

Feature Preserving Smoothing of Shapes using Saliency Skeletons

Alexandru Telea

Abstract We present a novel method that uses shape skeletons, and associated quantities, for feature-preserving smoothing of shapes in digital images. We preserve, or smooth out, features based on a saliency measure that relates feature size to local object size, both computed using the shape’s skeleton. Low-saliency convex features (cusps) are smoothed out, and low-saliency concave features (dents) are filled in, respectively, by inflating simplified versions of the shape’s foreground and background skeletons. The method is simple to implement, works in real time, and robustly removes large-scale contour and binary speckle noise whereas preserving salient features. We demonstrate the method with several examples on datasets the shape analysis application domain.

1 Introduction

Noisy shapes occur in a wide variety of applications and domains, such as image acquisition and processing in medical imaging and computer vision, object retrieval and matching from shape databases, image compression, and contour simplification in computer graphics. Considerable work has been invested in developing processing methods that are able to remove certain characteristics of a given shape, regarded to be *noise* from the perspective of the application at hand, and keep (or even enhance) other characteristics, known as *features*. In the following, we shall refer to such methods as feature-preserving smoothing methods.

Numerous feature-preserving smoothing techniques have been developed. They differ in several respects, *e.g.* the definition of what are features (to preserve) and noise (to be removed), shape representation (implicit or explicit), shape dimensionality (typically 2D or 3D), and space discretization (Lagrangian or Eulerian). If we

Alexandru Telea

Institute for Mathematics and Computer Science, University of Groningen, Nijenborgh 9, Groningen, the Netherlands e-mail: a.c.telea@rug.nl

consider shapes $\Omega \in \mathbb{R}^n, n = \{2, 3\}$, as having (closed) orientable and locally differentiable boundaries $\partial\Omega$, then a common definition of features and noise is based on the analysis of local perturbations of the surface $\partial\Omega$, typically in normal direction.

Traditionally, such perturbations are measured using various instruments based on higher-order surface derivatives, such as gradients, principal components or moments. Although a wealth of such techniques exists, curvature estimation on noisy shapes, discretized or not, is still a delicate process. This type of smoothing typically implies some form of signal filtering. By *e.g.* inversely correlating the filter width with the strength of the feature signal, smoothing can achieve a certain degree of feature preservation. Filtering can be applied on several scales, thereby smoothing the shape at several levels of detail. However, filtering of discrete signals usually cannot avoid some finite amount of undesired smoothing, *e.g.* in regions where one wants to preserve features. Figure 1 illustrates the idea: Typical curvature-based smoothing will produce the smooth shape (b) from the noisy shape (a), thereby removing cusps and dents but also smoothing out the perceptually important rectangle corners. The method proposed here produces image (c), which smooths out the noise but keeps the corners sharp at pixel level.

In this paper, we approach the goal of smoothing shapes in a feature-preserving manner from a different angle. First, we characterize both features and noise on a shape’s boundary $\partial\Omega$ by a new saliency metric computed on the shape’s skeleton or medial axis $S(\Omega)$. The saliency metric relates the size of a boundary perturbation, encoded by the skeleton’s so-called importance metric, to the local object size, encoded by the shape’s distance transform. Secondly, we prune the saliency-attributed skeleton by simple upper thresholding. Due to the properties of our saliency metric on the skeleton $S(\Omega)$, this effectively removes all features below a given saliency value but fully preserves features above that value, a property which is relatively difficult to achieve when using purely local techniques. Thirdly, we reconstruct the smoothed shape by inflating the pruned skeleton. By considering both the foreground and background skeletons of a digital image, we can smooth out cusps (protrusions) and fill dents (concavities) respectively.

The paper is structured as follows. Section 2 overviews related work in the area of feature-preserving shape smoothing and related medial axis methods. Section 3 introduces our new saliency metric. Section 4 describes the feature-preserving smoothing algorithm. Section 5 presents implementation details. Section 6 presents several results obtained on different datasets related to medical and life sciences, and discusses our method. Section 7 discusses various aspects of the presented method. Finally, Section 8 concludes the paper.

2 Related work

Shape smoothing methods can be classified along the way in which features and noise are defined, measured, and represented, as follows.




Fig. 1 a) Noisy shape. b) Local smoothing methods affect the corners. c) Feature-preserving smoothing removes noise but keeps corners sharp

2.1 Local methods

A first main class of methods models features and noise locally, using the surface curvature tensor. In this way, corners of 2D shapes and corners and edges of 3D shapes can be detected. Many methods exist for curvature evaluation on discrete surfaces, *e.g.* [Moreton and Séquin, 1992, Clarenz et al., 2004a, Desbrun et al., 1999, Peng et al., 2001]. Besides curvature, surface classifiers can be based on related integral quantities, such as moments [Clarenz et al., 2004b]. Gaussian scale-space representations of progressively smoothed shapes can be obtained by successive convolutions [Koenderink, 1984].

