

Shapes, Shocks and Wiggles

Kaleem Siddiqi

*Center for Computational Vision & Control
Yale University, New Haven, CT 06520-8285, USA*

Benjamin B. Kimia

*Division of Engineering
Brown University, Providence, RI 02912, USA*

Allen Tannenbaum

*Department of Electrical Engineering
University of Minnesota, Minneapolis, MN 55455, USA*

Steven W. Zucker

*Center for Computational Vision & Control
Yale University, New Haven, CT 06520-8285, USA*

Abstract

We earlier introduced an approach to categorical shape description based on the singularities (shocks) of curve evolution equations. The approach relates to many techniques in computer vision, such as Blum's grassfire transform, but since the motivation was abstract, it is not clear that it should also relate to human perception. We now report that this shock-based computational model can account for recent psychophysical data collected by Burbeck and Pizer. In these experiments subjects were asked to estimate the local centers of stimuli consisting of rectangles with "wiggles" (sides modulated by sinusoids). Since the experiments were motivated by their "core" model, in which the scale of boundary detail is proportional to object width, we conclude that such properties are also implicit in shock-based shape descriptions. More generally, the results suggest that significance is a structural notion, not an image-based one, and that scale should be defined primarily in terms of relationships between abstract entities, not concrete pixels.

Key words: Shocks, Shape Description, Shape Perception, Shape Significance.



Fig. 1. The small bumps significantly alter the percept of a straight rectangle (LEFT) to a “wiggle” (MIDDLE). However, when the boundaries are pulled apart (RIGHT) the shape is once again perceived as being straight, but now with wiggly sides. Adapted from [21,10].

1 Introduction

The notion of “scale” in computer vision is commonly formulated in *image-based* terms. Fine scale features are typically extracted using operators occupying a small spatial extent (in pixels), while coarse scale features are extracted using larger operators. This view has led to the development of scale-spaces in vision and image processing, where many mathematical properties of the fine-to-coarse simplification process have been formally motivated [22,23,48,1,37,28,47,14,13,15]. This effort is further supported by physiological evidence that receptive fields for visual neurons can be observed to be sensitive to a range of spatial support, a result that is normally interpreted to imply that receptive field size defines the spatial scale of a “biological” operator. The existence of many such receptive fields covering each retinotopic position suggests a “pyramid”-like organization of operators upon which to base image representations [11].

In contrast to this commonly-held view of scale, there are also structural considerations, which dictate that scale should be formulated in *object-based* terms and not image-based ones [4,29,50]: finger-nails are fine scale with respect to the fingers, fingers are fine scale with respect to the hand, and the hand with respect to the arm, *etc.* The commonly-held view implies that larger image structures should be more significant for shape; after all, these are the ones that will survive to the largest image scales. However, this is demonstrably *not* the case; see Figure 1. Here we show three shapes: a thin rectangle, a thin rectangle with bumps, and a thicker rectangle with bumps. Notice that the

bumps significantly alter the perception from a block to a “wiggle”, while for the last shape the bumps look like protrusions and the block is once again perceived as being straight. Clearly there is more to scale than size in pixels.

Secondly, the above example shows that context plays an important role; stated informally, note that a pimple on a teenager’s cheek is not quite the catastrophic event that the same pimple on her nose would be. Despite the emergence of an elaborate literature on scale-spaces in computer vision, to date there has been little insight into how such perceptual observations can be explained, or into their deeper implications for object recognition. The pimple on the nose is a simple counter-example to the image-based view that small operators will pick out structurally less significant features, and the wiggles illustration shows that perceptions change as a complex function of size.

In this paper we investigate the notion of scale for 2D shape from computational and from perceptual considerations. We claim, from both considerations, that *the significance of a structure depends on its size with respect to structures to which it is attached*. This would explain, for example, why a pimple is perceived to be more salient when situated on a (smaller) nose than when placed on a (larger) cheek. The wiggles, however, are more complex. In order to explain such percepts we first review the classification of shocks upon which our computational model is based, and briefly describe algorithms for computing such descriptions in Section 2. In Section 3 we summarize the recent psychophysical experiments of Burbeck et al. designed to measure the effect of boundary perturbations on the perceived centers of elongated rectangle shapes [10]. This provides a quantitative measure of how such shapes are represented by the human visual system. We then demonstrate that our shock-based model is consistent with the psychophysical data, and show that it leads to a natural structural scale for significance, as illustrated by numerous computational examples in Section 4. The key insight is to abstract the system of shocks into an *object-based* structure which, although derived from image properties, supports the computational significance of each component in relative terms (its size with respect to the components to which it is attached).

2 Shocks and the Shape Triangle

2.1 Shock Classification

Observing that shapes which are slight deformations of one another appear similar, Kimia *et al.* proposed the following evolution equation for visual shape

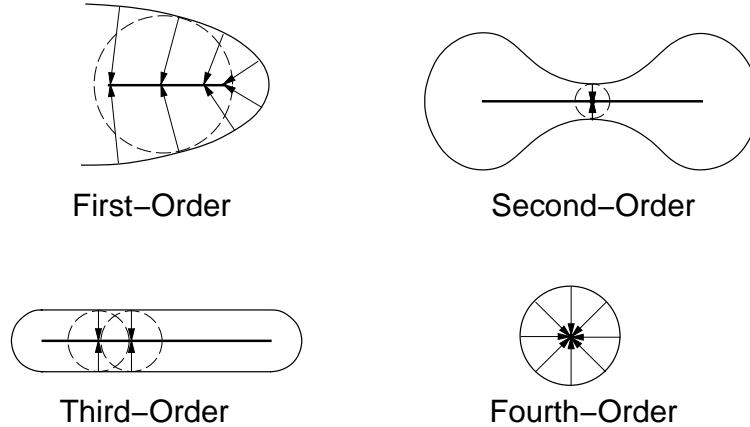


Fig. 2. A coloring of shocks into four types [20]. A 1-shock derives from a *protrusion*, and traces out a curve segment of adjacent 1-shocks. A 2-shock arises at a *neck*, and is immediately followed by two 1-shocks flowing away from it in opposite directions. 3-shocks correspond to an annihilation into curve segment due to a *bend*, and a 4-shock an annihilation into a point or a *seed*. The loci of these shocks gives Blum’s medial axis.

