A Shape Representation Based on Geometric Topology: Bumps, Gaussian Curvature, and the Topological Zodiac

A. Peter Blicher
IBM T.J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598

Abstract

We develop a discrete representation for smooth objects embedded in 3-space, which describes the nesting of bumps, depressions, saddles, and related features within each other. The representation is *intrinsic* and *stable* under perturbation of the surface shape and embedding. The final structure is a graph of level set graphs. Each level set graph represents a constant topology of level sets for a region of the *Gaussian projective plane* (obtained from the Gaussian sphere by identifying antipodal points) bounded by the images of parabolic curves, corresponding to a range of choices of height orientation. The graph of these graphs has the topology of the adjacency graph of the above Tegions on the Gaussian projective plane. We show what the topology of the graphs is, and specify what their bifurcations are. Subparts correspond to subgraphs or collapsed graphs in a simple way. Scale space transformations and smoothing correspond to simple bifurcations of the graph structure.

I. Introduction

Shape representation is crucial in image understanding, both for object recognition as well as signal representation. *Generalized cylinders* [e.g., Bin ford 1971, Agin and Binford 1973, Nevatia and Binford 1977, Brooks 1981] are one well-known representation with a notable history. More recently, there has been interest in using differential geometry in shape description [e.g., Haralick, Watson, Laffey 1983, Horn 1983, Nackman 1984, Brady et al. 1985, Besl and Jain 1985], as well as fractals [Pentland 1985].

Our present subject is a representation for shape based on tome ideas from differential topology, which turns out to be related to all 3 of the approaches mentioned above. The shapes we are interested in are smooth 2-dimensional surfaces embedded in 3-space, a common class for vision.

My starting point for this is that I want to represent an object as a bunch of bumps. The representation should have invariance properties: it should be independent of position or orientation of the object, i.e. it should be *intrinsic.* It should also have *stability* properties, i.e. it should be constant under small perturbations of most anything involved in generating or matching to the representation, like noise, viewpoint, and, less obviously, perturbations of the shape itself. And it should degrade well under partial information; in particular it should commute with occlusion: the representation of a subpart should be a subpart of the representation.

Anyway, the goals of invariance and stability, like spring, make a young man's thoughts turn to topology and geometry, so that is where we look for the tools of our representation.

Since space is limited here, many details, technical and otherwise, including how to make all this rigorous, are omitted. For a more complete description, consult [Blicher 1987].

II. Canonical graph structures

A. The level set graph for a height function

For a surface defined as a real function on the plane, the natural seroth order structure that is invariant under coordinate changes of the plane is the topology of level sets, i.e. the topology of contour lines. This generically leads to a binary tree structure where the leaves are extrema, and the branch nodes are saddles [Koenderink and van Doom 1979; Blicher 1983, 1984]. The effects of a scale spectrum can be expressed in the bifurcations of this tree's topology as the function changes. The effect of a generic smoothing, e.g. convolving with a gaussian, is to cause a sequence of such bifurcations in the tree structure. We will extend this structure to an arbitrary surface.

First, take some surface embedded in 3-spare, and choose a 2-direction (Figure 1). Let's look at the level sets, i.e. the intersections with horisontal planes. We can think of the surface embedding as a function which assigns to each point of the surface its *z* value. Then the level set structure will be very much like that of the function on the plane. The critical points (where the horisontal planes are tangent to the surface) are either extrema or saddles, and the component containing a regular point is a circle, that of an extremum is a point, and that of a saddle is a figure-8. So if we shrink each connected component of a level set to a single point, we can see that as the horisontal plane moves vertically, these points trace out a graph with the branch nodes at saddles, and the number of cycles of the graph is the number of holes in the surface. The gfaph can be thought of as the spine of a generalised cylinder representation of the surface. It has all the nice properties of the tree structure for functions on the plane, as well it should since the latter comes from specialising Morse theory for arbitrary functions to ones on the plane.

Figure 1: The canonical graph for a given z-direction.

B. What happens as the x-direction changes

We were able to describe our shape in terms of level sets by choosing a particular direction for the x coordinate as special Alas, if we were to choose some random other x-direction, we probably would get a different batch of critical points and a different graph structure.

