The Scale-Space Aspect Graph[1]

David W. Eggert* Kevin W. Bowyer*

Charles R. Dyer† Henrik I. Christensen‡ Dmitry B. Goldgof*

* Department of Computer Science and Engineering
University of South Florida
Tampa, Florida 33620
eggertd, kwb or goldgof@csee.usf.edu

† Department of Computer Science
University of Wisconsin
Madison, Wisconsin 53706
dyer@cs.wisc.edu

‡ Institute of Electronic Systems
Aalborg University
DK-9220 Aalborg East
hic@vision.auc.dk

Abstract

Currently the *aspect graph* is computed from the theoretical standpoint of perfect resolution in the viewpoint, the projected image and the object shape. This means that the aspect graph may include details that an observer could never see in practice. Introducing the notion of scale into the aspect graph framework provides a mechanism for selecting a level of detail that is "large enough" to merit explicit representation. This effectively allows control over the number of nodes retained in the aspect graph. This paper introduces the concept of the *scale space aspect graph*, defines an interpretation of the scale dimension in terms of the spatial extent of features in the image and presents a detailed example for a simple class of objects.

[1]This work was supported at the University of South Florida by Air Force Office of Scientific Research grant AFOSR-89-0036 and by National Science Foundation grant IRI-8817776.

The Scale-Space Aspect Graph

Abstract

Currently the *aspect graph* is computed from the theoretical standpoint of perfect resolution in the viewpoint, the projected image and the object shape. This means that the aspect graph may include details that an observer could never see in practice. Introducing the notion of scale into the aspect graph framework provides a mechanism for selecting a level of detail that is "large enough" to merit explicit representation. This effectively allows control over the number of nodes retained in the aspect graph. This paper introduces the concept of the *scale space aspect graph*, defines an interpretation of the scale dimension in terms of the spatial extent of features in the image and presents a detailed example for a simple class of objects.

Summary

1. What is this paper about?

   This research introduces the idea of incorporating the concept of scale into the aspect graph framework. One criticism of current aspect graph approaches mentioned at a recent workshop (IEEE Workshop on Directions in Automated CAD-Based Vision) was the lack of use of scale information, a problem which is actually prevalent in the entire computer vision field. It is important to include some quantitative information into the qualitative and idealized representation of the aspect graph. The general notion of a scale space aspect graph is presented and a detailed example of the use of image resolution as a scale parameter is given.

2. What is the original contribution of the work reported in this paper?

   This is the first and only research which incorporates the scale space concept into the aspect graph framework. A general definition of the scale space aspect graph is provided. Also, a particular interpretation of the scale parameter as image resolution is presented. This interpretation is explained through the complete analysis of this representation for the case of a nonconvex polygon in the plane.

3. What is the relationship of this work to related work of others?

   There has been a considerable volume of research on constructing aspect graphs in recent years as the domain of objects has grown from polygons to curved-surface objects. Also the application of scale space methods has evolved from 1-D curves to intensity images to 3-D object shape. However, there has not been any published work on applying scale space techniques to the aspect graph representation.

4. How can your contribution be used by others?

   This research is the first step in incorporating quantitative information into the aspect graph representation. The general theory presented discusses alternative interpretations of the scale parameter which are ready to be explored. Furthermore, the particular theory concerning the interpretation of scale as image resolution can be expanded to encompass broader classes of objects. Finally, the practical benefits of using a scale space aspect graph can be explored in the context of an actual object recognition system.

1 Introduction

The aspect graph is considered important because it provides a a complete view-centered representation of an object. Considerable research has been performed in recent years to generate algorithms that compute the aspect graph and its related representations [4, 7, 9, 14, 20, 21, 22, 24, 25, 26, 27]. However, the practical utility of the resulting aspect graphs has been questioned. A recent panel discussion on the theme "Why aspect graphs are not (yet) practical for computer vision" was held at the 1991 *IEEE Workshop on Directions in Automated CAD-Based Vision* [8]. One issue raised by the panel is that aspect graph research has not included any notion of *scale*. (Actually, the lack of knowledge concerning the use of the concept of scale was acknowledged to be a problem for computer vision in general rather than just aspect graphs in particular.) The use of scale can be seen as a method of making the representation more realistic.

Effectively, the aspect graph to date has been computed only for the ideal cases of perfect resolution in the viewpoint, the projected image and the object shape, leading to the following practical difficulties.

- A node in the aspect graph may represent a view of the object that is seen from such a small cell of viewpoint space that it is extremely unlikely to ever be witnessed,
- The views represented by two neighboring nodes in the aspect graph may differ only in some small detail that is indistinguishable in a real image, and
- A very small change in the detail of the 3-D shape of an object may drastically affect the number of nodes in the aspect graph.

Each of these factors contributes to a rather large overall size of the aspect graph representation. (For example, the worst-case node complexity is $O(N^9)$ for an N-faced polyhedra assuming a 3-D viewpoint space.) By introducing the concept of scale one hopes to reduce the large set of theoretical aspects to a smaller set of the "most important" aspects.

We begin by briefly reviewing the aspect graph representation and the scale space concept. Then a general definition of the *scale space aspect graph* is given. This is followed by a more detailed explanation of one promising interpretation of the scale dimension. This interpretation is supported by a complete example showing the scale space aspect graph for the case of a nonconvex polygon in the plane.

2 A Brief Review of the Aspect Graph Concept

A general definition of the aspect graph is that it is a graph structure in which

- there is a node for each *general view* of the object as seen from some maximal connected cell of *viewpoint space*, and
- there is an arc for each possible transition, called an *accidental view* or a *visual event*, between two neighboring general views.

A general viewpoint is defined as one from which an infinitesimal movement in each possible direction in viewpoint space results in a view that is equivalent to the original. In contrast, an accidental viewpoint is one for which there is at least one direction in which an infinitesimal movement results in a view that is different from the original.

Most algorithms for computing the aspect graph follow the same basic high-level approach:

1. Use the geometric definition of the object to enumerate the set of accidental views that occur, as well as the corresponding surfaces in viewpoint space from which they are seen.
2. Find the parcellation of viewpoint space defined by the set of accidental view surfaces.
3. Traverse the parcellation of viewpoint space to build the aspect graph and attribute the nodes with descriptions of the representative views.

The various algorithms that have been developed may be classified using three properties; the domain of objects, the view representation and the model of viewpoint space. The domain of objects has evolved from polygons [10], to polyhedra [9, 20, 25, 26, 27], to solids of revolution [7, 14], to piecewise-smooth objects [4, 21, 22, 24], to articulated assemblies [23]. Almost without exception, a view of the object is represented using a qualitative description of the line drawing, such as the *image structure graph* (ISG) [16]. The actual labeling of *contours* and *junctions* varies slightly among researchers. Distinctions between general and accidental views are usually based on isomorphism of the ISG. Lastly, two viewpoint space models are commonly used. The first is the 2-D *viewing sphere*, on which each point defines a viewing direction for *orthographic projection* [9, 14, 20, 21, 22, 24]. The other is 3-D space, in which each point is the focal point for a *perspective projection* [4, 7, 20, 23, 25, 26, 27]. (For greater detail on these algorithms, see [3].)

3 A Brief Review of the Scale Space Concept

In its strictest sense, the phrase "scale space of X" is taken to mean a parameterized family of X in which the detail of features in X is monotonically decreasing with increasing scale. Also, the qualitative features of X at a given scale can be traced back across all lower scales ("causality"). Research in this area was popularized by Witkin's scale space analysis of the inflections of a 1-D signal [28]. Since that time the scale space concept has been applied to the curvature of 2-D curves [5, 17], the curvature of 3-D curves [18], the 2-D intensity map [1, 12, 15, 29] and 3-D object shape [13]. In addition, a number of other researchers have described different "hierarchical" or "multi-resolution" representations, for example pyramids, that are similar to the scale space concept.

To better explain the approach we review the basics of Witkin's 1-D signal analysis. The qualitative (or symbolic) structure of a 1-D signal can be given in terms of the locations of its inflection points. The 2-D scale space of a 1-D signal is developed by introducing a second dimension, σ, that represents the the size of a Gaussian kernel used to smooth, or blur, the original signal. In this one parameter family of 1-D signals a value of $\sigma = 0$ yields the original signal, while $\sigma = \infty$ means the signal is reduced to a flat line. In the scale space, a particular inflection that exists at one value of σ may be traced over increasing values of σ until it is eventually annihilated (merged with a neighboring inflection). In keeping with the monotonicity requirement, inflection points can only be annihilated as σ increases, never generated. Thus the value of scale at which an inflection ceases to exist is a measure of its strength or importance.

4 The Scale-Space Aspect Graph

At this point, the high-level concept of a *scale-space aspect graph* should be rather apparent. The essence of the scale space concept is that a one parameter family of instances of some entity be created, and that changes in the qualitative complexity of the entity be apparent over changes in scale. The aspect graph is nothing more than a qualitative description of the underlying structure of the parcellation of viewpoint space into general views. Therefore it is appropriate to consider a parameterized family of these parcellations as the basis for the scale space aspect graph.

Scale space is then defined as a multi-dimensional space parameterized by the viewpoint location and the scale value. This means a 4-D (x, y, z, σ) space when the perspective projection viewing model is used, and a 3-D (θ, ϕ, σ) space in the case of orthographic projection. Each visual event surface is now also a function of both viewpoint and scale. Thus, at $\sigma = 0$ the parcellation of the viewpoint space, and so also the aspect graph, is exactly as computed by some known aspect graph algorithm. As σ increases, the parcellation of viewpoint space should deform in a way such that at certain discrete values of scale the aspect graph becomes simpler (has fewer nodes).

There are (at least) two alternative representations of the qualitative structure of scale space. In part (a) of Figure 1, the scale space aspect graph is depicted as an explicit sequence of aspect graphs, each representing a range of σ over which the aspect graph has a constant structure. This representation is more explicit and perhaps simpler conceptually, but there is potentially a great deal of redundancy in the multiple instances of the aspect graph. Part (b) of Figure 1 depicts a more compact representation that is more directly analogous to the typical form of the aspect graph. Each node represents a "volume" of the scale space for which the same general view exists. Each arc again represents a visual event, but the underlying event surface boundary is now parameterized by the scale dimension. In general this graph structure is not planar in nature. The structure of the complete aspect graph at any given scale is not explicitly represented in this second form.