analysis [19,20]:

$$\begin{aligned}
 \mathcal{C}_t &= (1 + \alpha\kappa)\mathcal{N} \\
 \mathcal{C}(p, 0) &= \mathcal{C}_0(p).
 \end{aligned}
 \tag{1}$$

Here $\mathcal{C}(s, t)$ is the vector of curve coordinates, $\mathcal{N}(p, t)$ is the inward normal, p is the curve parameter, and t is the evolutionary time of the deformation. The constant $\alpha \geq 0$ controls the regularizing effects of curvature κ . When α is large, the equation becomes a geometric heat equation; when $\alpha = 0$, the equation is hyperbolic and *shocks* [25], or entropy-satisfying singularities, can form. The locus of points through which the shocks migrate is related to Blum’s grassfire transformation [7,20], although significantly more information is available via a “coloring” of these positions. Four types can arise, according to the local variation of the radius function along the medial axis (Figure 2). Intuitively, the radius function varies monotonically at a type 1, reaches a strict local minimum at a type 2, is constant at a type 3¹ and reaches a strict local maximum at a type 4. The classification of shock positions according to their colors is at the heart of our result.

¹ Although this condition reflects the non-genericity of 3-shocks “bend”-like structures are abundant in the world and are perceptually salient.

2.2 Shock Detection

The numerical simulation of equation (1) is based on level set methods developed by Osher and Sethian [32,38]. The essential idea is to represent the curve $\mathcal{C}(p, t)$ as the zero level set of a smooth and Lipschitz continuous function $\Psi : \mathbf{R}^2 \times [0, \tau) \rightarrow \mathbf{R}$, given by $\{X \in \mathbf{R}^2 : \Psi(X, t) = 0\}$. Since $\mathcal{C}(p, t)$ is on the zero level set, it satisfies

$$\Psi(\mathcal{C}, t) = 0. \quad (2)$$

By differentiating equation (2) with respect to t , and then with respect to the curve parameter p , it can be shown that

$$\Psi_t = (1 + \alpha\kappa)\|\nabla\Psi\|. \quad (3)$$

Equation (3) is solved using a combination of discretization and numerical techniques derived from hyperbolic conservation laws. The curve \mathcal{C} , evolving according to equation (1), is then obtained as the zero level set of Ψ .

A convenient choice for Ψ is the signed distance function to the shape, which can be computed efficiently [12]. The real-valued embedding surface also provides high resolution information for detecting and localizing shocks. In particular, the four shock types can be classified using differential properties of Ψ (see Figure 3). The algorithm for shock detection, described in detail in [40], can now be summarized as follows:

ALGORITHM

- (1) For a small time step evolve the curve by equation (1) with $\alpha = 0$, using a shock capturing curve evolution scheme [39,33].
- (2) Obtain estimates of current shock locations and velocities by subpixel interpolation on Ψ [42,41].
- (3) Group current shocks together with past shocks with which they are consistent [41].
- (4) Repeat till the shape is annihilated.

The above algorithm provides accurate descriptions, but at significant computational cost. In particular, the subpixel interpolation algorithm requires multiple polynomial fits to portions of the surface Ψ . In ongoing work alternate methods for shock detection are being explored [46]; these are computationally more efficient and can be applied to fragmented contours as well.

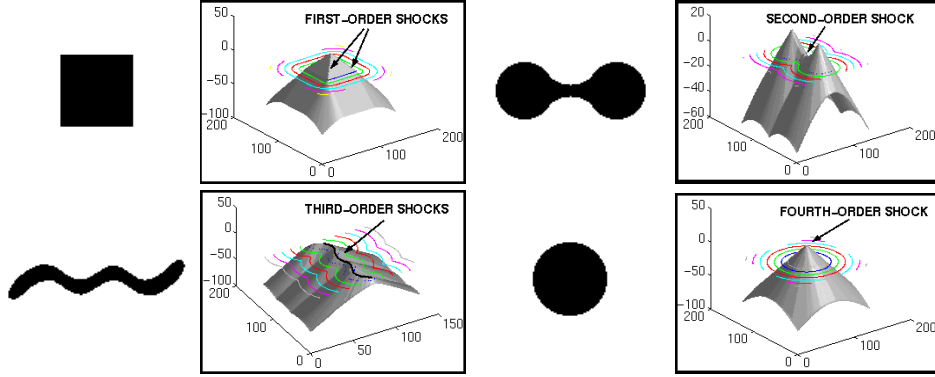


Fig. 3. A classification of shock types based on properties of the embedding surface Ψ (in this case the signed distance function, shown as a shaded surface). TOP LEFT: First-order shocks occur at corners of the square shape, corresponding to creases on the surface with $|\nabla\Psi| > 0$. TOP RIGHT: A second-order shock forms at the “neck” of the peanut shape, corresponding to a hyperbolic point with $|\nabla\Psi| = 0$. BOTTOM LEFT: A set of third-order shocks forms along the central axis of the worm shape, where $\kappa_1\kappa_2 = |\nabla\Psi| = 0$. BOTTOM RIGHT: A fourth-order shock forms in the center of the circle shape, where $\kappa_1\kappa_2 > 0$ and $|\nabla\Psi| = 0$. Adapted from [41].




