INTERPENETRATING WAVES AND MULTIPLE GENERATION SHOCKS VIA THE CEDT*

HÜSEYIN TEK
LEMS, Brown University, RI 02906, USA

FREDERIC LEYMARIE
Thomson-CSF/Syseca, S2IG/GIIS Dept., Malakoff 92247 Cedex, France

BENJAMIN B. KIMIA
LEMS, Brown University, RI 02906, USA

The extraction of symmetries as quench points of propagating orientation elements from edge maps of grey scale images for object recognition is faced with fundamental theoretical and computational challenges. The theoretical issues arise since object symmetries are drastically altered due to missing edges (gaps), introduction of new parts and occluders, and spurious edges. While the *full* symmetry set retains much of the original figure's symmetries and as a result is less sensitive to such changes, it brings to bear many unintuitive branches, thus requiring further selection for object recognition. In this paper, we view the full symmetry set as the superposition of shocks arising from *multiple generations of waves*: the quenching points of the waves from the initial edge map constitute the first generation of shocks. A second generation of waves initiated at these points, simulate interpenetrating waves and generate a second generation of shocks, and so on until no further shocks can be formed. This view of the full symmetry set supports a *selective* continuation of waves, *e.g.,* at shock loops to remove spurious edges, and at shock-hypothesized limbs to partition shape and close boundary gaps. This selective continuation of waves brings out relevant symmetries, but avoids the ambiguity of the full symmetry set. The computational challenge is addressed by a framework based on the Contour-based Euclidean Distance Transform (CEDT) for shock detection, classification, labeling, as well as for simulating interpenetrating waves and multiple generation shocks described above. The key feature of CEDT that makes this possible is the explicit simultaneous propagation of orientation and distance, as well as additional features, *e.g. labels*. In addition, CEDT is exact and of very low numerical complexity. The results for a number of illustrative examples indicate the suitability of this framework for the recovery of object symmetries from real imagery.

1 Introduction

The versatility of symmetry-based representations for object recognition from unsegmented *real* imagery is seriously challenged by at least three fundamental theoretical issues. First, the extraction of symmetries is ill-conditioned in that the removal of a small portion of the boundary introduces a large change in

*This research was supported by funding from NSF and the Whitaker Foundation.

the object symmetry, Figure 1a. Second, the changes induced on the object's visible boundary by an occluder, or as result of changes in the visibility of a part, induce large changes in the object symmetry, Figure 1b. Finally, spurious edges due to surface markings, texture edges, specularity highlights, noise, *etc.* alter the underlying object symmetry, Figure 1c. Thus, symmetries recovered in a straightforward fashion from edge responses from real images often contain *(i)* additional skeletal points due to gaps; *(ii)* missing skeletal points due to occluders and parts; and *(iii)* distorted symmetries due to spurious elements, such that the resulting skeletons are not recognizable, nor lead to figure/ground segmentation. These difficulties have prompted several approaches to extract the *full symmetry set* from unsegmented gray-scale images. Scott *et al.*[1] propagate waves to recover the full symmetries. They also suggest a convolution approach for implementing the full symmetry set. Pizer *et al.*[2] use a similar approach where by a voting scheme, edges measured at each scale vote for medialness at a point which is a constant proportion of scale away. The ridges of the resulting surface constitute the *core*, a skeleton in x, y and σ (scale). Kelly and Levine[3] use annular symmetry operators in a similar fashion to derive the full symmetry set. In related work, Tari and Shah[4] define symmetries as the curvature maxima of level sets constructed consistently with an edge strength functional. August *et al.*[5] use the notion of a gap skeleton to group certain nearby endpoints identified by a depth separation process.

It has been argued that since the full symmetry set represents all the symmetries of a shape, spurious elements and gaps affect the full symmetry set less than the SAT. We argue that while some of the full symmetry set is revealing, not all of it is useful, and some additional unintuitive branches can infact lead to ambiguities for object recognition, Figures 6 and 8e. Rather, we propose that only in select situations should further symmetries be recovered. First, we observe that the full symmetry set can be viewed as the union of the quench points of a series of waves: the first generation of waves is launched at the edges of the image, while the second is launched at the quench points of the first generation shocks. In general, the nth generation wave is launched at the quench points of the $(n-1)$th generation wave. The union of the multiple generations of shocks constitutes the full symmetry set. Second, we observe that the maze of unwanted symmetries can be avoided by launching the secondary and future generations of shocks *selectively*, *e.g.*, at loops, to bring out relevant symmetries, Section 4.