Distinction between features and noise are typically taken by analyzing the magnitude of the detector signal. Smoothing attempts then to preserve the former, and remove the latter, by *e.g.* convolving the classifier signal with filters correlated with the signal's strength [ter Haar Romeny, 1994, Weickert, 1997, Clarenz et al., 2004a]. When normals are present in the input data, they can be used to efficiently perform feature-preserving filtering [Osher and Sethian, 1988, Fleishman et al., 2003]. However, if normals lack, their computation from position data involves a finite amount of smoothing, similar to curvature estimation. All in all, since differential classifiers decide whether a certain signal variation at point $x \in \partial\Omega$ is a feature or noise based on the analysis of a small neighborhood $N(x) \subset \partial\Omega$ or size, or diameter, δ , around x , they will typically smooth both features and noise details below scale δ . Also, discrete curvature computations need to be regularized, typically by local integration over a neighborhood of finite size smaller or equal to the noise size δ , which introduces some artificial smoothing. An early example of the challenges involved in curvature-based salient feature detection in a multiresolution setting is offered by [Fernmuller and Kropatsch, 1992]. A good overview of the above-mentioned challenges involved in discrete curve curvature computations is given in [Jalba et al., 2006].

2.2 Global methods

A second class of smoothing methods models features and noise in a more global manner, using the so-called *skeleton*, or medial axis, of a shape, defined as follows. For a shape boundary $\partial\Omega$, we first define the distance transform $D : \Omega \rightarrow \mathbb{R}_+$ and the feature transform $F : \Omega \rightarrow \mathcal{P}(\partial\Omega)$ as

$$D(x \in \Omega) = \min_{y \in \partial\Omega} \|x - y\| \quad (1)$$

$$F(x \in \Omega) = \{y \in \partial\Omega \mid \|x - y\| = D(x)\} \quad (2)$$

The skeleton of $\partial\Omega$ is defined as the locus of centers of maximally inscribed balls. Since these balls touch $\partial\Omega$ in at least two feature points [Kimmel et al., 1995], we have

$$S(\Omega) = \{x \in \Omega \mid |F(x)| \geq 2\} \quad (3)$$

It is well known that the terminations of the skeleton branches (which are curves in 2D and manifolds in 3D) map, via the feature transform, to curvature maxima on $\partial\Omega$ [Kimmel et al., 1995]. Several methods exploit this property to smooth a shape by pruning, or regularizing, its skeleton and then inflating it back. Pruning a so-called skeleton terminal branch $b \in S$ effectively corresponds to replacing the points $F(x \in b) \subset \partial\Omega$ by a circle arc (in 2D) or spherical sector (in 3D). Skeleton-based shape simplification is intimately related to differential shape processing, *e.g.* curvature flow: The so-called skeleton scale-space obtained by computing skeletons of increasingly smoothed versions of a given shape Ω , corresponds to increasingly pruning the skeleton of Ω from its endpoints to its center [Bai et al., 2007, Ogniewicz and Kubler, 1995, Pizer et al., 1987].

Hisada *et al.* detect salient features (edges) of polygonal surfaces by extracting the 3D skeleton, detecting the terminations (edges) of the separate 3D skeletal sheets, and mapping these back to the shape [Hisada et al., 2002]. Since both the shape and skeleton are represented as a non-uniformly sampled point set, problems arise with the density and continuity of the detected salient features. The skeleton computation, based on Voronoi techniques and related to the power crust [Amenta et al., 2001], is extremely noise sensitive. Robustness is achieved by Laplacian smoothing of both the surface and its skeleton, but this actually removes salient details one wants to find.

Computing robust, exact, and connected skeletons is, however, perfectly doable for noisy 2D and 3D shapes. Among others, Telea *et al.* achieved this by defining the importance $\rho(x)$ of each skeleton point x as the longest shortest-path length between any two feature points of x . Intuitively, ρ equals the boundary length subtended by a skeleton point's features. Efficient implementations are provided for 2D images [Telea and van Wijk, 2002] and 3D voxel volumes [Reniers et al., 2008]. A key observation is that ρ is minimal at skeleton endpoints and increases monotonically towards the skeleton's center or root [Ogniewicz and Kubler, 1995, Reniers et al., 2008]. Hence, a regularized robust skeleton S_ρ can be obtained by upper-thresholding ρ with a desired value ρ_{min} . Using this property, Reniers *et al.* proposed a skeleton-

based classifier for 3D shapes, which detects salient features such as valleys and ridges simply by computing the image of $F(x \in S, \rho(x) > \rho_{min})$ for a given saliency value ρ_{min} . This classifier showed to be more noise-resistant than curvature-based classifiers *e.g.* [Taubin, 1995]. However, the classifier is not further used to smooth the shape.

Tek and Kimia presented probably the earliest result in boundary smoothing using skeletons [Tek and Kimia, 2001]. They iteratively smooth a 2D shape by applying a so-called splice transform, which removes terminal branches from both the inner and outer skeletons. Removals are ordered by saliency measure, which is defined as the area difference between the smoothed and original shapes divided by skeleton branch length. The splice transform, however, needs to manipulate a relatively complex graph representation of the skeleton, and maintain explicit connections between Lagrangian curve representations of boundary segments and Eulerian representations of the distance transform.