Fig. 4. The detection of shocks for a dumbbell shape undergoing constant inward motion. Each sub-figure is a snapshot of the evolution in time, with the outline of the original shape shown in black, the evolved curve overlaid within, and the arrows representing velocity vectors for the current first-order shocks. Adapted from [41].

To get an intuitive feel for the shock detection process, consider the numerical simulation of a dumbbell shape evolving under constant inward motion, Figure 4. Note the emergence of a qualitative description of the shape as that of two parts separated at a “neck” (second-order shock), with each part consisting of three “protrusions” (first-order shock groups) merging onto a “seed” (fourth-order shock). The notion of parts for 2D shape, popularized by Hoffman and Richard’s work [17], has been further investigated in [45].

It can be proven that the locus of positions covered by all of the above shocks is formally equivalent to the Blum skeleton [5], which raises the question of why we choose to work with shocks. The examples in Figures 10, 11 and 12 provide the answer: with shocks there are two additional factors labeling the skeleton: (i) the type of shock which “colors” the skeletal points; and (ii) the significance of the shock, related to its time of formation. As we shall later see, type and significance provide an enriched description that is better suited




Fig. 5. LEFT: The sides of the shape triangle represent continua of shapes; the extremes correspond to the “parts”, “protrusions” and “bends” nodes [21]. We plot the time till formation of the second-order shock (for constant inward motion), normalized by the lifetime of the shape, for samples of the “bow-tie” stimuli used in [21]. Observe the increase in this ratio in moving from the *parts* node (shape 10) to the *protrusions* node (shape 1).

to recognition than simply the locus of shock points.

2.3 The Shape Triangle

The generic perceptual shape classes discussed above have lead to the metaphor of a shape triangle, where three distinct processes compete with one another to explain shape [21], Figure 5 (left). The sides of the triangle reflect this competition and capture the tension between object composition (parts), boundary deformation (protrusions) and region deformation (bends). When one considers the *parts-protrusions* continuum and plots the time till formation of the second-order shock under constant inward motion, normalized by the time till annihilation of the shape, Figure 5 (right), a curious connection is unearthed. The ratio is a measure of local width versus global width, and provides a metric for the perceptual distance of the shape from the “parts” node along this axis.

Now, note that intermediate shapes along the *bends-protrusions* continuum closely resemble wiggles. The term “wiggle” appropriately describes not only the sinusoidal nature of the boundary modulation, but also its perceived effect on the object’s central axis [10]. In the context of the shape triangle, for thin objects placed close to the “bends” node the axis is seen to wiggle, for thicker

objects placed close to the “protrusions” node the axis is perceived to be straight. Within the shock-based framework, such effects are reflected in the geometry of the high-order shocks that arise, Figure 8. We are formally led to the following prediction:

Prediction 1 *The perceived center of a “wiggle” along a horizontal line in alignment with a sinusoidal peak coincides with a high-order shock² that forms under constant inward motion.*

The psychophysical experiments of Burbeck *et al.*, introduced next, involved precisely the above class of shapes.

3 Wiggles

Pizer *et al.* [30,9,10] have developed an alternative approach to visual shape analysis called the *core* model. Underlying the formulation of the core model is the hypothesis that the scale at which the human visual system integrates local boundary information towards the formation of more global object representations is proportional to object width. Psychophysical examinations of Weber’s Law for separation discrimination support this proposal [8]. Arguing that the same mechanism explains the attenuation of edge modulation effects with width, Burbeck *et al.* have recently reported on an elegant set of psychophysical experiments where subjects were required to bisect elongated stimuli with wiggly sides. In the following we present their main findings. As predicted above, we will later demonstrate a strong correlation between these findings and computational results obtained from a shock-based description of the wiggles.

The stimuli consisted of rectangles subtending 4 degrees of visual arc in height, with sinusoidal edge modulation, Figure 6 (left). Two widths were considered (0.75° and 1.5°) and for each width there were 6 edge modulation frequencies (0.25, 0.5, 1, 2, 4, 8 cycles/o) and 2 edge modulation amplitudes (20% and 40% of object width). A black probe dot appeared near the center of each stimulus, in horizontal alignment with a sinusoidal peak. The subject was asked to indicate “whether the probe dot appeared to be left or right of the center of the object, as measured along a horizontal line through the dot.”³ As a sample experiment view the stimulus on the right of Figure 6 for a period of one second from a distance of 1.5 meters. You are likely to judge the dot to be to the right of the object’s center. It may surprise you to find that it actually

² Type 3 or 4 for in-phase sinusoids, and type 2 for out-of-phase ones (mirror symmetric shapes).

³ See [10] for further details.




Fig. 6. LEFT: The geometry of a “wiggle” stimulus. RIGHT: Is the dot to the left or to the right of the object’s center?

lies midway between the boundaries on either side, as can be verified by placing a ruler across the figure. In fact, despite instructions to make a local judgment your visual system is biased towards acquiring edge information across a more global spatial extent.

Burbeck and Pizer quantified this effect of edge modulation on the perceived center by varying the horizontal position of the probe dot and subjecting the data to probit analysis. The center of the object was inferred as the 50% point on the best-fitting probit function⁴, and the *bisection threshold* was defined as the variance of this function. The *perceived central modulation* was then obtained as the horizontal displacement between the perceived centers in alignment with left and right sinusoidal peaks. The main findings were:

Result 1 *For a fixed edge modulation frequency the perceived central modulation decreases with increasing object width.*

Result 2 *For a fixed object width the perceived central modulation decreases with increasing edge modulation frequency.*

These results appear to be consistent with our earlier prediction. Specifically, if the perceived centers of the wiggle stimuli inferred by Burbeck *et al.* coincide with high-order shocks, *the central modulation computed as the horizontal displacement between fourth-order shocks in alignment with successive left and right sinusoidal peaks, Figure 8, should agree with the psychophysical data.* Thus in the following section we compare computational results obtained using the shock detection algorithm (Section 2.2) with Burbeck *et al.*’s data.

