International Journal of Computer Vision 55(2/3), 139–153, 2003

Untangling the Blum Medial Axis Transform

ROBERT A. KATZ AND STEPHEN M. PIZER
University of North Carolina at Chapel Hill
katz@cs.unc.edu
smp@cs.unc.edu

Received July 19, 2001; Revised December 23, 2002; Accepted December 26, 2002

Abstract. For over 30 years, Blum's Medial Axis Transform (MAT) has proven to be an intriguing tool for analyzing and computing with form, but it is one that is notoriously difficult to apply in a robust and stable way. It is well documented how a tiny change to an object's boundary can cause a large change in its MAT. There has also been great difficulty in using the MAT to decompose an object into a hierarchy of parts reflecting the natural parts-hierarchy that we perceive. This paper argues that the underlying cause of these problems is that medial representations embody both the substance of each part of an object and the connections between adjacent parts. A small change in an object's boundary corresponds to a small change in its substance but may involve a large change in its connection information. The problems with Blum's MAT are generated because it does not explicitly represent this dichotomy of information. To use the Blum MAT to it's full potential, this paper presents a method for separating the substance and connection information of an object. This provides a natural parts-hierarchy while eliminating instabilities due to small boundary changes. The method also allows for graded, fuzzy classifications of object parts to match the ambiguity in human perception of many objects.

Keywords: medial, medial analysis, form analysis, Blum MAT, object part hierarchies

1. Introduction

The lure of a form representation that inherently reflects the perceptual qualities of form has long attracted scientists. Blum proposed his Medial Axis Transform (MAT) as a representation that embodies the skeleton of an object as well as the width of the object at every point on the skeleton (Blum, 1967). This work has spawned a tremendous amount of research into the use of the Blum MAT and other skeleton representations (Blum and Nagel, 1978; Brady and Asada, 1984; Bruce et al., 1985; Ogniewicz and Ilg, 1992; Pizer et al., 1987; Pizer et al., 1998; Szekely, 1996). A goal of much of this work has been to create a form representation that defines a natural decomposition of an object into a set of basic parts that mirrors the object parts we perceive. At the same time, these representations are intended to allow easy access to the full form information about each part and about the object as a whole to support form analysis and computation.

A driving consideration faced by all skeleton methods is the condition put forth by Marr and Nishihara that small changes to a boundary's form should cause only small changes to the form's representation (Marr and Nishihara, 1978). This rule insists that skeletons remain stable under boundary change, and considerable effort has been spent on modifying Blum's MAT to exhibit this kind of stability (August et al., 1999; Szekely, 1996). Complicating this rule is the problem that some boundary perturbations actually are important; the growth of a new tail or limb from a body, for instance, can be a small boundary change that has a large perceptual importance. Significant boundary perturbations correspond to the parts of an object; examples are the fingers on a hand, legs on a table and limbs on a body. In order to decompose an object into its parts,

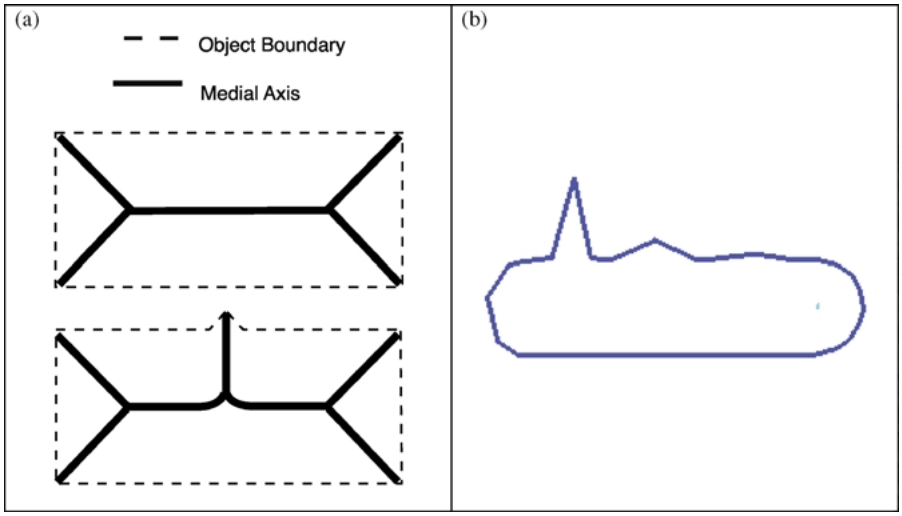

Fig. 1. Challenges with using the Blum MAT. (a) MAT instabilities: A tiny change in the boundary produces a large change in the MAT. (b) Part/protrusion ambiguity: Classifying a protrusion as a bump or separate part can be ambiguous.

any general form description needs to be able to make a stable and clear distinction between an insignificant bump on the boundary and a separate and significant part of the object.

Used as originally proposed, the Blum MAT falls short of each of these goals (see Fig. 1). It is notoriously unstable since even tiny perturbations can cause the creation of very long segments with prominent radius values to be added to the MAT. In fact, instabilities in the MAT can arise even under boundary smoothing (August et al., 1999). Furthermore, the MAT's network of branching axis segments quickly grows as object complexity grows, causing difficulty in creating a MAT hierarchy that reflects an object's perceptual parts. Without a clear hierarchy, there is no distinction between object parts and insignificant object protrusions.

This paper offers a new method for using the Blum MAT that does not suffer from these problems. We argue that objects are an assemblage of solid parts and that form representations should explicitly separate the substance of each part from the connections between parts. The method proposed here adds a component to the MAT that describes for every point on the MAT the degree to which the point embodies object substance as opposed to object connection. We call this measure the *substance* value. The measure easily reveals the natural distinction between the parts of an object and mere bumps, and it is a continuous measure that reflects the naturally ambiguous part classification found in many objects. The new method also reduces small boundary perturbations to small substance changes in the new weighted representation, maintaining a stable form description even with changes to the boundary. The new method permits a substance-weighted application of the MAT where the MAT pieces that cause instabilities in the traditional Blum representation are weighted by small values of the new substance measure. In this way, the method presented here maintains a stable form description even with changes to the boundary.

While this paper describes how to instrument the Blum MAT to separate substance and connection, we anticipate that these ideas are applicable to the general class of skeletons and medial form representations.

2. Objects are Substance and Connections

Objects in our physical world are solid, tangible entities. They have mass and fill space. Physical objects are real substance and not mathematical abstractions. Complex objects are perceived as collections of many solid parts. A boulder is a round-ish blob, and hands are a blob (the region between the wrist and the fingers) with five elongated protrusions (the fingers).

For multi-part objects, we perceive not only the substance of each part but also the connections between parts. A hand is a hand not just because it has a palm and five fingers, but because of the way that the parts are connected. The same five parts can be connected in ways that are not at all hand-like. Clearly, both the

substance information in each part of an object and the connection information between parts are crucial components of form description.

Another type of connectivity is also implicit in skeleton methods. If we consider each part of an object as a series of tiny slices, then each slice is connected to its neighboring slices to form the continuous substance of the part. This *continuity* information is usually built into the representation used for a skeleton; for the piecewise linear boundaries used in this work, Blum's MAT uses line segments and parabolic segments to represent the skeleton, and the continuity along each of these segments is implicitly defined. The distinction between continuity along a single part and connections between parts is exploited in this paper.