The "obvious" mathematical solution to the choice of in-direction is to consider *all* x-directions, and *all* possible graph structures—then things won't depend on choice of coordinates, since we use them all! The space of all x-directions is ... the Gaussian sphere. So now, every choice of x-direction, hence every point on the Gaussian sphere, has a graph associated with it that tells us the level set structure of the surface for that particular x-axis.

Let's consider stability. Look at what happens at an atomic bump (Figure 2). Looked at with a slightly different vertical, it*s still a bump. In particular, we can't get rid of the critical point (the tangency with the plane). It may move (and will), but it still exists somewhere, if the change in orientation is small enough, and if we started at a typical orientation. This tells us that things are not all that bad: for typical (all but measure sero) points on the Gaussian sphere, the graph structure is constant in some neighborhood of the point.

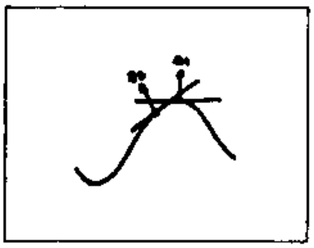

Figure 2: Stability of bumps.

When does it change? Consider what happens as we keep rotating the x-direction (Figure 3). As the extremum from our bump moves to the left, a saddle point moves toward it to the right, inexorably. Once they meet, poof! No more critical point, and our graph has changed. This is a saddle-node bifurcation, since a saddle and an extremum have annihilated each other. Running our rotation backwards would result in the birth of a saddle-node pair. Where do they meet? At an *inflection,* i.e. at a point where the curvature is sero. For a surface, this is where the Gaussian curvature is sero. Such points are called *parabolic points,* and they generically group to form *parabolic curvet* (which are not parabolas!). That means that if we rotate the x-direction somewhat differently, we can't get rid of the parabolic point-well Just hit the parabolic curve in a slightly different place.

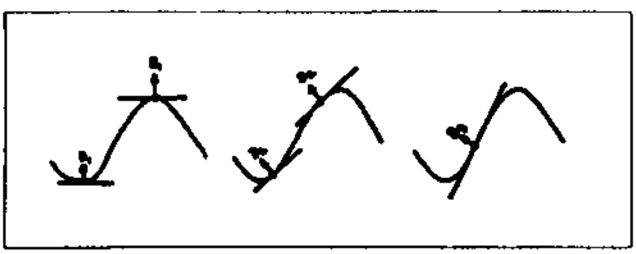

Figure S: x-directions where the canonical graph changes.

The other type of bifurcation, a saddle connection, can also occur. This will happen when changing the x-direction results in a saddle that was higher than another becoming lower with respect to the changing equal-height planes. From now on, we will only describe the part of the structure that comes from the saddle-node bifurcations.

C. Parabolic curves and the Gaussian curvature graph

For a compact surface, the parabolic curves—the points of tangency where the saddle-node bifurcations happen—generically are closed curves that don't intersect [Koenderink and van Doom I960]. Now both the surface and the Gaussian sphere are divided into regions by the parabolic curves. On the Gaussian sphere, the parabolic curves (or really their images under the Gauss map) are the x-directions where the level set graph bifurcates (changes its

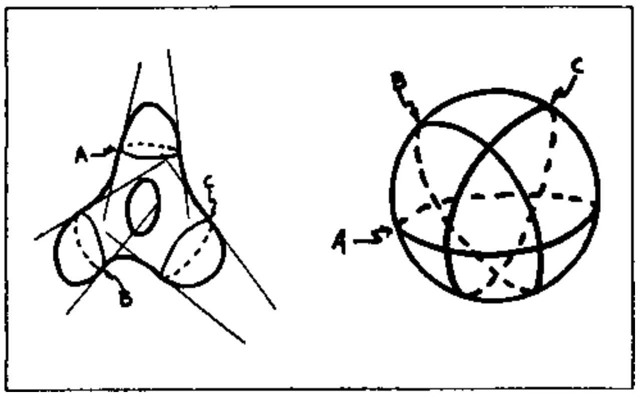

Figure 4: Parabolic lines on a surface and their Gaussian images. In general, the Gaussian image curves may have cusps.

