Singularity Theory in Computer Vision

Peter Giblin

Department of Mathematical Sciences

The University of Liverpool

Since 1985 (a visit to University of Massachusetts, Amherst) I have been involved in four lines of enquiry with computer scientists and engineers at

University of Cambridge (Roberto Cipolla)

University of Massachusetts (Richard Weiss)

University of Minnesota (Guillermo Sapiro)

Brown University (Ben Kimia, David Mumford)

IT University of Copenhagen (Mads Nielsen)

UNC Chapel Hill (Steve Pizer, James Damon)

he's a mathe- matician

and at Liverpool (Bill Bruce, Farid Tari), also mathematicians

Four broad areas:

- 1. Recovery of shape and motion from profiles (Weiss, Cipolla, Äström: rather a long time ago—problems got too hard!)
- 2. Medial axes, local symmetry, either euclidean or affine (Kimia, Sapiro, Nielsen, Zakalyukin: still ongoing)
- 3. Moving views of illuminated surfaces (Damon, ongoing)
- 4. Salient features of surfaces (crest lines etc.) (Mumford, Bruce, Tari: related to (2), no direct work recently)

(and in all cases students/postdocs at Liverpool)

The common thread is **geometry** in 2D or 3D

In addition **singularity theory** comes in (basically this deals with classification of functions and mappings)

1. Recovery of shape from profiles (outlines, apparent contours)

folds

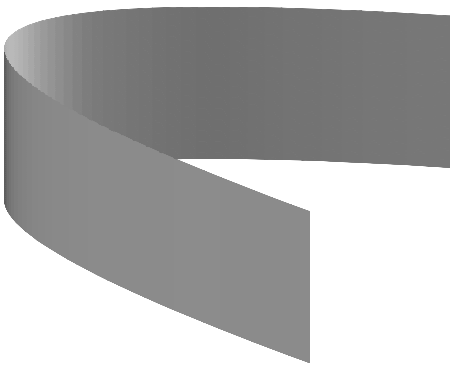

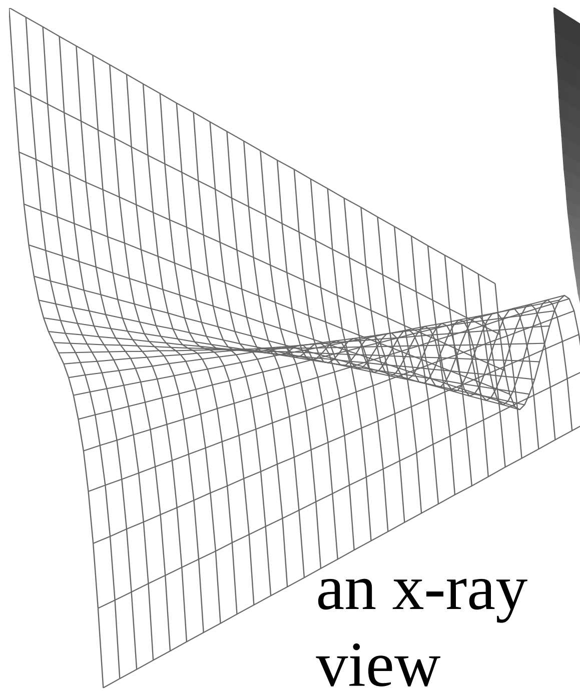

an x-ray view

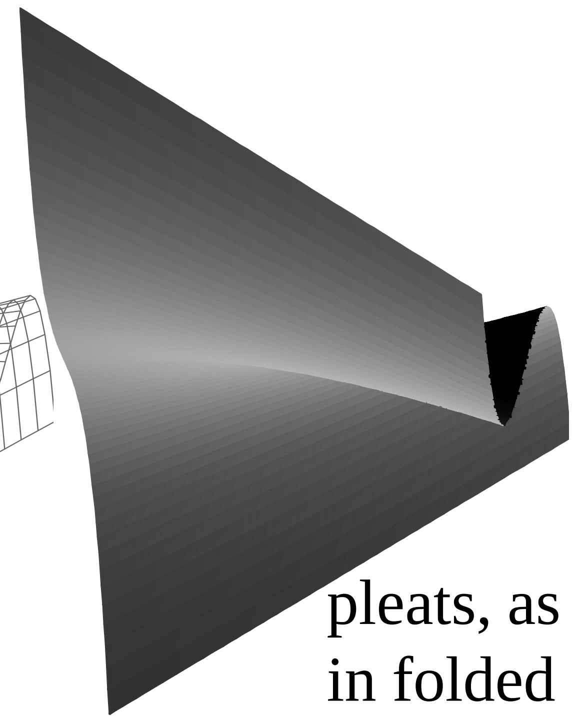

pleats, as in folded cloth

The **singularities** here are those which arise from projection of a surface into a viewplane.

Moving viewplanes and x-ray images of the contours (profiles)

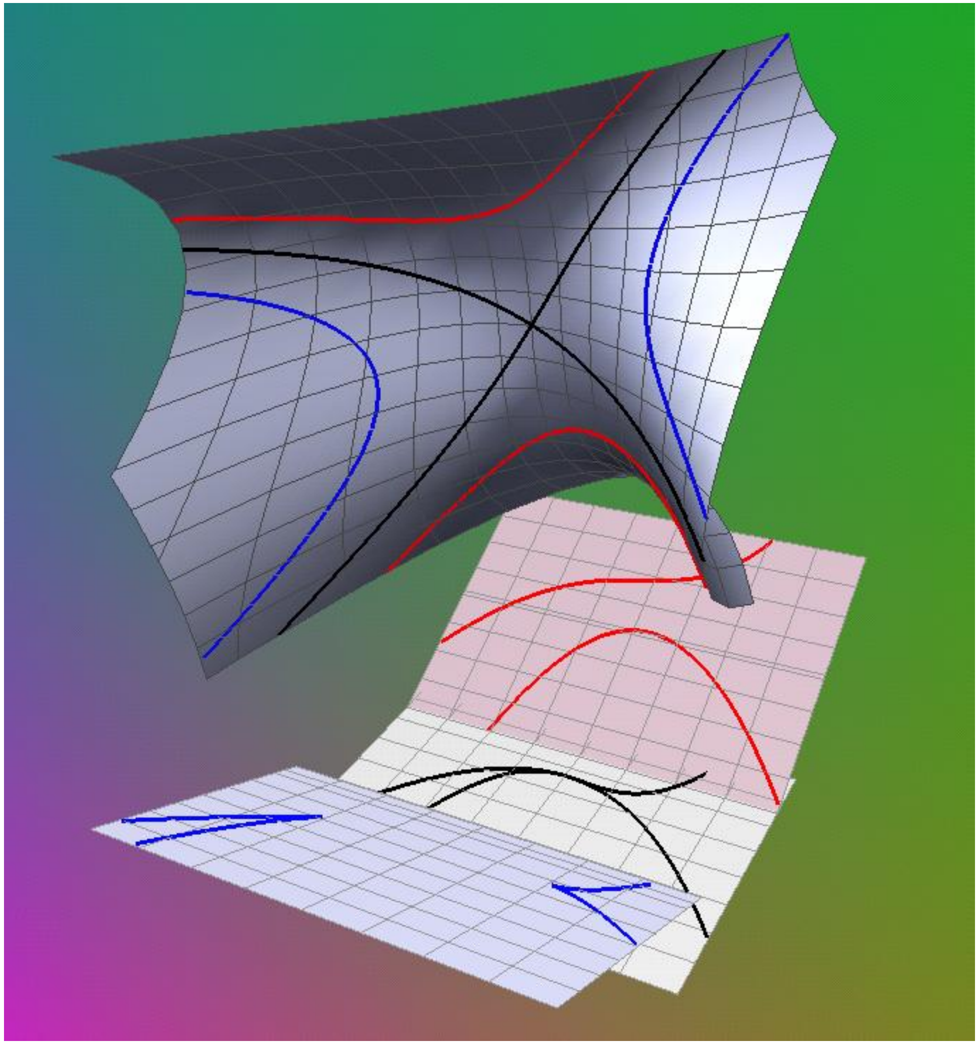

The key idea is that a surface may be reconstructed as an envelope of cones centred on the camera positions **c** and containing the ‘visual rays’ which are tangent to the surface along a ‘contour generator’ and which therefore contribute to the profile.

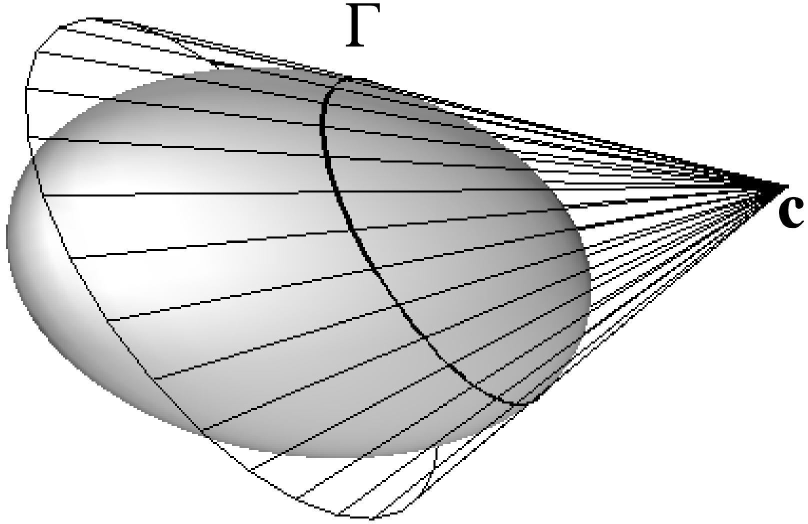

as **c** moves on a known path, the contour generator Γ sweeps out (part of) the surface.