Analyzing and computing with form requires access to substance information and connection information. Form researchers want to be able to work explicitly with the substance information for calculations involving the perceptual properties of an object such as saliency, and they want to work explicitly with the connection information for analyzing the decomposition of objects into parts and analyzing the relationships between parts. Medial form representations naturally include substance information. By representing objects as "middle" points and the "widths" at those points, medial representations define an object from it's inside outward thus implicitly defining a solid object. Medial representations also can include connection information. In the same way that we can trace our own bones to know that our legs are connected to the bottom of our torso, medial representations can naturally define the connectivity of an object.

While medial representations are natural carriers of substance and connection information, many implementations do not harness the full potential of the two components. For example, medial methods such as SLS (Brady and Asada, 1984) and Cores (Pizer et al., 1987) provide substance information but have no explicit connection information. Fig. 2 shows how the Blum MAT contains both substance and connection information, but they are not easily separable for creating parts-hierarchies or for form computation. The next section describes what it is about the standard approaches to using the Blum MAT that fails to discriminate between substance and connection, and Section 4 presents a method that successfully separates the two form components.

3. Parts Hierarchies and Form Information in the Blum MAT

In order to solve the problems inherent in the Blum MAT, we must understand their causes. A complex object with many parts and many bumps generates a MAT with a complex web of axes. There are far more branching axis segments than the number of parts in the object, and creating a parts-hierarchy requires deciding which axis branches correspond to part substance and which to part connections. Previous hierarchy methods have attempted to impose a binary hierarchy, examining every axis point to determine if it belongs with its neighbors in a single object part, either with its neighbors in a single part or on a new object part that is different from its neighbors, or if it should be pruned away and ignored as merely connection information. Shaked and Bruckstein give a good overview of pruning methods (Shaked and Bruckstein, 1998). Fig. 3 shows how using binary decisions to create parts-hierarchies is highly unstable. A small change to the boundary can cause a drastic change in the topology of the hierarchy,

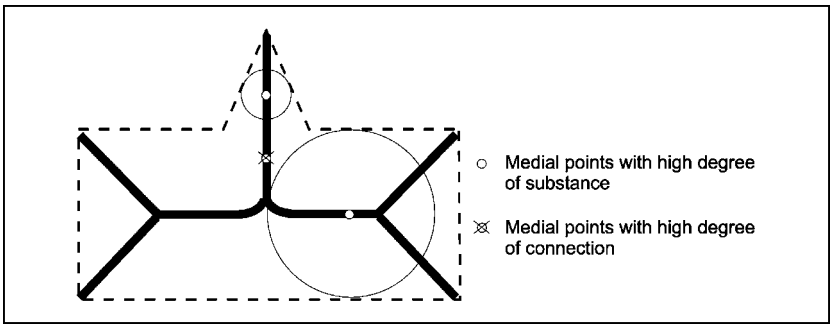

Fig. 2. Examples of medial points with high degree of substance information and with high degree of connection information. Also shown is a perceptual aperture around the substance points, showing the aperture within which computations are performed.

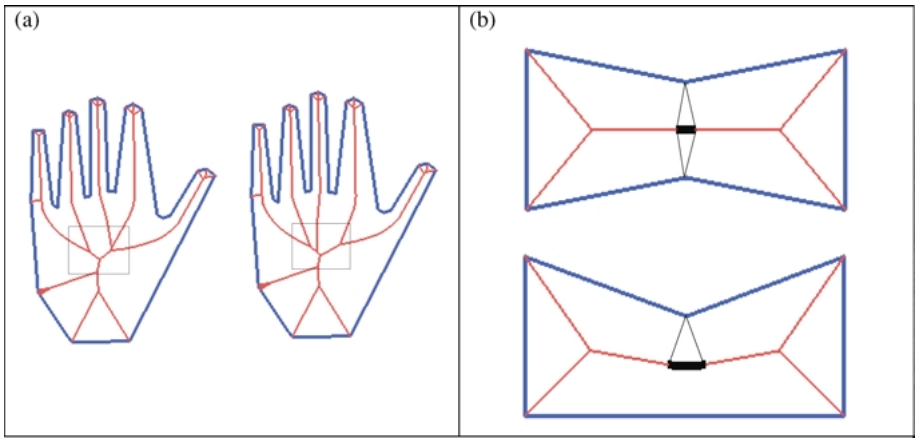

Fig. 3. (a) Unstable part-hierarchies caused by binary hierarchy decisions. A slight change to the boundary between the index and middle fingers causes a drastic change to the MAT topology in the middle of the palm (within the rectangles). (b) The thick lines are examples of ligature (top) and semi-ligature (bottom).

and with binary classifications, insignificant boundary bumps can change to significant parts as a consequence of the slightest boundary perturbation.

The instabilities in the Blum MAT are due to the fact that the MAT is not a one-to-one correspondence; As Fig. 3(b) shows, a single point on the boundary can generate many medial axis points. Thus, changing a single boundary point can cause many axis points to change, creating the instability. These many-to-one axis points were identified by Blum as *semi-ligature* and *ligature* (see Fig. 3). August et al. nicely describe the instabilities caused by ligature (August et al., 1999). Their work reveals that some ligature points are purely connection information and some contain substance information as well. They use a tertiary classification based on the geometric heat equation to identify axis points with substance-like information, and then they cull the ligature that roughly corresponds to what we call connection information; they propose that the culled ligature are the points that cause instabilities. However, their classification still requires explicit decisions to be made at every point, and in this way they encounter the same problems as a binary part-hierarchy.

Furthermore, some objects have a parts-hierarchy that is naturally ambiguous. Fig. 4 shows progressive deformations of an object that starts with one clearly defined bending part and ends with three clearly defined parts. However, there is no exact point in the deformation where the object changes from one part to three. Instead, as the deformation progresses, the single main part exhibits less and less continuity in its substance and more and more connection among distinct parts.

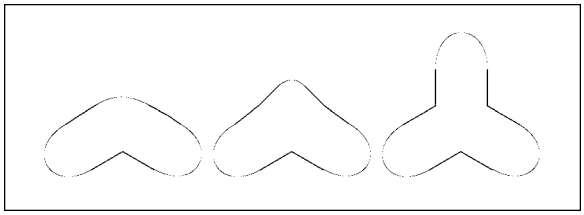

Fig. 4. Part classification that is ambiguous under deformation.

The actual part classification is perceptually ambiguous for the middle configurations of the deformation. This example demonstrates the perceptual continuum from substance to connection that must be addressed in any part-classification approach.

Based on this line of reasoning, the creation of perceptual parts-hierarchies can be resolved by first extracting explicit substance and connection information from the MAT and then using these components to create a fuzzy parts-hierarchy that allows for ambiguity. This allows the inherent form information to be used for parts-hierarchies and form computation. By using these measures of substance and connection, the underlying instabilities of the MAT are not changed but their effects are eliminated.

4. Calculating Substance and Connection Components

In order to create a natural parts-decomposition, this method defines a measure that separates substance and connection components for the Blum MAT. Every axis

point will be considered to have these two components, and a measure between 0 and 1 is calculated for each. Neighboring axis points that have a substance component that is much greater than their connection component are continuous points on the same part of the object. Axis points with a high connection component and low substance component are connection points that separate different parts of the object. Axis points with similar substance and connection measures represent regions of the object that are perceived to have an ambiguous part classification.