topology) in a saddle-node. On the surface, they are the inflections. Also, they are level sets of the Gaussian curvature function, for the level value of sero. Like other level sets, they exhibit a nesting structure, so that they, too have a graph structure associated with them. The Gaussian curvature is a smooth function defined on the surface, so we can play the same level set graph game with it. This yields the Gaussian curvature graph of the surface, which has the same topological features as the other level set graphs. This graph is intrinsic, though, since there is only one Gaussian curvature function for (a given embedding of) the surface. As we smoothly change the surface (say by deforming it), both the graph structure of the Gaussian curvature map (hence the parabolic lines) and of the surface level sets in each domain of the Gaussian sphere will remain constant, while the shape and locations of the domains will change. This will continue until we reach a bifurcation of the "parabolic" graph, at which time a domain may be created or annihilated, signifying a qualitative change in appearance for some view aspect.

There is a notion of scale for each of the graph structures, analogous to that for the level set tree of a function on the plane, defined by the nesting of the respective level sets. "Smoothing" can be effected by lopping off leaf nodes, or whole terminal subgraphs.

D. The topological iodise

We've seen so far that the (images of the) parabolic curves on the Gaussian sphere mark the normal directions where the level set graph undergoes a saddle-node bifurcation. Equivalently, they divide the Gaussian sphere into regions where the topology of the level set graph is constant (modulo saddle connections). We've also seen that distortions of the surface move these parabolic curves around on the Gaussian sphere. To capture the topology of these regions, we can use the adjacency graph of the regions.

Actually, we were not completely candid about the construction of the regions on the Gaussian sphere. You may have noticed that it's not entirely clear which of the 2 possible normals to a level plane is the right one for mapping to the Gaussian sphere at any given time. The answer is that really they both are, and the only solution is not to use the Gaussian sphere at all, but to use the *Gaussian projective plane* obtained by identifying antipodal (diametrically opposite) points of the Gaussian sphere to single points. To make our pictures completely precise, we have to map the parabolic curves to this *Gaussian projective plane* by the antipodal identification, and consider when the normal line to a level plane crosses one of these curves, i.e. when *either* normal of the level plane crosses the parabolic curve on the Gaussian sphere. Figure 5a shows the result of applying the above mapping. We are representing the projective plane by a hemisphere, with features from the rear hemisphere of the Gaussian sphere mapped to the front hemisphere (which we show) by the antipodal map. In such a picture, the points on the circle at the boundary of the hemisphere are understood to be identified with their antipodal points (on the same circle). The dashed lines represent the parabolic curves that were on the rear of the sphere in Figure 4, and which now have been mapped to the front of the sphere by the antipodal map, so that all the features are now on the front of the sphere, and we therefore only depict a hemisphere. This projective plane represents all the *lines* through the center of the sphere—the normal vectors with their sense of direction stripped away. Note that rotating a line through only 180 degrees results in the same line again; that's why we have to identify antipodal points on the sphere.

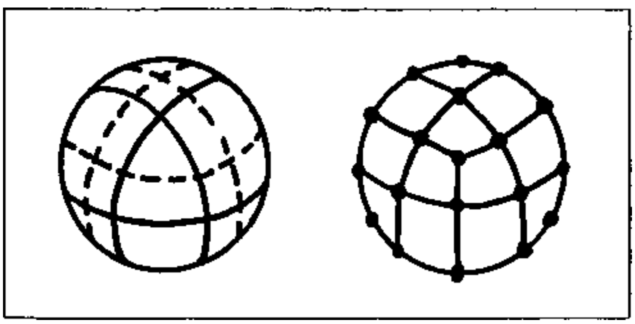

Figure 5: The Gaussian projective plane, (a) The images of parabolic curves, (b) The adjacency graph of the regions. Note that antipodal points on the periphery circle are identified.

It's the regions in the Gaussian projective plane that really are the regions of constant level set graph. In Figure 6 we show some of the level set structures for our example surface. We've depicted the surface rotating relative to fixed level planes for clarity. The relative direction of the normal line for each set of level planes lies on the associated region of the Gaussian projective plane. The adjacency graph associated with these regions is shown in Figure 5b. Keep in mind that antipodal nodes on the outside circle are identified, even though each one has been shown twice to allow a comprehensible graph.