The computational challenges posed by the ideas of *inter-penetrating waves* and *multiple generation shocks* can be met by a framework based on the contour based distance transform (CEDT). Previously, many approaches were developed to obtain skeletons from the distance transform. However, such schemes

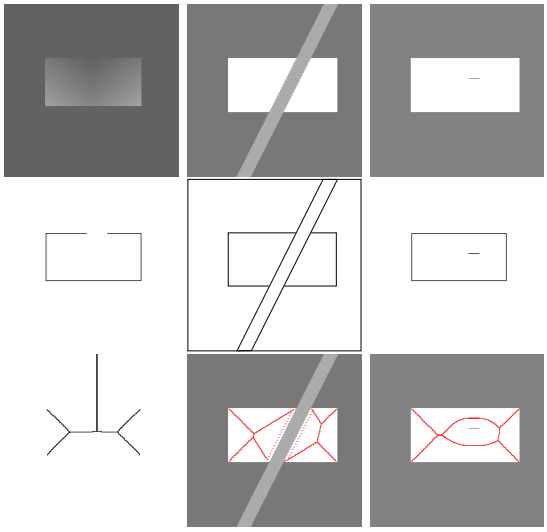

Figure 1: Gaps (a), occluders (b) spurious edges (c) drastically alter object symmetries. Top row: Original image; middle row: edge map; bottom row: skeletons pertaining to the object.

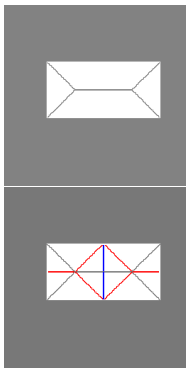

Figure 2: (a) SAT and (b) Full symmetry set from CEDT

do not make propagation of orientation and additional labels explicit. The CEDT not only propagates *distances* but also *orientation* and shock labels such that it is immediately clear at each point on the wavefront not only how far a wave has traveled, but also what the *direction* of propagation is, which point on the original boundary gave rise to it, and whether it is a regular or a rare-faction wavefront. This is the fundamental advantage of the CEDT which we utilize for the recovery of shocks, shock classification, and shock-labeling as an alternative to curve evolution. Two additional key advantages of CEDT over traditional raster-scan based schemes are its lower numerical complexity and its accuracy [6,7].

2 Wave Propagation for Skeleton Computations

The classic paper by Blum[8] has motivated a number of approaches for extracting the symmetries of binary segmented shape in the form of a "skeleton" or "stick figure", including extracting the center of maximally inscribed circles[9], fitting generalized cylinders, computing mid chords of double-tangent circles[10],

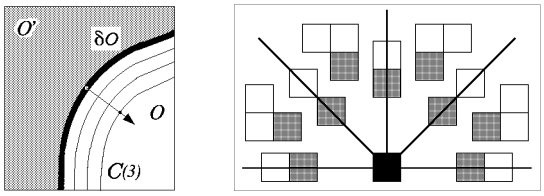

Figure 3: (a) Three iterations of the propagation of a boundary δO at one pixel per iteration. (b) The minimal set of masks (only a mirror image constructs the bottom half of the full picture): The black pixel represents the source while the grey and white pixels represent subsequent propagation of waves. Black lines represent directly supported directions of propagation. for the intermediate directions of propagation, larger masks are required. (Adapted from [6])

extracting ridges of distance transforms [11,12,13], Voronoi diagram methods [14], and thinning algorithms [9]. Blum's original idea was based on a "grass fire" initiated at the contours of the shape which quenches at the skeleton. The reaction-diffusion space and the formation of shocks (singularities) implements and generalizes this idea [15,17,18]. The reaction process can be simulated much more efficiently, however, by the distance transform methods which map a binary image into an image where the value at each point is the Euclidean distance from the object [11]. Each constant distance sets then represents a front where distance represents "time".