In this work, the two components are defined to be directly related; if ω is the connection measure, then substance measure ψ is $1 - \omega$. In this way every point in the object is either substance, connection or a blend of the two. In the remainder of this section, we will develop the calculation for the connection measure ω, but we will refer to both the substance and the connection components of an object. The end result will be a new component added to the Blum MAT representation. Instead of every medial point having just location and scale components, (x, y, σ), every point will be a 4-tuple, (x, y, σ, ψ).

As described in Katz (2002), perceptually based form calculations at a medial point should be integrated over an aperture centered at that point. This idea captures the effects of the receptive fields in our visual systems that operate at a scale that depends on the medial width of the object at the point of interest (Burbeck et al., 1996; Kovacs and Julesz, 1994).

According to this model, the size of the aperture is directly related to the radius of the maximal circle at the medial point, and calculations within the aperture are weighted with a Gaussian fall-off to give points closer to the center point a higher influence than more distal points. All calculations to produce the connection measures for medial points are integrated over and performed within such an aperture. Fig. 2 shows the perceptual aperture with RMS width that is equal to width of the maximal circle around a medial point.

The key to calculating the connection measure for the Blum MAT is noticing that branches are the crucial points in determining the connection and substance components of an object. A single axis representing a simple object would be fully substance with no connection information. This reflects the intra-part continuity along the axis; there are no other connected parts and thus there is no connection information. However, when there is a branch point with three incident axes, we must face the question: "which two axes combine to form the 'main body' of the object, and which remaining axis reflects a bump on the object or an entirely new part?" Or, in other words, "what are the connection and the substance components of the axes around this branch point?"

Any answer to this question must reflect the notion that our visual systems will tend to follow the main medial path of an object while ignoring less significant off shoots. This continuity of a visual path is determined by many factors; two of the most important factors are the continuity of the direction of the medial axes and the continuity of scale along the axes. In this work, *visual conductance* is developed as a measure of how likely any two of the axes at a branch point will be perceived as the single, main perceptual path. In the remainder of this section, the properties required for useful perceptual metrics such as visual conductance are presented, and then functions, sometimes ad-hoc ones, are offered that satisfy these properties.

Visual Conductance

To measure the likelihood that we will perceptually connect two axes at a branch point, visual conductance is defined to be a relative measure that compares, at a branch point, each pair of the branch's axes. When one pair has a high conductance while the other two pairs are low, the first pair is perceived as the main visual path through the branch, and therefore that path is perceived as mostly substance while the third axis is perceived as mostly connection information. When two or all three pairs have similar values of visual conductance through the branch point, there is perceptual ambiguity about which branches are the main visual path, and therefore all branches may reflect a significant amount of both substance and connection information.

Visual conductance needs to capture the visual continuity provided by the continuity of direction along two branching axes, and it should capture the visual continuity provided by similar medial widths across the medial axes. In order to calculate visual continuity through a branch according to these notions, a "visual vector" is first calculated for each axis at the branch (see Fig. 5), and then these vectors are compared to produce the visual conductance measure at the branch. Visual conductance is defined to be a value between 0.0 and 1.0, with all non-branch points defined to have a conductance of 1.0. The following development, therefore, is applicable to branch points.

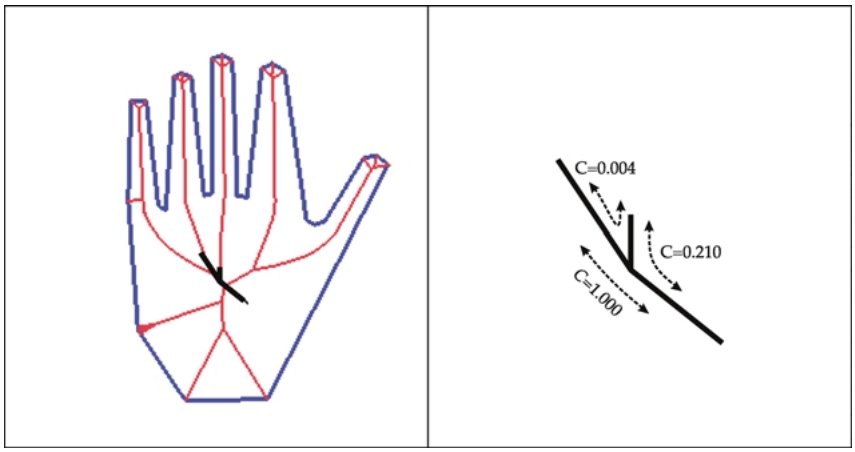

Fig. 5. (a) Hand object with its Blum MAT and visual vectors at the branch where the middle finger's axis joins the other axes. Longer vectors represent greater accumulated visual scale in that direction. (b) Visual conductance values for each pair of axes represented by the visual vectors in (a).

The visual vector for an axis emanating from a branch point is defined to be a weighted directional average of the tangent of the medial axes A_i starting at the branch point and following that branch axis outward through any neighboring branches. To be consistent with the scale-based medial view, this directional average is weighted by the width $\sigma(s)$ of every axis point s on the averaged path. In this work, the medial width $\sigma(s)$ is defined to be the Blum MAT medial radius R, that is, the radius of the maximal disk centered at the medial point s. The final magnitude of the vector reflects the accumulated scale of the object along the vector.

The vector includes contributions from all axis points that fall within an aperture centered at the initial branch point with the aperture size equal to the medial radius of the initial branch point. The calculation is performed within a perceptual aperture of radius $\sigma(s_b)$ around the initial branch point b, and closer axis points are given a higher influence then more distant ones using a Gaussian fall-off $G(s)$ with standard deviation $\sigma(s_b)$. Distance from the branch point is calculated as the shortest distance along the medial axes. Finally, since the scale of all axis points around a branch point are nearly the same (and are exactly equal at the branch point), a term $\Phi(s)$ is added to lessen the influence of axis points very close to the branch.

For a generic branch point with three intersecting axes, there is a visual vector associated with each of the three axes. Each vector includes in its calculation all the axes that can be reached by traversing the MAT, starting along axis A_i and emanating away from the branch point b, traveling along the axes in a depth-first manner up to a distance of $\sigma(s_b)$.

To incorporate all of these requirements, the visual vector $\bar{\mathbf{v}}$ along axis A_i emanating from branch point b is defined as follows. Here, the functions are parameterized as the arc-length distance from branch point b.