⁴ The location at which a subject is statistically equally likely to judge the probe dot to be to the left or to the right of the object’s center.




Fig. 7. The shock-based description of selected 40% amplitude modulation stimuli used in [10]. a) 0.75 degree width, 0.25 cycles/degree edge modulation; b) 0.75 degree object, 2.0 cycles/degree edge modulation; c) 1.5 degree object, 0.25 cycles/degree edge modulation; d) 1.5 degree object, 2.0 cycles/degree edge modulation.




Fig. 8. Shape (d) from Figure 7 is rotated and second-order and fourth-order shocks are labeled (all other shocks are first-order). Note that the fourth-order shocks are in alignment with the sinusoidal peaks.

4 Shock-Based Results

4.1 Psychophysics

We computed shock-based representations for all 24 wiggles using the algorithm in Section 2.2. Results for selected stimuli are shown in Figure 7, with the geometry of the high-order shocks explained in Figure 8. As evidence in support of our hypothesis, consider the computed central modulations overlaid as solid lines on the original observer data taken from [10], in Figure 9. We note that the predicted central modulations are consistent with an “aver-

age” observer’s data.

Whereas the core and the shock-based representation are motivated from quite different points of view, the strong overlap between computational and psychophysical results for each model points to an intimate connection between them. A precise mathematical comparison is beyond the scope of this paper. Instead, we identify the qualitative connections that have emerged. In particular, the “fuzziness” of the core model, whereby the width of the core scales with object width, is paralleled by the ratio of a shock’s formation time to the lifetime of the shape, a measure of local width/global width. This property is also reflected in the “bisection-threshold” or variance of the perceived centers in Burbeck *et al.*’s psychophysical experiments. Underlying this notion is the concept that the scale at which boundaries should interact to form more global object models is proportional to the spatial extent across which they communicate. Both models support medial-axis like representations for shape [5,2,3,6,27,35,26,31,36,24,18]. It is indeed gratifying that they are further quantifiably consistent with human performance involving shape.

4.2 *Recognition and Structural Significance*

We now expand the previous ideas to show how they are relevant for the recognition of generic objects. To begin, in Figures 10, 11 and 12 we show examples of a collection of different objects with their labeled shocks. Recalling that the Blum skeleton is the locus of positions through which the shocks migrate, notice that from a topological perspective several very different objects have almost identical Blum skeletons. (To do this, just ignore the colors in the diagrams that follow; the skeleton is the locus of shock positions.) The problem of shape matching using medial-axis based representations, attempted by [31,49,16] and others, has proved difficult precisely due to this reason. Changes in the representation due to noise are often more extensive than inherent structural differences. To make matters worse, even perfect skeletons may not suffice for recognition: different objects may give rise to the same skeleton, as exemplified by the pliers in Figure 10 and the skeleton of an elementary rectangle. Of course, there are quantitative geometric differences, but it is these differences that are most sensitive to noise.

Labeling the skeleton with the type of the shock from which each skeletal point derives provides a much richer foundation for recognition. In effect, this puts the “color” back into our illustrations and, in particular, now reveals that the pliers, the brush, and the hammer all have very different *labeled skeletons*. Moreover, the labeling can direct the matching. For example, it would be implausible to match a 3-shock branch in the nose of the pliers (Figure 10) to a 1-shock branch in the corner of the brush (Figure 11), although each of