Our proposal is related to [Tek and Kimia, 2001] as follows. We exploit the same principle of smoothing a shape by pruning its internal and external skeletons. However, our saliency metric is different (see Sec. 3), and so is the skeleton pruning order, algorithm, and results (Secs. 5 and 7). We work fully in a pixel-based (Eulerian) setting, without the need of maintaining a skeleton graph representation or to explicitly manipulate boundary curve segments represented in a Lagrangian setting. In particular, we obtain the reconstructed (smoothed) shape by inflating the regularized skeleton using its distance transform, rather than explicitly editing the boundary to replace fragments by circular arcs. This yields a much simpler overall implementation with arguably lower complexity, which enables us to smooth considerably more complex shapes (see Sec. 6).

3 Skeleton-based Saliency Metric

Our general aim is similar to that of Tek and Kimia [Tek and Kimia, 2001]: We want to build a multiscale shape representation so that *perceptually* salient features, such as sharp corners, are retained on coarse scales. Consider Figure 1: On a coarse scale, we arguably see a rectangle with sharp corners but no noise on the edges (Fig. 1 c), and not a rectangle with rounded corners (Fig. 1 b). Hence, the saliency of a feature relates not just to its size, but to whether that feature is relevant for the local interpretation of the object (see also [Dudek and Tsotsos, 1997]). In the following, we consider first convex salient features. Concave features are treated similarly (see Sec. 4.4).

The first step in our saliency metric design is to measure the size of a feature. For this, we use its boundary length, which is exactly the importance metric ρ introduced in Sec. 2 [Telea and van Wijk, 2002]. Consider the case of the noisy rectangle in Fig. 1. The color-coded importance ρ is shown in Fig. 2 a using a rainbow colormap. Terminal skeleton branch pixels have low ρ values (blue), whereas pixels closer to the skeleton center have high values (red).

Let us compare the skeleton branches corresponding to the upper-eight rectangle corner, respectively the neighboring small-scale noise in this figure. Along the corner branch, ρ increases steadily. Along the two cusp branches, ρ increases until the branches leave the cusp and enter the rump of the shape. After that, ρ stays constant on that branch, indicating the presence of so-called *ligature points* [August et al., 1999], *i.e.* skeleton points that connect a branch fragment, corresponding to a small protrusion, with the skeleton's main part.

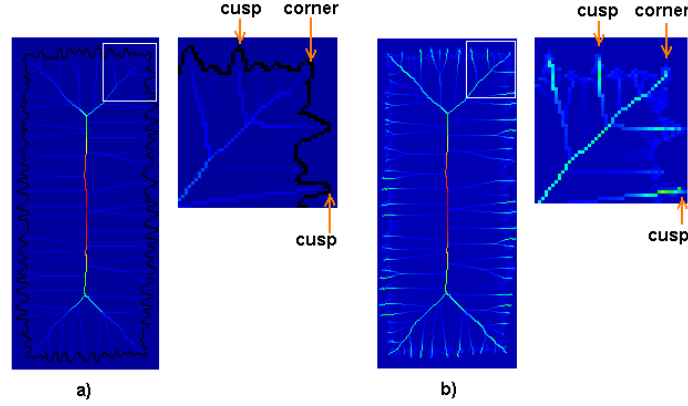


Fig. 2 Comparison of skeleton importance ρ (a) with saliency metric σ (b). While ρ stays constant along ligature branches, σ decreases markedly

By upper thresholding ρ with some desired value ρ_{min} , and inflating back the regularized skeleton S_ρ , we achieve smoothing which essentially replaces all boundary features shorter than ρ with circle arcs (Fig. 1 b). However, what we want is to keep the rectangle's corners sharp and smooth out the cusp.

We first make two observations concerning the perceptual saliency of a shape detail:

- saliency is proportional with *size*, which can be measured by boundary length. Longer features are more salient than shorter ones [Ogniewicz and Kubler, 1995];
- saliency is inversely proportional with local object *thickness*. A feature located on a thick object is less salient than the same feature located on a thin object [Tek and Kimia, 2001].

Hence, we can define a saliency metric σ on the skeleton of a shape as

$$\sigma(x \in S(\Omega)) = \frac{\rho(x)}{D(x)} \quad (4)$$

where $\rho(x)$ is the skeleton importance defined as in [Telea and van Wijk, 2002] and $D(x)$ is the distance transform (Eqn. 1).

Figure 2 b shows the saliency computed for the noisy rectangle shape. We see that σ stays constant along the branches corresponding to the rectangle corners, but decreases rapidly along ligature portions of the small-scale cusp branches, as indicated from the light-blue to dark-blue color transitions. Following Eqn. 4, σ is 0 for all non-skeletal points, has a constant value of $2/\tan(\alpha)$ for points along the skeleton branch of an angular cusp of aperture angle α , and has a theoretical maximal value of $|\partial\Omega|/\Phi$, where $|\partial\Omega|$ is the shape's perimeter and $\Phi = \min_{x \in S(\Omega)} D(x)$ is the shape's minimal local thickness.

Computing σ (Eqn. 4) for any pixel or voxel-based shape is immediate. Here, we use the AFMM Star implementation [Reniers and Telea, 2007, Telea and van Wijk, 2002] which delivers us the skeleton S , importance ρ , and distance and feature transforms D and F .