$\bar{\mathbf{u}}(s)$ = tangent vector at a medial point s
$\sigma(s)$ = radius of maximal circle at a medial point s
s_b = arc length parameter at branch point b
$\sigma(s_b)$ = radius of maximal circle at branch point b
A_i = the i'th axis emanating from the branch point b

$$G(s) = e^{-\frac{1}{2}(\frac{s}{\sigma(s_b)})^2}$$
$$\Phi(s) = 1 - e^{-(\frac{s}{\sigma(s_b)})^2}$$
$$\bar{\mathbf{v}} = \int_{\substack{t=-\sigma(s_b) \\ A_i}}^{t=\sigma(s_b)} \bar{\mathbf{u}}(t)\sigma(t)[G(t)][\Phi(t)]dt$$

With visual vectors defined for each axis emanating from a branch, the vectors can then be analyzed in pairs to determine the visual conductance between each pair of branch axes. Our perceptual systems use a combination of factors to determine how we perceive an object at its medial branching points. Visually, we will tend to follow along medial axes that provide a straight path rather than a curved path; at the same time we will tend to follow axes that have the most similar scales. For a branch point b, the visual conductance measure incorporates these phenomena; visual conductance C

between branching axes A_i and A_j at branch b is defined as follows.

$$C_b^{continuity}(A_i, A_j) = \left[\left(1 - \frac{\bar{\mathbf{v}}_\mathbf{i}}{|\bar{\mathbf{v}}_\mathbf{i}|} \cdot \frac{\bar{\mathbf{v}}_\mathbf{j}}{|\bar{\mathbf{v}}_\mathbf{j}|}\right)\right]^2 \left[\min\left(\frac{|\bar{\mathbf{v}}_\mathbf{i}|}{|\bar{\mathbf{v}}_\mathbf{j}|}, \frac{|\bar{\mathbf{v}}_\mathbf{j}|}{|\bar{\mathbf{v}}_\mathbf{i}|}\right)\right]$$

The first term captures the continuity of direction using the dot product function, and the value is squared as an ad-hoc means to emphasize the continuity along straight paths and decrease continuity when the path bends. The second term captures the continuity of scale by comparing the magnitude of the two branches under consideration and then keeping the comparison value less than one.

There are three measurements $C_b^{continuity}(A_i, A_j)$ for every generic branch point, one for each pair of the three branching axes. If we had been required to deal with the non-generic case of more than three axes meeting at a branch point, we could have handled that by dividing the branch into two or more very nearby generic branches.

To use this measure, the visual conductances for each pair of axes at a branch point are compared. The two axes that have the highest conductance are chosen as the "main visual path", and their conductance is scaled to be 1. The conductance of the other pairs are then scaled by the same amount to maintain their relative magnitudes but to have a value less than or equal to 1. In this way, the main visual path through a branch point fully "conducts" perceptually while other paths with conductances less than 1 "attenuate" visual flow.

For example, in Fig. 5, the conductance between the left-most and right-most vectors is the highest, so it is scaled to be equal to 1. The conductance between the right-most and the middle vector are both scaled by the same amount. As Fig. 5 shows, while there is full visual conductance between the axes represented by the left-most and right-most vectors, the conductance is attenuated between each of those axes and the axis represented by the middle vector.

Part-End Adjustment

While this formulation for visual conductance applies to general branch points, it fails when an object part ends in a curved tip (see Fig. 6). The MAT in these regions consists of a large scale axis segment branching into segments with much smaller scales. As defined so far, visual conductance would recognize the small scale segments at such branches as most similar in scale and connect them as the main visual path unless they meet at a very sharp angle. Clearly, this is not perceptually correct. Visually, at such a branch point the large scale segment appears to connect with one of the small scale ones to become the main visual path through the

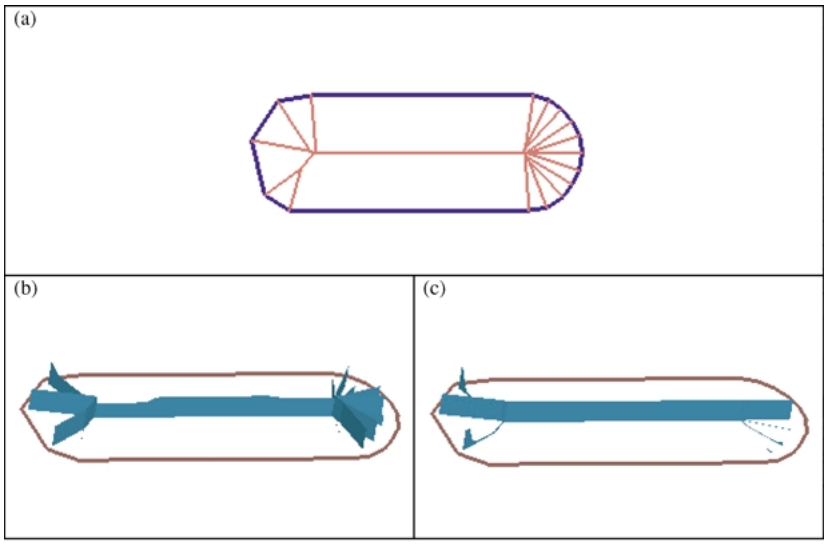

Fig. 6. (a) Rounded-end object with its Blum MAT. (b) Substance measure of an object with only continuity component of visual conductance. Height at each medial point is the substance measure at that point. (c) Substance measure with endness adjustment added to visual conductance in order to capture the visual effects at the tip of a part.

branch, and the remaining segment is seen as a connecting bump.

In order to implement this, we leverage the understanding in human vision that ends of object parts are identified by end-stopped cells independently from the identification of other features (Orban et al., 1979a, 1979b). In this work, an endness factor is calculated separately for each axis at every branch point. *Endness* measures the discontinuity in scale between one axis and its other two branch axes; in other words, the endness factor represents how much larger in scale each axis is compared to the other two neighboring axes. The largest endness factor of the three axes becomes $endness_b$, the degree to which the branch point b will be perceived as an ending of an object part.

Endness is used in formulating visual conductance to signal branch points in which an axis with a high degree of endness should be kept as part of the main visual path through the branch even when the previous conductance calculation indicates that it should not be. When this is the case, the large scale axis is combined with one of the other two smaller scale axes to form the visual path, and is given an endness-conductance $C_b^{endness}$ that is equal to 1. Then the remaining axis is considered an offshoot of that path, and its endness-conductances are calculated as in the non-endness case but they are left unscaled.

If the degree of endness is very low at a branch, the visual conductances at the branch should not be changed by the endness calculations; if *endness* is very high, the branch's conductances should become these alternative end-condition conductances. For cases in between where there is perceptual ambiguity about whether a part is ending at that medial point, the visual conductance should be a blended combination of the two conductance values.

The *endness* measure and final visual conductance are calculated at branch point b, following the path between branch axes A_i and A_j as follows. The new variables and parameters are described in the next paragraph.

$$\lambda = \max\left[\min\left(\frac{|\bar{\mathbf{v}}_\mathbf{i}|}{|\bar{\mathbf{v}}_\mathbf{j}|}, \frac{|\bar{\mathbf{v}}_\mathbf{i}|}{|\bar{\mathbf{v}}_\mathbf{k}|}\right), \min\left(\frac{|\bar{\mathbf{v}}_\mathbf{j}|}{|\bar{\mathbf{v}}_\mathbf{i}|}, \frac{|\bar{\mathbf{v}}_\mathbf{j}|}{|\bar{\mathbf{v}}_\mathbf{k}|}\right), \min\left(\frac{|\bar{\mathbf{v}}_\mathbf{k}|}{|\bar{\mathbf{v}}_\mathbf{i}|}, \frac{|\bar{\mathbf{v}}_\mathbf{k}|}{|\bar{\mathbf{v}}_\mathbf{j}|}\right)\right]$$

$$endness_b = \left(1 - e^{-(\lambda)^p}\right)^n$$

$$C_b(A_i, A_j) = endness_b C_b^{endness}(A_i, A_j) + (1 - endness_b) C_b^{continuuity}(A_i, A_j)$$

Recall that the endness adjustment is intended to modify the visual conductances at MAT branches at the end of a part. The initial factor λ is used to detect the part-end situation by comparing branch-axis scales at a branch to determine the greatest mismatch. The ad-hoc formulation for the final $endness_b$ factor is designed to create a high threshold that the end condition must meet before it influences visual conductance. The values of p and n are used to adjust this threshold, with p typically set to 4 and n typically set to 2. When a part-end situation is indicated by a non-zero value of the *endness* measure, the alternative visual conductance $C_b^{endness}$, which is always 1.0, is factored in to the total visual conductance between two branch axes.