Distance transform methods may be classified into those based on raster scans or those that are contour-based. Danielsson showed that Euclidean Distance Transform (EDT) can be computed by comparing neighborhood pixels with vector-valued masks used to evaluate distance steps along the axis of the supporting grid. A raster scan version for EDT (REDT) was described in [19] and later extended to signed EDT [20] and then to contour-based EDT (CEDT) [6,7]. Previously, REDT and CEDT have been considered of similar usefulness when used for wave propagation and symmetry computation mainly because REDT is relatively simpler to implement, due to independence of embedded shapes. In addition, CEDT implementations for symmetry set elicitation have relied upon dilation/erosion of chain-coded representations of contour [21,22], requiring a pre-processing of contour features. Thus, CEDT has received relatively little attention. However, for simulating wave propagation, we propose that CEDT has a key advantage over REDT in that it provides an explicit representation of *orientation* in addition to *distance* of propagation. This is particularly attractive since CEDT can be initiated at points, open and closed contours, and surface patches, without requiring chain-code pre-processing, leading to a more efficient and direct simulation of wave propagation.

The basic design of CEDT follows ideas originated in Montanari [23], brought

to the foreground by more recent work, *e.g.*, [24,22] and, Ragnemalm [6,7]. While Montanari had the key insight that wave propagation from boundary feature was potentially more efficient than raster-scan sequential DT, Ragnemalm imported the idea of using vectored values for DT [19] and studied different masks and their properties for propagating various metrics from contours. Figure 3 illustrates the minimal complete set of masks required for Euclidean metric propagation on a 2D rectangular distance grid from a point source. Each masks maintains a distance vector (L_x, L_y) from its origin, thus *explicitly* representing the direction and distance of the propagating wavefront. Note that the distance vector values may in addition propagate other features, *e.g.*, a labeling of original front waves into regular or rarefaction, Section 3. The metric $L^2 = (L_x^2 + L_y^2)$ is also carried to optimize operations. A further optimization uses *buckets* to store wavefront distance values in order of metric L^2, thus leading to constant speed propagation. Other advantages of CEDT include: CEDT is exact, is nearly optimal in terms of numerical complexity when compared to raster-scan based DT, and is easily extended to 3D by defining additional masks.

The extraction of skeleton from distance transform is faced with a number of difficulties. First, most approaches for computing skeletons from DTs rely on non-Euclidean metrics giving rise to highly inaccurate results, and in particular failing to provide rotation invariance [11,25]. Second, approaches relying on a raster-scan implementation to compute DTs, extract skeleton by some post-processing of the ridges of the computed distance map [11,26]. Third, approaches based on the retrieval of a smooth contour representation, such as splines, to compute skeletons from a derived distance map [27] suffer from the two major difficulties: (i) an additional complexity due to the contour modeling (finding good nodes) and (ii) the creation of artifacts due to non-smooth contour features. Finally, extensions to the third dimension usually leads to high numerical complexity. We now consider an alternative approach by tailoring a previous framework based on curve evolution to use CEDT.

3 Shock Detection and Classification by CEDT

Shape can be completely described as the collection of four types of shocks which form in the course of deformations of shape in the reaction-diffusion space [15,17,18] $\frac{\partial \mathcal{C}}{\partial t} = (\beta_0 - \beta_1 \kappa)\vec{N}$. The four types of shocks correspond to intuitive elements of shape, namely, parts, protrusions, and bends [16]. The deformations are implemented via the *curve evolution* paradigm by embedding the curves $\mathcal{C}(s,t)$ as the level set of a surface $\{\psi(x,y,t) = 0\}$ evolving by $\frac{\partial \psi}{\partial t} = (\beta_0 - \beta_1 \kappa)|\nabla \psi|$. Table 1 shows a classification of shocks from ψ which

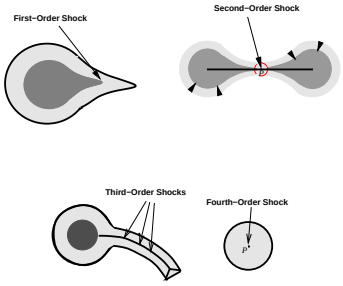

Shock Type	Orientation	Curvature
First−Order	non− vanishing $\nabla\phi$	high level set curvature
Second−Order	isolated vanishing $\nabla\phi$	$\kappa_1\kappa_2 < 0$
Third−Order	non−isolated vanishing $\nabla\phi$	$\kappa_1\kappa_2 = 0$
Fourth−Order	isolated vanishing $\nabla\phi$	$\kappa_1\kappa_2 > 0$

Table 1: This table depicts the classification of shock types based on the the gradient level set curvature and the principal curvatures of the surface.