4 Saliency-preserving smoothing

Now that we have a saliency indicator, we can use it to selectively smooth out low-saliency features. In the following, we consider low-saliency convex features, or *cusps* (see *e.g.* Fig. 1). Concave features are treated analogously (see Sec. 4.4). The entire smoothing pipeline, illustrated in Fig. 3, is described below.




Fig. 3 Feature-preserving smoothing pipeline. See Section 4) for details

4.1 Saliency-based pruning

The first smoothing step is to prune the skeleton S to S_σ retaining all points where σ exceeds a minimal saliency σ_{min}

$$S_\sigma = \{x \in S \mid \sigma(x) \geq \sigma_{min}\} \quad (5)$$

In Eqn. 5, S denotes the *full* skeleton of a given shape Ω . In practical terms of the implementation used (see Sec. 5), this means all image points where $\rho > 1$. Using the full skeleton, thus having a method able to compute it with pixel precision, is essential, as we later inflate this skeleton and wish to exactly reconstruct salient features (Sec. 4.3).

4.2 Connected component detection

Since the signal σ is not monotonic over Ω , S_σ will be a collection of disconnected, compact, skeletal components. For 2D shapes, disconnection immediately follows the thresholding in Eqn. 5, since the skeleton S computed by the AFMM is always one pixel thick [Reniers and Telea, 2007]. Note, however, that this is not always valid for 3D shapes (Sec. 7). These skeletal components correspond to multiple small fragments located inside low-saliency features, and one large 'core' fragment containing the skeleton rump. Note that this holds also for higher-genus shapes, *i.e.* shapes with holes.

The second step isolates the skeleton rump component, $S'_\sigma \in S_\sigma$, defined as the connected component of S_σ which passes through $\max_{x \in \Omega} \rho(x)$. This is easily done *e.g.* by performing a 8-connected flood-fill on S_σ from the maximum of ρ . S'_σ is a different regularization of S than S_ρ introduced in Sec. 3. While S_ρ eliminates both ligature branches corresponding to small-scale noise, it also shortens branches corresponding to salient features by the same factor ρ . In contrast, S'_σ eliminates *entire* branches that have ligature components, but does not touch salient feature branches.

4.3 Reconstruction

The third and last step inflates S'_σ back to reconstruct a smoothed shape Ω' . For this, we use the well-known Fast Marching Method [Sethian, 1999]: We evolve each point $x \in S'_\sigma$ until it reaches its distance transform value $D(x)$. To simplify the stopping criterion, we actually solve

$$\begin{aligned} \nabla T &= 1 \\ T|_\Omega &= -D \end{aligned}$$

on the image domain, and obtain the reconstructed shape as the level set $T(0)$. This effectively removes all shape details which correspond to the pruned skeleton branches.

4.4 Concave features

The result of the three-step procedure outlined above effectively eliminates small-scale *convex* features (cusps), while it preserves the coarse-scale convex features intact. However, we would like to treat *concave* features (dents) similarly. For this, we cannot use the shape's skeleton directly. While convexities map to terminal branches, which we can easily prune as described in Sec. 4.1, concavities map to so-called inner skeleton branches [August et al., 1999], which cannot be edited with the same ease, as they also describe other non-noise shape points.

An elegant solution is to use the background skeleton \bar{S} , *i.e.* the skeleton of all points $\bar{\Omega}$ located outside the shape Ω . While S describes (and allows the selective smoothing of) the shape's convexities, \bar{S} describes, and allows the selective smoothing of, the shape's concavities. Figure 3 (bottom row) illustrates this by showing the entire process described in Secs. 4.1-4.3 for the background of the rectangular shape in the image. The background skeleton \bar{S} , saliency $\bar{\sigma}$ and simplified skeleton \bar{S}' are computed using the same settings, *e.g.* σ_{min} , as their foreground counterparts. Note that the background saliency image in Fig. 3 is almost everywhere dark blue. This is correct, since $\bar{\sigma}$ is very low everywhere for this shape. Indeed, the shape has no salient *dents* which are to be preserved, like it would be the case, for example, with the four concave corners of a cross shape.

When computing \bar{S} , we must define a closed shape $\bar{\Omega}$ from the background pixels. One border of this shape will be the actual foreground shape Ω . We create the other, outer, border by artificially adding a thin one-pixel border to the image. This explains the form of the background skeleton \bar{S} in Fig.3 bottom.

4.5 Convex and concave smoothing combination

If we simply apply the pipeline described above separately to the foreground, respectively background, of a given image, we obtain two smoothed shapes Ω' and $\bar{\Omega}'$ from which the cusps, respectively dents of the original shape Ω have been removed. However, how to combine these two shapes in one final result? Using *e.g.* a distance-based interpolation between Ω' and $\bar{\Omega}'$ would be doable, but it would have the undesired effect of reducing the elimination of the cusps and dents to roughly half.

A better solution is to apply the concave smoothing pass (Sec. 4.4) on the background of the output of the convex smoothing pass, *i.e.* the background of Ω' (Sec. 4.3) rather than on the background of the original image Ω . This effec-

tively combines the dent and cusp removal in one single result. Intuitively, the foreground skeleton simplification pulls cusps inwards, whereas the background skeleton simplification pulls dents outwards.