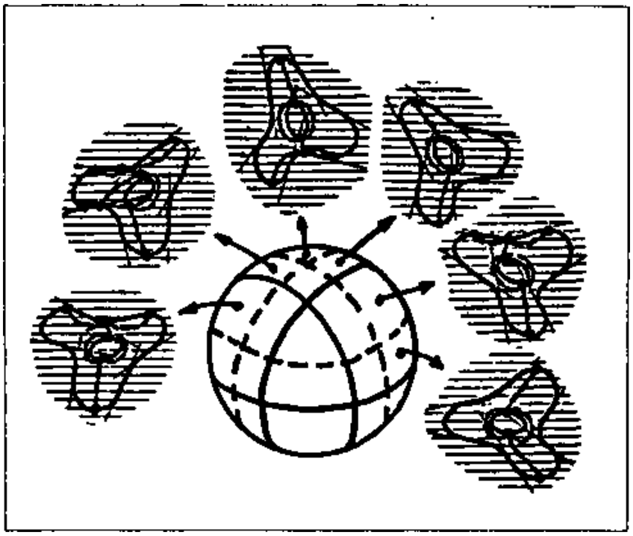

Figure 6: Level set structures for regions on the Gaussian projective plane. (Not all the possibilities are illustrated.)

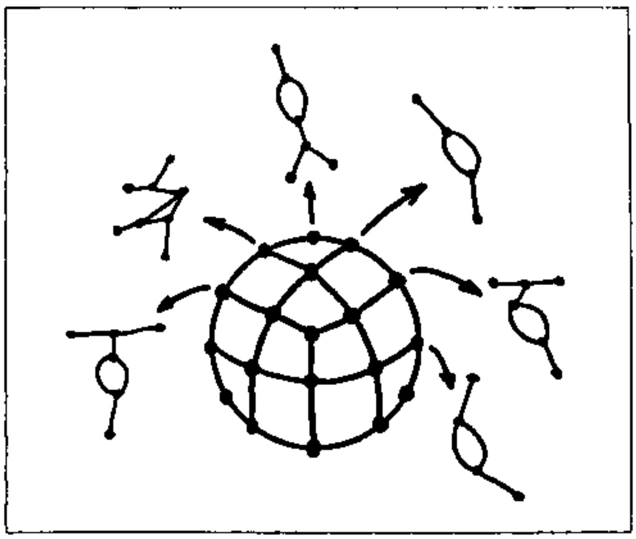

Figure 7: The topological zodiac of the surface: The graph of level set graphs for all choices of z-axis. (Not all the possibilities are illustrated.)

Then each node in the adjacency graph has associated with it the level set graph that describes the bump structure for z-axes within its region (Figure 7). We call this the *topological zodiac* since it is reminiscent of some astrological charts. This graph of graphs completely specifies all the level set structures possible for the surface, and all the possible transitions between them. Of course further information could be appended to the structure, e.g., metric information about the sise of regions, their actual orientation, etc.

The topological sodiac is related to the *visual potential* of [Koenderink and van Doom 1979] in that both representations are graphs of topological types, based on analysis of singularities. The former, however, represents solid shape; the latter, shape of outlines. The full details of the relationship remain to be worked out.

III. Summary and Conclusions

We sought to represent protrusions with seroth order structure, and found that this is organised by 1st order singularities: critical points. The changes in this structure, in turn, are governed by 2nd order singularities: parabolic curves. The parabolic curves embed in the graph structure of the level sets of Gaussian curvature. The resulting representation provides a stable discrete structure which captures some (but by no means all) of the intuitive notion of shape. It allows restriction to a range of scale, a characterisation of substructure, and thus permits a coarse-to-ftne matching strategy.

Since all the information we have used is really contained in the Gaussian curvature graph and the Gaussian projective plane adjacency graph, it should be possible to dispense with the topological sodiac (the graph of level set graphs). It seems that the latter should be deducible by very simple operations directly from the above graphs. We haven't figured out how to do that yet, though.