Figure 4: Each of the four types of shocks each is correlated with a perceptual/semantic category, *i.e.*, *protrusion, part, bend*, and *seed*.

has been implemented to sub-pixel accuracy[18]. While the derived shock structure when $\beta_1 = 0$ is related to skeletons, several properties of shocks relating to the notions of type, velocity, grouping, salience, and hierarchy are significant. (*i*) certain deformations, *e.g.*, bending, affect shock types selectively, *e.g.*, third-order shocks of a rectangle; (*ii*) the skeletal representation lacks sufficient explicit dimensions for qualitative approximation, *e.g.*, the addition of shock type substantially narrows down the range of shapes it can generate; (*iii*) topological and differential properties of shocks, *e.g.*, velocity, directly reflect boundary properties; (*iv*) the notion of time of formation induces a hierarchy on shocks; and (*v*) a notion of shock grouping and salience based on the diffusion process ($\beta_1 \neq 0$) leading to a stability of representation with small changes.

We now argue that the explicit and simultaneous propagation of orientation and distance in the CEDT leads to an alternate and more efficient framework for the detection and classification of shocks. Observe that each point on the original shape (or even partial contour) can be considered a source for propagation of waves. Since both orientation and distance are available, the source of each point on the front can be determined. Thus, waves arriving at a point from two distinct sources can be identified and distinguished. Since shocks are the quench point of such waves, all image points receiving two or more waves can be easily detected. In addition, since the direction of propagation of each wave is explicit the shock type can be determined. Specifically, the CEDT approach updates the wavefront by propagating each point in the direction where the point had propagated from. In the discrete domain, the propagation continues until the propagating wave meets an incoming wave, as determined by the minimal distance carried by the waves. The collision signals the formation of a shock. Observe that two distinct directions must necessarily

have arrived at the shock points. If these two directions are not aligned, a first order shock has formed. The velocity of the shock is the vector sum of the two front velocities; *i.e.*, shock speed is $\frac{1}{sin\frac{\theta}{2}}$ times front speed where θ is the angle between the two fronts. If the two directions are aligned, a higher-order shock has formed. Table 1 shows a classification based on an embedding surface, which in this case we take to be the CEDT generated surface: negative, zero and positive Gaussian curvatures gives second-, third, and fourth-order shocks, respectively. Observe that the crucial advantage of CEDT in representing the angle at which wavefronts meet or cross each other, in contrast to REDT which necessitates a cumbersome and inaccurate post-processing of the distance map [11,26].

The removal of some portions of shape's complete boundary does not simply only lead to the removal of portions of its symmetries, *e.g.*, as represented by shocks, but also to generation of seemingly un-intuitive symmetries, Figure 5. These newly formed "spurious" shocks must be distinguished from shocks common with the previous case. Tek *et al.* [28] suggest that such a distinction should be based on whether propagating waves carry "true" orientation information as supported by the original boundary or carry "bogus" orientation arising from rare-fraction waves, *e.g.*, as arising from a concave corner. This distinction between two types of waves leads to labeling of shocks into three classes [28]:

Definition: Contour points with regular tangent give rise to *regular* wavefronts. Contour points without a uniquely defined tangents give rise to *degenerate* wavefronts. A shock point arising from the interaction of two regular wavefronts is *regular*. A shock point arising from one regular and one degenerate wavefront is *semi-degenerate*. A shock point formed from the interaction of two degenerate wavefronts is *degenerate* [28].

It is suggested that (i) regular shocks represent the only symmetries arising from partial contour segments, (ii) the semi-degenerate shocks are altered form of the underlying symmetries, and (iii) the degenerate shocks arise represent potential candidates for contour continuity and grouping, as in-partitioning shape [29], or arising in completing gaps [30]. The latter statement (iii) has similarities to August *et al.* [5] who use a notion of gap skeleton to group certain nearby endpoints in the edge map segregated by a depth process. However, there are several important distinctions, as (i) they operate on depth-segregated not full edge maps, (ii) our approach is motivated by the notion of orientation propagation and rare-faction waves [28]. Thus, waves simulated by CEDT carry not only orientation and distance but also a rarefaction/regular label which forms the basis of subsequent shock labeling by CEDT, Figures 5. In summary waves can be labeled and propagated, and shocks can be detected, classified, and