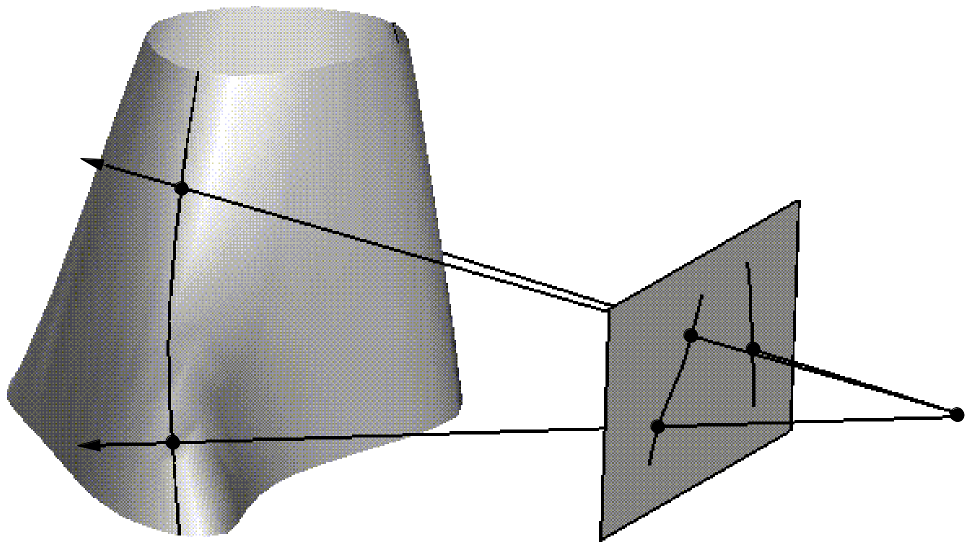

The profiles (apparent contours) are obtained by intersecting the visual rays with a plane or sphere.

Things become much more interesting when the path of the ‘camera centre’ is unknown. I did some work on ‘circular motion’ with Cipolla and Äström (1995).

WONG AND CIPOLLA: RECONSTRUCTION OF SCULPTURE FROM ITS OUTLINES
2004

Fig. 14. Eight images of an outdoor sculpture acquired by a hand-held camera under approximate circular motion.

Disclaimer! This is not my work but is an extension of work I did with Cipolla.

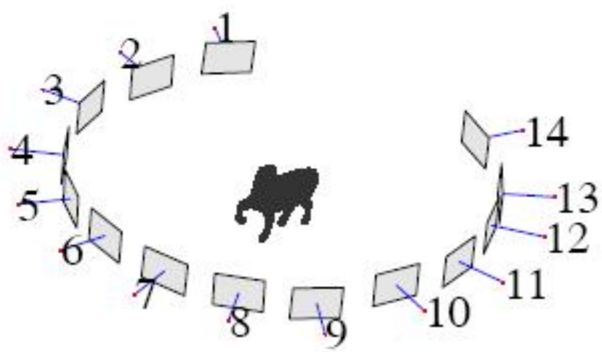

Recover the shape from unknown but roughly circular motion, using the outlines (profiles) of the shape.

Fig. 15. Camera poses estimated from the outdoor sculpture sequence.

Fig. 16. Triangulated mesh of the reconstructed outdoor sculpture. This model was composed of 29,672 triangles.

2. Medial axes, local symmetry in 2D and 3D

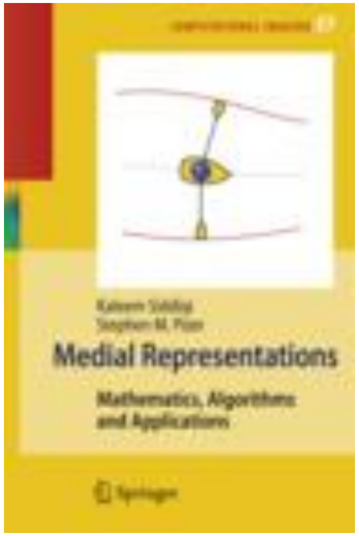

Medial axes (skeletons) have applications in recognition, classification and comparison of shapes in 2D and 3D.

There is a recent book edited by Steve Pizer and Kaleem Siddiqi, called **Medial Representations** (Kluwer Publishers)

Medial Representations
Mathematics, Algorithms and Applications
Series:Computational Imaging and Vision, Vol. 37
Siddiqi, Kaleem; Pizer, Stephen (Eds.)
2008, XVI, 413 p. 204 illus., 84 in color., Hardcover
ISBN: 978-1-4020-8657-1

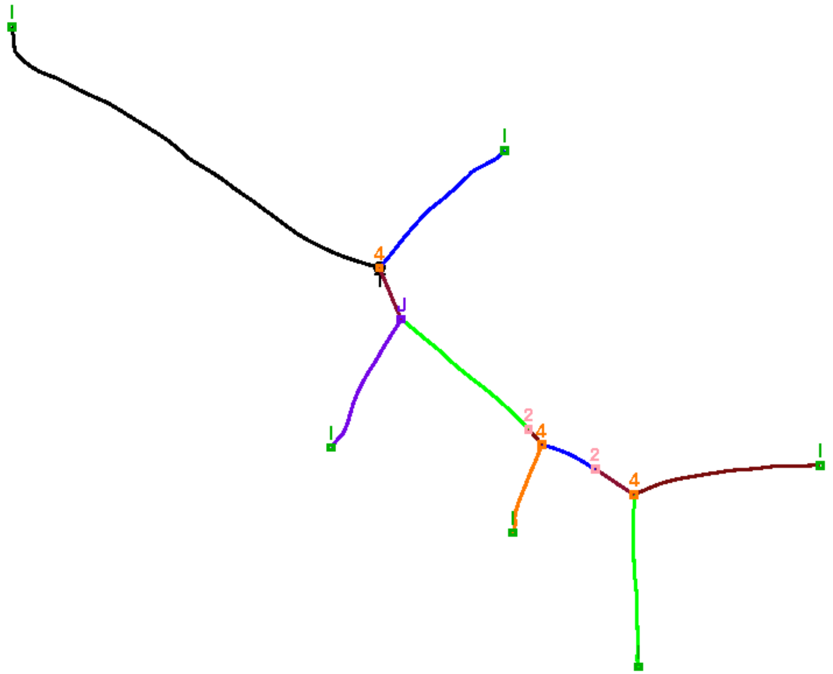

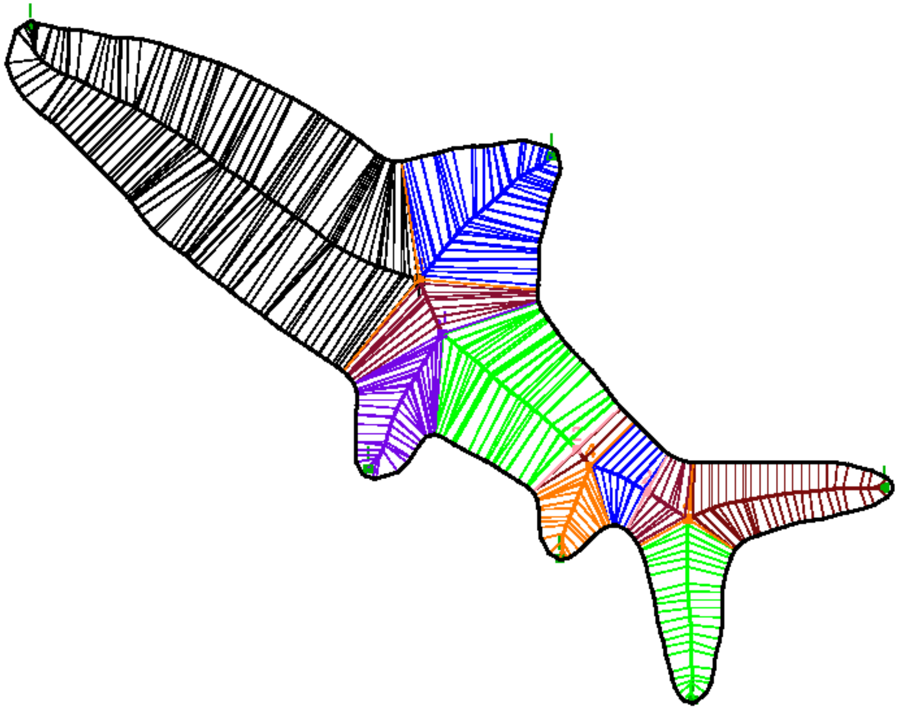

A typical **medial axis** (MA, skeleton) in 2D showing the tree structure (branches but no loops) which ‘encodes’ the shape

not **medical axis** though I understand there are medical applications!