Acknowledgments

I would like to thank Ted Selker for some very helpful discussions, Danny Sabbah for suggesting applying these ideas to shape representation, Sanjay Addanki for his critical appraisal, Gabi Kuper for bringing TEX to Yorktown, Scott Penberthy for mating machines in captivity, and Sesh Murthy for countless favors.

References

[Agin and Binford 1973] Agin, Gerald A. and Thomas O. Binford, "Computer Description of Curved Objects," *Proceedings of Third International Joint Conference on Artificial Intelligence,* Stanford, California, August 1973, 629-635.

[Besl and Jain 1985] Besl, Paul and Ramesh Jain, "Intrinsic and Extrinsic Surface Characteristics," *Proceedings of IEEE Computer Vition and Pattern Recognition Conference,* San Francisco, June 1986, 226-233.

[Binford 1971] Binford, Thomas O., "Visual Perception by a Computer,*[1] *IEEE Conference on Systems and Controls,* Miami, December 1971.

[Blicher 1983] Blicher, A. Peter, "The Stereo Matching Problem from the Topological Viewpoint," *Proceedings of the Eighth International Joint Conference on Artificial Intelligence,* 1983, 1046-1049.

[Blicher 1984] Blicher, A. Peter, "Edge Detection and Geometric Methods in Computer Vision," Ph.D. thesis, Mathematics Department, University of California, Berkeley, December 1984. Also Stanford University Computer Science Department technical report STAN-CS-85-1041 and AI memo AIM-352, February 1985.

[Blicher 1967] "A Shape Representation Based on Geometric Topology: Bumps, Gaussian Curvature, and the Topological Zodiac," IBM T.J. Watson Research Center Technical Report, to appear.

[Brady et al. 1985] Brady, Michael, Jean Ponce, Alan Yuille, and Haruo Asada, "Describing Surfaces," in Robotics Research, MIT Press, Cambridge, 1985.

[Brooks 1981] Brooks, Rodney A., "Symbolic Reasoning Among 3-D Models and 2-D Images," *Artificial Intelligence,* 17, 1981, 265-348.

[Bruce 1981] Bruce, J.W., "The Duals of Generic Hypersurfaces," *Mathematica Scandinavica,* 49, 1981, 36-60.

[Haralick, Watson, LarTey 1983] Haralick, Robert M., Layne T. Watson, and Thomas J. LarTey, "The Topographic Primal Sketch," *The International Journal of Robotics Research,* vol. 2, no. 1, spring 1983.

[Horn 1983] Horn, Berthold K.P., "Extended Gaussian Images," MIT AI Memo 740, Cambridge, July 1983.

[Koenderink and van Doom 1976] Koenderink, J.J. and A.J. van Doom, "The Singularities of the Visual Mapping," *Biological Cybernetics,* 24, 1976, 51-59.

[Koenderink and van Doom 1979] Koenderink, J.J. and A.J. van Doom, "The Internal Representation of Solid Shape with Respect to Vision," *Biological Cybernetics,* 32, 1979, 211-216.

[Koenderink and van Doom 1980] Koenderink, J.J. and A.J. van Doom, "Photometric Invariants Related to Solid Shape," *Optica Acta,* vol. 27, no. 7, 1980, 981-996.

[Koenderink and van Doom 1984] Koenderink, J.J. and A.J. van Doom, "What Does the Occluding Contour Tell Us About Solid Shape?," *Perception,* vol. 13, 1984, 321-330.

[Koenderink and van Doom 1979] Koenderink, J.J. and A.J. van Doom, "The Structure of Two-dimensional Scalar Fields with Applications to Vision," *Biological Cybernetics,* 33, 1979, 151-168.

[Nackman 1984] Nackman, Lee R., "Two dimensional Critical Point Configuration Graphs," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. PA MI-6, no. 4, July 1984, 442-450.

[Nevatia and Binford 1977] Nevatia, Ramakant and Thomas O. Binford, "Description and Recognition of Curved Objects," *Artificial Intelligence,* vol. 8, no. 1, February 1977, 77-98.

[Pentland 1965] Pentland, Alex P., "On Describing Complex Surface Shapes," *Image and Vision Computing,* vol. 3, no. 4, November 1985, 153-162.