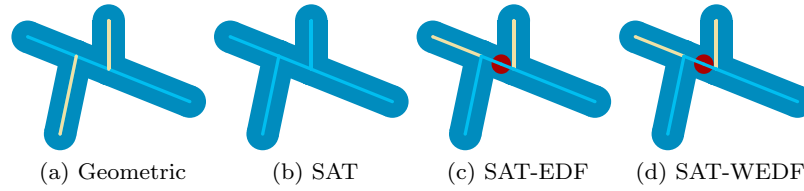
(a) Geometric   (b) SAT   (c) SAT-EDF   (d) SAT-WEDF

Fig. 16 Comparison of Algorithm Results. (a) The Geometric Algorithm, (b) The SAT, (c) EDF-Adjusted SAT with EMA in red, (d) WEDF-Adjusted SAT with WEMA in red.

feature, the SAT and SAT-EDF algorithms do not. Eventually, for larger scaling, both the SAT and SAT-EDF algorithms identify the same part as a feature.
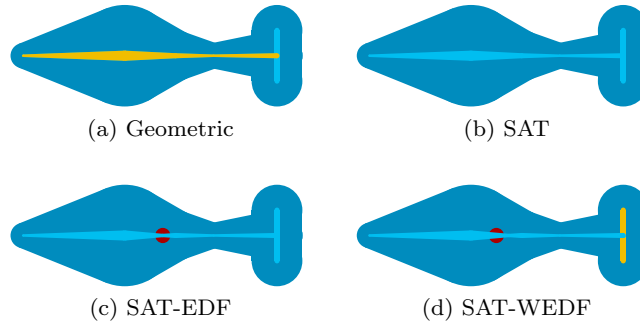

(a) Geometric   (b) SAT

(c) SAT-EDF   (d) SAT-WEDF

Fig. 17 Comparison of Algorithm Results. (a) The Geometric Algorithm, (b) The SAT, (c) EDF-Adjusted SAT with EMA in red, (d) WEDF-Adjusted SAT with WEMA in red.

In Figure 18, the results are shown for two shapes with repeated features. The SAT algorithm is unable to detect shape features due to the nearly-constant radius of the medial balls for these two shapes. However, the geometric algorithm as well as SAT-EDF and SAT-WEDF both detect the same shape features. For the shape with less protruding shape features, these results are consistent with the user study. For the shape in the first row, the end features are not detected. None of these algorithms currently detect repeated shape features. In the first row of Figure 18, the two end segments may not be considered shape features without the middle three features. This leads us to a possible extension of the current methods, based on the following statement:
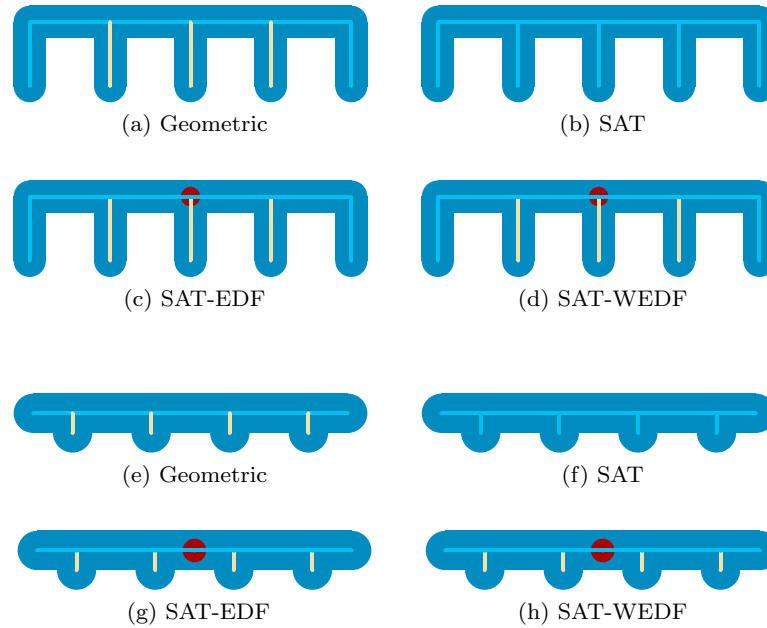
Fig. 18 Comparison of Algorithm Results on two shapes. (a) and (e) The Geometric, (b) and (f) The SAT, (c) and (g) EDF-Adjusted SAT with EMA in red, (d) and (h) WEDF-Adjusted SAT with WEMA in red.