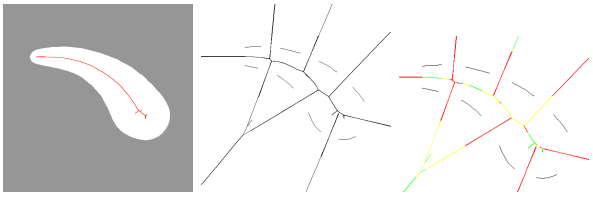

Figure 5: Extracting partial shocks from partial contours. (b) original image with complete boundary and its shocks, (c) partial boundary and the introduction of "spurious" shocks, (d) a labeling of shocks into regular (green, semi-degenerate (yellow) and degenerate (red). Observe that the regular shocks are the partial shocks extracted from partial contour (black).

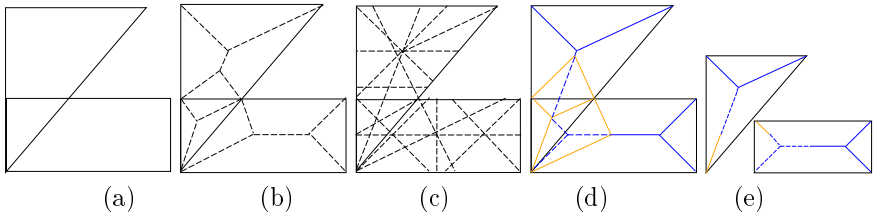

Figure 6: Edge evidence in real images often suggests a multiplicity of groupings, e.g., (a) suggesting a rectangle and a triangle. While the SAT representation does not capture triangular and rectangular symmetries (b) the full symmetry set is capable of bringing out the appropriate symmetries, but also introduces superfluous ones (c). Shocks of secondary waves exclusively initiated at loops (d), however, bring out only the missing pieces leading to two appropriate groupings (e).

labeled via the CEDT framework.

4 Inter-penetrating Waves and Multiple Generation Shocks

The sensitivity of SAT and the ambiguity of the full symmetry set prompts us to propose an alternative which is based on a view of the full symmetry set as the union of quench points of a series of waves:

Definition: The *first generation wave* is the wavefront initiated at the edge map of an image. The *first generation shock set* is the set of quenched points (shocks) arising from the propagation of first generation wave. The n^{th} *generation wave* is the wavefront initialized at the $(n-1)^{\text{th}}$ generation shocks at the time indicated by its formation. The n^{th} *generation shocks* are the shocks corresponding to the n^{th} generation wavefront.

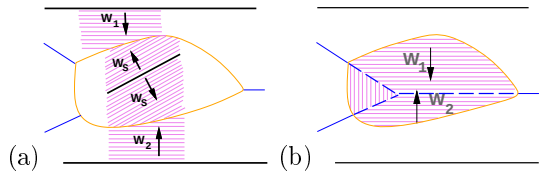

Figure 7: (a) waves W_1 and W_2, are quenched by waves from the spurious boundary, W_s, resulting in a loop in the shock structure. The shocks on the loop can now simulate the passage of the original flow of waves W_1 and W_2 via secondary waves, (b), which are initiated at the shock in a delayed manner, resulting in the formation of shocks due to the top and the bottom boundaries.

Proposition: The union of all generations of shocks is the full symmetry set.
Proof: Each point in the full symmetry set can be viewed as the quench point of two waves traveling without interruption from two boundary segments. Since wave initialized at quench points are the continuation of waves quenched at these points all two boundary segments eventually interact. Conversely, each multiple generation shock is clearly a point of the full symmetry set by the same argument, Figure 7.