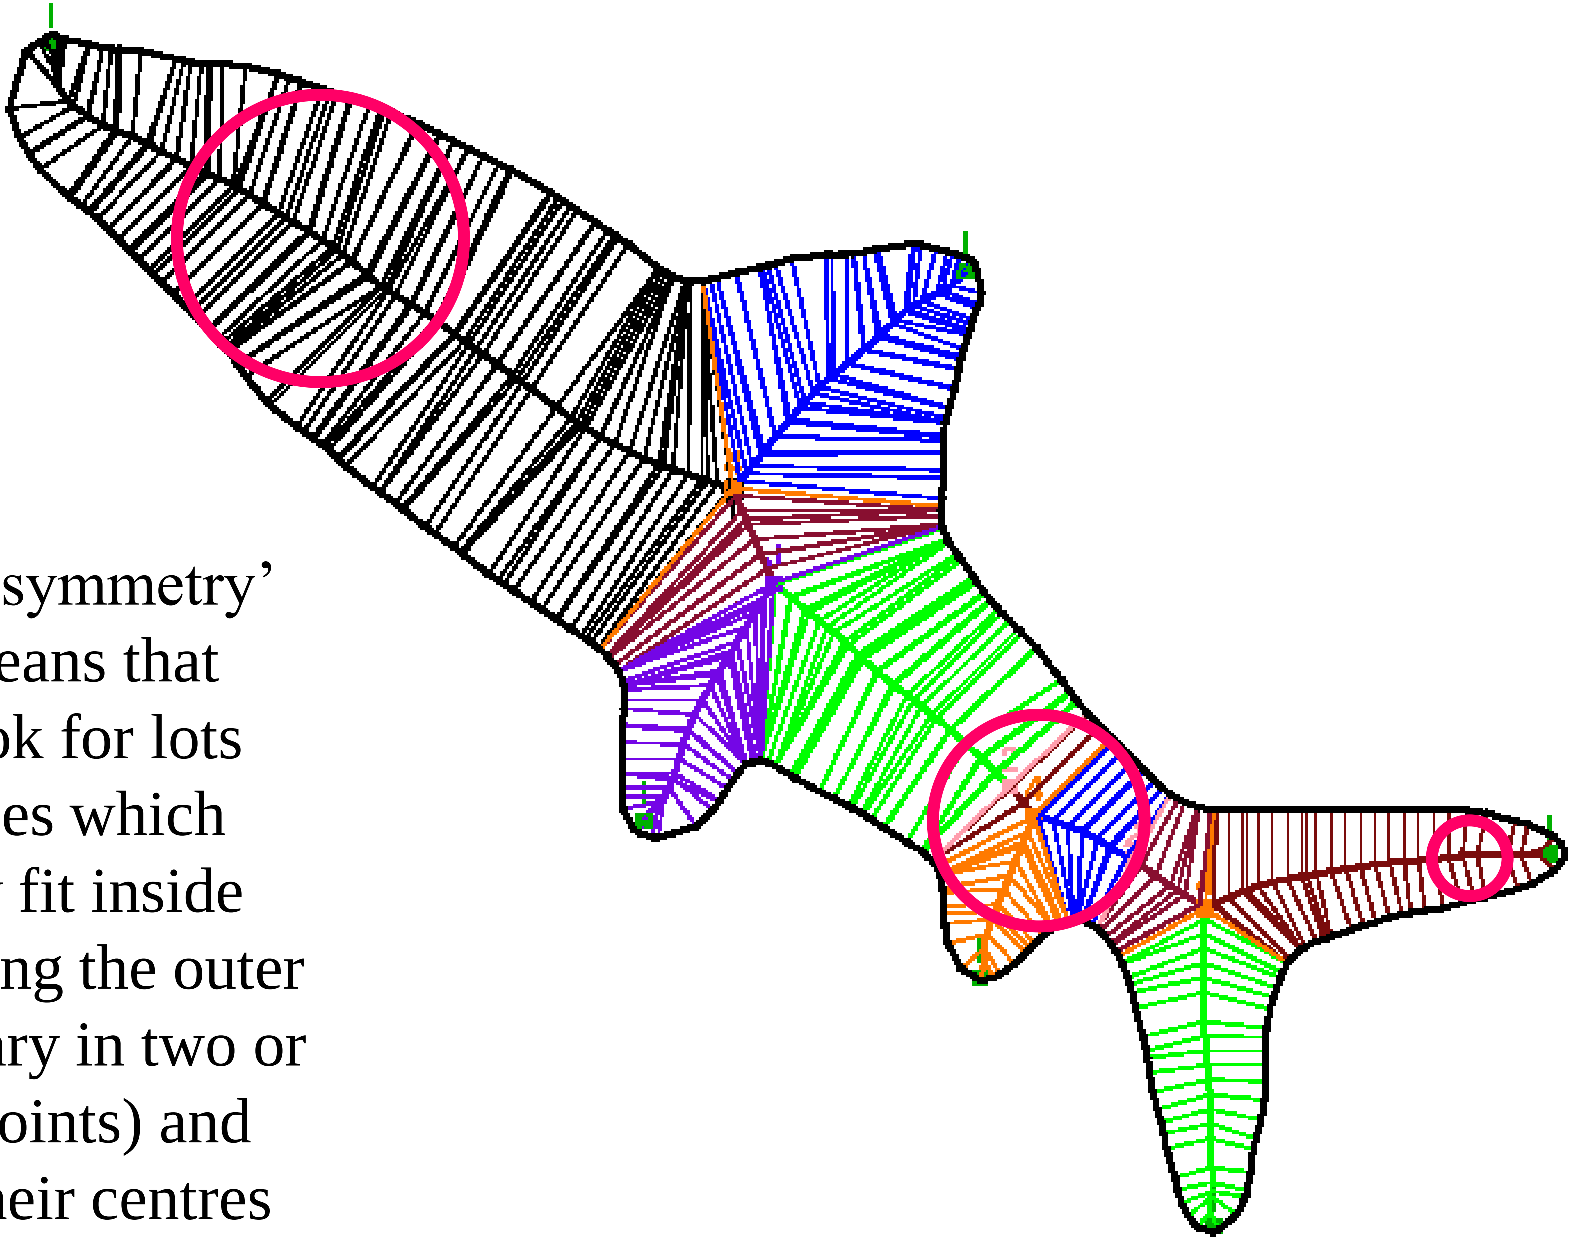

‘Local symmetry’ here means that you look for lots of circles which exactly fit inside (touching the outer boundary in two or more points) and trace their centres

The singularities which arise in this situation come from a function (the distance from a point in the plane to a point on the curve) which is **has two (equal) global minima** at the centre of one of the 'bitangent circles'

singularity theory is well set-up for analyzing situations like this

sometimes three global minima

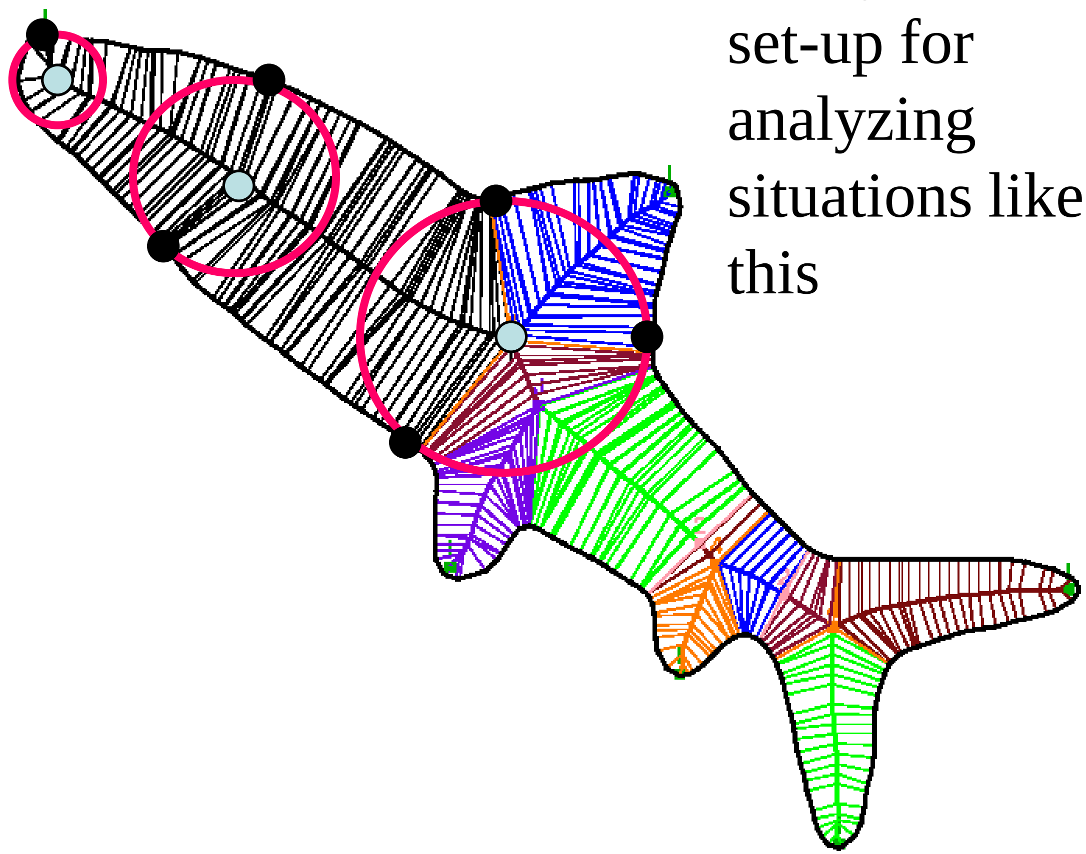

or a degenerate global minimum, where two ordinary minima have coalesced

For a simple shape like an ellipse (in red) the medial axis is just the blue line

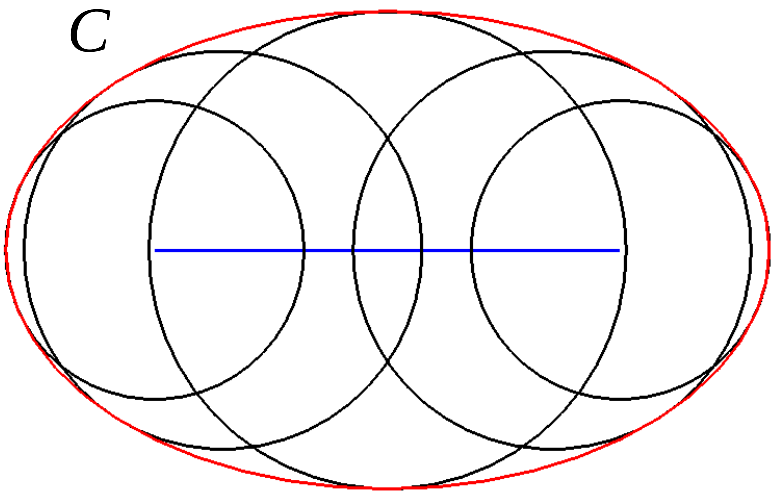

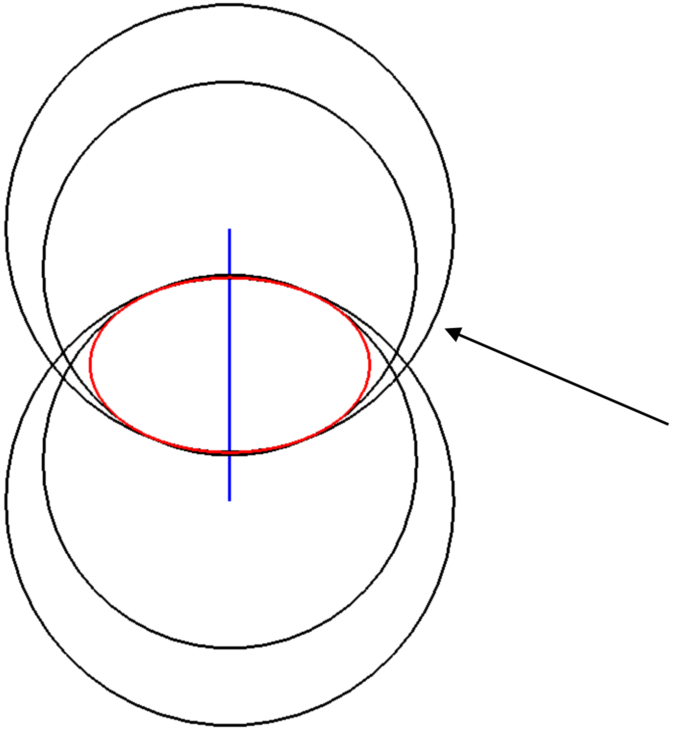

If we allow circles which are not necessarily inside the shape *C* then their centres trace the so-called ***symmetry set*** (SS) of *C*.

Note the SS is *straight* indicating global symmetry about an axis.