Repeated parts of a shape are likely to be perceived as shape features.

In Figure 19, an example is shown where SAT-WEDF outperforms SAT-EDF. In this case, the geometric algorithm, the SAT algorithm, and SAT-WEDF all correctly identify the long, thin segment as the shape feature. For SAT-EDF, the algorithm considers this long thin feature to be the main shape, whereas the nob at the bottom was identified as a feature. As this was considered to be part of the main shape in the user study, these results show the sensitivity of SAT-EDF to long, thin shape parts.

These results demonstrate the strengths and weaknesses of four algorithms we tested for detecting shape features. Our conclusions are given next.

7 Discussion and Conclusion

In this paper, the perception of 2D shape features was measured via a user study. This information was then used to construct three algorithms to identify shape features. The first algorithm was a geometric algorithm based on
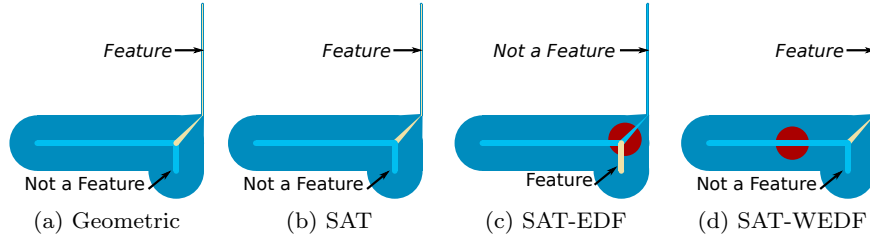
Fig. 19 Comparison of Algorithm Results. (a) The Geometric Algorithm, (b) The SAT, (c) EDF-Adjusted SAT with EMA in red, (d) WEDF-Adjusted SAT with WEMA in red.

change in radius and orientation at junctions of the medial axis. This method offered results that were consistent with the user study on most examples, but may fail as it is only based on local criteria (see Figure 17). The Scale Axis Transform uses global criteria and performs well on many examples, but cannot identify features correctly on shapes with near-constant thickness, whereas the user study shows very clear identification of features on such shapes. We thus extended the SAT algorithm to incorporate branch length (based on EDF) in the scaling. This method—SAT-EDF—is able to identify features on a shape with near-constant thickness. However, the lengthier branches are always considered the main shape, which is not consistent with the user study. The last algorithm—SAT-WEDF—was able to handle cases with long and thin features, as it uses both length and thickness information in the SAT scaling. In this case, the last algorithm reproduced well the results of the user study. However, the results of these algorithms on a set of test shapes indicated that all algorithms should be coupled with a mechanism for detecting repeated shape features.

The benefit of using the SAT-based algorithms is that they immediately offer a hierarchy of shape features by tuning the scaling threshold s. The shapes in this paper were chosen to have only two level of details (main shape and features) in order to identify perceptual features as defined by a tractable user study. These two levels are adapted for handling shape deformation operations: properties of the features may be preserved, while the main shape undergoes the deformation. However, in future work we also intend to benefit from the multiresolution inherent to the SAT based algorithms. By choosing an appropriate scaling factor, s, these algorithms automatically find the features detected under that scaling. Determining the scale of details on shapes is well-suited to be investigated in an expanded perceptual study.
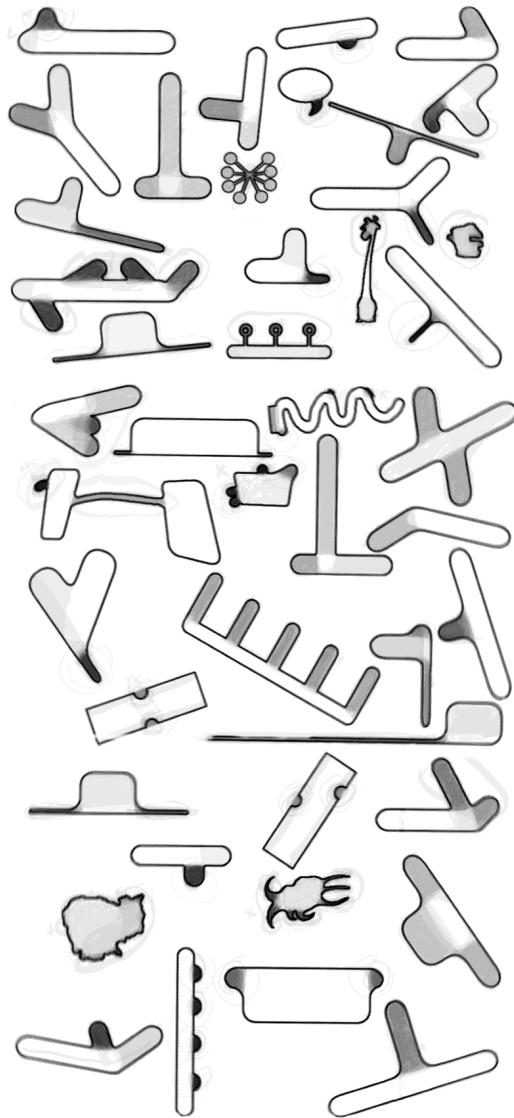
The Results of the User Study