The shock-based representation can implement such a process since complete information about the incoming waves is stored as shock location, time of formation, and velocity. The second observation is that multiple generation waves and shocks can recover the distorted or missing symmetries. The idea is to launch second and further generation of waves only at *select groups of shocks* as indicated by special properties of the shock itself. For example, an isolated spurious edge or equivalently a hole in the object interferes with the formation of appropriate symmetries, Figure 6, but also always leads to a *loop* in the shocks arising from it. Thus, selectively launching second generation waves at shock loop effectively removes this element Figure 7 and recovers appropriate symmetries, Figure 8f without generating additional symmetries, Figure 8e. Figure 9 illustrates two interacting spurious shocks requiring a second generation of waves. As a second example consider how symmetries are affected by the appearance of a newly visible part, Figure 10a. The partitioning theory of shape [29] cast in the language of shock labels [28], proposes that salient semi-degenerate and degenerate shocks signal parts. This provides a second criterion for selectively launching a new generation of shock, namely, at these semi-degenerate shocks and at a hypothesized limb part-line, Figure 10b to recover each object's part symmetry axes. A similar argument holds for gaps, which are viewed as "null parts" [30], Figure10e-g. Observe the need for multiple generation shocks to lead to an appropriate grouping essential for segmentation and recognition tasks, *e.g.* in indexing into image databases, Figure 11.

5 References:

1. Guy L. Scott *etal.* *IVC*, 7(1):63–70, 1989.
2. S. M. Pizer and C. A. Burbeck. *Vision Research*, 35,13:1917-1930, 1995.
3. M. F. Kelly and M. D. Levine. ICCV, pages 1016–1021, 1995.
4. S. Tari, J. Shah, and H. Pie. In *IEEE Work. Math. Meth. Bio. I. A.*, 1996.
5. J. August, K. Siddiqi, and S. W. Zucker. In *ICPR*, 1996.
6. I. Ragnemalm. Licenciate Thesis No 206, Linkoping University, Sweden, 1990.
7. I. Ragnemalm. *CVGIP: Image Understanding*, 56(3):399–409, 1992.
8. Harry Blum. *J. Theor. Biol.*, 38:205–287, 1973.
9. C. Arcelli and G. Sanniti di Baja. *IVC*, 11(3):163–173, 1993.
10. M. Brady and H. Asada. *Robotics Research*, 3(3), 1984.
11. F. Leymarie and M. D. Levine. *PAMI*, 14(1):56–75, January 1992.
12. Carlo Arcelli and Gabriela Sanniti di Baja. *PAMI*, 11(4):411–414, 1989.
13. C. Arcelli and G. Sanniti di Baja. *Patt. Rec. Let.*, 13(4):237–243, 1992.
14. R. L. Ogniewicz and O. Kübler. *Pattern Recognition*, 28(3):343–359, 1995.
15. B. Kimia, A. Tannenbaum, and S. W. Zucker. In *ECCV I*, pages 402–407.
16. B B. Kimia, A. Tannenbaum, and S W. Zucker. In IWWF 1993.
17. B. Kimia, A. Tannenbaum, and S. W. Zucker. In *IJCV I*, 15:189–224, 1995.
18. K. Siddiqi and B. Kimia. In *CVPR*, 1996.
19. P.E. Danielsson. *CVGIP*, 14:227–248, 1980.
20. Q.-Z. Ye. In *ICPR*, volume 1, pages 495–499, Bejing, China, 1988.
21. Y. Xia. *PAMI*, 11(10):1076–1086, 1989.
22. Luc Vincent. In *CVPR*, pages 520–525, Maui HI, 1991.
23. U. Montanari. *JACM*, 15(4):600–624, October 1968.
24. B.J.H. Verwer. In *ICPR*, pages 137–142, 1988.
25. F. Leymarie and M. D. Levine. *CVGIP*, 55(1):84–94, January 1992.
26. B. Kruse. Tech. Rep. LiTH-ISY-I-1116, Linkoping U., Sweden, 1990.
27. M. W. Wright, R. Cipolla, and P. J. Giblin. *PAMI*, 13(5):367–375, 1995.
28. H. Tek, P. Stoll, and B. B. Kimia. In CVPR 1997.
29. K. Siddiqi and B. Kimia. *PAMI*, 17(3):239–251, March 1995.
30. K. Siddiqi, K. Tresness, and B. Kimia. In *IWVF*, pages 507–521, Italy, 1994.