Pre-symmetry set

Parametrization of ellipse

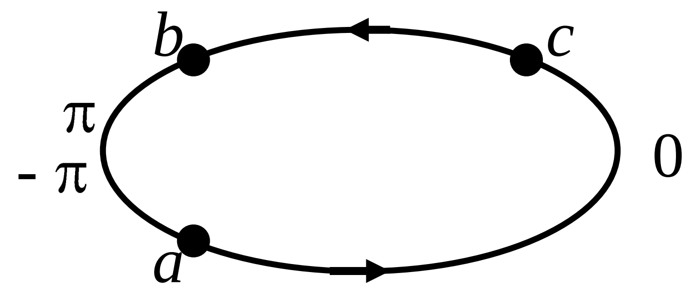

Green circle contributes two green dots (a,b) and (b,a) to the pre-SS

Diagram for all pairs of points (s,t) on the ellipse

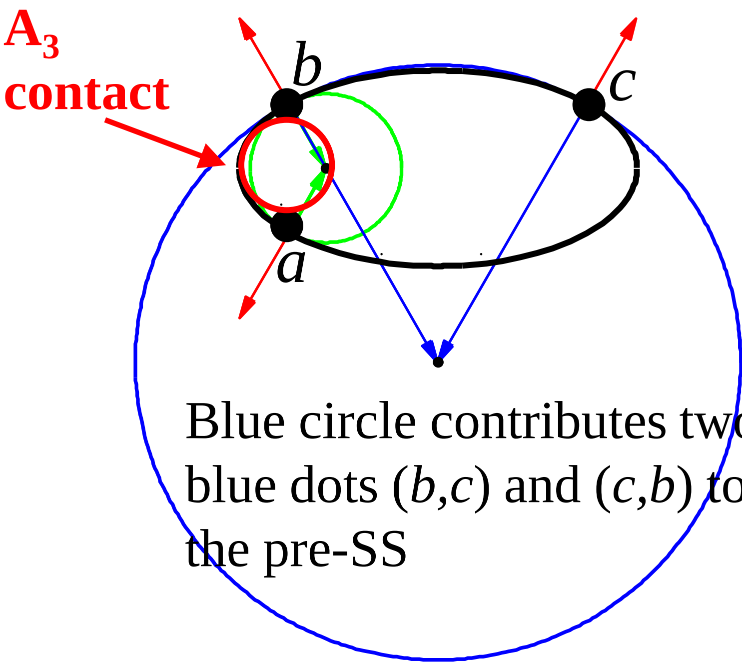

Blue circle contributes two blue dots (b,c) and (c,b) to the pre-SS

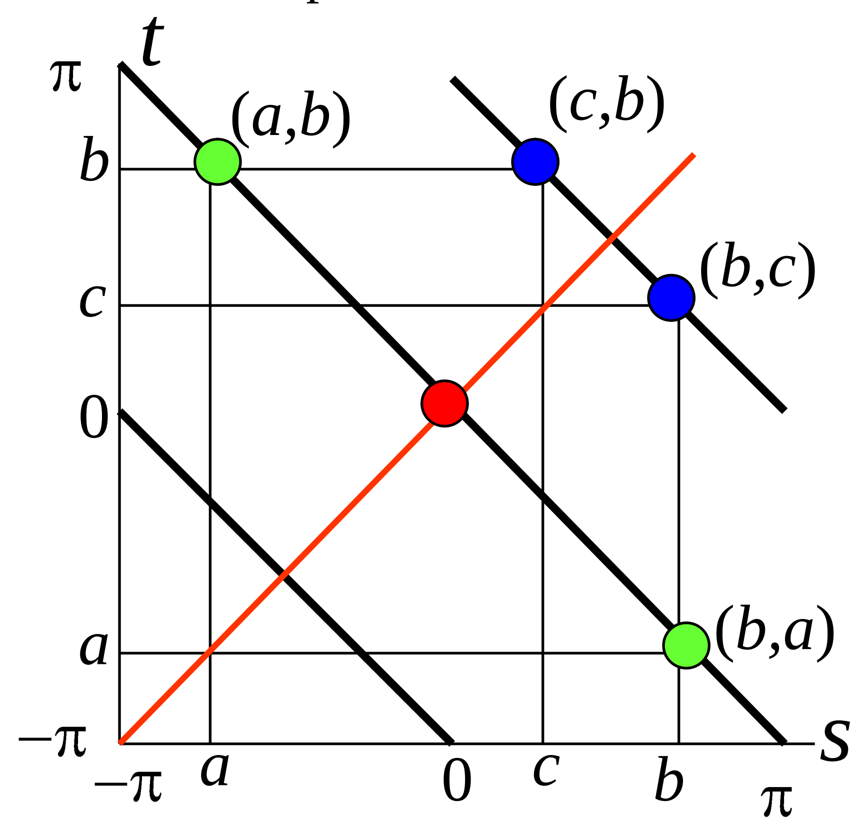

All bitangent circles give the full pre-SS, symmetric about the diagonal $s=t$

Each diagonal point gives a circle where the points of contact have *coincided* at a **max** or **min** of curvature (here a **max**).

This is called A_3 contact.

Topologically the rectangle of pairs (s,t) is a **torus** and the two loops are **essential loops**, not contractible to points.