Fig. 20 Results of the user study. This shows the average pixel value for all users on all 44 shapes administered.

References

1. N. Amenta, S. Choi, and R. K. Kolluri. The power crust. In *Proceedings of the sixth ACM symposium on Solid modeling and applications*, pages 249–266. ACM, 2001.
2. D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *Computer Vision and Image Understanding*, 67(3):261–273, 1997.
3. H. Blum et al. A transformation for extracting new descriptors of shape. *Models for the perception of speech and visual form*, 19(5):362–380, 1967.
4. H. Blum and R. N. Nagel. Shape description using weighted symmetric axis features. *Pattern recognition*, 10(3):167–180, 1978.
5. F. Chazal and A. Lieutier. The λ-medial axis. *Graphical Models*, 67(4):304–331, 2005.
6. J. Giesen, B. Miklos, M. Pauly, and C. Wormser. The scale axis transform. In *Proceedings of the 25th annual symposium on Computational geometry*, pages 106–115. ACM, 2009.
7. B. Kimia, A. Tannenbaum, and S. Zucker. Toward a computational theory of shape: An overview. In O. Faugeras, editor, *Computer Vision ECCV 90*, volume 427 of *Lecture Notes in Computer Science*, pages 402–407. Springer Berlin Heidelberg, 1990.
8. B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. Shapes, shocks, and deformations I: the components of two-dimensional shape and the reaction-diffusion space. *International journal of computer vision*, 15(3):189–224, 1995.
9. M. Leyton. Symmetry-curvature duality. *Computer vision, graphics, and image processing*, 38(3):327–341, 1987.
10. L. Liu, E. W. Chambers, D. Letscher, and T. Ju. Extended grassfire transform on medial axes of 2D shapes. *Computer-Aided Design*, 43(11):1496–1505, 2011.
11. B. Miklos, J. Giesen, and M. Pauly. Discrete scale axis representations for 3D geometry. *ACM Transactions on Graphics (TOG)*, 29(4):101, 2010.
12. L. Prasad. Rectification of the chordal axis transform and a new criterion for shape decomposition. In *Discrete Geometry for Computer Imagery*, pages 263–275. Springer, 2005.
13. T. B. Sebastian, P. N. Klein, and B. B. Kimia. Recognition of shapes by editing shock graphs. In *ICCV*, volume 1, pages 755–762, 2001.
14. D. Shaked and A. M. Bruckstein. Pruning medial axes. *Computer vision and image understanding*, 69(2):156–169, 1998.
15. A. Sud, M. Foskey, and D. Manocha. Homotopy-preserving medial axis simplification. *International Journal of Computational Geometry & Applications*, 17(05):423–451, 2007.
16. A. Tagliasacchi. Skeletal representations and applications. arXiv preprint arXiv:1301.6809, 2013.
17. R. Tam and W. Heidrich. Feature-preserving medial axis noise removal. In *Computer Vision-ECCV 2002*, pages 672–686. Springer, 2002.
18. J. Zeng, R. Lakaemper, X. Yang, and X. Li. 2D shape decomposition based on combined skeleton-boundary features. In *Advances in Visual Computing*, pages 682–691. Springer, 2008.