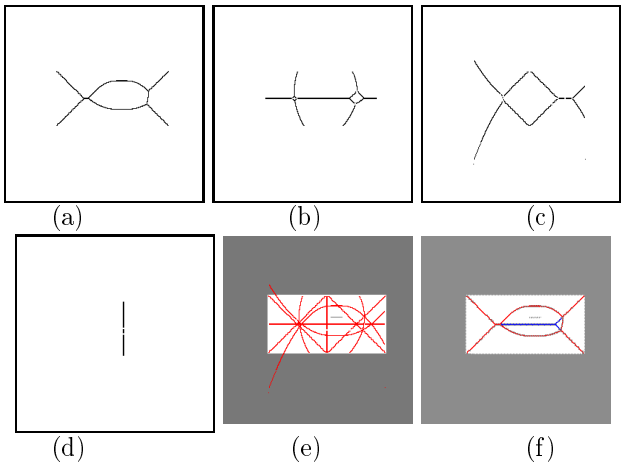

Figure 8: Multiple generation shocks of the image is Figure 1c; (a), (b), (c), and (d) depict first-, second-, third-, and fourth-generation shocks, respectively. (e) the superposition of all generations of shocks constitutes the full symmetry set. Observe that un-intuitive nature of the full symmetry set; (f) second generation of waves exclusively initiated at the loops gives rise to shocks which complete the rectangular symmetry set.

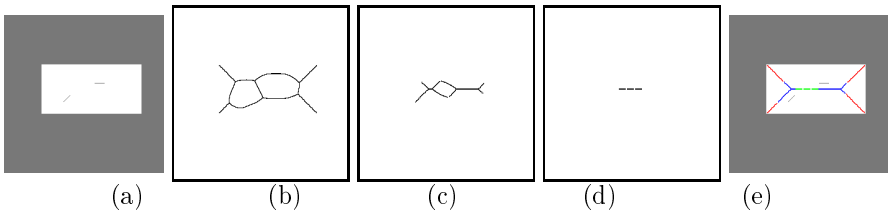

Figure 9: (a) original image; (b) first generation of shocks: (c) and (d) loop transformations to remove the effect of spurious edge elements. (e) the superposition of all generations of shocks.

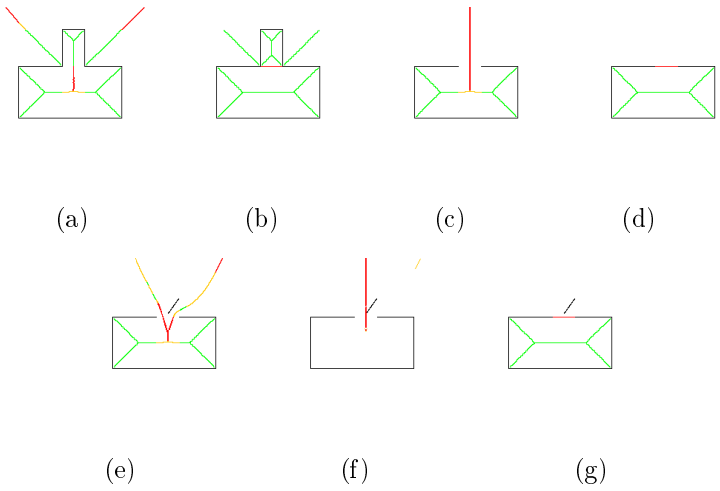

Figure 10: (a) A shape with two parts and its shocks. (b) The recovery of individual part's shocks by shock labeling followed by a second generation shocks of waves initiated at limbs. (c) a shape with a gap (null part) and its shocks. (d) Second generation waves initiated at salient limbs remove the degenerate shocks and correct for the distortion at semi-degenerate shocks. (e) a spurious edge element's interference with contour grouping can be removed by considering multiple generation of shocks as shown in (f) where second generation shocks arising from the shocks loops (completed by the image boundary) generate a new grouping hypothesis thus completing the rectangle symmetries (g).

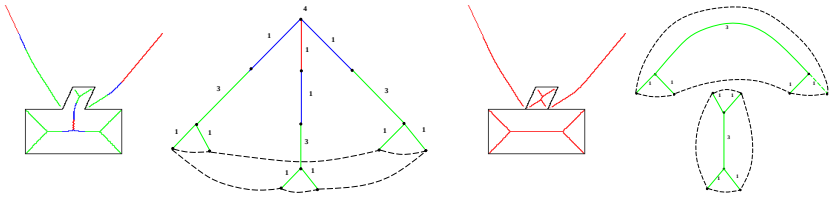

Figure 11: Role of interpenetrating wave in the segmentation of shape and the resulting graphs.