In addition to the endness adjustment, the measures developed here must also account for multiple branches, as developed in Section 3. For a simple object with a single branch, if a point lies close to the branch and is on the main visual path through the branch, then that point has a high substance component. Similarly, if the point lies on an axis whose conductance to the other two axes is much less than the conductance of the main path, then that point has a high connection component. However, this classification fails when multiple branches are close together. Each branch point may influence the substance/connection classification of neighboring axis points. For example, in Fig. 7(c) the axes labeled A_2 and A_4 around branch b_2 would have a substance measure of 1 if b_2 were the only branch in the object. However, because of b_2's proximity to branch b_1, A_2 and A_4 are low in substance and are mostly connection information as shown in Fig 7(b). When many branches are close together, each branch casts a shadow of influence around itself that affects the calculation of nearby connection measures.

Connection Shadows

Since the receptive fields in our visual systems process visual input from a region of an object and not at a single point, there may be many medial branches within this aperture of viewing, and each branch can affect the determination of connection and substance components at the focal point. A branch that is close to the point of interest may indicate that the point is on a high conductance visual path, but a branch farther upstream may indicate a much lower conductance path at the same point. More branch points may each indicate a different conductance along the point's path. To combine the effects of many branch points, there needs to be a way

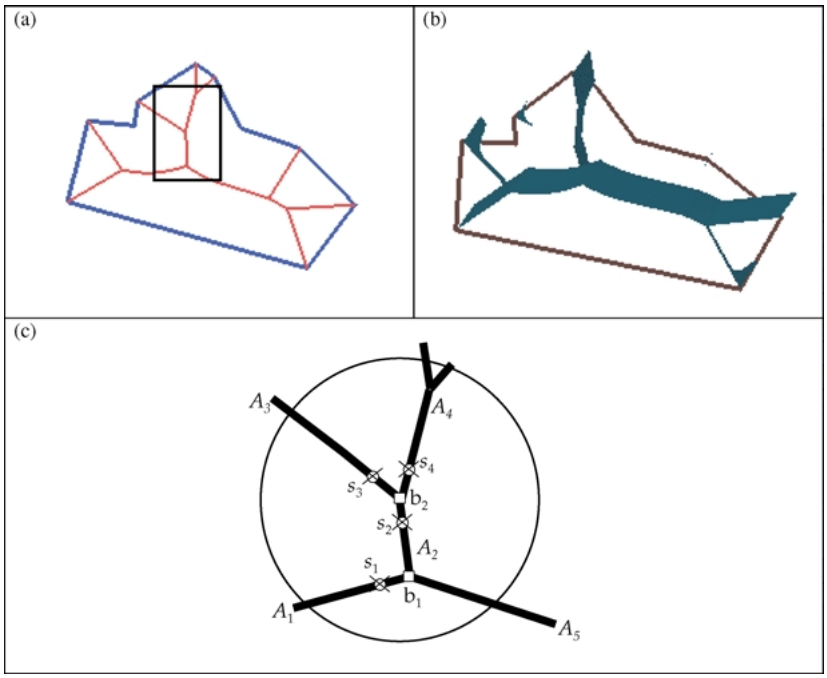

Fig. 7. (a) Object with its Blum MAT. (b) The substance measure of the object. (c) A blow-up of the highlighted region from (a) showing a connection shadow and the axes that it covers. See the text for more details.

to propagate the influence of a branch through other branches and combine the many influences into a single measure. At the same time, any solution to this problem must insulate the main visual paths through an object from branches that are on branches with low visual conductance to a main path. The solution presented in this work is to allow branch points to cast a shadow that covers the local region around the branch. These *connection shadows* propagate the substance-attenuating effect of their branch, where this attenuation results in increased connection measure (and reduced substance measure) at the points their shadow affects. The large circle in Fig. 7(c) illustrates a shadow at branch b_2 and the following paragraphs describe how the attenuation is calculated.

To implement connection shadows, the attenuating effect cast by a branch along one of its branch axes is 1 minus the largest of the conductances of that axis connecting with any of the other branch axes. In this way, a branch will cast a highly attenuating shadow along its branch axes that are not on a main visual path, and it will cast no attenuation along its main visual axes. For the example in Fig. 7, axis A_3 has low conductance with both A_2 and A_4 so the shadow from b_2 creates high connection measure on A_3. At the same time, since the conductance between A_2 and A_4 is 1, there is no substance attenuation cast along those axes by b_2. The shadow is adjusted by a fall-off from the branch point casting the shadow, with the aperture of the falloff determined by the scale at the branch point. This provides a high shadowing effect within most of the shadow with a rapid fall-off at the edges of the shadow.

The connection shadow is calculated as follows, where n is a parameter used to adjust the rate of fall-off from the branch.

If s is a point on axis A_i, where A_i ends at branch points b_1 and b_2, and b_β of A_i has neighboring axes $A_{\beta,1}$ and $A_{\beta,2}$, then the shadow from a branch point b cast on an axis point s that lies on axis A_i is

$$S_{i,b}(s) = \left[1 - 2\left(\operatorname{erf}\left[\frac{|s - s_b|}{\sigma(s_b)}\right]^n - 0.5\right)\right] \times \left[1 - \max_{j=1,2} C_\beta(A_i, A_{\beta,j})\right]$$

This value is clamped to zero if the result becomes negative.

Following the principle that our visual processing occurs within an aperture and not at a single point,

connection shadows are propagated out from a branch along each branch axis. A high value for the shadow, indicating high connection and low substance components, will flow along the axis and along all branching axes downstream from the initial branch point. This in effect casts a shadow along those axes that is used to determine the connection measure at points along the axes that it covers. With this construction, a branch's effect is felt by all points within its shadow even when separated by other branches.

However, as a connection shadow flows through other branch points its effect should be attenuated by every branch that it flows through. This can be understood by remembering that visually there are the main medial paths through an object and there are sub-parts that branch from the main paths. These sub-parts are separated from the main parts by axes with a high connection measure, and the medial configuration of the sub-parts have little visual impact on the main part. For example, in Fig. 7(c), point s_1 is on the main visual path of the object and its substance measure should not be attenuated by the shadow cast from b_2. In general, every branch propagates a shadow according to the conductance of the path the shadow follows through the branch. In this way, branches that are located on a path that is already highly connection information will not cast influence on neighboring branches that are on the main visual path attached to the connection path. In other words, main visual paths through a branch are insulated from connection shadows on attached low conductance paths.

In order to compute shadow propagation, the accumulated conductance $V_{i,b}$ between a branch point b and a point s on axis A_i must be determined. This is the product of conductances encountered through each branch point on a path starting at branch point b and traversing the medial axes to reach axis A_i. These conductances potentially attenuate the effects of branch point b's shadow on the points on a distant axis.

Accumulated conductance $V_{i,b}$ is computed as follows.