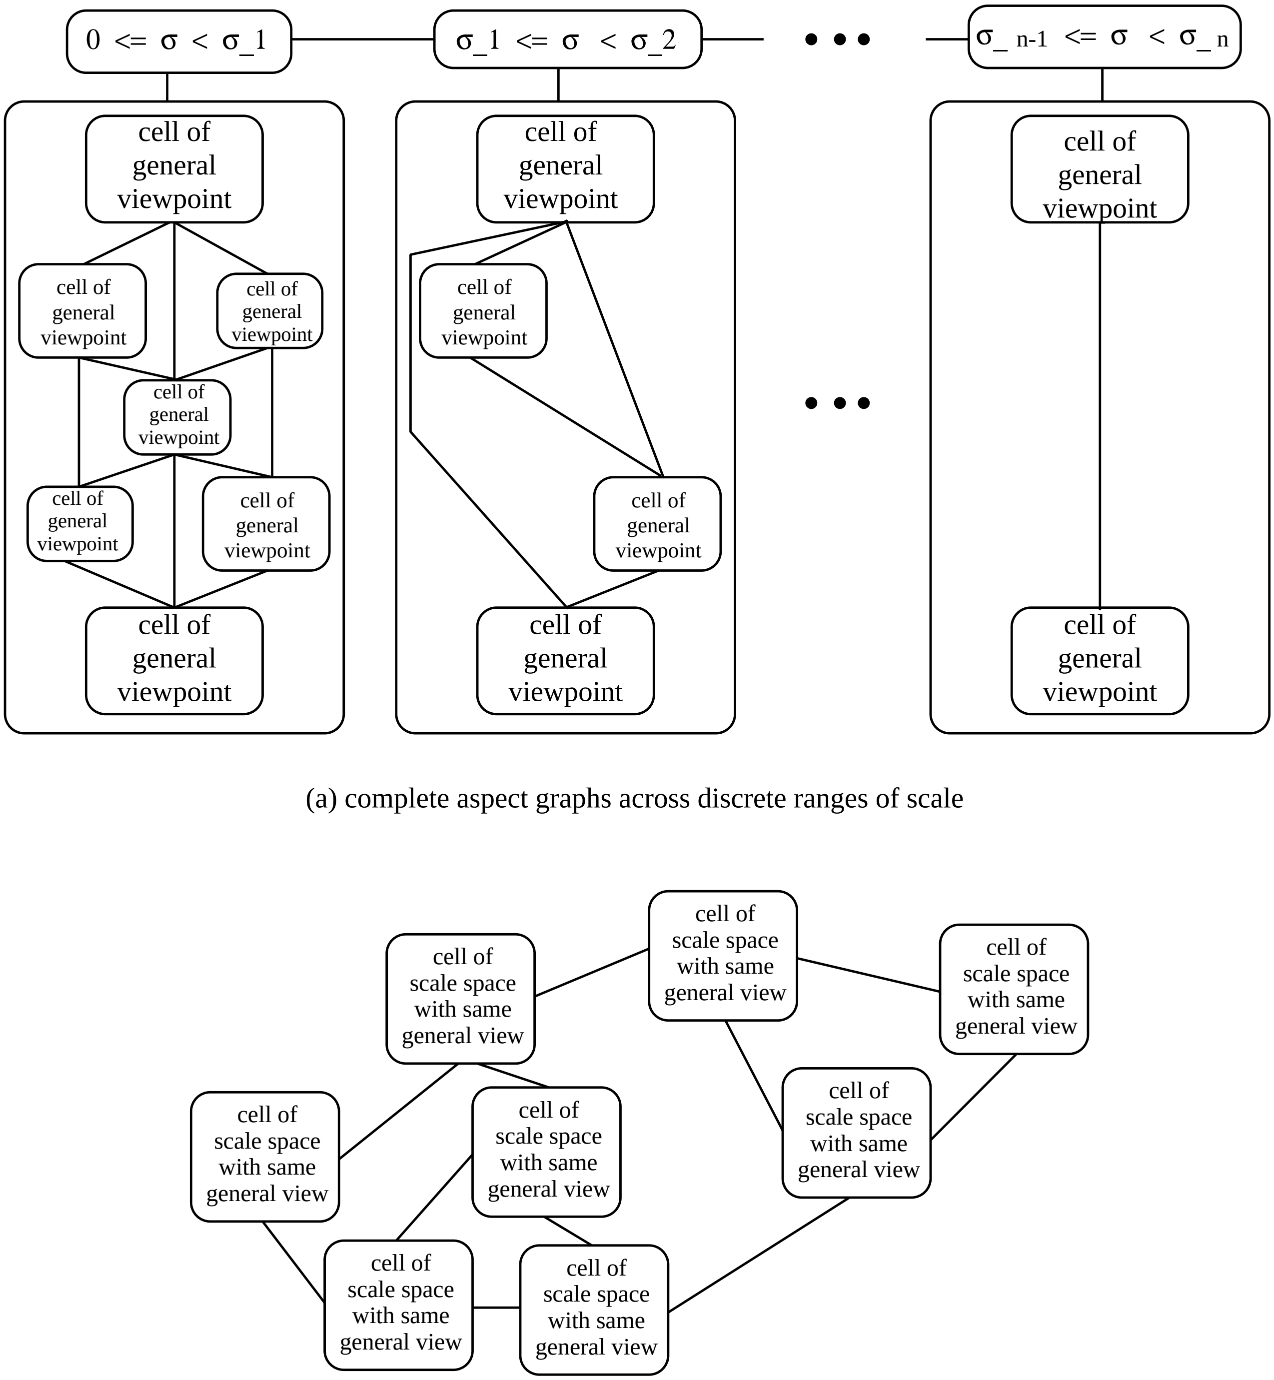

(a) complete aspect graphs across discrete ranges of scale

(b) interconnected cells representing general views in scale space

Figure 1: Conceptual Depictions of the Scale Space Aspect Graph.

However, it should be obvious that one can convert back and forth between the two representations to present whatever information is necessary.

The above representations bear resemblance to other recent forms. The more explicit of the two is analogous to the *visual potential* of Sallam *et. al.* [23]. In their representation separate instances of the aspect graph are recorded for varying articulation parameter values of an object. Here scale can be thought of in a similar manner. The more implicit of the two forms corresponds most closely to the *asp* of Plantinga and Dyer [20]. In their representation the aspect graph was formed as the projection of certain higher-dimensional "volumes", representing particular feature configurations, into the viewpoint space. This is essentially the conversion process used between the two forms given here. Other representations such as extensions of the *interval tree* concept [15, 28] may exist depending on the interpretation of the scale parameter, which is the topic of the next section.

4.1 Interpretations of Scale in the Scale Space Aspect Graph

We have proceeded this far without assigning any particular meaning to the scale dimension and without saying how the scale parameter might be used to create a family of parcellations of the viewpoint space. This is a question which has neither an obvious nor a unique answer. Previous scale space representations have been applied to 1-D, 2-D and 3-D intensity functions by interpreting the scale parameter in terms of the solution to the diffusion equation [12] (or more specifically, as the variance of a Gaussian kernel used to blur the function). It has been proven that only under this interpretation will the qualitative features of the function disappear and not be created as the scale value is increased [12]. However, the entities on which the aspect graph concept is based (such as visual events, projected line drawings, and 3-D shape) are not intensity functions. As such, it is hard to define what one means by "blurring" the parcellation of viewpoint space, and so the requirement that the quantity of features monotonically decrease in size may have to be relaxed.

In order to define a set of interpretations of "blurring" the parcellation, we return to the three areas of concern for the aspect graph structure.

- One interpretation would be to examine the relative sizes of the cells of viewpoint space. This is similar to past approaches which have considered the probability of certain views based on cell size [2, 11, 26, 27]. This interpretation may correspond most intuitively to blurring the parcellation until only the large cells are left. In terms of a physical phenomenon, this idea relates to changing the relative sizes of viewer and object, from say a gnat orbiting about a large building to a Cyclops observing a doll house.

- A second interpretation would examine the relative sizes of the various features in the image. This corresponds to examining the view under varying levels of image resolution. As the pixel size increases, the necessary size of a feature in the image must increase to maintain its detectability. As certain features are lost from view the descriptions of neighboring views may become the same, reducing the overall number of views. It is also possible to incorporate the changes in size due to varying distance into this approach.

- The third interpretation involves examining the relative sizes of features of the object. This corresponds to smoothing the features of the object surface, such as knocking down small bumps and filling in small holes. Continual smoothing should lead to a featureless appearance of the object in the image and thus to fewer visual event surfaces in the parcellation.

It seems obvious that each of the above approaches offers some form of reduction in the size of the aspect graph as a function of scale, and that each corresponds to a slightly different visual phenomenon. Ideally we would like to incorporate all of these effects in one solution, but due to the complexity it is necessary to more fully understand each one individually first. This includes exploring alternative methods that account for the above descriptions, as it is still not precisely clear how scale should parameterize the parcellation under each approach. Thus this paper focuses on what currently appears to us to be the most promising of the three ideas – the use of image resolution as a scale measure. The other possible interpretations of scale are left for future study.

4.2 The Scale of Features in the Projected Line Drawing

In current aspect graph analyses, a projected line drawing is constructed assuming an infinite resolution image plane. Also, the exact dimensions of the lines are lost as the drawing is abstracted to the image structure graph, where all arcs are treated equally. Both of these assumptions ease the computations that are made to calculate visual event boundaries, but limit the practical usefulness

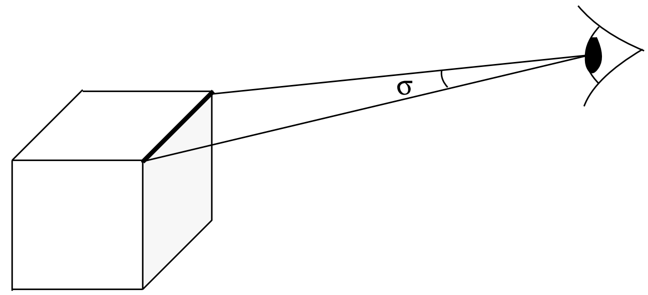

Figure 2: The degree of visual arc occupied by an edge feature.

of the result. Thus it is proposed that the scale of the features in the view be reintroduced as a function of image resolution. Implicit in this method is the fact that the size effects due to viewing distance will also be accounted for. These ideas are similar in nature to those used by researchers determining visibility constraints for automatic sensor placement [6].

First one must determine which features should be measured. In order to be measured a feature must have some spatial extent in the image. This means that a junction, which occurs at a single point, should not be one of the features we now concentrate on. Alternatively, both edges (limbs) and object faces (portions of surface patches) generally have measurable extent in a view. How does one quantify the size of a feature? It is not sufficient to measure the length of an edge or the area of a face on the object. It is the projection of these features which is of interest.

The first solution which comes to mind is then to measure the dimensions of the features in an image coordinate system, the resolution of which is based on our scale parameter. Then the length along a projected edge, the perimeter or area of a face, or possibly the radius of the sphere which circumscribes the feature would be quantified in terms of a number of pixels. Unfortunately, this approach implicitly requires a more detailed camera model. It would be necessary to know such parameters as the focal distance, the image plane size (field of view), the particular viewing direction and the viewing position. While such a sophisticated model would undoubtedly be more realistic, it is unlikely to be described using a single scale parameter. And this is a necessary requirement to make this refinement to the aspect graph concept manageable.

An alternative measurement, used commonly by psychologists and biologists, is the degree of visual arc σ, or field of view, occupied by the feature. (See Figure 2.) Furthermore, if this is combined with the natural perspective viewing model[2] [19] implicitly used by many aspect graph researchers, every feature's size can be described by only a single parameter value in the range $0° - 360°$. Exactly how this value is measured depends on the feature.