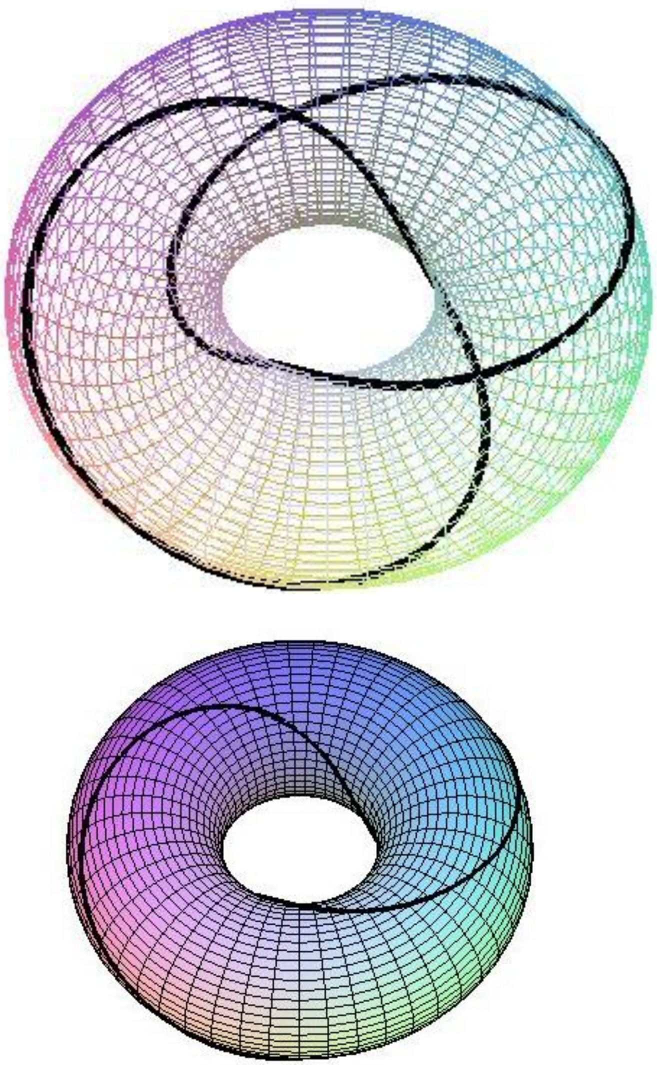

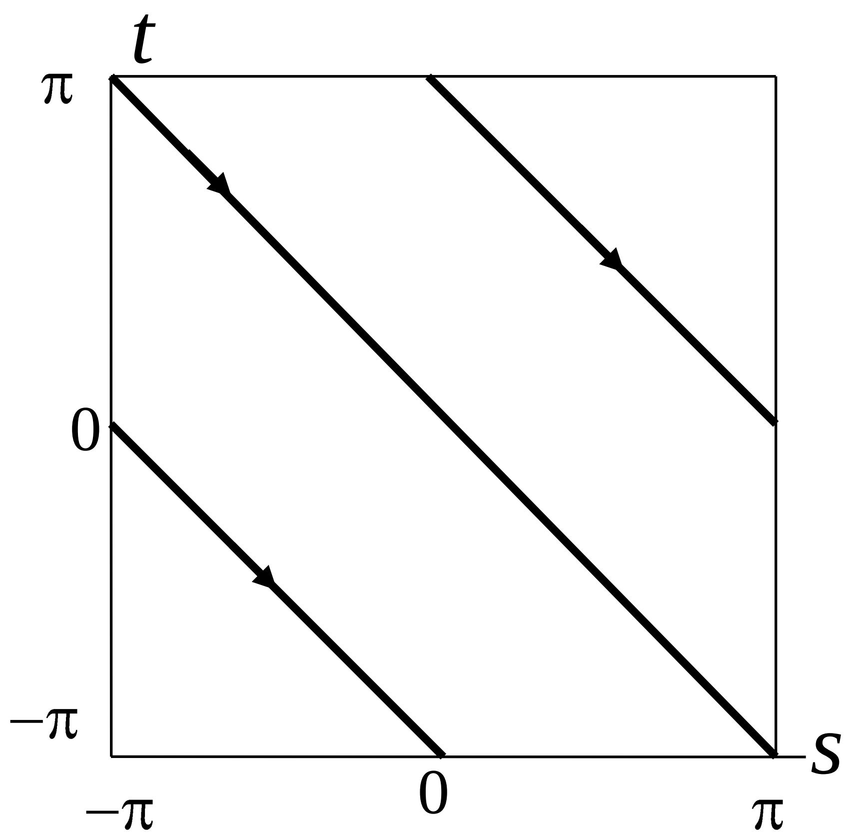

−π and π give the same point

Bitangent circles

Branch of symmetry set

Pre-symmetry set

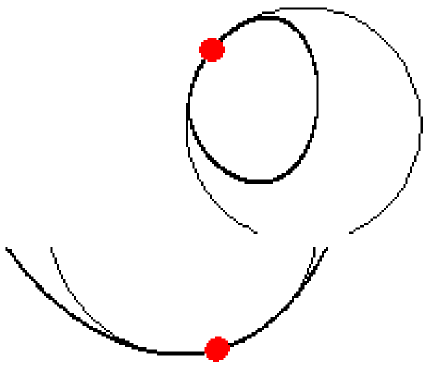

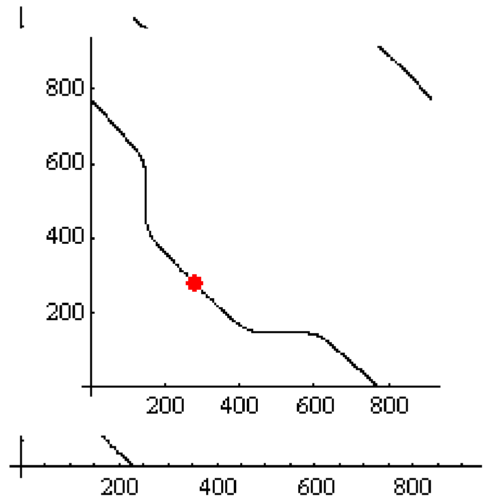

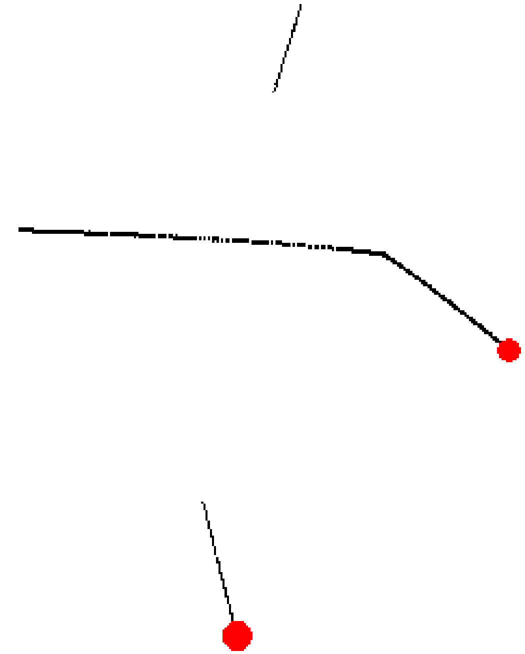

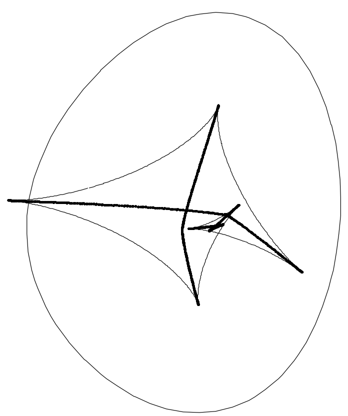

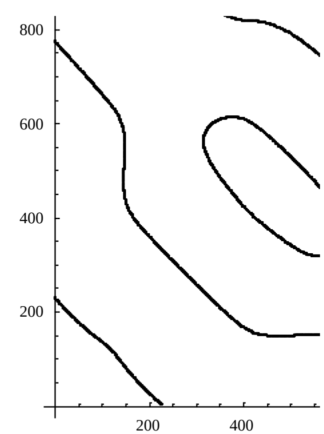

Full symmetry set

Animation by Arjan Kuijper (Copenhagen)

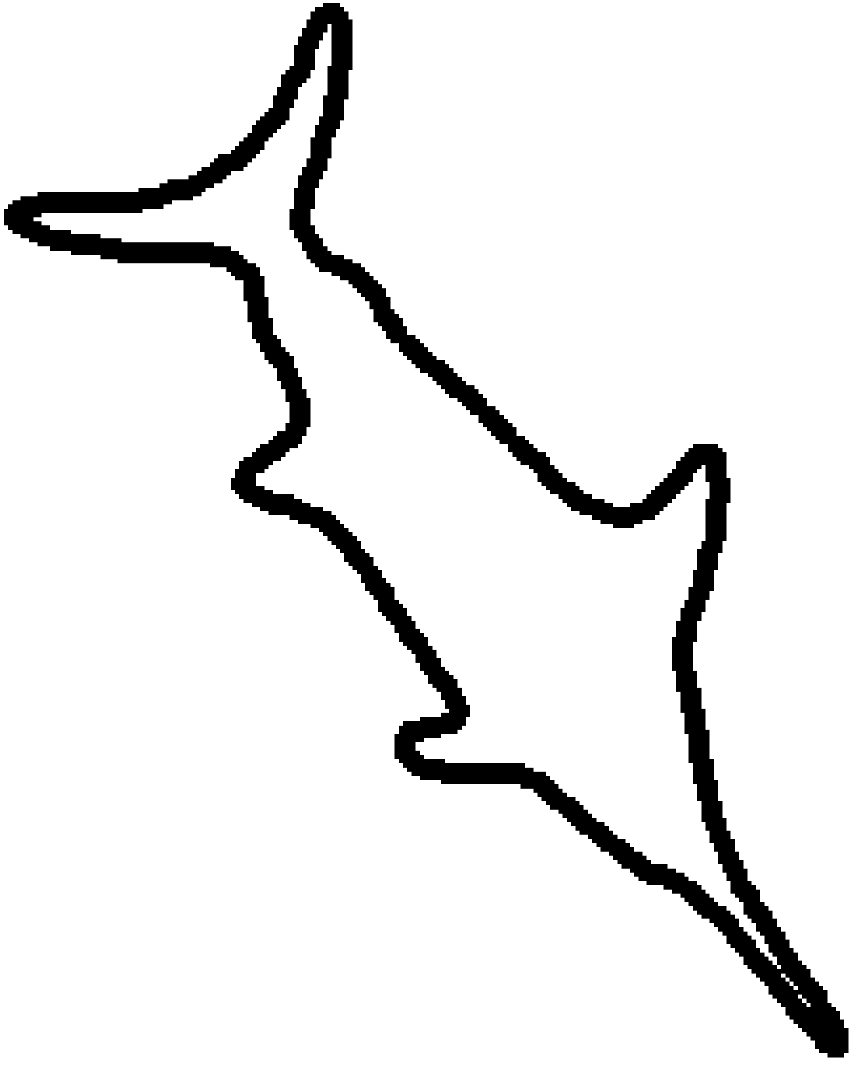

Animation by Arjan Kuijper

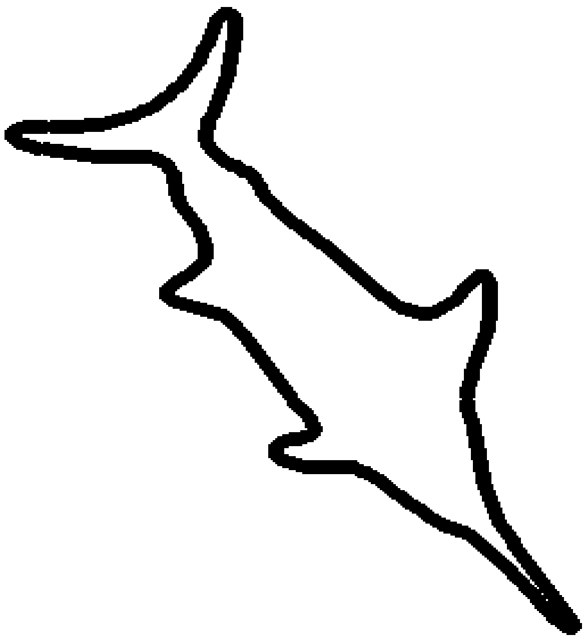

Animation by Arjan Kuijper

All the ‘transitions’ on the pre-SS, SS and medial axis have been classified (in 2D and 3D) using singularity theory.

Some very interesting problems remain when we consider instead of evolving families of *smooth* curves, families of curves one member of which is *singular*, for example planar sections of a surface close to the tangent plane.

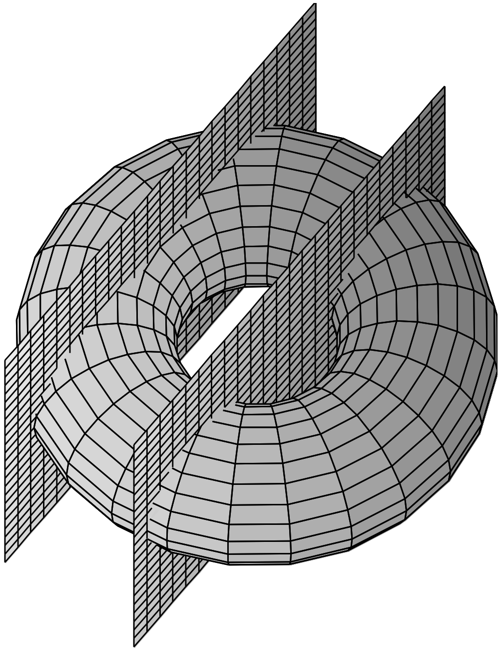

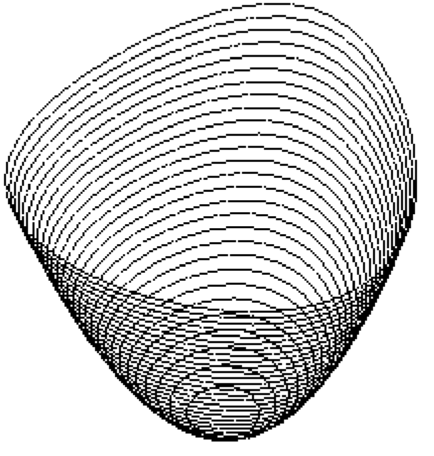

When we consider **surfaces in 3-space** and their SS, MA and pre-SS we get similar but more complex behaviour.

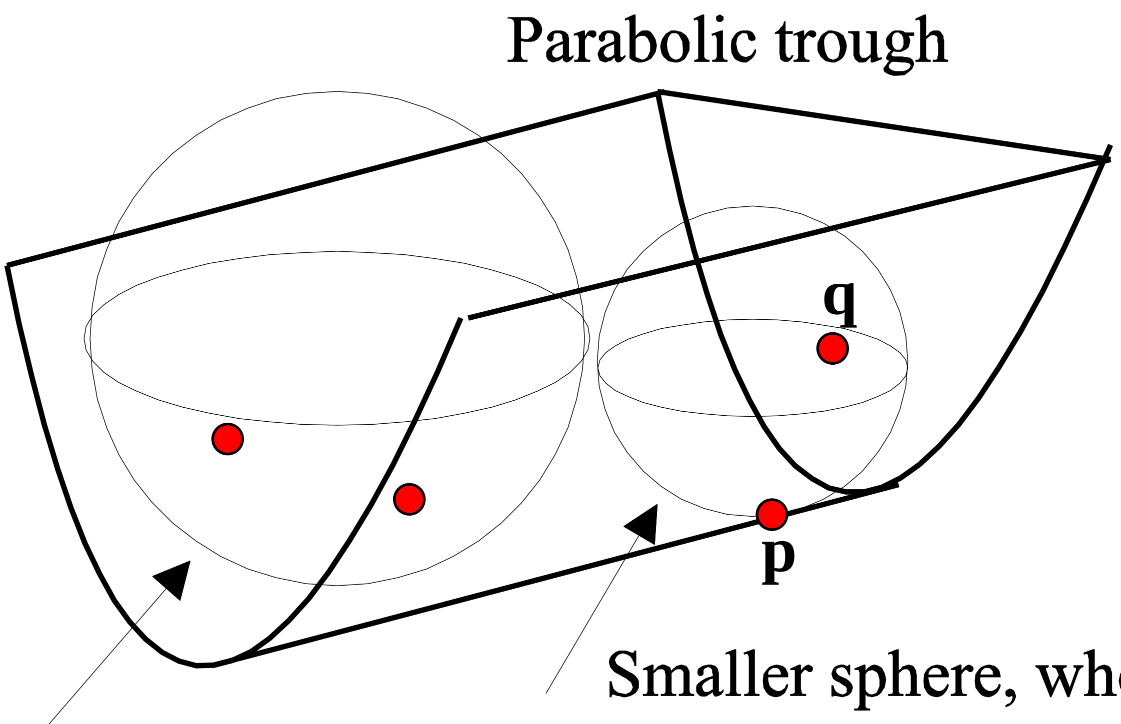

p is called an A_3 point and **q** an A_1 point; the sphere is an A_1A_3 sphere

Big sphere, tangent at 2 points, front and back

Smaller sphere, where two tangency points have coincided at a *ridge point* **p.** This sphere is also tangent at another point, **q**