$\{p_{i,b,m}\}$ for a fixed axis A_i and branch point b is the sequential ordered list of branch points numbered $1..m_{\max}$ on the path from axis A_i to branch point b.

$\{k_{i,b,m}\}$ is the sequential ordered list of axes numbered $1..m_{\max}$ on the path from axis A_i to branch point b that is associated with $\{p_{i,b,m}\}$.

The accumulated conductance $V_{i,b}$ between a point on axis A_i and a branch point b through the path of branch points $\{p_{i,b,m}\}$ and axes $\{k_{i,b,m}\}$ is the product of conductances encountered at each of the m branch points along the path,

$$V_{i,b} = \prod_m C_m(k_{i,b,m}, k_{i,b,m+1})$$

Finally, the shadowing effect of any given branch point b onto a point s on axis A_i is the attenuation of the connection shadow from b multiplied by the accumulated conductance between the branch point $b$and the axis A_i: $V_{i,b}S_{i,b}(s)$.

Combining all of these developments, the final measure $\omega(s)$ of the connection component (and the related substance measure $\psi(s)$) at every point on a Blum medial axis is the combined effects of all branch points on the point of interest, s. That is, the connection measure $\omega(s)$ is the sum of the effects of all branch points on point s. While every branch point is considered, only those whose shadow overlaps the point of interest will have any effect on s. This is defined as follows, with a summary of the terms involved.

$$\text{Connection measure: } \omega(s) = \sum_b V_{i,b}S_{i,b}(s)$$

Substance measure: $1 - \omega(s)$

Accumulated conductance:

$$V_{i,b} = \prod_m C_m(k_{i,b,m}, k_{i,b,m+1})$$

$\{p_{i,b,m}\}$ for a fixed axis A_i and branch point b is the sequential ordered list of branch points numbered $1..m_{\max}$ on the path from axis A_i to branch point b.

$\{k_{i,b,m}\}$ is the sequential ordered list of axes numbered $1..m_{\max}$ on the path from axis A_i to branch point b that is associated with $\{p_{i,b,m}\}$.

The accumulated conductance $V_{i,b}$ between a point on axis A_i and a branch point b through the path of branch points $\{p_{i,b,m}\}$ and axes $\{k_{i,b,m}\}$ is the product of conductances encountered at each of the $m_{\max}$ branch points along the path.

$$\begin{aligned}\text{Connection Shadow} : S_{i,b}(s) &= \left[1 - 2\left(\operatorname{erf}\left[\frac{|s - s_b|}{\sigma(s_b)}\right]^n - 0.5\right)\right] \\ &\times \left[1 - \max_{j=1,2} C_\beta(A_i, A_{\beta,j})\right]\end{aligned}$$

This is the connection shadow cast by branch point b on a point s on axis A_i, where A_i ends at branch

points b_1 and b_2, and b_β of A_i has neighboring axes $A_{\beta,1}$ and $A_{\beta,2}$.

$$\begin{aligned}
&\text{Visual Conductance} : C_b(A_i, A_j)\\
&\quad = endness_b C_b^{endness}(A_i, A_j)\\
&\qquad + (1 - endness_b) C_b^{continuuity}(A_i, A_j)\\
&\lambda = \max\left[\min\left(\frac{|\bar{\mathbf{v}}_\mathbf{i}|}{|\bar{\mathbf{v}}_\mathbf{j}|}, \frac{|\bar{\mathbf{v}}_\mathbf{i}|}{|\bar{\mathbf{v}}_\mathbf{k}|}\right), \min\left(\frac{|\bar{\mathbf{v}}_\mathbf{j}|}{|\bar{\mathbf{v}}_\mathbf{i}|}, \frac{|\bar{\mathbf{v}}_\mathbf{j}|}{|\bar{\mathbf{v}}_\mathbf{k}|}\right),\right.\\
&\qquad\qquad \left.\min\left(\frac{|\bar{\mathbf{v}}_\mathbf{k}|}{|\bar{\mathbf{v}}_\mathbf{i}|}, \frac{|\bar{\mathbf{v}}_\mathbf{k}|}{|\bar{\mathbf{v}}_\mathbf{j}|}\right)\right]\\
&endness_b = \left(1 - e^{-(\lambda)^p}\right)^n\\
&C_b^{continuity}(A_i, A_j)\\
&\quad = \left[\left(1 - \frac{\bar{\mathbf{v}}_\mathbf{i}}{|\bar{\mathbf{v}}_\mathbf{i}|} \cdot \frac{\bar{\mathbf{v}}_\mathbf{j}}{|\bar{\mathbf{v}}_\mathbf{j}|}\right)\right]^2 \left[\min\left(\frac{|\bar{\mathbf{v}}_\mathbf{i}|}{|\bar{\mathbf{v}}_\mathbf{j}|}, \frac{|\bar{\mathbf{v}}_\mathbf{j}|}{|\bar{\mathbf{v}}_\mathbf{i}|}\right)\right]
\end{aligned}$$

The *endness* measure and final visual conductance are calculated at branch point b, following the path between branch axes A_i and A_j.

$$\text{Visual Vector: } \bar{\mathbf{v}} = \int_{\substack{t=-\sigma(s_b)\\ A_i}}^{t=\sigma(s_b)} \bar{\mathbf{u}}(t)\sigma(t)[G(t)][\Phi(t)]\,dt$$

$\bar{\mathbf{u}}(s)$ = tangent vector at a medial point s
$\sigma(s)$ = radius of maximal circle at a medial point s
s_b = arc length parameter at branch point b
$\sigma(s_b)$ = radius of maximal circle at branch point b
A_i = the i'th axis emanating from the branch point b

$$G(s) = e^{-\frac{1}{2}(\frac{s}{\sigma(s_b)})^2}$$
$$\Phi(s) = 1 - e^{-(\frac{s}{\sigma(s_b)})^2}$$

In the example given in Fig. 7(c), the shadow cast from b_1 has no effect on s_1 because A_1 and A_5 constitute the main pathway through b_1. The shadow does cast its attenuation on s_2, and this is propagated to s_3. The shadow is blocked by b_2 from affecting s_4 because A_2 and A_4 are the main path through that branch. The shadow cast from b_2 has no effect on s_2 or s_4 because they are on the main pathway through b_2, and there is no effect on s_1 because it is on the main visual path from A_2 to A_1. The shadow from b_2 does cast its full attenuating effect on s_3. Summing all of these effects produces no substance attenuation at s_1, but much attenuation at s_2, s_3, and s_4. The results are seen in Fig. 7(b), showing the final substance measure $\psi(s)$ for the object. This measure is added to the Blum MAT representation, giving a value of ψ at every medial point and leading to an augmented representation (x, y, σ, ψ) at every medial point.

The Blum MAT produces $O(n)$ axes for n lines on a polygonal boundary. This meas that the calculation of the substance measure is an $O(n^2)$ process. A first pass over all branches is required to compute the visual vectors, visual conductances and connection shadows. In a second pass over all the axes, the substance measure is actually computed.

5. Results

Figs. 8–10 show the substance measure of several objects. The medial points with a high substance measure correspond to the parts of the objects. Points where substance measure is low correspond to axes that serve mainly to connect the object parts. Points where the substance measure is between 0.0 and 1.0 reflect regions of the object where a parts-classification is perceptually ambiguous.