For a straight edge, the projected distance between its endpoints will cover a particular visual arc as shown in Figure 2. For a curve, the maximum projected distance between any two points, with respect to the viewpoint of course, indicates its visual extent. For a face, the maximum projected distance between any two points on its boundary curve could describe the overall area occupied by it in the image. For the square face of the cube seen in Figure 2, the distance between opposing corners would define its extent. For a curved-surface patch the boundary curve includes both surface discontinuities and apparent contours. (It is possible to consider additional measurements that describe when a portion of a face is no longer distinguishable as well.)

The model above accounts for the quantitative size of a feature in both an absolute and relative sense. In the absolute sense, the actual effects of viewing distance are considered. Given a particular visual arc threshold, a feature cannot be distinguished after a certain distance away, as one would realistically expect. In a relative sense, a feature can be considered as insignificant if it is much smaller than the neighboring features, even when clearly visible. Again this is a function of viewpoint, as two features may change their relative size as the viewpoint moves. Unfortunately, this model does not directly correspond to Gaussian blurring of the features. Therefore, there may not be a strictly monotonic decrease in the size of the scale space aspect graph as a function of scale. This will become apparent in our later case study.

So how is the above interpretation used? It should be obvious that image resolution can be defined in terms of degree of visual arc. The size of a pixel in the image directly corresponds to the minimum visual arc necessary to distinguish a feature. At a value of $0°$ the camera has infinite resolution. At a value of $360°$ there is only a single pixel in the image and everything projects to it. For a given scale, any feature that maps to a size less than one pixel is considered as not

[2]Under this model points in the scene are projected onto a sphere of infinitesimal radius about the viewpoint rather than onto an image plane at a fixed distance from the focal point, pointed in a particular direction. Each point's image coordinates are defined by the spherical angles of the ray from the focal point to it.

observable. In this way, for a given resolution, two views that were previously different may now appear equivalent.

More exactly, the image resolution has a direct effect on the shape of the visual event boundaries. In the ideal case ($\sigma = 0°$), a visual event that denotes the alignment of some set of features is seen from a ruled surface. Once the scale becomes nonzero, the features do not have to align exactly. As long as the points in question map to within one pixel of each other in the image, they can be considered as aligned. Thus the restriction of a ruled surface no longer holds, and the distortion from the ideal surface is increased as the scale parameter (visual arc required to distinguish the features) grows.

The next general question is how does one go about constructing the scale space aspect graph? Basically there is an approach for each of the two forms mentioned earlier. The easiest conceptually, but the more difficult practically, is to actually construct a parcellation of scale space using visual event boundaries that are functions of position and scale. Since these surfaces are not ruled, the reliable calculation of the intersections between them can be difficult. Alternatively, one can analyze the overall partition in viewpoint space as scale is varied and detect those values at which the structure of the subdivision is altered. This technique lends itself well to sampling methods used to locate these values, as has been done with articulated assemblies [23].

Finally, how does this method lead to a selection of the "most important" aspects. First, if a particular image resolution is chosen, a more accurate depiction of the viewpoint space parcellation is available. This in conjunction with techniques based on relative cell size could order the aspects. Second, in keeping with Witkin's model, a particular aspect's strength could be measured according to the range of visual arc values over which it exists, just as an inflection on a 1-D curve was rated. This approach does not directly take into account the relative cell size of each aspect at the different scale values. However, the use of the scale space "volume" of an aspect as an indication of its importance seems both reasonable and feasible.

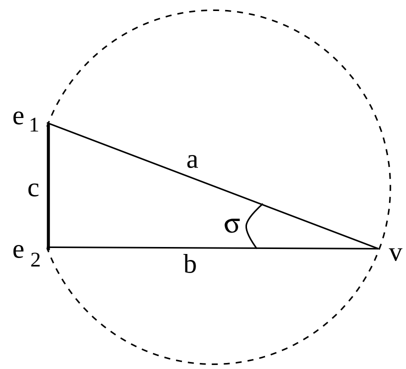

Figure 3: Visual event curve defined for an edge.

4.2.1 Case study : A nonconvex polygon in a plane

We now examine how the above theory can be applied to a concrete example. Following the tradition of the first aspect graph algorithms [10], we begin our study for nonconvex polygons in a plane. In this case there exists a 3-D scale space (x, y, σ). Also, the only feature of interest is an edge. An edge in isolation is considered visible if its two endpoints are distinguishable in the image. The degree of visual arc covered by the projected line segment connecting the endpoints will measure the feature. Thus we study the visibility of a single edge by constructing the equation for the visual event surface in the following manner.

Assume the endpoints of the edge are given by $e_1 = < x_1, y_1 >$ and $e_2 = < x_2, y_2 >$. Then given a viewpoint defined by $v = < x_v, y_v >$, the visibility curve can be defined as the set of viewpoints such that the angle formed by the lines drawn from the viewpoint to each endpoint is constant. This is merely the set of vertex points of a set of triangles formed by the endpoints and the viewpoint, each triangle having the edge as one side and the chosen visual arc angle opposite it. See Figure 3. Using the law of cosines this curve can be expressed by the equation:

$$c^2 - a^2 - b^2 + 2abcos\sigma = 0$$

where the lengths of the sides of the triangle shown in Figure 3 are given by:

$$c = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2} \qquad b = \sqrt{(x_2 - x_v)^2 + (y_2 - y_v)^2} \qquad a = \sqrt{(x_1 - x_v)^2 + (y_1 - y_v)^2}$$

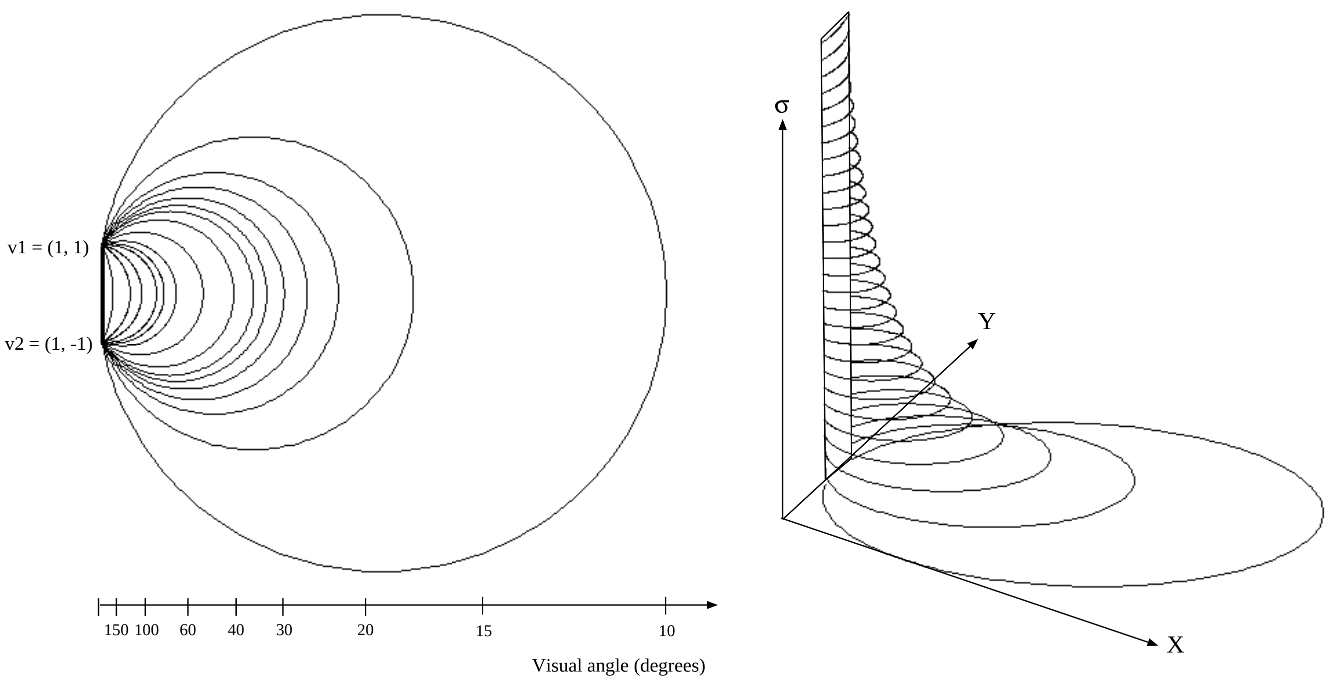

(a) instances of visibility boundary curves in viewpoint space

(b) visibility boundary surface for edge in scale space

Figure 4: Visibility ranges of an edge for varying visual arc angles.

and σ is the visual arc angle. A plot of this curve for varying values of σ for a particular edge is shown in part (a) of Figure 4. The shape of the surface in scale space can be seen in part (b).

The nature of this surface is that for the ideal case of infinite camera resolution ($\sigma = 0°$) the curve in the plane degenerates to the line containing the edge. As aspect graph theory would predict, in this case the edge is visible from the infinite half-space on the proper side of this line. At any finite resolution ($\sigma > 0°$) the curve bounds a finite region of the plane, outside of which the projection of the edge is too small to distinguish. The distortion of the original line into the circular curve is an instance of relaxing the ruled event surface restriction mentioned earlier.

Next we must consider the interaction of a pair of edges, A and B, that are joined at a common vertex as shown in Figure 5. The particular interaction varies slightly based on whether the edges form a convex or concave angle. For a convex-angled pair, a particular visual arc angle will define a visibility region bounded by the above event curve for each edge. If the given visual angle is not large there may be an area of overlap between these two regions. For viewpoints in the overlap

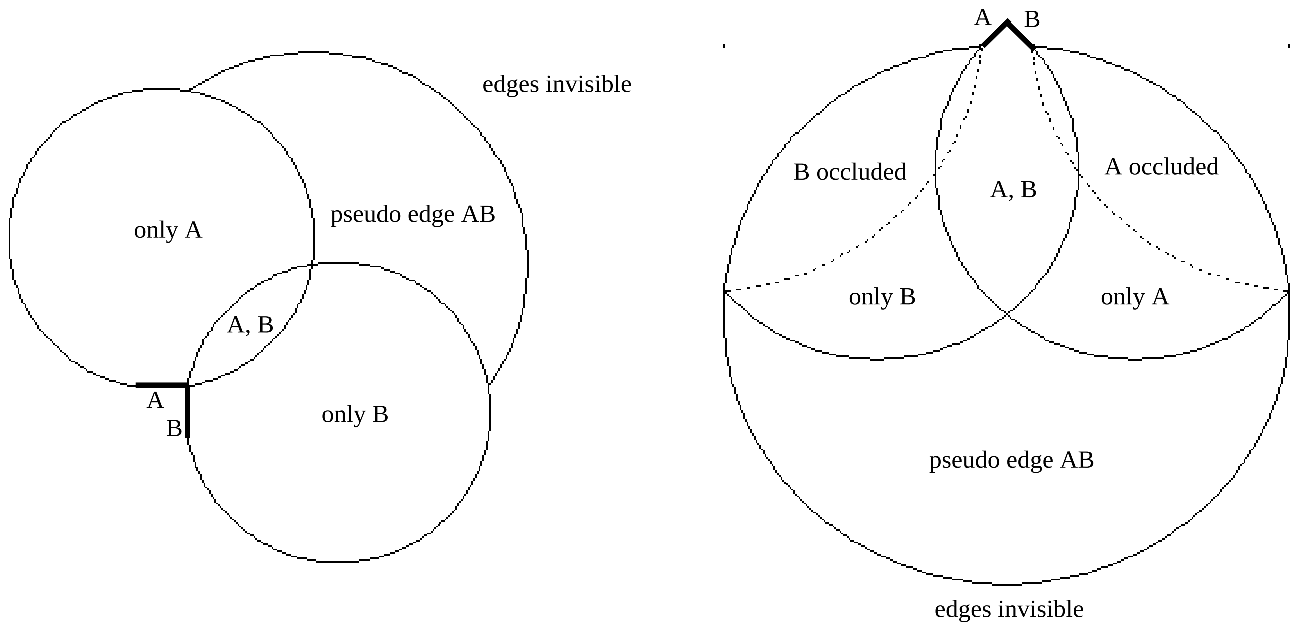

(a) joint visibility of convex-angled edge pair (b) joint visibility of concave-angled edge pair

Figure 5: Joint visibility ranges of edge pairs.

portion both edges are visible and seen joined to one another. The points of intersection between the two curves fall along the line that bisects the angle between the edges.