A typical medial axis (MA) of a surface (an elliptical bin with a lid)

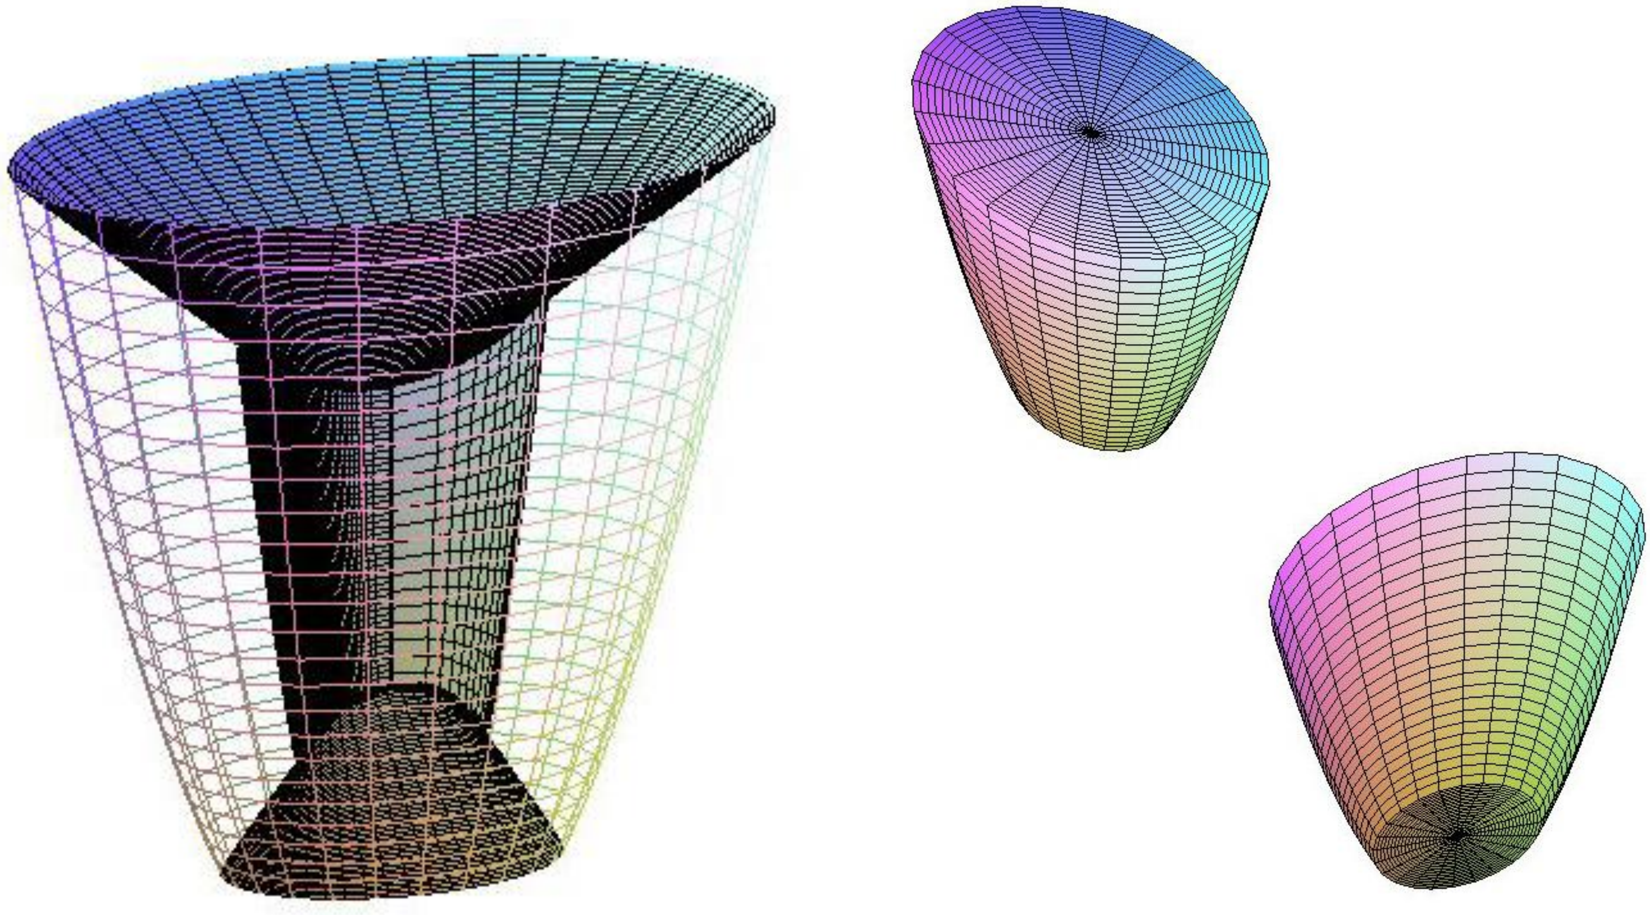

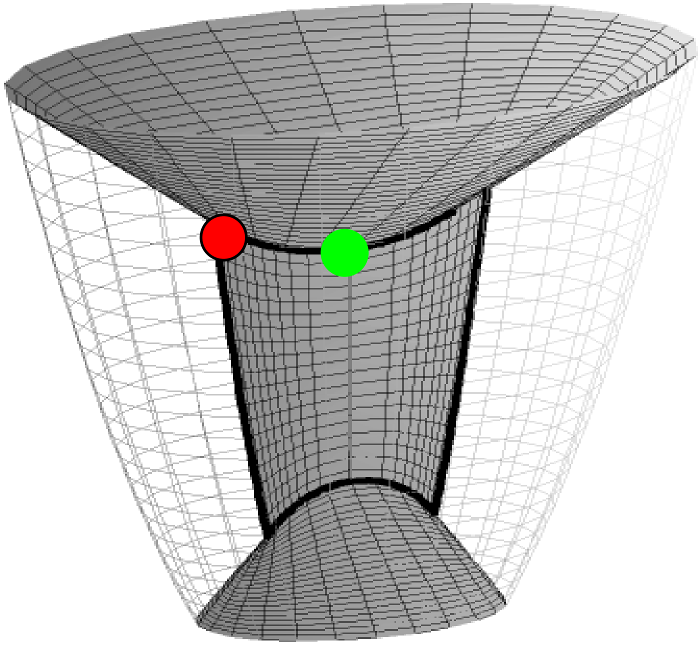

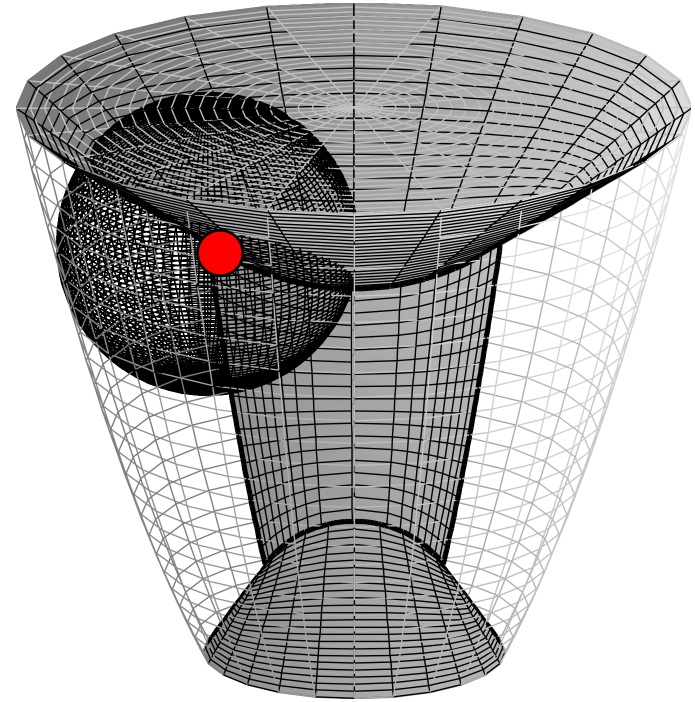

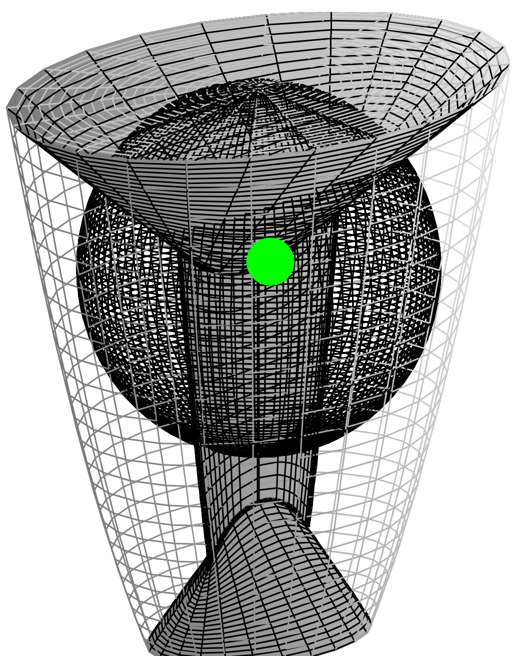

Sphere centred at red ‘fin point’ is tangent to the top and at a ridge point on the left (an A_1A_3 sphere)

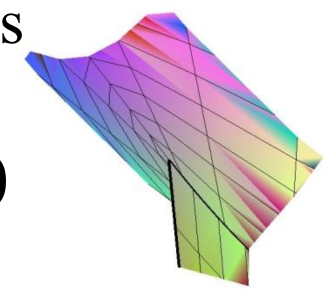

Sphere centred at green ‘Y-junction point’ is tangent in three places: top, front, back. (An $A_1A_1A_1$ or A_1^3 sphere.)