The bottom-right image in Fig. 3 shows the final result. Note that the foreground skeleton simplification can also affect branches which describe the convex borders of a dent. For example, the smoothed Ω' in the same figure shows a rounded dent on the left side. However, this is not a problem: This dent will be further removed by the background skeleton simplification pass.

5 Implementation

The entire smoothing method is implemented in under 300 lines of C++, based on the AFMM Star method, which provides skeletons, skeleton importance values, and feature transforms of digital images [Reniers and Telea, 2007]. As expected, the method is very efficient: Images of roughly 1000^2 pixels can be processed in under one second on a 2.5 GHz PC computer. For an image of N pixels, the complexity of the algorithm is $N \log N$, determined by the underlying Fast Marching Method [Sethian, 1999].

6 Results

Figure 4 shows feature-preserving smoothing applied to several datasets.

The first two rows (Fig. 4 a-h) show pairs of noisy shapes and their respective smoothings. The shapes are used as test cases for several papers in the image processing community (*e.g.* [Jalba et al., 2006, Ogniewicz and Kubler, 1995]). Small-scale noise is eliminated, and sharp coarse-scale corners are kept unsmoothed, *e.g.* the tails of the leafs. The method is able to handle relatively large-scale noise, as illustrated by the star figure.

Multiscale smoothing: By successively increasing the salience threshold σ_{min} , increasingly more sharp features are removed, as illustrated by Fig. 4 i-l, where σ_{min} was successively increased by 10%. The notion of scale corresponds to the definition of σ (Eqn. 4), *i.e.* feature size as a fraction of local object thickness. Fine scales capture small wiggles at thick object parts, coarse scales capture large boundary perturbations at thinner object parts.

The simplification of the foreground and background skeletons by increasing σ_{min} successively removes terminal branches, and hence only replaces (but does not add) sharp shape features by circle arcs with radii higher than $\min_{p \in \partial\Omega, q \in S_\sigma} \|p - q\|$. As such, the so-called 'scale-space causality' is respected, so the space of progressively simplified shapes under σ has the necessary properties of a scale space. A similar observation was made by [Tek and Kimia, 2001] for their skeleton-based