Now consider the "pseudo edge" formed by chaining together the two edges. This pseudo edge has a wider defining angle than either of the others for certain vantage points, given by the distance between the noncommon endpoints of the two. Thus it can be expected that from certain viewpoints neither of the two edges is individually distinguishable in the image, but the combined length of the pseudo edge is. This pseudo edge is seen as a blending of the two underlying edges, with the vertex that connects them considered as indistinguishable. Of course, even this pseudo edge has a visibility limit. Part (a) of Figure 5 shows the range of the visual event boundaries for each of the edges. The meaningful portion of the curve representing the pseudo edge's visibility ends at the intersections with the curves for each edge. These intersection points occur along the lines containing the edges. The pseudo edge is not considered to exist if either edge is distinguishable. Thus the event curve cannot enter the visibility regions. The intersection points occur where they do because at least some portion of both edges must be visible in order to define the pseudo edge.

The existence of this third event curve is something that the original aspect graph theory does not predict. This is because the assumption of infinite resolution allows every edge to be perceived regardless of distance. The concept of a feature which is a blending or blurring together of other features is unique to the scale space approach.

For a concave-angled edge pair a somewhat similar situation occurs, as given in part (b) of Figure 5. Again there are three event curves derived, which yield regions from which either (1) a single edge is visible, (2) both edges are visible, or (3) the pseudo edge is visible. In this case, the visibility boundary curve for the pseudo edge encompasses the individual curves because it is not possible to see one of the edges alone if you cannot see the pair. Thus the meaningful portions of the individual curves end at the intersection point with the pseudo edge curve. Again the intersection points occur along the line extensions of the edges. Now the additional type of interaction centered about this type of pair involves occlusion. The pseudo event curve, which marks the apparent merging of the noncommon vertices, also marks a limit to the occlusion of one of the edges by another neighboring edge. The visibility limits of the neighboring edges, which are also shown by dotted curves in the figure, form the other boundaries to the regions in which the edges are occluded. The change from occlusion limit to pseudo edge visibility limit occurs at the point of intersection with the individual edge event curve.

When an entire polygon is considered, the set of event curves is derived from the set of pseudo edges. Each pseudo edge can be modeled as coming from a convex or concave corner of the object, no matter how many edges are actually involved. It may be the case that this corner is only implied as the junction of two smaller pseudo edges. An example of this will be seen shortly. From this set of pseudo edges the proper event curves as well as their meaningful ranges are established.

The remaining task is to compute the structure of the scale space aspect graph. As mentioned before, calculating the event surfaces and their intersections in scale space can be difficult. Therefore we examine the structure of the aspect graph as a function of scale and discuss the different types of *scale events* which can alter its structure. These basically fall into three categories:

1. *Begin or end overlap of two curves.* Two curves, either event curves or actual edges, may initially have their meaningful portions touch at a single point as they start to overlap, or their meaningful portions touch at a single point as they cease to overlap. If two event curves are involved, either a new region is created or one is deleted to represent the overlap area. If an event curve and an edge are involved, generally it is only the boundary of an existing cell which is altered.

2. *Triple point.* The meaningful portions of three curves, either event curves or actual edges, intersect at a single point in the parcellation. If three event curves are involved, both before and after the critical scale these curves bound a region. The region that existed before will no longer be in the parcellation and it will be replaced by the new region. If two event curves and an edge interact, either a new region will be created or one will be removed. Because the object exists on one side of the edge, only one of the regions is meaningful. If one event curve and two edges are involved, again a region may be created or removed. In general this creation or deletion coincides with the insertion or removal of a meaningful segment of the event curve.

3. *Curve coincidence.* The meaningful portions of two curves, either event curves or actual edges, coincide along their length in the parcellation. In the case of two event curves, before the critical scale they intersect in the parcellation and are part of the boundaries of a set of regions. At the critical scale the definitions of the two curves are the same, and any regions that existed between them no longer do. Beyond the critical scale the curves again intersect one another and bound another set of new regions. When an event curve and an edge are involved, the region which is between them is deleted from the parcellation. In addition the event curve is removed, since it has passed into the object where it has no meaning.

In addition to the above changes to the parcellation, there is the initial transformation from the ideal parcellation ($\sigma = 0$) to a scale space parcellation. This only happens once. In truth, the ideal parcellation is a degenerate form in which many of the event curves collapse to lines and others are

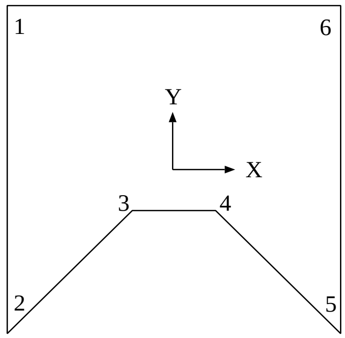

(a) nonconvex polygon

Vertex	Coordinates (x,y)
1	(-1, 1)
2	(-1, -1)
3	(-0.25, -0.25)
4	(0.25, -0.25)
5	(1, -1)
6	(1, 1)

(b) vertex coordinates

Real edges	Vertex endpoints
A	6-1
B	5-6
C	4-5
D	3-4
E	2-3
F	1-2

(c) actual edge definitions

Pseudo edges	Vertex chain	Corner type	Subedges
G	2-3-4	concave	E D
H	3-4-5	concave	D C
I	2-3-4-5	concave	G H, G C, E H
J	1-2-3	convex	F E
K	4-5-6	convex	C B
L	5-6-1	convex	B A
M	6-1-2	convex	A F
N	1-2-3-4	convex	F G
O	1-2-3-4-5	convex	F I
P	3-4-5-6	convex	H B
Q	2-3-4-5-6	convex	I B

(d) pseudo edge definitions

Figure 6: Definition of example nonconvex polygon.

simply not defined. Thus it is either derived separately using a conventional aspect graph algorithm or from the more general form, rather than trying to derive the more general parcellation from the ideal representation.

4.2.2 An example

We now examine how the scale space theory can be applied to an example polygon. Figure 6 defines a six-sided nonconvex polygon. In part (c) a list of the real edges is presented, while in part (d) the list of pseudo edges is defined. Each pseudo edge is defined by the string of actual edges it represents. Also given is its interpretation as a convex or concave corner, and the pair of edges that make up the corner. Notice that pseudo edge I can be considered as a concave corner for three different edge pairs in the concavity.

There are seventeen event curves which serve as region boundaries somewhere in the scale space aspect graph. These are labeled in Figure 7 for a sample parcellation at $\sigma = 5°$. The curves $A - F$ represent visibility limits of the actual edges, while $J - Q$ represent the visibility limits of the pseudo edges across convex corners. The pseudo edges across concave corners, $G - I$, generate their respective visibility limits, as well as six curve segments that mark the occlusion limits of the edges in the concavity by the side edges. Notice that the curve generated by vertices two and five serves

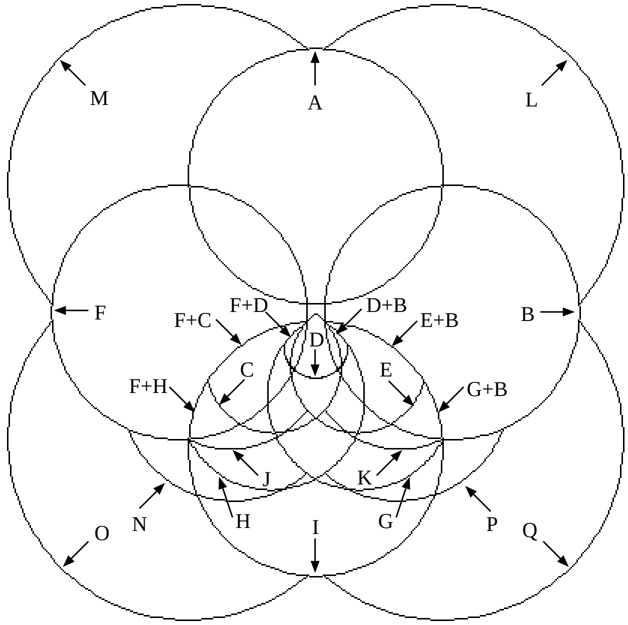

Figure 7: Labeling of meaningful portions of event curves for polygon.

as the termination limit for the majority of the event curves since the elements of the concavity are not visible if the two vertices are seen as one.

If one constructs the parcellation of the 3-D scale space, a total of fifty aspects are determined. These are listed in Table 1 and described according to the range in scale over which they exist and the view that is seen in terms of the edges and the vertices between them. The vertices have been broken down into four categories; T junctions, real vertices on the object, the implied vertex that is the merging of the endpoints of an indistinguishable edge, and an implied vertex in the overlap range of two pseudo edges. The structure of the scale space aspect graph, which corresponds to the parcellation of scale space, is given in Figure 8.

It is perhaps more instructive to examine the second form of the scale space aspect graph, which is the series of distinct parcellations existing over the changes in scale. There are 23 different angles at which scale events occur that modify the parcellation, excluding the initial transformation at $\sigma = 0°$. Each of these events is described in Table 2. At several of the critical angles it is the case that multiple events occur. The representative parcellations at intermediate values of σ are drawn

Table 1: Index of polygon aspects with visual arc existence ranges.