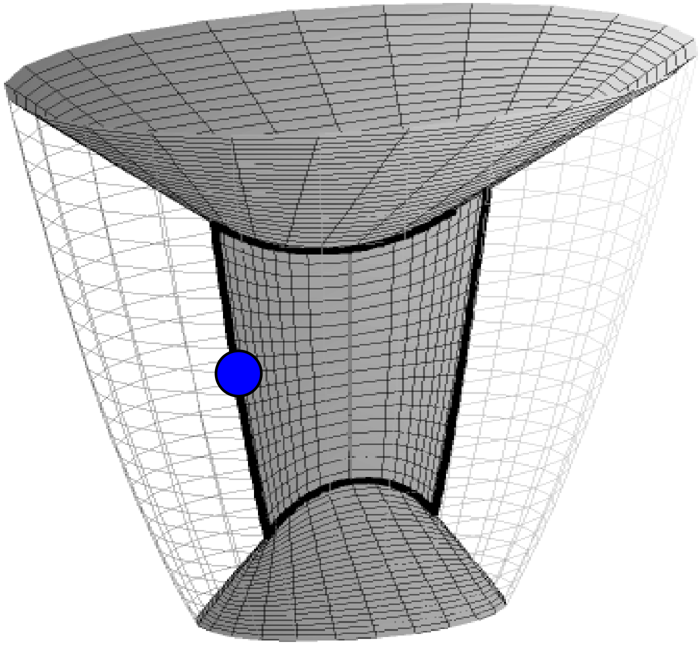

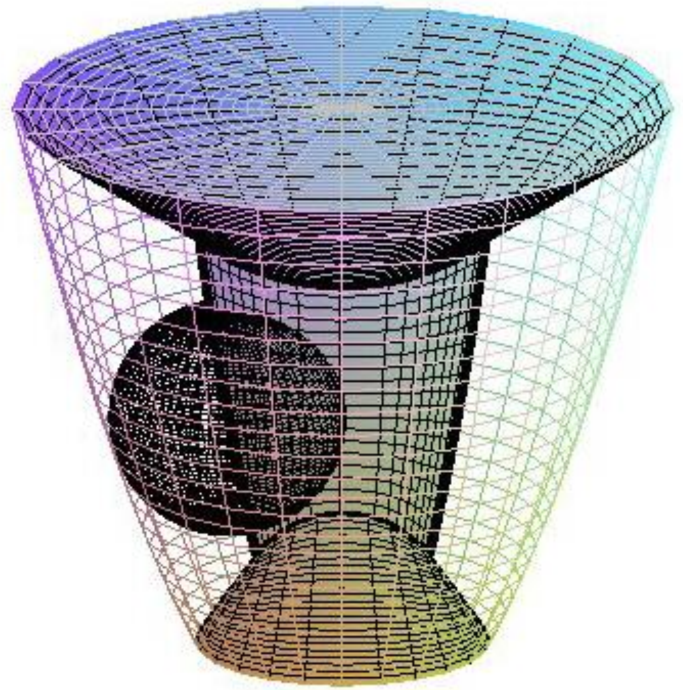

A sphere centred at a blue ‘edge point’, tangent to the bin at a ridge point but not elsewhere. An A_3 sphere.

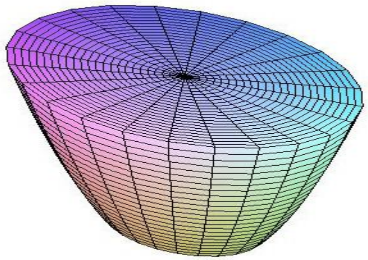

If we squash the bin…so that a sphere can fit inside touching top, bottom, front and back....

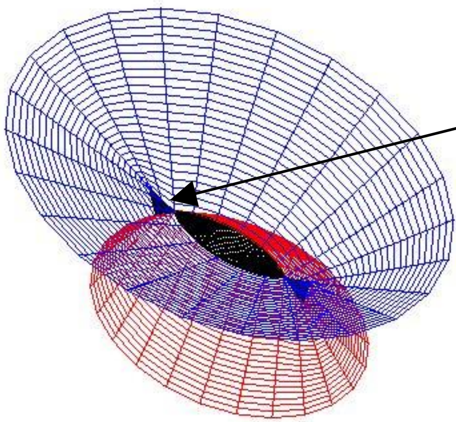

Centre of a sphere tangent in 4 places (an A_1^4 sphere)

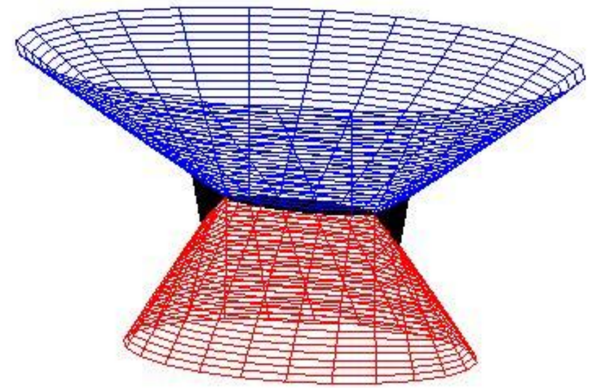

The new medial axis

This is one of 7 ‘transitions’ on the medial axis in 3D, classified initially by Ilya Bogaevsky, subject also of a recent paper (TPAMI 2009) by Kimia and myself.

Another example: squeezing a tube. Only the curves of the medial axis (Y-junction = A_1^3 red, A_3 = blue) shown here.