These results show how extraneous axes are removed and an object's parts are naturally extracted. In the cat example, the most salient parts of the cat are signaled by the visual path that runs from the head through the body and down to the end of the tail. The ears and front paws are clearly distinct perceptual parts with the connecting axes exhibiting very low substance measure (and therefore a high connection measure). In a similar manner, the lizard's legs are perceptually distinct from its body, and the stem of the leaf is distinct from the body of the leaf.

The examples also demonstrate how perceptual ambiguity about an object's parts is reflected in the substance measure. The five points of the maple leaf have no clear hierarchy of parts and substance measures in the region reflect this. Here, the branching axes reflect the medial axis filling a region of similar saliency rather than connecting parts.

The object with protrusions in Fig. 10 shows how the substance measure captures a broad range of perceptual ambiguity. Some protrusions are clearly separate parts while others have some tendency to been seen as an extension of the main object. The bottom-most protrusions demonstrate this; the lower-right one is clearly an independent part, a property highlighted by the high connection measure of its connecting axis. On the other hand, the lower-left protrusion can be perceived as an extension of the central body, and the lower-middle one can be perceived as an extension of the upper-left protrusion. Perceptually there is no clear hierarchy in

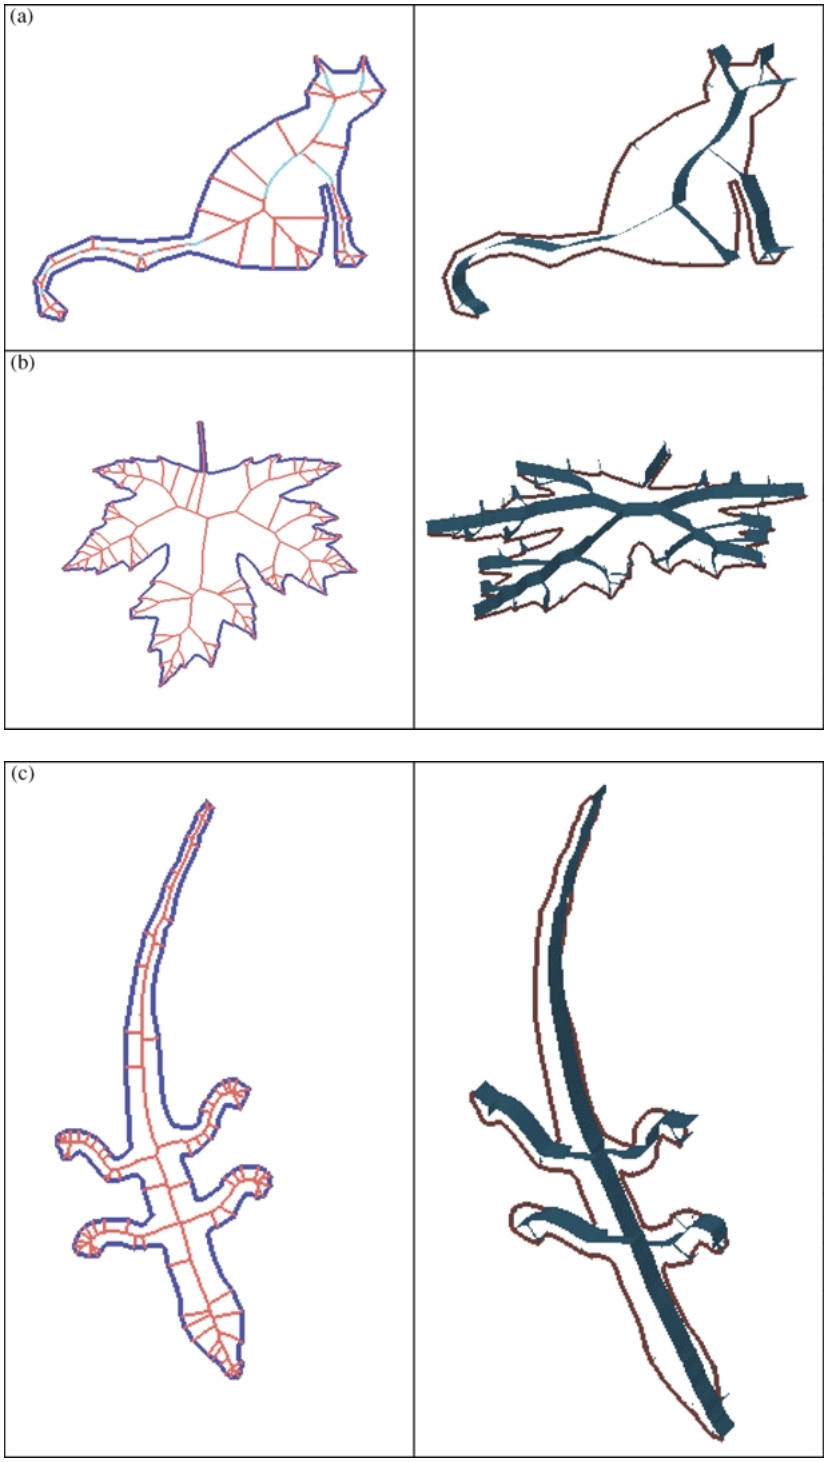

Fig. 8. The subtance measure of various objects, where the height of each medial point reflects the substance measure. (a) Cat object. (b) Maple leaf object. (c) Lizard object.

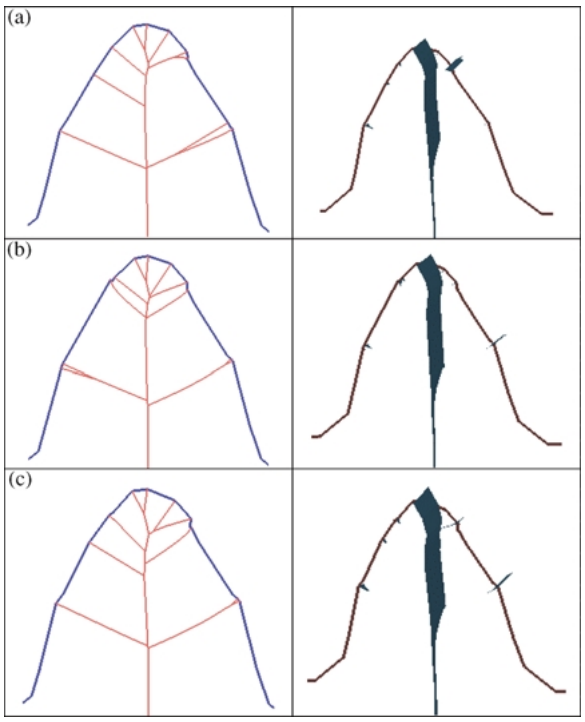

Fig. 9. (a) The MAT and substance measure of an object's protrusion. (b) and (c) Two deformed versions of the object, showing significantly different MATs but very similar substance measures.

this region, and the substance and connection measures reflect this.

A subtle effect of performing perceptual calculations within a local aperture can be seen in the object with protrusions example of Fig. 10, with the right-most protrusion and the smaller one underneath it. While the larger, right-most protrusion can be perceived clearly as being a more salient part in a global sense, at the medial junction of the parts there is no clear part hierarchy. Within the local region around those branch points, this can be seen to be valid. The substance and connection measures are not intended to be used directly for high-level global object interpretation.