ASPECT	EDGES IN VIEW	VISUAL ANGLE RANGE (degrees)	ASPECT	EDGES IN VIEW	VISUAL ANGLE RANGE (degrees)
1	A	0* - 180	26	F + H	0 - 45
2	B	0* - 180	27	G + B	0 - 45
3	C	45 - 180	28	G • C	0 - 127.98
4	D	36.87 - 180	29	G × H	0 - 216.87
5	E	45 - 180	30	G • K	0 - 45
6	F	0* - 180	31	G × P	0 - 45
7	G	0 - 225	32	I × P	0 - 59.04
8	H	0 - 225	33	J − C	0 - 30.96
9	I	0 - 255.96	34	J • H	0 - 45
10	L	0 - 90	35	N • C	6.1 - 24.75
11	M	0 - 90	36	N × H	0 - 45
12	N	0 - 59.04	37	N × I	0 - 59.04
13	O	0 - 90	38	E − C • B	0 - 22.5
14	P	0 - 59.04	39	E • D + B	0* - 14.04
15	Q	0 - 90	40	E • D • C	0* - 75.4
16	A • F	0* - 45	41	E • H • B	0 - 29.5
17	B • A	0* - 45	42	E • H × K	0 - 30.96
18	D • C	30.96 - 112.5	43	F + D • C	0* - 14.04
19	E + B	0* - 30.96	44	F • E − C	0 - 22.5
20	E − C	0 - 77.65	45	F • G • C	0 - 29.5
21	E • D	30.96 - 112.5	46	G × H × K	0 - 6.1, 24.75 - 45
22	E • H	0 - 127.98	47	J × G • C	0 - 30.96
23	E − K	0 - 30.96	48	J × G × H	0 - 6.1, 24.75 - 45
24	E • P	6.1 - 24.75	49	E • D • C • B	0* - 14.04
25	F + C	0* - 30.96	50	F • E • D • C	0* - 14.04

* indicates aspect exists for visual arc angle of 0°
\+ indicates edges joined by T junction vertex
− indicates edges joined by implied vertex that is collapsed edge
• indicates edges joined by actual vertex on object
× indicates edges joined by implied vertex due to overlap of edges

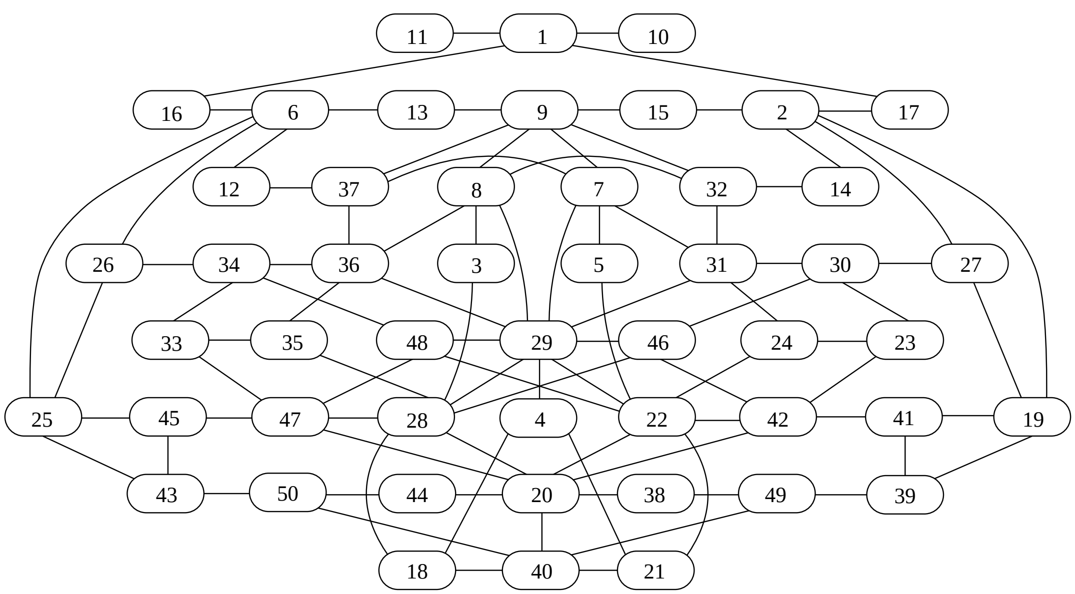

Figure 8: Scale space aspect graph of example polygon.

Table 2: List of changes in parcellation due to varying visual arc limits.

EVENT ANGLE	EVENT DESCRIPTION	INSERTED ASPECTS	REMOVED ASPECTS
0.0°	infinite lines become finite curves (cA - cC, cE - cH), new finite curves (cD, cI - cQ)	7-15, 20, 22, 23, 26-34, 36-38, 41, 42, 44-48	–
6.1°	triple points (cJ, cG, cC), (cK, cH, cE)	35, 24	48, 46
14.04°	triple points (cD, cE, cG), (cD, cE, cF), (cD, cC, cH), (cD, cC, cB) begin overlap (cD, eE), (cD, eC)	– 20^{+}, 40^{+}	43, 50, 39, 49 –
22.5°	end overlap (cE, cF), (cC, cB)	–	44, 38
24.75°	triple points (cJ, cG, cC), (cK, cH, cE)	48, 46	35, 24
29.5°	end overlap (cF, cG), (cB, cH)	–	45, 41
30.96°	triple points (cC, cF, cI), (cC, cF, cG), (cE, cB, cI), (cE, cB, cH) triple points (cC, eD, eE), (cE, eC, eD) coincide (cC, cE)	– 21, 18 22, 28	25, 33, 19, 23 – 22, 28, 42, 47
35.26°	triple points (cC, cD, eE), (cE, cD, eC)	22^{+}, 28^{+}	–
36.87°	triple point (cC, cE, eD) end overlap (cC, eE), (cE, eC)	4 –	– 22^{+}, 28^{+}
45.0°	triple points (cH, cE, eE), (cG, cC, eC) end overlap (cD, eE), (cD, eC), (cE, eD), (cC, eD) end overlap (cA, cF), (cA, cB), (cF, cI), (cB, cI) triple point (cJ, cF, cG), (cJ, cE, cG), (cJ, eE, eF) triple point (cK, cB, cH), (cK, cC, cH), (cK, eB, eC) coincide (cG, cH)	5, 3 – – – – 7, 8	– 21^{+}, 22^{+}, 18^{+}, 28^{+}, 4^{+} 16, 17, 26, 27 34,48 30, 46 7, 8, 31, 36
59.04°	triple point (cN, cF, cI), (cN, cG, cI), (cN, eF, eG) triple point (cP, cB, cI), (cP, cH, cI), (cP, eB, eH)	– –	12, 37 14, 32
75.4°	triple point (cC, cD, cE)	29^{+}	40
75.96°	end overlap (cH, eE), (cG, eC)	–	5^{+}, 22^{+}, 3^{+}, 28^{+}
77.65°	end overlap (cC, cE)	–	20, 29^{+}
90.0°	triple points (cL, cA, cB), (cL, eA, eB), (cM, cA, cF), (cM, eA, eF) triple points (cO, cF, cI), (cO, eF, eI), (cQ, cB, cI), (cQ, eB, eI)	– –	10, 11 13, 15
112.5°	end overlap (cD, cE), (cD, cC)	–	21, 18
127.98°	end overlap (cE, cH), (cC, cG)	–	22, 28
180.0°	coincide (cA, eA), (cB, eB), (cC, eC), (cD, eD), (cE, eE), (cF, eF)	–	1, 2, 3, 4, 5, 6
194.04°	begin overlap (cG, eE), (cH, eC)	7^{+}, 8^{+}, 9^{+}	–
210.96°	begin overlap (cG, eD), (cH, eD)	7^{+}, 8^{+}	–
216.87°	triple point (cG, cH, eD)	–	29
225.0°	begin overlap (cI, eE), (cI, eC) triple points (cG, eD, eE), (cH, eD, eC)	9^{+} –	– 7, 8
253.74°	begin overlap (cI, eD)	9^{+}	–
255.96°	triple points (cI, eD, eE), (cI, eD, eC)	–	9

+ indicates existence of aspect not changed, but boundary description is altered

cX indicates event curve X, eX indicates actual edge X

and labeled in Figure 9. We now discuss some of the scale events that occur for this object.

- *triple point between three event curves.* At angles 6.1° and 24.75° typical examples of this event occur, with one cell being created and another deleted. What is interesting about these cases is that they are exact opposites of one another. Cells 24 and 35 exist over the finite scale range between the events, while cells 46 and 48 exist before and after. The scale event angles were found using a search over the scale angle to find the value at which the intersection points between curves converged.

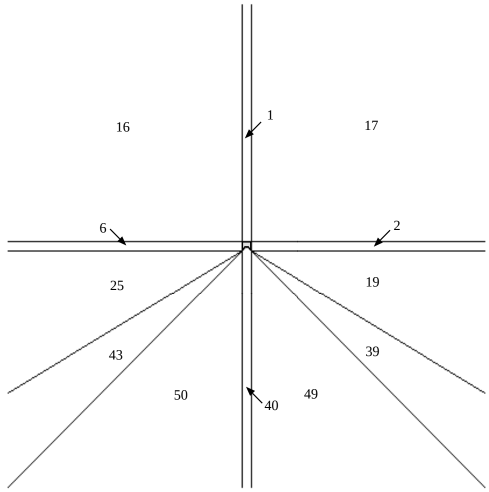

(a) visual angle = 0.0

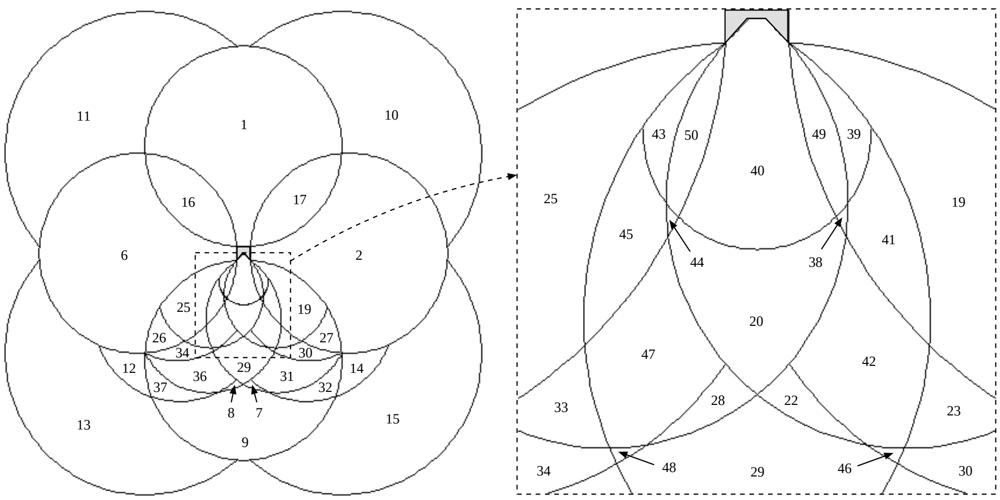

(b) visual angle = 4.0

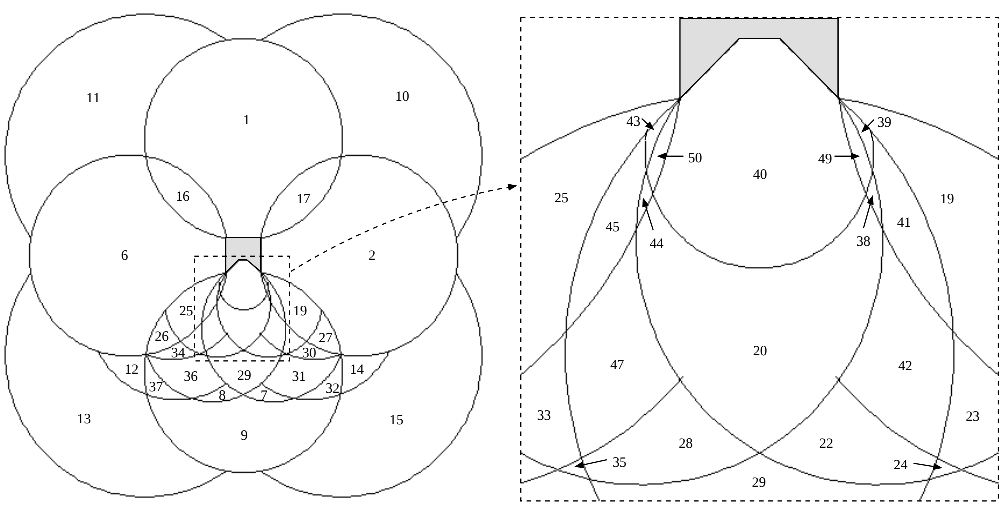

(c) visual angle = 10.0

Figure 9: Representative aspect graphs over changes in scale.