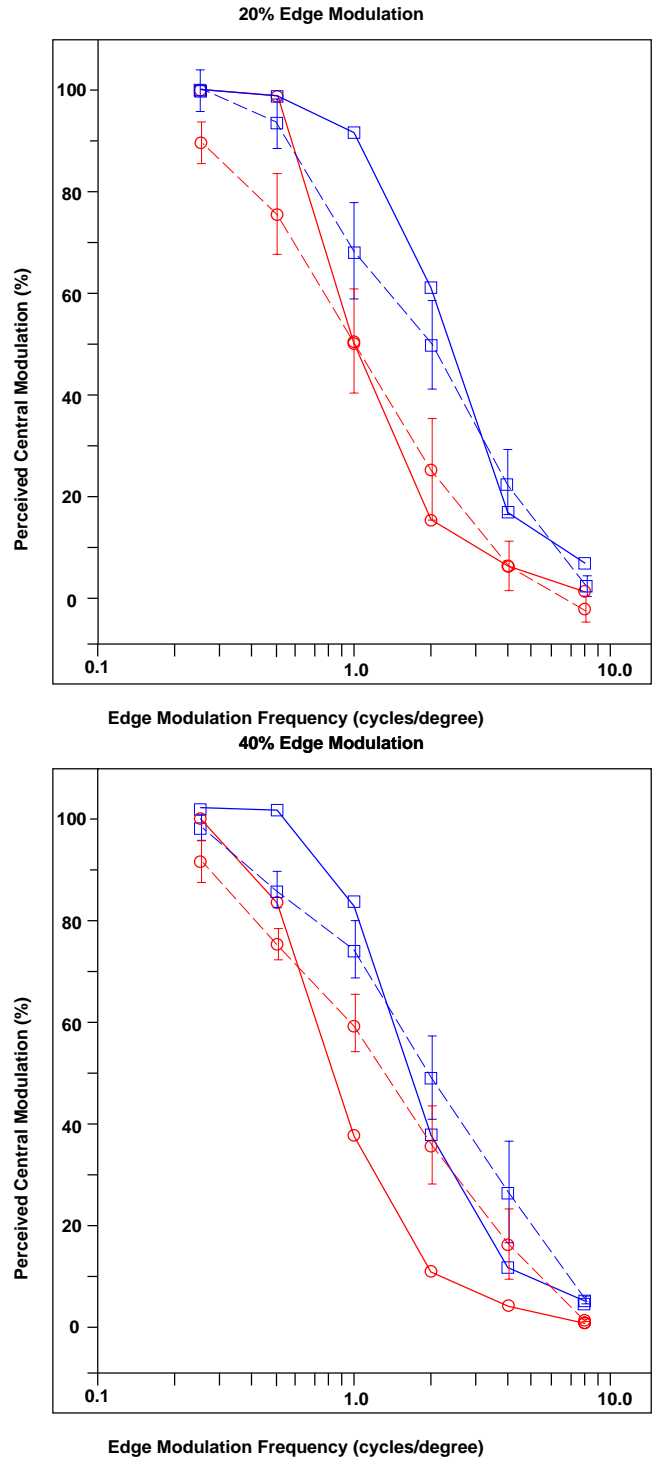


Fig. 9. Central modulations computed from shock-based descriptions (solid lines) are overlaid on the average observer data (dashed lines) from [10]. The central modulations are expressed as a percentage of the edge modulation amplitude and are plotted against edge modulation frequency for amplitudes of 20% of the object width (top) and 40% of the object width (bottom). Results for the wider 1.5° object are depicted by the circles and for the narrower 0.75° object by the squares.

these is simply a skeletal branch in the Blum sense. Such notions have been formalized via an abstraction of the shock-based system into a *shock graph*, and algorithms have been developed for shape matching using graph-theoretic concepts [44,34]. The topological structure of the graph provides a powerful means for indexing into database to find structurally similar shapes, and the matching algorithms are very efficient, typically taking only a few seconds to converge.

More is required for matching than simple labeling, however, as the psychophysical results in this paper have shown. The perceptual effect of a “wiggle” was a function of the geometry of the shocks as well as the ratio of local width to global width, as measured by the ratio of 2-shock to 4-shock formation times. This ratio is an example of a structural significance measure, which extends to far more complicated shapes than the symmetric wiggles. For example, the significance of any component of the hand shape in Figure 10, such as a finger, can be measured with respect to the formation time of the central 3-shock. This is an example of an object-based scale, like our earlier example of a nose pimple as compared to a pimple on the cheek. Much remains to be done with respect to object-based scales, and we hope that the examples in this paper provide a motivation for considering them more fully in the future.

Acknowledgments

This work was supported by grants from the Natural Sciences and Engineering Research Council of Canada, from the National Science Foundation, from the Air Force Office of Scientific Research and from the Army Research Office. We thank Christina Burbeck, Steve Pizer and Xiaofei Wang for fruitful discussions and for supplying us with the wiggle stimuli, and Sven Dickinson for the use of shapes from the Rutgers Tools database.

References

- [1] L. Alvarez, F. Guichard, P. L. Lions, and J. M. Morel. Axioms and fundamental equations of image processing. *Archive of Rational Mechanics and Analysis*, 123, 1993.
- [2] C. Arcelli, L. P. Cordella, and S. Leviadi. From local maxima to connected skeletons. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(2):134–143, March 1981.
- [3] C. Arcelli and G. S. di Baja. Euclidean skeleton via centre-of-maximal-disc extraction. *Image and Vision Computing*, 11(3):163–173, 1993.