Figure 9 shows a protruding part of an object and two deformed versions created with tiny perturbations to the object's boundary. The MAT is significantly different in each of the three versions, yet the substance measures remain very similar. Only the most distal edges of the MAT pieces created by deforming the object show any change in substance. This example demonstrates how the substance-weighted MAT reflects only the minor changes to the boundary even when the underlying MAT shows drastic instabilities.

Finally, Fig. 10 shows how the adjustment for ends of figures works and where it fails. At the tip of each protrusion, only one of the small-scale branches remains and it is connected to the main visual path that passes through the MAT to the tip; the other small scale branches are identified as connection information with only their very ends possibly showing any substance. However, Fig. 10(b) shows where the ad-hoc endness function is applied in a way that generates a substance measure that has no perceptual basis. This object reflects the case of a large-scale body of an object with two small wings off of one end. Clearly, both of the wings are separate parts and neither is a continuation of the main body. In this implementation, however, two of the axes at a branch are always connected and given a visual conductance of 1.0. While this choice was suitable for using these ideas to develop a visual saliency measure and an associated object simplification algorithm (Katz, 2002), it probably will not support other perceptual analysis. For such work, the previous endness function could still be used to identify MAT points where an adjustment should be applied and a new adjustment created that does not connect any axes at that point. Alternatively, a new endness function also could be developed for these cases.

6. Conclusions

This paper describes a method for instrumenting the Blum MAT to perform perceptual part-decompositions of objects and to support stable and robust form calculations. The method considers objects to be a collection of solid parts and a set of connections among those parts. A measure is created at every point on the medial axes to grade the amount of substance information at each medial point and the amount of connection information. This measure is then used to decompose the object into a set of perceptually defined parts. A key aspect of this classification is that this metric gives a fuzzy measure of "part-ness", and the parts of an object are not forced into a binary classification scheme.

Using the substance measure and complementary connection measure as well as the resulting fuzzy classification of parts, form calculations remain stable under MAT instabilities caused by continuously deforming boundaries. This stability has been demonstrated by extensive tests and is detailed in (Katz, 2002). Even when drastic changes to the MAT result from tiny changes to the boundary, these instabilities are confined

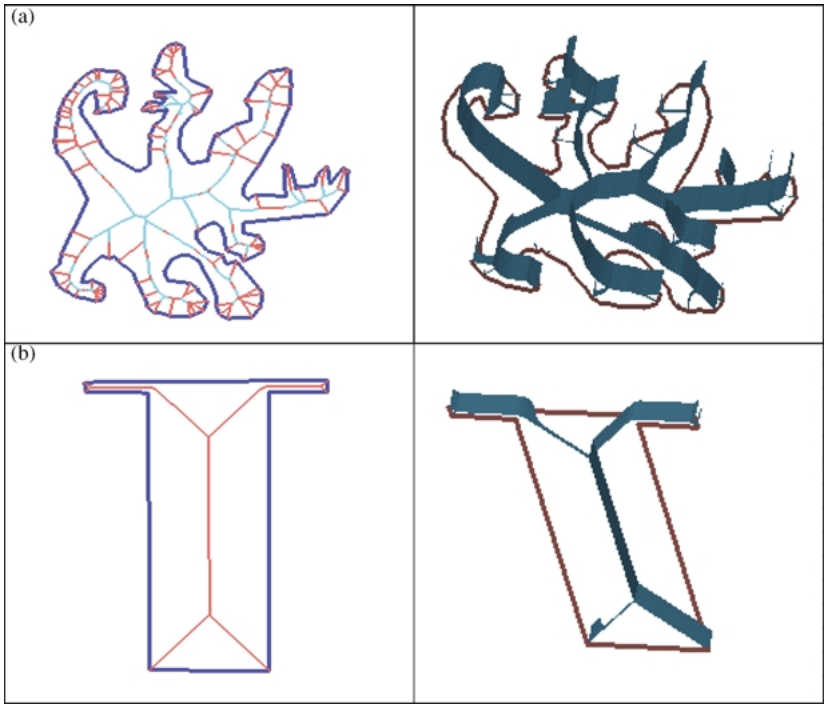

Fig. 10. Objects and their substance measure. (a) Object with protrusions. The protrusions of the object show how perceptual uncertainty about the part hierarchy is reflected in the substance measure. (b) This example shows how the implementation of the endness adjustment gives unintuitive results.

to the connection component of the object and the substance-weighted MAT shows only the small deformation. In this way, form analysis can track the minor boundary changes and remain immune to the sensitivities of the underlying MAT.

Acknowledgments

We thank Steven Zucker for a discussion that became the catalyst for the ideas developed in this work.

References

August, J., Siddiqi, K., and Zucker, S. 1999. Ligature instabilities in the perceptual organization of form. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*. Fort Collins, Colorado, pp. 42–48.

Blum, H. 1967. A transformation for extracting new descriptors of form. In *Models for the Perception of Speech and Visual Form*, W. Whaten-Dunn (Ed.). MIT Press: Cambridge, MA, pp. 362–380.

Blum, H. and Nagel, R. 1978. Form description using weighted symmetric axis features. *Pattern Recognition*, 10: 167–180.

Brady, M. and Asada, H. 1984. Smoothed local symmetries and their implementation. *International Journal of Robotics Research*, 3(3):36–61.

Bruce, J., Giblin, P., and Gibson, C. 1985. Symmetry sets. *Proc. Royal Society Edinburgh*, 101A:163–186.

Burbeck, C., Pizer, S., Morse, B., Ariely, D., Zauberman, G., and Rolland, J. 1996. Linking object boundaries at scale—A common mechanism for size and form judgements. *Vision Research*, 36(3):361–372.

Katz, R. 2002. *Form Metrics for Interactive Rendering via Figural Models of Perception*. Dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC.

Kovacs, I. and Julesz, B. 1994. Perceptual sensitivity maps within globally defined visual forms. *Nature*, 370:644–646.

Marr, D. and Nishihara, K. 1978. Representation and recognition of the spatial organization of three dimensional structure. *Proc. of the Royal Society of London*, 200:269–294.

Ogniewicz, R. and Ilg, M. 1992. Voronoi skeletons: Theory and applications. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 63–69.

Orban, G., Kato, H., and Bishop, P., 1979a. End-zone region in receptive fields of hypercomplex and other striate neurons in the cat. *Journal of Neurophysiology*, 42:818–832.

Orban, G.A., Kato, H., and Bishop, P.O. 1979b. Dimensions and properties of end-zone inhibitory areas in receptive fields of hypercomplex cells in cat striate cortex. *Journal of Neurophysiology*, 42:833–849.

Pizer, S., Oliver, W., and Bloomberg, S. 1987. Hierarchical form description via the multiresolution symmetric axis transform. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 9(4):505–511.

Pizer, S., Eberly, D., Fritsch, D., and Morse, B. 1998. Zoom-invariant vison of figural form: The mathematics of cores. *Computer Vision and Image Understanding*, 69(1):53–71.

Shaked, D. and Bruckstein, A. 1998. Pruning medial axes. *Computer Vision and Image Understanding*, 69(2):156–169.

Szekely, G. 1996. *Form Characterization by Local Symmetries*. Habilitationsschrift: Institut fur Kommunikationstechnik, Fachgruppe Bildwissenschaft, ETH, Zurich.