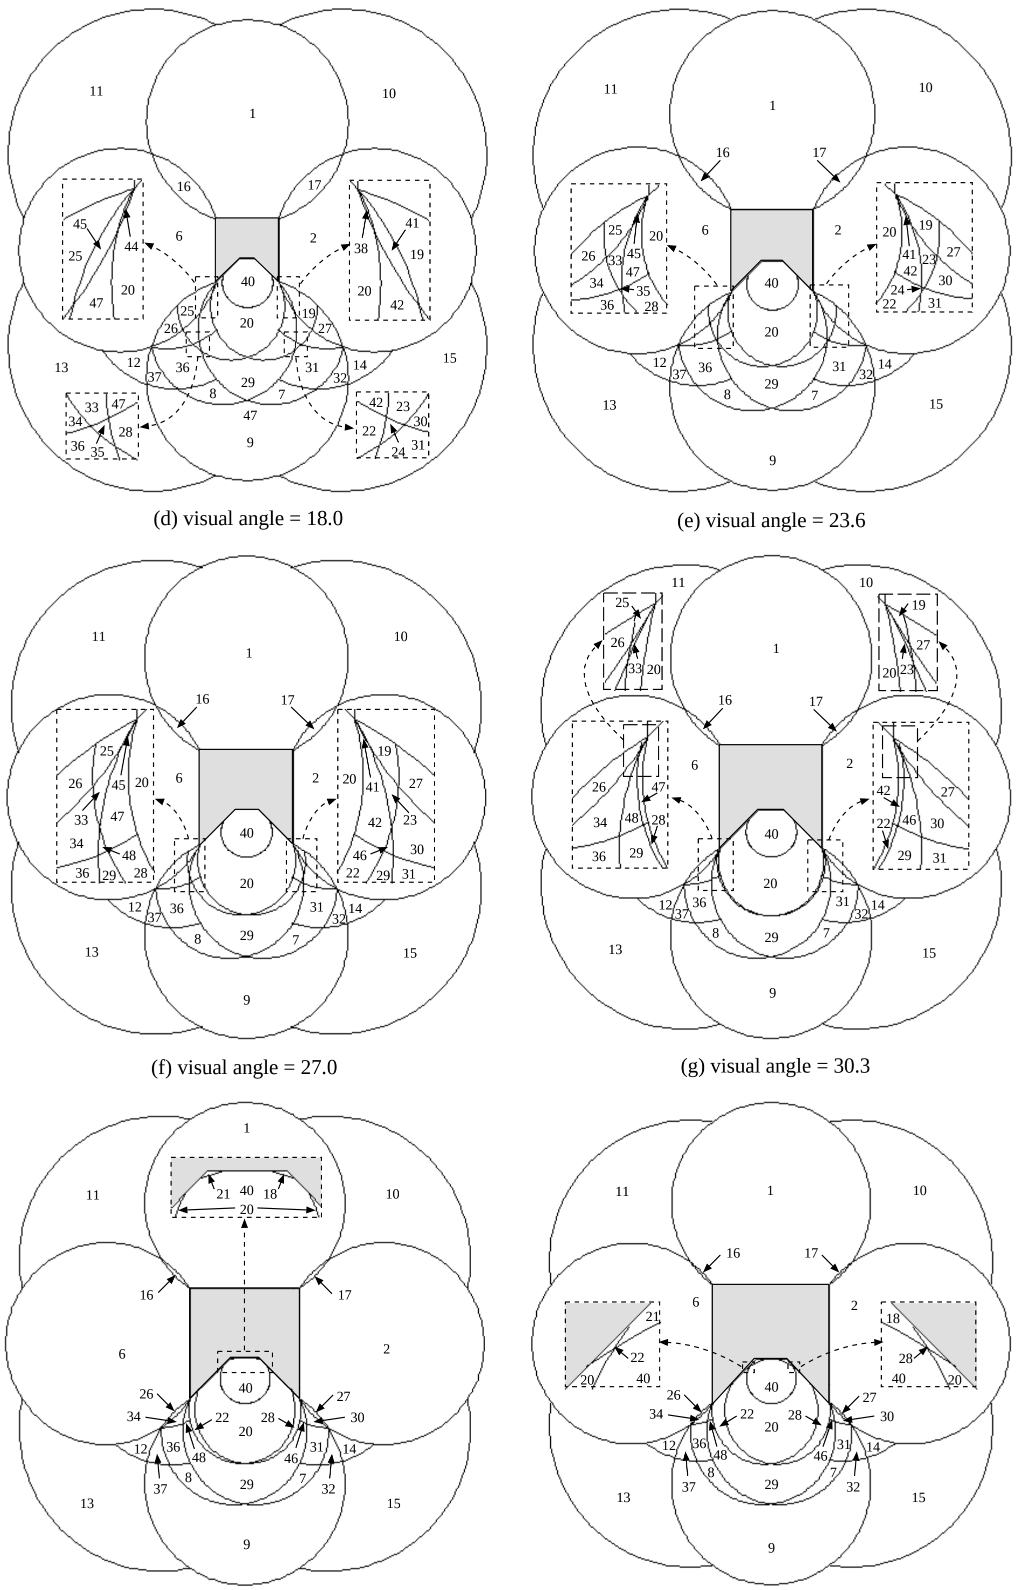

(d) visual angle = 18.0

(e) visual angle = 23.6

(f) visual angle = 27.0

(g) visual angle = 30.3

(h) viewing angle = 33.3

(i) viewing angle = 36.0

Figure 9: continued.

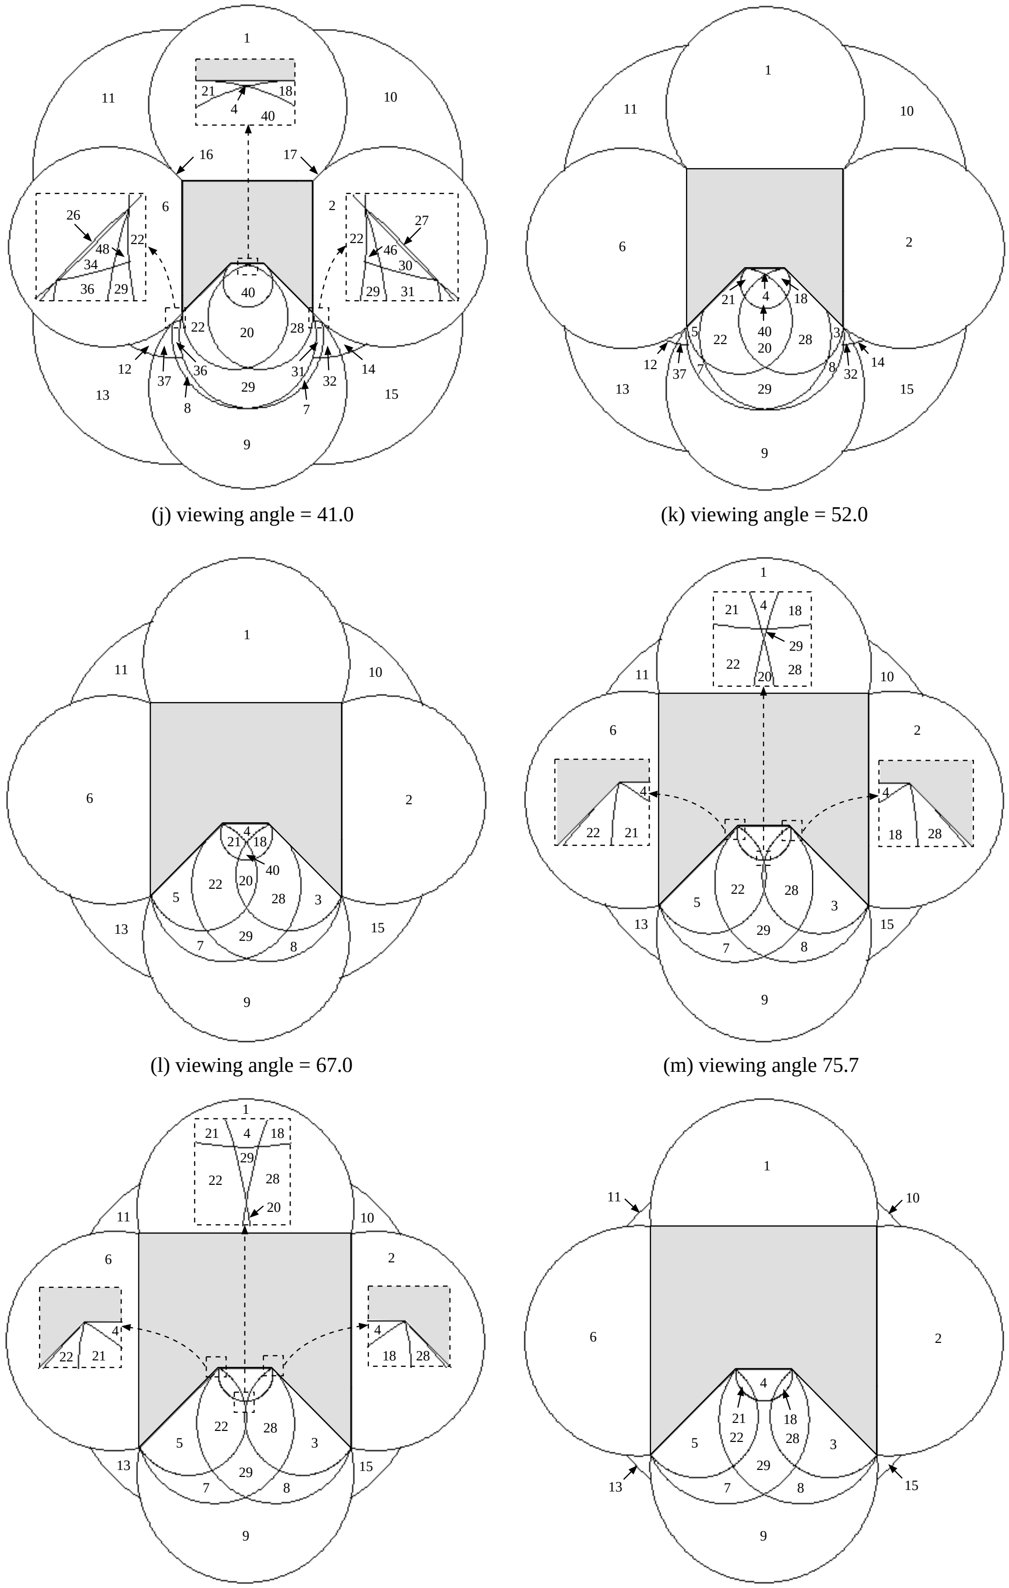

(j) viewing angle = 41.0

(k) viewing angle = 52.0

(l) viewing angle = 67.0

(m) viewing angle 75.7

(n) viewing angle = 76.8

(o) viewing angle = 84.0

Figure 9: continued.

