Scott McCloskey COMP 766

Paper Critique, Jan J Koenderink

"What does the occluding contour tell us about solid shape?"

1 The Big Picture

The main point of this paper is the presentation of a simple yet counterintuitive analysis of what the occluding contour of an object tells us about the surface that it describes. Specifically, convexities in the occluding contour correspond to convexities of the surface around that position. Concavities in the contour correspond to saddle-shaped patches of the surface, and inflections correspond to what we've described as hyperbolic points on the surface. This contradicts previous work in the field, which had suggested that other cues were used to judge surface characteristics from the contour.

2 The Gory Details

To begin the paper, the author reviews the basic constructs that are needed to prove the main point. We start with the *rim*, which is the set of points on an object's surface where a ray coming from the viewpoint is tangent. This is a smooth curve which isn't generally planar and is not self-intersecting. The *contour* is, in a certain sense, the opposite: it isn't necessarily smooth, it may intersect itself, and is planar. It is, in fact, the projection of the rim on a plane perpendicular to the viewing axis.

The points on the rim, and their corresponding points on the contour, are only of two different types. The first group are those where the perspective ray, which is a tangent, grazes the surface and heads away from the object. If we imagine a projection of a red object on a blue background, these points are the boundary between red and blue regions. The second type of point is where the perspective ray grazes the surface and then, at some greater depth, intersects it again.

We now consider an arbitrary point R on the rim of the object viewed from direction ζ. Assuming that the surface is smooth, we have a well-defined normal at R which, together with the ray ζ, gives us a normal plane. That plane cuts the surface, giving us the curve λ. Secondly we have the plane π, which is normal to ζ at R, and the curve μ that is cut from the surface by π.

With these constructs in mind, we can talk about the *radial curvature* of the surface at point R. The radial curvature is the curvature along the line of sight, and can be expressed as $K_r = \frac{1}{r_t}$ where r_t is the radius of curvature of λ. We have a similar expression using r_t, the curvature of μ, the projection. The *transverse curvature* at R is $K_t = \frac{1}{r_t}$. The importance of these values becomes apparent in Appendix I, where the author proves that the *Gaussian curvature* at R is $K = K_t K_r$.

While this is an interesting result, it isn't terribly useful to us because we don't measure any of these values directly. What we can measure is K_{app}, the *apparent curvature* from a given perspective projection. It turns out that $K_{app} = dK_t = \frac{dK}{K_t}$, which is the main result of the paper. For the record, d is the distance between the point R and the viewpoint.

In order to analyze this relationship, we have to consider the signs of the quantities involved. Clearly d is positive since it is a distance. K_r is a non-negative quantity, though this is less obvious. If K_r were negative, the object must obscure point R. This contradicts the fact that R is on the rim, so we can say that K_R is non-negative.

What we're left with, after the consideration of d and K_r is that the sign of K_{app} is the same as the sign of K. This is remarkable because of what we already know about the sign of K and the properties of the surface at the point. Namely, if K_{app} is positive then the surface is convex at R. If K_{app} is negative, then R is a saddle point of the surface. If K_{app} is 0 then R is an inflection point of the surface.

3 The Scott Opinion

The most remarkable thing about this paper, on the first read, is the fact that it was published in 1984. The main result is fairly simple, and isn't based on late-breaking developments in the field, so it's strange that it went unproven until the mid-80's. The author addresses the point in the introduction, and it's hard to improve upon that wording:

> Oddly enough, our visual systems automatically interpret some complicated spacial configurations correctly, despite the fact that we may hold primitive and even incorrect intellectual notions about geometrical relations.

It's almost a funny point - our brains are smarter then we are - but it makes us wonder how these incorrect notions came about, and how they came to afflict everyone who had previously considered the problem. One might look to blame David Marr, as his importance to computer vision probably gave undue weight to his - incorrect, as it turns out - observations. Of course, the real problem is that people believed these incorrect assertions; perhaps because of Marr's name, but perhaps because they didn't look into the problem.

The author points out the fact that artists hadn't let this alleged inability prevent them from conveying surface features in their artwork. It seems that vision scientists were likely to claim that some prior information was used to interpret such drawings, which is faintly reminiscent of the topic of depth perception. In the latter topic, Bela Julesz demonstrated our ability to perceive depth absent any familiar cues by presenting images that produced stereo perception without using anything familiar. Koenderink disproves the idea of outside information in a more satisfying way: without example, by explaining the causal relationships.

The main points of the paper are hard to dispute because of this. The relationships are clearly explained, and the methodology is solid. One could object to the use of non-standard notation, but this is a minor point. What's more is that the author explicitly mentions the conventional notation, to give us a decoder ring for interpreting the paper.

Organization seems to be something of a weak point for this paper. It's unclear why the author decided to prove results in the two appendices rather then the body of the paper, where it would seem to fit more naturally. More importantly is the fact that the paper's main result isn't clearly stated in one place. On page 323 the author handles the case where $K = 0$, but doesn't reiterate it on the next page when the non-zero cases are handled. It would have been clearer had the $K = 0$ case been held off for a page.

The style of the paper is effective, and one could imagine that writing it was a stylistic challenge. How do you tell all of your peers that they're wrong without being too heavy-handed? The author chooses to highlight Marr's text because it is so well-respected that his reputation could withstand this critique. It's my understanding that Marr had previously been contradicted on other points, so this probably wasn't a big shock to the community at large. In any event, the author of this paper demonstrated both good writing style and tact.