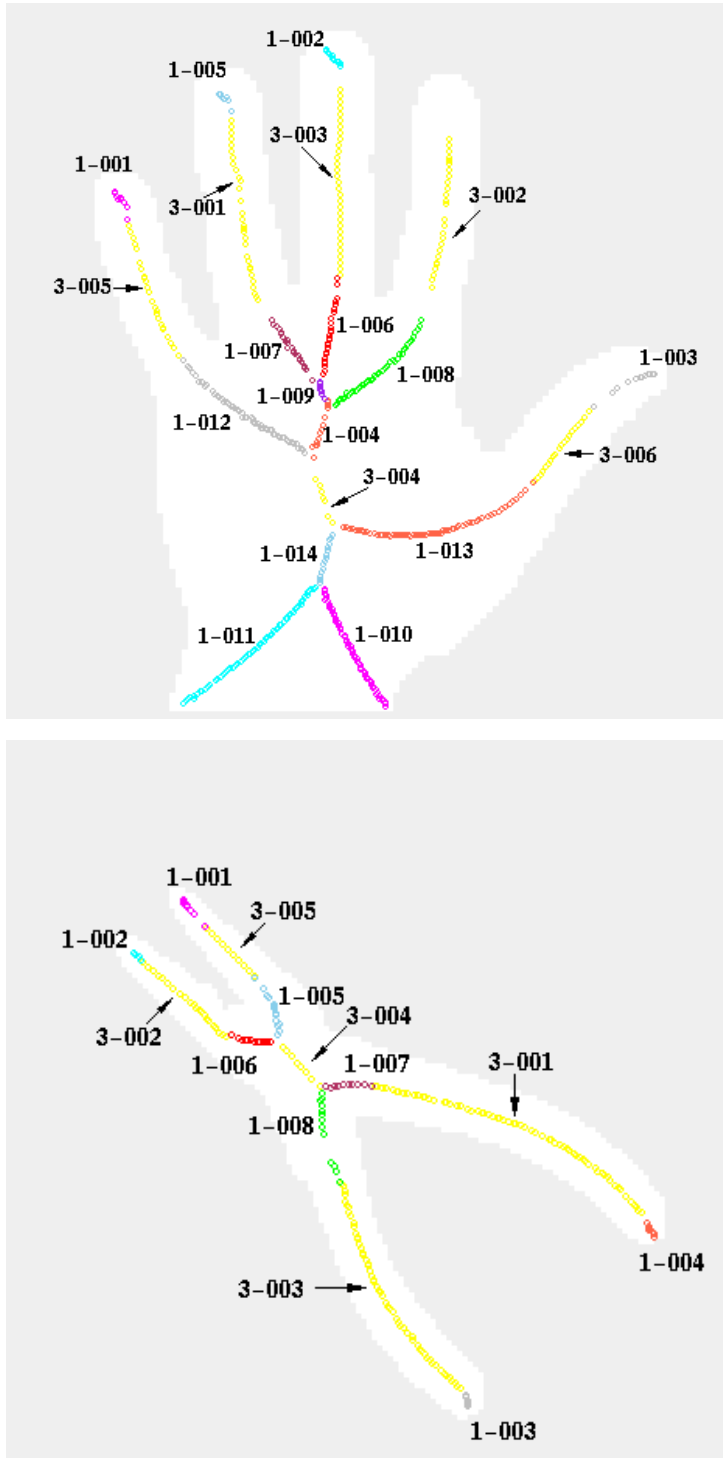


Fig. 10. The shocks computed for a hand (from the NRCC database), and a plier (from the Rutgers University Tool Database) using the algorithm in Section 2.2. The shapes were segmented from intensity images using the active contour developed in [43]. Each distinct group of shocks is given a unique label, with the first digit denoting the “type”. Adapted from [44].

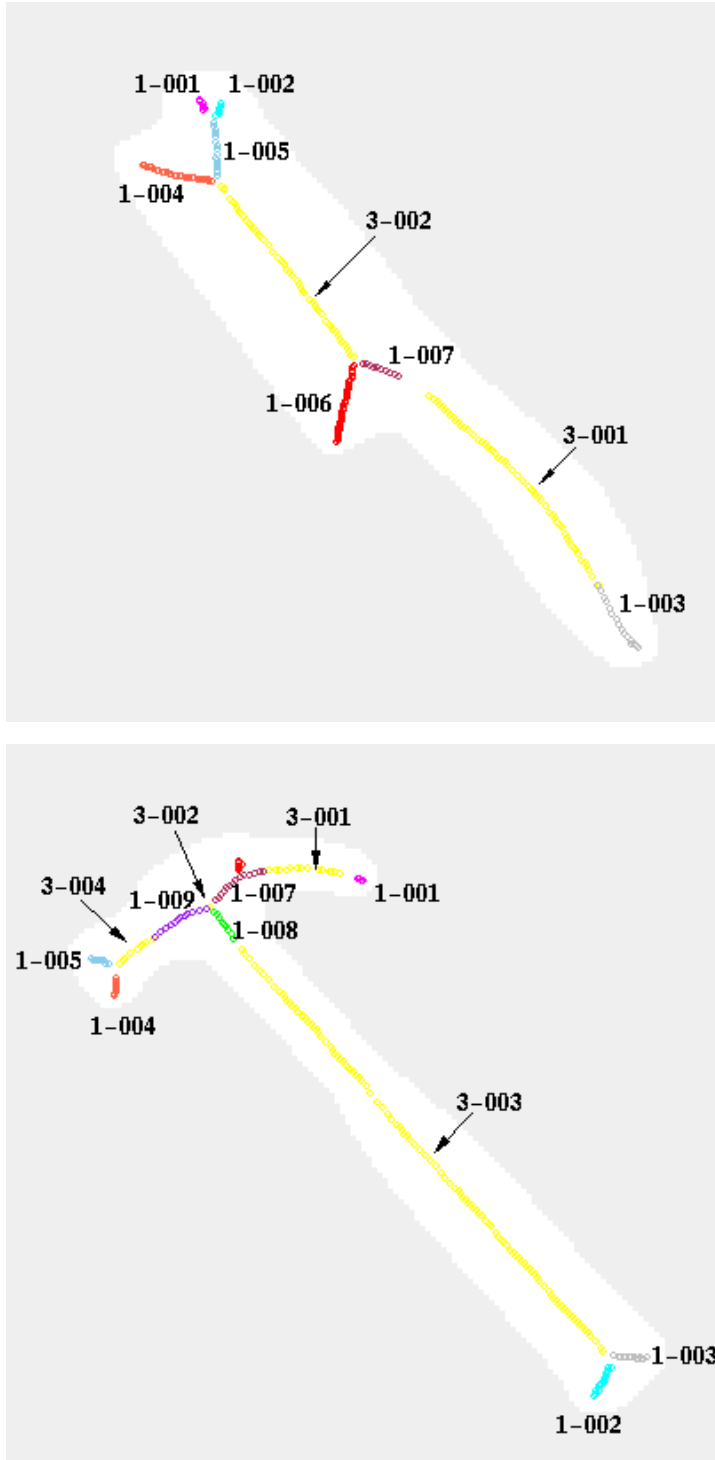


Fig. 11. The shocks computed for a brush and a hammer (from the Rutgers University Tool Database) using the algorithm in Section 2.2. The shapes were segmented from intensity images using the active contour developed in [43]. Each distinct group of shocks is given a unique label, with the first digit denoting the “type”. Adapted from [44].