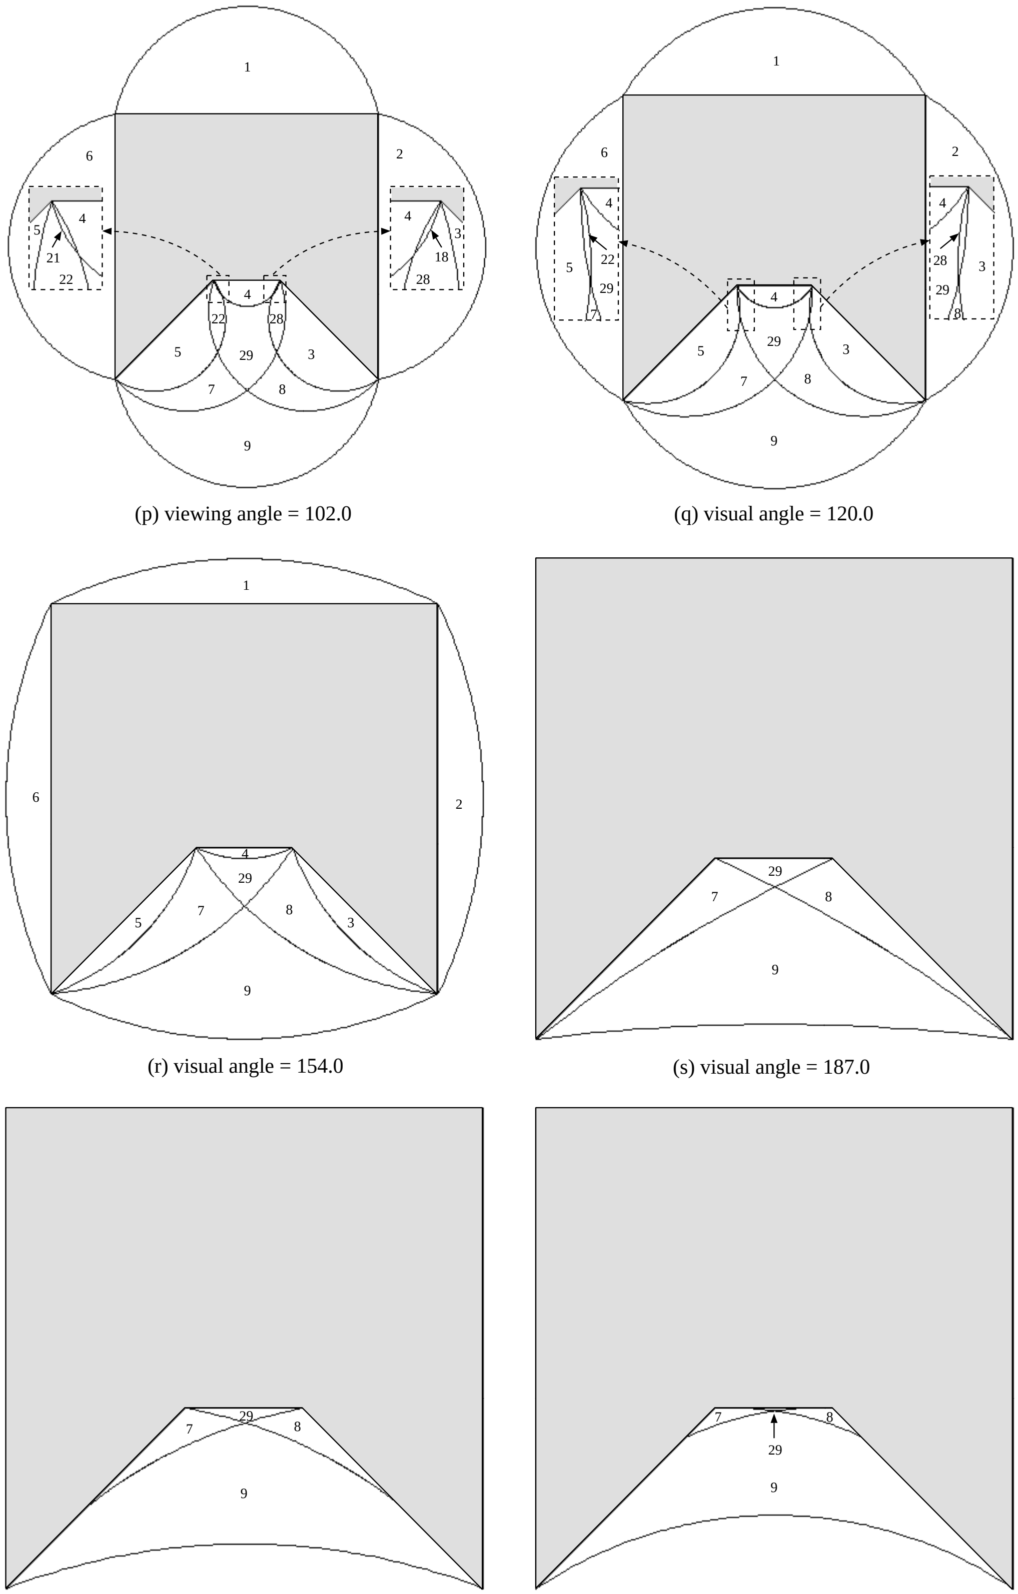

(p) viewing angle = 102.0

(q) visual angle = 120.0

(r) visual angle = 154.0

(s) visual angle = 187.0

(t) visual angle = 201.0

(u) visual angle = 214.0

Figure 9: continued.

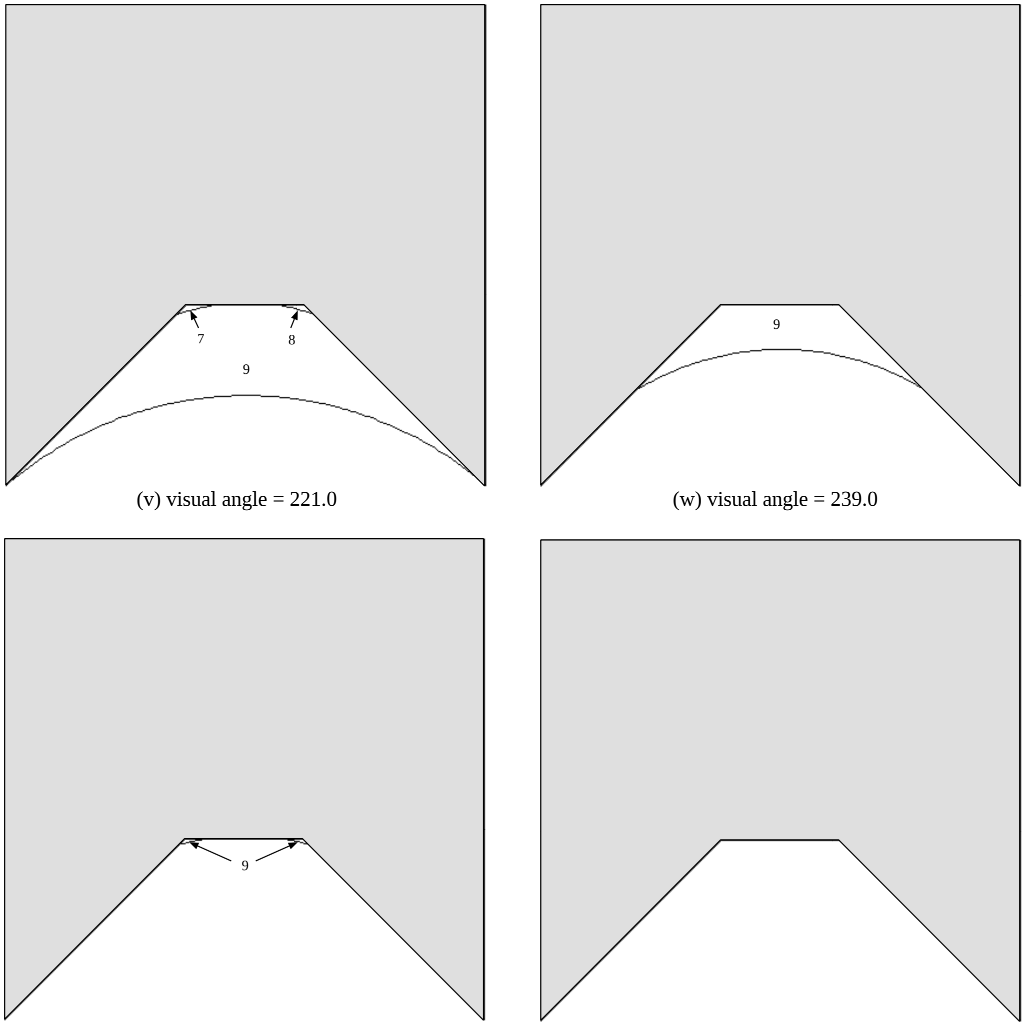

(v) visual angle = 221.0
(w) visual angle = 239.0
(x) visual angle = 254.8
(y) visual angle = 260.0

Figure 9: continued.

- *triple point between two event curves and one edge.* The types of events for this group vary. At angle 36.87° cell 4 can be seen to come into existence, while at 216.87° the opposite occurs as cell 29 disappears. At 35.26° an interesting event occurs in which a new cell in the parcellation appears (cell 22), but it is merely another area in which that aspect can be seen. These two areas are later merged as the result of another event. The angles at which these events occur can again be calculated by solving for the point on the edge at which the two visible edges are seen to be the same size.

- *triple point between one event curve and two edges.* The emergence of a cell can occur when a curve passes over a vertex and emerges from the object, as for cell 21 at 30.96°. The angle at which this occurs is found by solving for σ after substituting the vertex position into the event curve equation. It is also possible for cells to disappear as event curves are absorbed into the object after passing over the vertex. In these cases the event curves that represent a pseudo edge about that vertex no longer exist. For a convex corner this coincides with a triple point of three event curves, as occurs at 90° for pseudo edge AF. For a concave corner there is only the curve and edges involved, as happens for pseudo edge G at 225°. These events occur at angles equal to that between the participating edges.