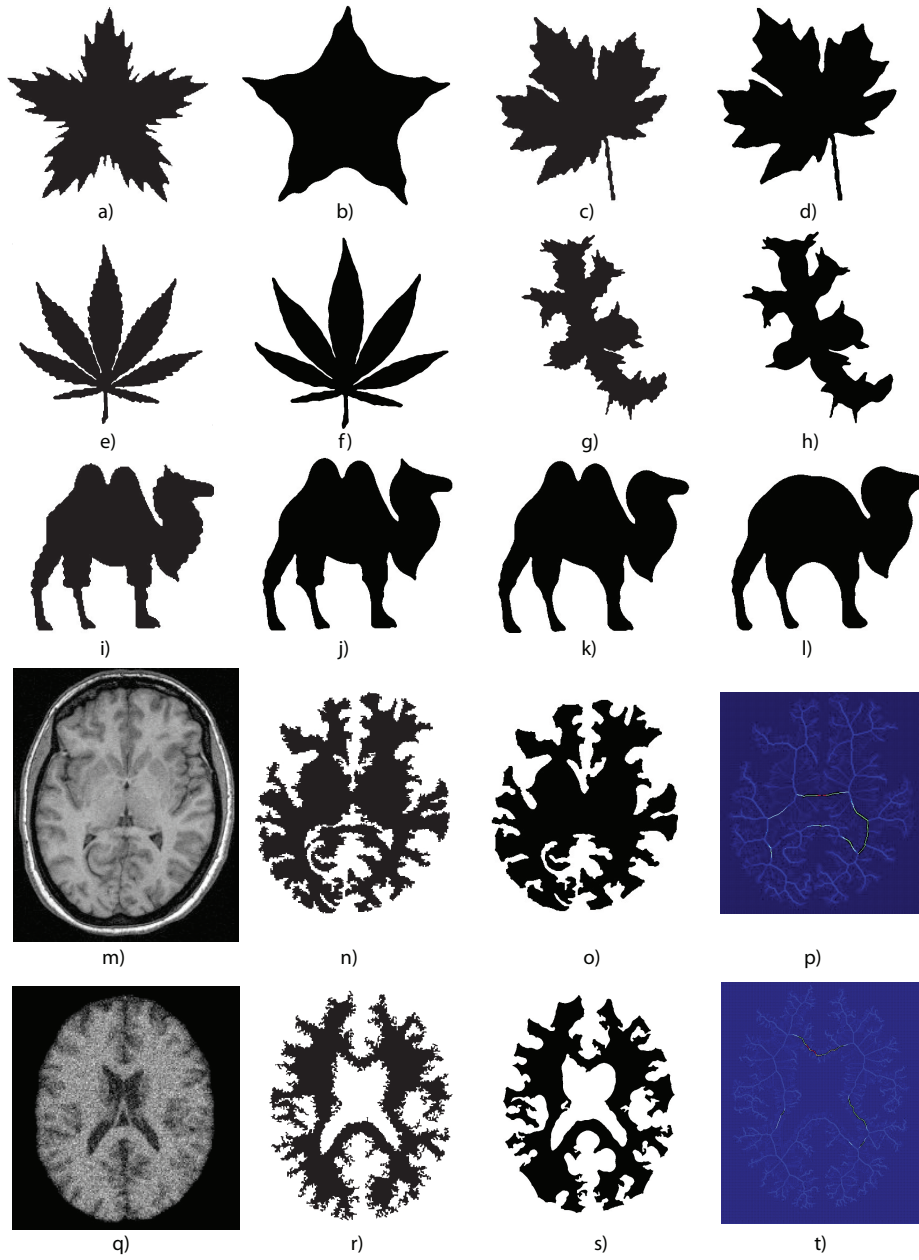


Fig. 4 Feature-preserving shape smoothing examples. See Section 6 for explanations.

simplification method.

Complex shapes: We have tested our method also on more complex shapes. Figure 4 m-p and q-t show the smoothing of two value-based segmentations of white matter in MRI grayscale images, taken from [Shah, 2009, Rehm et al., 2004]. Since no preprocessing of the segmented data was done, apart from a simple morphological pass to fill in tiny holes and remove specks, the shapes are quite noisy and complex (Fig. 4 n,r). They exhibit both segmentation noise, visible as fine-level 'wiggles' and filament-like structures along the boundary, but also actual shape information, visible as undulations of the boundary at a level of detail slightly coarser than the noise. In certain areas, drawing a border between the actual shape boundary and noise is quite challenging. The saliency values σ are shown in Figs. 4 p,t using a blue (low) to red (high) colormap. The final results show that small-scale noise has been eliminated, but the boundary has not been excessively rounded and smoothed (Fig. 4 o,s).

As a final example of the ability of the presented method to handle highly noisy shapes, we show a smoothing of a value-based segmentation of the MRI image from Fig.4 q *prior* to performing morphological hole filling and small-scale island removal. The input image in Fig. 6 a, which is actually a set of many high-genus shapes with complex boundaries, is smoothed to yield the output in Fig. 6 b. The simplification of the foreground and background skeletons achieves similar effects to the morphological operators used to create Fig. 4 r, at no added cost. The explanation for this capability resides in the definition of σ : small-scale 'islands' in the image have a low eccentricity (perimeter to maximal thickness ratio), *i.e.* a low saliency, thus are readily removed by low σ_{min} thresholds.




Fig. 5 Smoothing of highly noisy images achieves the effect of morphological erosion and dilation operators

7 Discussion

Below we discuss a number of aspects of our proposed method.

Comparison with [Tek and Kimia, 2001]: Figure 6 compares our method with the one proposed in [Tek and Kimia, 2001]. The image in Figure 10 b taken from the respective publication. The initial noisy image is shown in black outline; the smoothed result is filled green. Two insets show zoomed-in details. Looking at these, we see that our method constructs the smoothed edge roughly *between* the cusp and dent extremas, thereby behaving like an approximator of these points. In contrast, the method of Tek and Kimia tends to pull the smoothed contour outwards in most places, thereby preferring to fill the dents rather than remove the cusps. This is also visible in other images from the respective paper. The explanation lies in the different type of saliency metric used and performed optimization: While we consider feature size related to object thickness, Tek and Kimia remove branches in increasing order of affecting the overall shape’s area. Also, we apply the dent smoothing pass after the cusp filling pass, whereas Tek and Kimia interlace the two passes.




Fig. 6 Feature-preserving shape smoothing. Tek and Kimia (a) and current method (b). Original shape in black outline. Smoothed shape is shown filled. Insets show selected zoom-in details

Parameter setting: The method has one adimensional parameter: σ_{min} , with a geometric meaning. For example, setting σ_{min} to 0.1 means removing all wiggles smaller than 10% of the local object thickness. The images in this paper have been produced with $\sigma_{min} \in [0.2, 0.3]$. Setting σ_{min} is quite robust to small variations: In contrast to *e.g.* curvature-based methods, skeleton-based removal of noise is a *discrete* process, as noise details go out in a successive sequence as skeleton branches get pruned out. For all shapes we experimented with, we noticed there are only a few critical values of σ_{min} where notable removal events happen. Similar observation were made for a different type of shape descriptor, shock graphs, by [Sebastian et al., 2004, Siddiqi et al., 2004].

Saliency measure: The saliency metric proposed here essentially favors sharp corners which are far away from locally thick shape parts. As such, smoothed shapes will exhibit long straight edges, without wiggles, but keep convex or concave corners which connect coarse-scale edge-like structures. This is in line with the perceptual saliency model proposed by [Alter and Basri, 1995]. If desired, other saliency measures can be easily incorporated. For example, given a skeleton point $x \in S(\Omega)$, the A of the feature corresponding to x can be computed by cutting away the feature from the shape's rump using the line segment delimited by the furthest two feature points of x (see Fig. 7)

$$\{p_1, p_2\} = \arg \max_{p_1, p_2 \in F(x)} \|p_1 - p_2\| \quad (6)$$




Fig. 7 Definition of additional skeleton-based saliency measures (see Sec. 7)

Since x is a skeleton point, it is the center of a maximal ball C touching $\partial\Omega$ at p_1 and p_2 , so the segment p_1p_2 is fully enclosed in Ω . Defining a saliency based on the so-called feature aperture angle $\alpha = \widehat{p_1xp_2}$ can be done analogously.