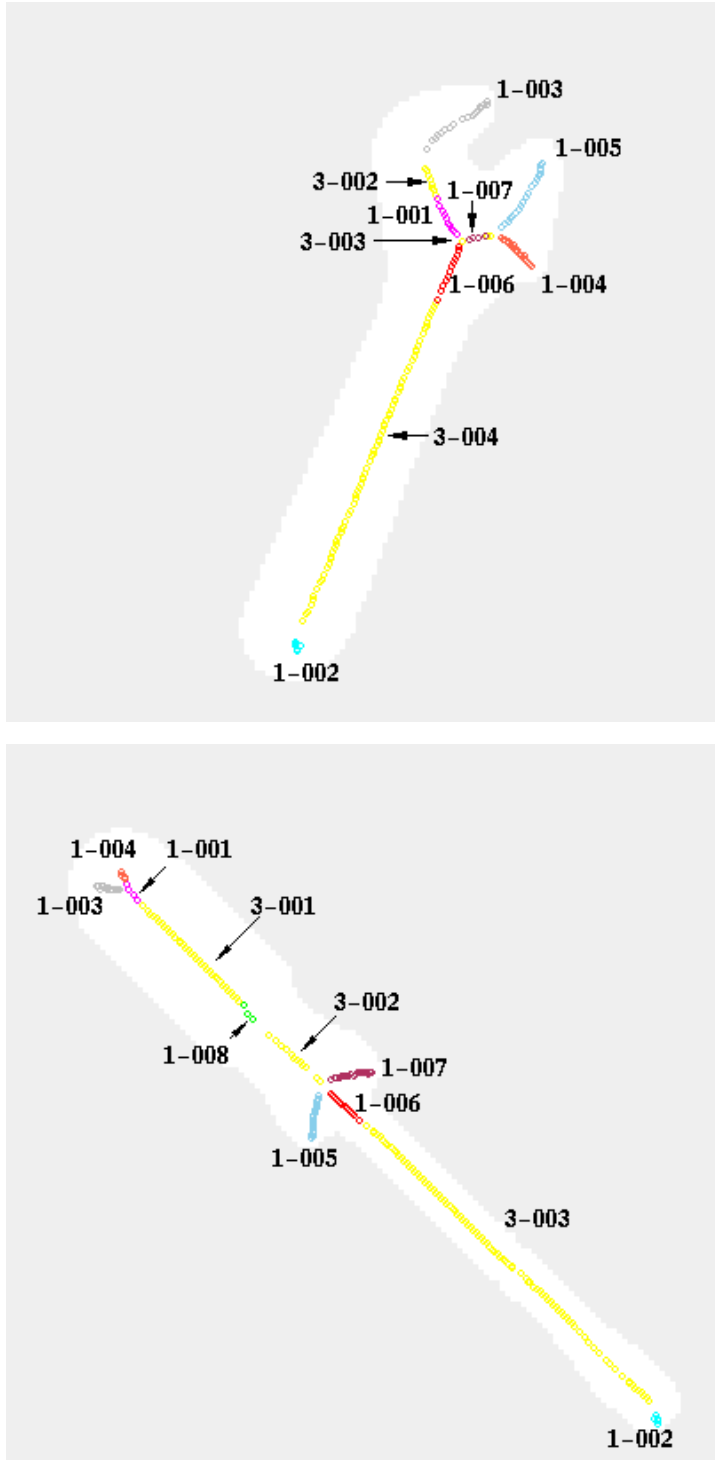


Fig. 12. The shocks computed for a wrench and a screwdriver (from the Rutgers University Tool Database) using the algorithm in Section 2.2. The shapes were segmented from intensity images using the active contour developed in [43]. Each distinct group of shocks is given a unique label, with the first digit denoting the “type”. Adapted from [44].

- [4] T. O. Binford. Visual perception by computer. In *IEEE Conference on Systems and Control*, December 1971.
- [5] H. Blum. Biological shape and visual science. *J. Theor. Biol.*, 38:205–287, 1973.
- [6] M. Brady and H. Asada. Smoothed local symmetries and their implementation. *International Journal of Robotics Research*, 3(3):36–61, 1984.
- [7] R. Brockett and P. Maragos. Evolution equations for continuous-scale morphology. In *Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing*, San Francisco, CA, March 1992.
- [8] C. Burbeck and S. Hadden. Scaled position integration areas: Accounting for weber’s law for separation. *Journal of the Optical Society of America A*, 10(1):5–15, 1993.
- [9] C. A. Burbeck and S. M. Pizer. Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research*, 35:1917–1930, 1995.
- [10] C. A. Burbeck, S. M. Pizer, B. S. Morse, D. Ariely, G. S. Zauberman, and J. Rolland. Linking object boundaries at scale: A common mechanism for size and shape judgements. *Vision Research*, 36(3):361–372, 1996.
- [11] P. J. Burt and E. H. Adelson. The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*, 31(4):532–540, April 1983.
- [12] P. Danielsson. Euclidean distance mapping. *Computer Graphics and Image Processing*, 14:227–248, 1980.
- [13] L. M. J. Florack, A. H. Salden, B. M. t. Haar Romeny, J. J. Koenderink, and M. A. Viergever. Nonlinear scale-space. *Image and Vision Computing*, 13:279–294, May 1995.
- [14] L. M. J. Florack, B. M. t. Haar Romeny, J. J. Koenderink, and M. A. Viergever. Linear scale-space. *Journal of Mathematical Imaging and Vision*, 4(4):325–351, 1994.
- [15] L. M. J. Florack, B. M. t. Haar Romeny, J. J. Koenderink, and M. A. Viergever. The gaussian scale-space paradigm and the multiscale local jet. *International Journal of Computer Vision*, 18:61–75, April 1996.
- [16] A. Francois and G. Medioni. Generic shape learning and recognition. In *International Workshop on Object Recognition*, 1996.
- [17] D. D. Hoffman and W. A. Richards. Parts of recognition. *Cognition*, 18:65–96, 1985.
- [18] M. Kelly and M. D. Levine. Annular symmetry operators: A method for locating and describing objects. In *Fifth International Conference on Computer Vision*, pages 1016–1021, June 1995.