- *end overlap of two event curves.* This event marks the end of existence of a cell from which multiple edges are seen. It may be the case that when overlap ends there is still a common point between the curves, for instance vertex 1 when cell 16 disappears at 45°, or they may be completely separated, as when cell 20 collapses at 77.65°. In the former case the event angle corresponds to the half angle between the edges at the vertex, while in the latter one must search for the final common point to calculate the angle.

- *end overlap of event curve and edge.* This event alters the boundary of a cell without changing its existence. Again there may be a single common point left after contact, as for curve D and edge E at 45°, or not, as when curve C leaves edge E at 36.87°. This latter event merges the two cells corresponding to aspect 22. The event angle in the first case is found as the angle between the intersecting edge and the edge corresponding to the curve, less 180° since this occurs at a concave corner. In the second case a search is done to find the final point of contact on the edge to calculate the angle.

- *begin overlap of event curve and edge.* In this case the boundary of a cell is again altered without changing the existence of the cell, in essence performing the inverse operation of ending overlap. Before overlapping, the edge and curve may have a common point at a vertex or not, and if not may begin intersecting at an endpoint of the edge or in the middle. Instances of these are given for curve G and edge E at 194.04°, curve D and edge E at 14.04°, and curve I and edge D at 253.74°, respectively. The visual angles at which the events occur can be calculated in the same manners as the ending of overlap.

- *coincidence of two event curves.* There are two occurrences of event curve coincidence, curves C and E at 30.96° and curves G and H at 45.0°. In each instance four cells are removed and two inserted into the parcellation. What is interesting is that two of these cells correspond to the same aspect. For instance the aspect which sees only pseudo edge G exists between the curves G and H. Before coincidence this area is in the right half of the parcellation, while after coincidence it is in the left half. The angle at which this event occurs is found by solving the event curve equation with the location of the other curve's defining point inserted.

- *coincidence of event curve and edge.* For this event the curves merge with their respective edges, and both the curve and cell between curve and edge no longer exist. This occurs at an angle of 180° for all six real edges and their curves.

One phenomenon noticed in the series of parcellations, that does not occur for an ideal aspect graph, is that there is no longer a 1-1 correspondence between aspects and cells. This is because the cells in scale space are not "convex" and a given projection into viewpoint space may result in separate cells. It is also the case that the cell in scale space representing a single aspect need not be a single volume as we think of it. For instance, there are two distinct cells for the view of pseudo edge G, which only touch at a single point in the scale space. Thus it is necessary to loosen the restriction that the points in the space which correspond to an aspect be from a maximally connected region.

As a final conclusion, one might wonder what the "important" aspects are for this polygon. If one uses range of scale as a measure, then aspect 9, which sees the view of the concavity, is the strongest. A complete volumetric analysis in scale space has not been performed for this polygon. However, from examining the representative parcellations in Figure 9 it seems that those aspects corresponding to views of edges A, B and F as well as the pseudo edges I, L, M, O and Q occupy the larger areas in each parcellation. The majority of the other views exist in near the object's concavity. It is interesting to note these views are not among the larger sized cells of the ideal parcellation in part (a) of Figure 9.

5 Discussion

We have presented the idea of incorporating a quantitative measure of scale into the qualitative representation of the aspect graph. This was done in an effort to address the limitations recently expressed about the aspect graph. We began by defining what a scale space aspect graph should be. Then several alternative interpretations for scale were proposed, with the interpretation of scale as feature size in the image explored most fully. We now need to examine how closely this interpretation follows our original definition.

The main criterion of the scale space aspect graph was that it should be monotonically reduced in complexity for changes in scale. This is not the case for the feature size interpretation. Most disconcerting is the large increase in aspects as a nonzero value of scale is assumed. However, it does not seem so bad when one realizes that the ideal case (which could never happen in practice) is merely a highly degenerate instance of the more general parcellation. Even barring this, there are still scale events in which new aspects are generated without the accompanying removal of aspects (for instance cell 4 at 36.87°). But it is also true that the overall size of the parcellation will eventually reach zero as the scale parameter approaches 360°. Thus while there may be local increases in the parcellation complexity, the global trend is a decrease in size.

Regardless of the above goal, it was also desired to account for certain real visual phenomena in the aspect graph representation in order to get an indication of the practically "important" aspects. It is definitely the case that the scale space aspect graph is more realistic. If one has a camera with a given resolution, a much truer description of the parcellation is available. In terms of important aspects, examining the various parcellations in Figure 9 leads one to believe that some use of the "volume" of the aspect's cell in scale space would be the best indicator of its importance. Exactly how this volume is calculated can be viewed as another interpretation of the scale parameter. It was also apparent that those aspects which are important when measured in this way, generally are not those that would be labeled as important in the ideal form of the aspect graph.

Several possible scale interpretations have been mentioned here, and others can be envisioned. Exploring each of these requires specifically defining the organization of scale space and how the scale interpretation will be measured. Such alternatives were presented as we reached our final measurement of scale as the visual arc spanned by a feature.

The most direct extension of this work would be to generalize the theory presented here for polygons to broader classes of objects. We are currently considering extensions for polyhedra. To see how this extension begins, think of the visibility of, say, a triangular face of an object. In order for the face to project some visible area, those two points on the face boundary which are farthest apart should be distinguishable, otherwise the face is observed to collapse and project only as a line. These two points could either be the endpoints of an edge, or a vertex and some point on the edge not containing it.

The first of these cases generates an event surface that is formed by rotating the event curves described in 2-D about the edge in 3-D space. The other event surface is more complex, composed of the curves marking the visibility of each the quasi-edges formed between points along the edge and the vertex. The complete visibility boundary for the triangular face is formed by the portions of these six event surfaces which are nearest the face in every direction. Using this approach, an extension to convex polyhedra can be easily made. However, developing similar surfaces for nonconvex polyhedra requires considering the concepts of occlusion and the visibility of nonconvex polygons, both in whole and in part.

References

[1] Babaud, J., Witkin, A. P., Baudin, M. and Duda, R. O. 1986. "Uniqueness of the Gaussian kernel for scale-space filtering", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, pp. 26-33, 1986.

[2] Ben-Arie, J. "Probabilistic models of observed features and aspects with application to weighted aspect graphs", *Pattern Recognition Letters*, vol. 11, pp. 421-427, 1990.

[3] Bowyer, K. W. and Dyer, C. R. "Aspect graphs: An introduction and survey of recent results", *International Journal of Imaging Systems and Technology*, vol. 2, pp. 315-328, 1990.

[4] Chen, S. and Freeman, H. "On the characteristic views of quadric-surfaced solids", *Proceedings of the IEEE Workshop on Directions in Automated CAD-Based Vision*, pp. 34-43, 1991.

[5] Clark, J. J. "Singularity theory and phantom edges in scale space", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, pp. 720-727, 1987.

[6] Cowan, C. K. "Automatic camera and light-source placement using CAD models", *Proceedings of the IEEE Workshop on Directions in Automated CAD-Based Vision*, pp. 22-31, 1991.

[7] Eggert, D. and Bowyer, K. "Perspective projection aspect graphs of solids of revolution: An implementation", *Proceedings of the IEEE Workshop on Directions in Automated CAD-Based Vision*, pp. 44-53, 1991.

[8] Faugeras, O., Mundy, J., Ahuja, N., Dyer, C., Pentland, A., Jain, R., Ikeuchi, K. and Bowyer, K. "Panel theme: Why aspect graphs are not (yet) practical for computer vision", *Proceedings of the IEEE Workshop on Directions in Automated CAD-Based Vision*, pp. 98-104, 1991.

[9] Gigus, Z. and Malik, J. "Efficiently computing and representing aspect graphs of polyhedral objects", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, pp. 542-551, 1991.

[10] Gualtieri, J. A., Baugher, S. and Werman, M. "The visual potential: One convex polygon", *Computer Vision, Graphics, and Image Processing*, vol. 46, pp. 96-130, 1989.

[11] Kender, J. R. and Freudenstein, D. G. "What is a 'degenerate' view?", *Proceedings of the ARPA Image Understanding Workshop*, pp. 589-598, 1987.

[12] Koenderink, J. J. "The structure of images", *Biological Cybernetics*, vol. 50, pp. 363-370, 1984.

[13] Koenderink, J. J. Solid Shape, (MIT Press, Cambridge, Mass.), 1990.

[14] Kriegman, D. and Ponce, J. "Computing exact aspect graphs of curved objects: Solids of revolution", *International Journal of Computer Vision*, vol. 5, pp. 119-135, 1990.

[15] Lindeberg, T. and Eklundh, J. "Scale detection and region extraction from a scale-space primal sketch", *Proceedings of the 3rd International Conference on Computer Vision*, pp. 416-426, 1990.

[16] Malik, J. "Interpreting line drawings of curved objects", *International Journal of Computer Vision*, vol. 1, pp. 73-103, 1987.

[17] Mokhtarian, F. and Mackworth, A. K. "Scale-based description and recognition of planar curves and two-dimensional shapes", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, pp. 34-43, 1986.

[18] Mokhtarian, F. "Multi-scale description of space curves and three-dimensional objects", *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 298-303, 1988.

[19] Nalwa, V. S. "Line-drawing interpretation: A mathematical framework", *International Journal of Computer Vision*, vol. 2, pp. 103-124, 1988.

[20] Plantinga, H. and Dyer, C. R. "Visibility, occlusion and the aspect graph", *International Journal of Computer Vision*, vol. 5, pp. 137-160, 1990.

[21] Ponce, J. and Kriegman, D. "Computing exact aspect graphs of curved objects: Parametric surfaces", *Proceedings of the 8th National Conference on Artificial Intelligence*, pp. 340-350, 1987.

[22] Rieger, J. "The geometry of view space of opaque objects bounded by smooth surfaces", *Artificial Intelligence*, vol. 44, pp. 1-40, 1990.

[23] Sallam, M. Y., Stewman, J. S. and Bowyer, K. W. "Computing the visual potential of articulated assemblies", *Proceedings of the 3rd International Conference on Computer Vision*, pp. 636-643, 1990.

[24] Sripradisvarakul, T. and Jain, R. "Generating aspect graphs for curved objects", *Proceedings of the IEEE Workshop on Interpretation of 3D Scenes*, pp. 109-115, 1989.

[25] Stewman, J. H. and Bowyer, K. W. "Creating the perspective projection aspect graph of convex polyhedra", *Proceedings of the 2nd International Conference on Computer Vision*, pp. 494-500, 1988.

[26] Wang, R. and Freeman, H. "Object recognition based on characteristic view classes", *Proceedings of the 10th International Conference on Pattern Recognition*, pp. 8-12, 1990.

[27] Watts, N. "Calculating the principal views of a polyhedron", *Proceedings of the 9th International Conference on Pattern Recognition*, pp. 316-322, 1988.

[28] Witkin, A. P. "Scale-space filtering", in From Pixels to Predicates, (Ablex Publishing Corp., Norwood, NJ), pp. 5-19, 1986.

[29] Yuille A. and Poggio, T. "Scaling theorems for zero crossings", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, pp. 15-25, 1986.