The line segment p_1p_2 defined by Eqn. 6 represents the thickness of the feature at its base, *i.e.* where it joins the main shape rump. If we want to smooth features based on their thickness, one metric which can be used is $\varepsilon = \lambda\sigma$, where $\lambda = |p_1p_2|$. Figure 8 a-d illustrates this possibility. Here, the noisy brain segmentation contour from Fig. 4 n was smoothed progressively for increasing values of ε . Comparing with Fig. 4 s, we see now that features which are thin at the base *and* have a relatively small length, are removed. Thereby, relatively narrow and short dents and cusps disappear more rapidly, but without the effect of excessive overall rounding.

Combined smoothing: In this paper, we have used the *full*, unsimplified, skeleton as a starting point for the saliency computation (Sec. 4). This guarantees that sharp corners are precisely kept. However, if desired, one can combine our saliency-based smoothing, based on σ metric, with classical skeleton pruning based on the im-




Fig. 8 Progressive smoothing based on feature saliency and thickness (Sec. 7)

portance ρ , *e.g.* by first performing the ρ -based pruning, followed by our σ -based smoothing. This allows combining smoothing of all corners with removal of noise details. Studying the scale-space properties of the two-dimensional space created by ρ and σ is a potentially interesting avenue.

Selective processing: Since cusps and dents are detected separately, they can be processed accordingly with no additional effort. For example, one can imagine applications where cusps are to be smoothed more aggressively than dents, or conversely.

3D extension: Computing the skeletons, distance and feature transforms, and ρ and σ metrics is immediate for 3D shapes, using *e.g.* the approach in [Reniers et al., 2008]. The only step of our smoothing method which does not immediately generalize to 3D shapes is the connected component detection following the saliency thresholding (Sec. 4.2). Thresholding σ will, in general, not disconnect the 3D skeleton. Think for example of a cubic shape having a small-scale noisy cusp that extends from the middle of a face to one of the cube’s edges. In 2D, this problem does not appear: Removing a pixel from a skeleton branch corresponding to an arbitrary size cusp or dent always disconnects the skeleton. We are currently investigating an extension of the current analysis to handle 3D shapes.

Limitations: It is important to acknowledge several limitations of our method. First and foremost, this method is *not* to be used as a universal replacement for other denoising methods such as local curvature-based ones. It performs best on shapes which have inherent angular features, which cannot be distinguished locally from high-amplitude boundary noise. Noisy shapes without such features can be better smoothed by local methods. It would be interesting to study different application-specific, quasi-global, saliency metrics in this context, as our framework should accommodate a wide range hereof. Secondly, our current implementation only handles shapes described as boundaries, not full grayscale images. For the latter, skeletons are not (yet) effective descriptors.

8 Conclusion

We have presented a method for feature-preserving smoothing of digital shapes using shape skeletons. Using a saliency metric that relates the importance of a shape feature to its size and the local object thickness, this method is able to smooth relatively large-scale boundary noise and fully preserve sharp convex and concave corners. Another application is creating compact shapes from images corrupted by high amounts of binary speckle noise. The method uses skeletons at various places: to compute feature sizes, local object thickness, and to perform the actual feature smoothing. A simple and efficient implementation is presented here that can robustly handle complex shapes.

In the future, we aim to extend this method to efficiently and effectively handle 3D shapes, as outlined in Sec. 7. Another extension is the ability to handle full grayscale images, *e.g.* by representing these as a stack of skeletons for different isophotes.