- [19] B. B. Kimia, A. Tannenbaum, and S. W. Zucker. Toward a computational theory of shape: An overview. *Lecture Notes in Computer Science*, 427:402–407, 1990.
- [20] B. B. Kimia, A. Tannenbaum, and S. W. Zucker. Shape, shocks, and deformations I: The components of two-dimensional shape and the reaction-diffusion space. *Int. J. Computer Vision*, 15:189–224, 1995.
- [21] B. B. Kimia, A. R. Tannenbaum, and S. W. Zucker. The shape triangle: Parts, protrusions, and bends. In *Aspects of Visual Form Processing*, pages 307–323. World Scientific, June 1994.
- [22] J. J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.
- [23] J. J. Koenderink and A. J. van Doorn. Dynamic shape. *Biological Cybernetics*, 53:383–396, 1986.
- [24] I. Kovacs and B. Julesz. Perceptual sensitivity maps within globally defined visual shapes. *Nature*, 370:644–646, 1994.
- [25] P. D. Lax. Shock waves and entropy. In *Contributions to Nonlinear Functional Analysis*, pages 603–634, New York, 1971. Academic Press.
- [26] F. Leymarie and M. D. Levine. Simulating the grassfire transform using an active contour model. *IEEE Transactions On Pattern Analysis and Machine Intelligence*, 14(1):56–75, January 1992.
- [27] M. Leyton. Symmetry-curvature duality. *Computer Vision, Graphics, and Image processing*, 38:327–341, 1987.
- [28] T. Lindeberg. *Scale-Space Theory In Computer Vision*. Kluwer Academic Publishers, 1994.
- [29] D. Marr and K. H. Nishihara. Representation and recognition of the spatial organization of three dimensional structure. *Proceedings of the Royal Society of London*, B 200:269–294, 1978.
- [30] B. S. Morse, S. M. Pizer, and C. A. Burbeck. General shape and specific detail: Context-dependent use of scale in determining visual form. In *Aspects of Visual Form Processing*, pages 374–383. World Scientific, June 1994.
- [31] R. L. Ogniewicz. *Discrete Voronoi Skeletons*. Hartung-Gorre, 1993.
- [32] S. J. Osher and J. A. Sethian. Fronts propagating with curvature dependent speed: Algorithms based on hamilton-jacobi formulations. *Journal of Computational Physics*, 79:12–49, 1988.
- [33] S. J. Osher and C.-W. Shu. High-order essentially non-oscillatory schemes for Hamilton-Jacobi equations. *SIAM J. Numer. Anal.*, 28:907–922, 1991.
- [34] M. Pelillo, K. Siddiqi, and S. W. Zucker. Matching hierarchical structures using association graphs. In *Proceedings of the Fifth European Conference on Computer Vision*, June 1998.

- [35] S. M. Pizer, W. R. Oliver, and S. H. Bloomberg. Hierarchical shape description via the multiresolution symmetric axis transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 9(4):505–511, July 1987.
- [36] H. Rom and G. Medioni. Hierarchical decomposition and axial shape description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(10):973–981, October 1993.
- [37] G. Sapiro and A. Tannenbaum. Affine invariant scale-space. *International Journal of Computer Vision*, 10:25–44, 1993.
- [38] J. A. Sethian. *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Materials Science*. Cambridge University Press, 1996.
- [39] C.-W. Shu and S. J. Osher. Efficient implementation of essentially non-oscillatory shock-capturing schemes. *Journal of Computational Physics*, 77:439–471, 1988.
- [40] K. Siddiqi and B. B. Kimia. A shock grammar for recognition. Technical Report LEMS 143, LEMS, Brown University, September 1995.
- [41] K. Siddiqi and B. B. Kimia. A shock grammar for recognition. In *Conference on Computer Vision and Pattern Recognition*, pages 507–513. IEEE Computer Society Press, June 1996.
- [42] K. Siddiqi, B. B. Kimia, and C. Shu. Geometric shock-capturing eno schemes for subpixel interpolation, computation and curve evolution. *Graphical Models and Image Processing*, 59(5):278–301, September 1997.
- [43] K. Siddiqi, Y. B. Lauzière, A. Tannenbaum, and S. W. Zucker. Area and length-minimizing flows for shape segmentation. *IEEE Transactions on Image Processing*, 7(3):433–443, 1998.
- [44] K. Siddiqi, A. Shokoufandeh, S. J. Dickinson, and S. W. Zucker. Shock graphs and shape matching. In *Proc. Int. Conf. Comp. Vision*, pages 222–229, Bombay, India, 1998.
- [45] K. Siddiqi, K. J. Tresness, and B. B. Kimia. Parts of visual form: Psychophysical aspects. *Perception*, 25:399–424, 1996.
- [46] H. Tek, F. Leymarie, and B. B. Kimia. Interpenetrating waves and multiple generation shocks via the cedt. In C. Arcelli, L. P. Cordella, and G. S. di Baja, editors, *Proceedings of the Third International Workshop on Visual Form: Advances in Visual Form Analysis*, pages 582–593. World Scientific, 1997.
- [47] B. M. ter Haar Romeny, editor. *Geometry-Driven Diffusion in Computer Vision*. Computational Imaging and Vision. Kluwer Academic Publishers, Dordrecht, 1994.
- [48] A. P. Witkin. Scale-space filtering. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 1019–1022, Karlsruhe, West Germany, Aug. 1983.

- [49] S. Zhu and A. L. Yuille. Forms: a flexible object recognition and modelling system. *International Journal of Computer Vision*, 20(3):187–212, 1996.
- [50] S. W. Zucker. Structural scales in computational vision. In K. Bowyer and N. Ahuja, editors, *Advances In Image Understanding: A Festschrift for Azriel Rosenfeld*, pages 130–141. IEEE Computer Society Press, 1996.