References

- [Alter and Basri, 1995] Alter, T. and Basri, R. (1995). Extracting salient curves from images: An analysis of the saliency network. *Intl. J. of Computer Vision*, 27(1):51–69.
- [Amenta et al., 2001] Amenta, N., Choi, S., and Kolluri, R. (2001). The power crust. In *Proc. Solid Modeling*, pages 249–260. IEEE.
- [August et al., 1999] August, J., Siddiqi, K., and Zucker, S. W. (1999). Ligature instabilities in the perceptual organization of shape. *Comput. Vis. Image Underst.*, 76(3):231–249.
- [Bai et al., 2007] Bai, X., Latecki, J., and Liu, W. Y. (2007). Skeleton pruning by contour partitioning with discrete curve evolution. *IEEE TPAMI*, 29(3):449–462.
- [Clarenz et al., 2004a] Clarenz, U., Diewald, U., and Rumpf, M. (2004a). Processing textured surfaces via anisotropic geometric diffusion. *IEEE Transactions on Image Processing*, 13(2):248–261.
- [Clarenz et al., 2004b] Clarenz, U., Rumpf, M., and Telea, A. (2004b). Robust feature detection and local classification for surfaces based on moment analysis. *IEEE TVCG*, 10(5):516–524.
- [Desbrun et al., 1999] Desbrun, M., Meyer, M., Schröder, P., and Barr, A. H. (1999). Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. Intl. Conf. on Computer Graphics and Interactive Techniques*, pages 317–324.
- [Dudek and Tsotsos, 1997] Dudek, G. and Tsotsos, J. K. (1997). Shape representation and recognition from multiscale curvatures. *Computer Vision and Image Understanding*, 68:170–189.
- [Fermuller and Kropatsch, 1992] Fermuller, C. and Kropatsch, W. (1992). Multiresolution shape description by corners. In *Proc. CVPR*, pages 271–276. IEEE.
- [Fleishman et al., 2003] Fleishman, S., Drori, I., and Cohen-Or, D. (2003). Bilateral mesh denoising. *ACM Trans. on Graphics*, 22(3):950–953.
- [Hisada et al., 2002] Hisada, M., Belyaev, A., and Kunii, L. (2002). A skeleton-based approach for detection of perceptually salient features on polygonal surfaces. *Computer Graphics Forum*, 21(4):689–700.
- [Jalba et al., 2006] Jalba, A., Wilkinsoon, M., and Roerdink, J. (2006). Shape representation and recognition through morphological scale spaces. *IEEE Trans. Image Processing*, 15(2):331–341.
- [Kimmel et al., 1995] Kimmel, R., Shaked, D., and Kiryati, N. (1995). Skeletonization via distance maps and level sets. *Comp. Vision and Image Understanding*, 62(3):382–391.
- [Koenderink, 1984] Koenderink, J. J. (1984). The structure of images. *Biological Cybernetics*, 50:363–370.

- [Moreton and Séquin, 1992] Moreton, H. and Séquin, C. (1992). Functional optimization for fair surface design. In *Proc. ACM SIGGRAPH*, pages 167–176.
- [Ogniewicz and Kubler, 1995] Ogniewicz, R. and Kubler, O. (1995). Hierarchic Voronoi skeletons. *Patt. Recogn.*, 28(3):343–359.
- [Osher and Sethian, 1988] Osher, S. and Sethian, J. (1988). Fronts propagating with curvature-dependent speed. *J. of Computational Physics*, 79:12–49.
- [Peng et al., 2001] Peng, J., Strela, V., and Zorin, D. (2001). A simple algorithm for surface denoising. In *Proc. IEEE Visualization*, pages 107–112.
- [Pizer et al., 1987] Pizer, S. M., Oliver, W. R., and Bloomberg, S. H. (1987). Hierarchical shape description via the multiresolution symmetric axis transform. *IEEE TPAMI*, 9(4):505–511.
- [Rehm et al., 2004] Rehm, K., Schaper, K., Anderson, J., Woods, R., Stoltzner, S., and Rottenberg, D. (2004). Putting our heads together: a consensus approach to brain/non-brain segmentation in T1-weighted MR volumes. *NeuroImage*, 22(3):1262–1270.
- [Reniers and Telea, 2007] Reniers, D. and Telea, A. (2007). Tolerance-based distance transforms. In *Advances in Computer Graphics and Computer Vision (eds. J. Braz, A. Ranchordas, H. Araujo, J. Jorge)*, pages 187–200. Springer LNCS.
- [Reniers et al., 2008] Reniers, D., van Wijk, J. J., and Telea, A. (2008). Computing multiscale curve and surface skeletons of genus 0 objects using a global importance measure. *IEEE TVCG*, 14(2):355–368.
- [Sebastian et al., 2004] Sebastian, T. B., Klein, P. N., and Kimia, B. (2004). Recognition of shapes by editing their shock graphs. *IEEE TPAMI*, 26(5):550–571.
- [Sethian, 1999] Sethian, J. A. (1999). *Level set methods and fast marching methods: Evolving interfaces in computational geometry, fluid mechanics, computer vision, and materials science*. Cambridge Univ. Press.
- [Shah, 2009] Shah, M. (2009). MRI brain image segmentation. www.81bones.net/mri/mri_segmentation.pdf.
- [Siddiqi et al., 2004] Siddiqi, K., Shokoufandeh, A., Dickinson, S., and Zucker, S. W. (2004). Shock graphs and shape matching. *Intl. J. of Computer Vision*, 35(1):13–32.
- [Taubin, 1995] Taubin, G. (1995). Estimating the tensor of curvature of a surface from a polyhedral approximation. In *Proc. Intl. Conf on Computer Vision*, pages 902–917. IEEE.
- [Tek and Kimia, 2001] Tek, H. and Kimia, B. (2001). Boundary smoothing via symmetry transforms. *J. of Math. Imaging and Vision*, 14:211–223.
- [Telea and van Wijk, 2002] Telea, A. and van Wijk, J. J. (2002). An augmented fast marching method for computing skeletons and centerlines. In *Proc. EG/IEEE Symp. on Data Visualization (VisSym)*, pages 251–259.
- [ter Haar Romeny, 1994] ter Haar Romeny, B. (1994). *Geometric-Driven Diffusion in Computer Vision*. Kluwer.
- [Weickert, 1997] Weickert, J. (1997). A review of nonlinear diffusion filtering. In *Proc. Intl. Conf. on Scale Space*, pages 3–28. Utrecht, the Netherlands.