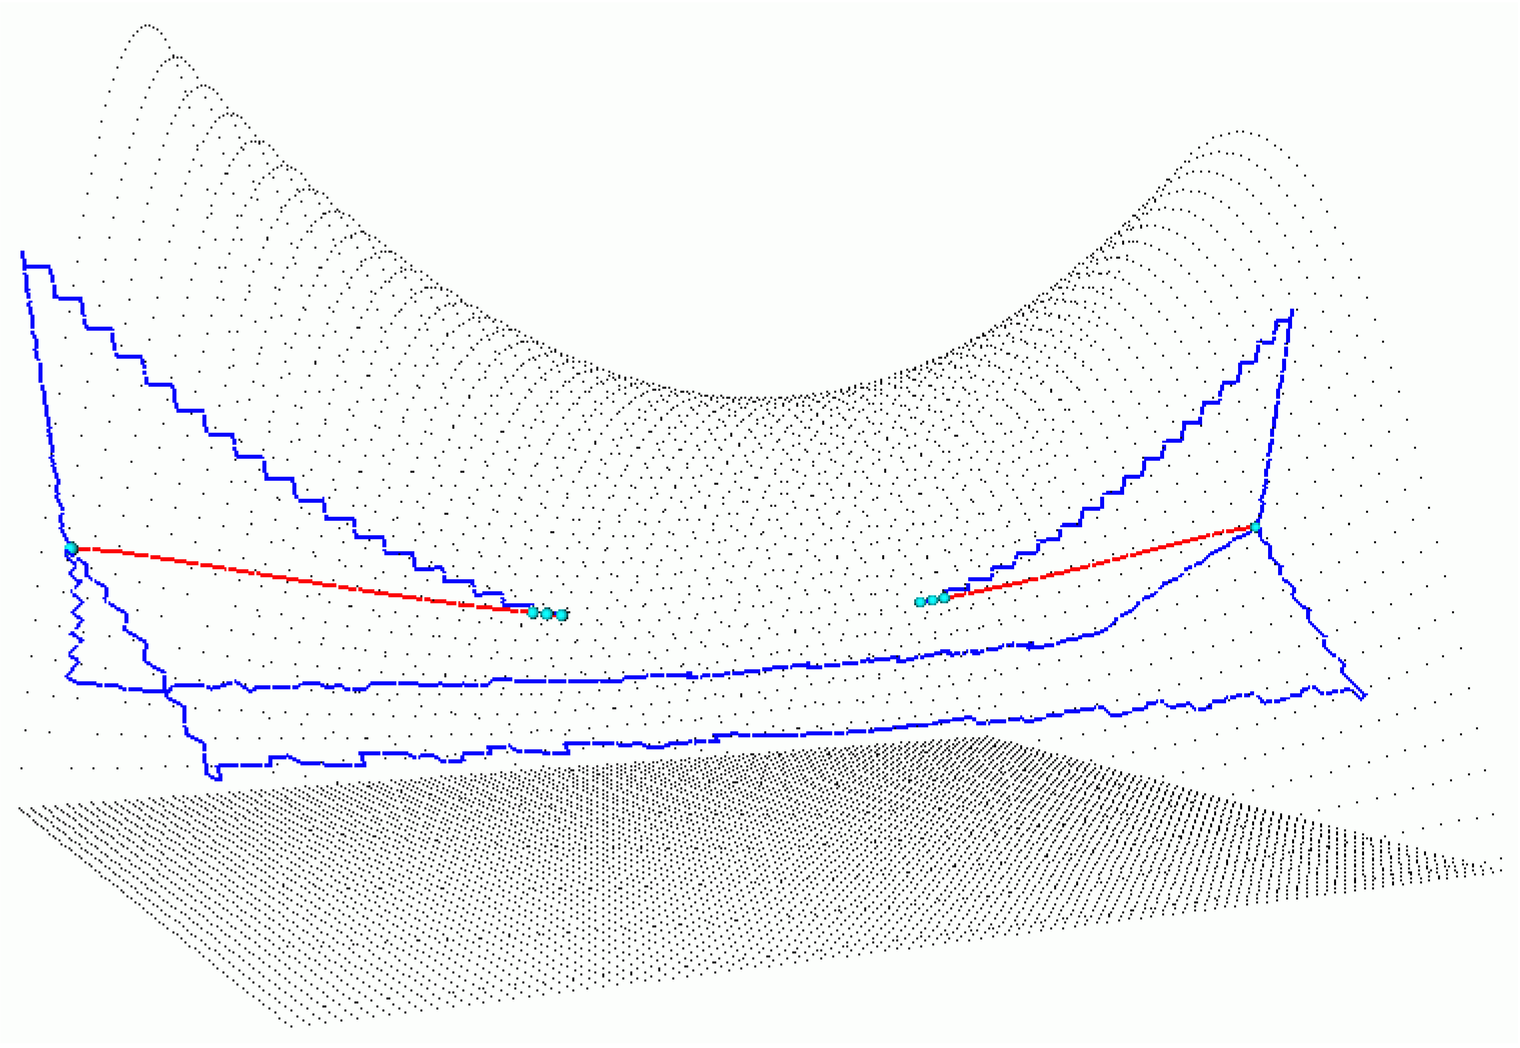

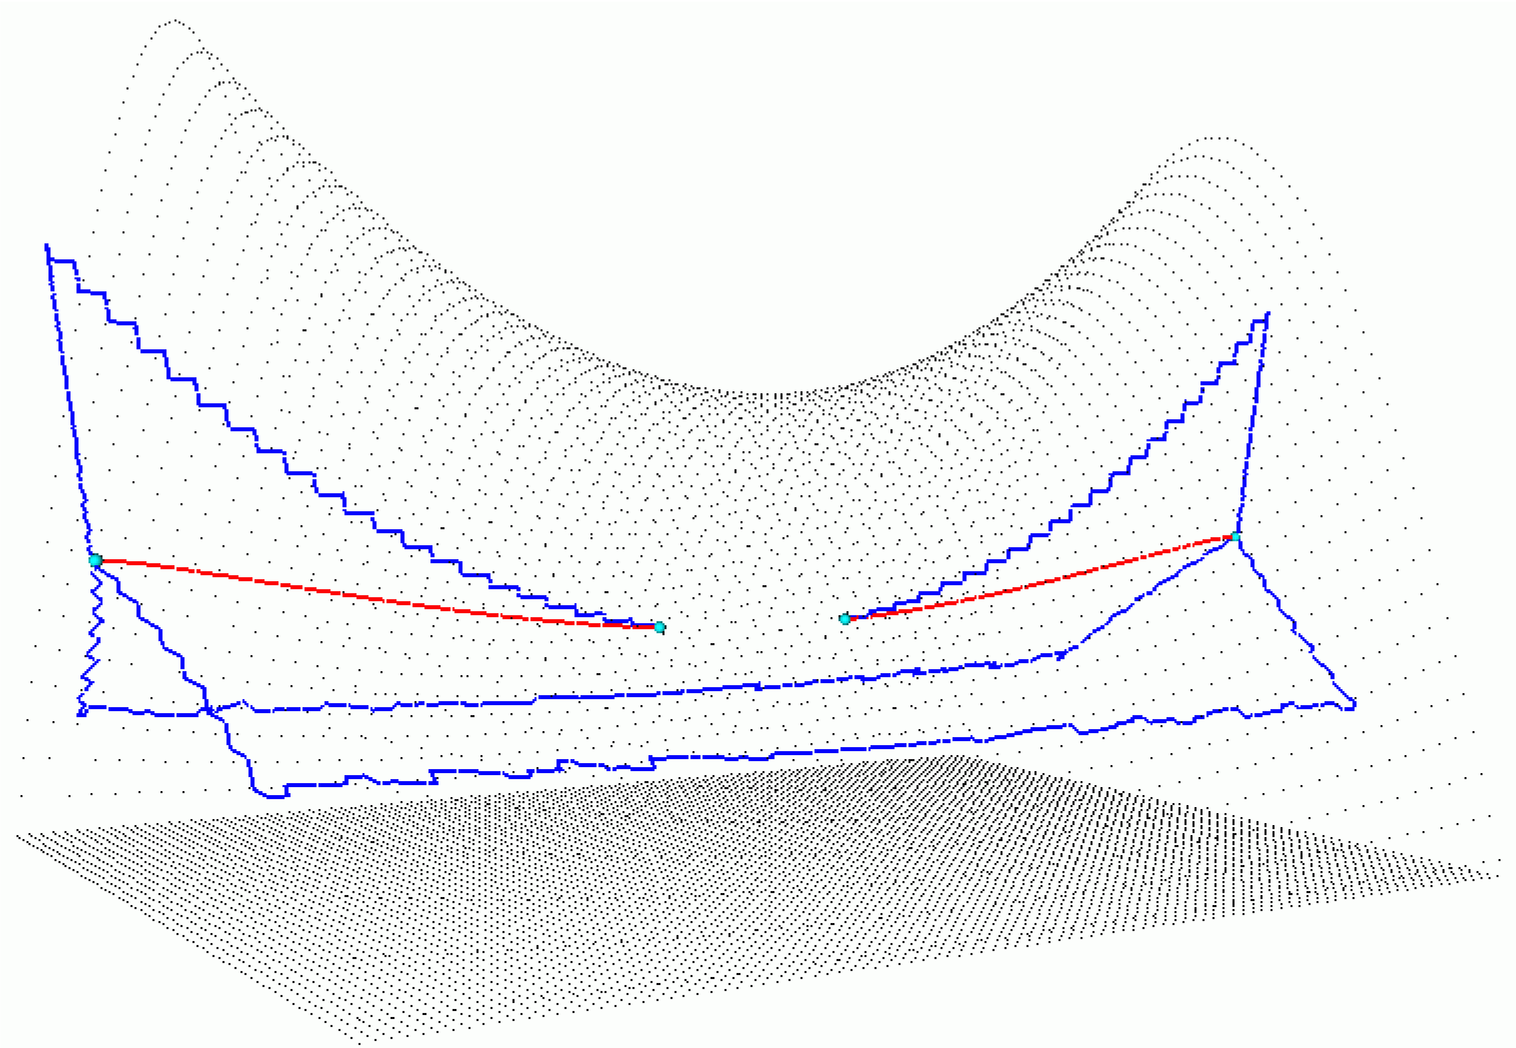

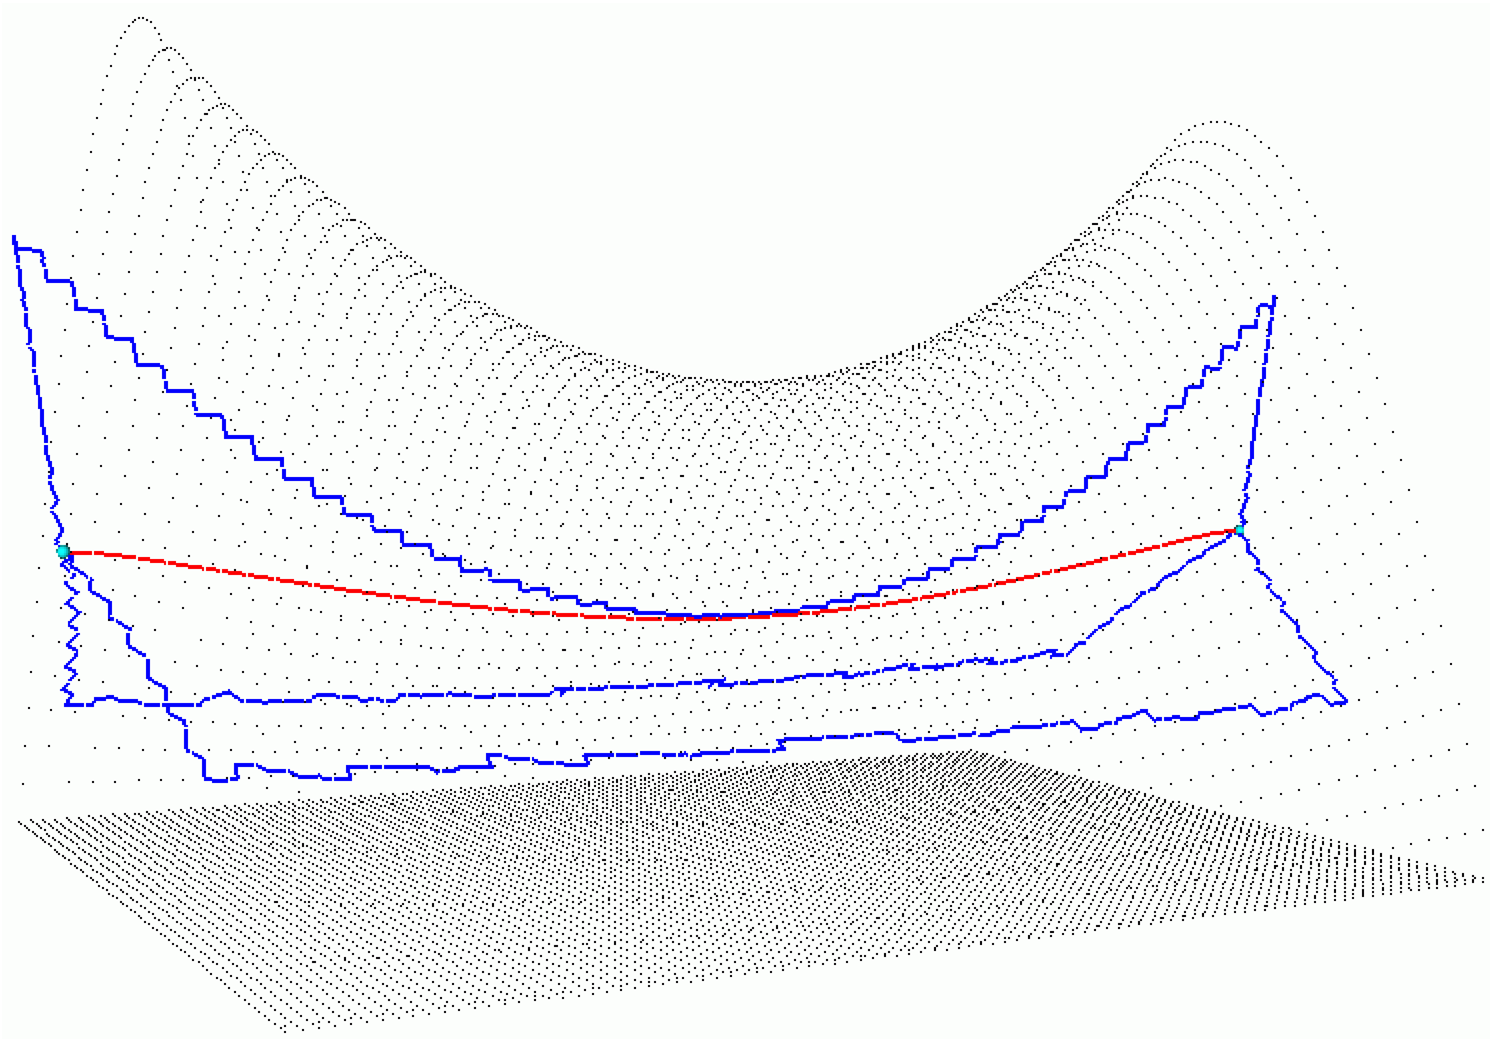

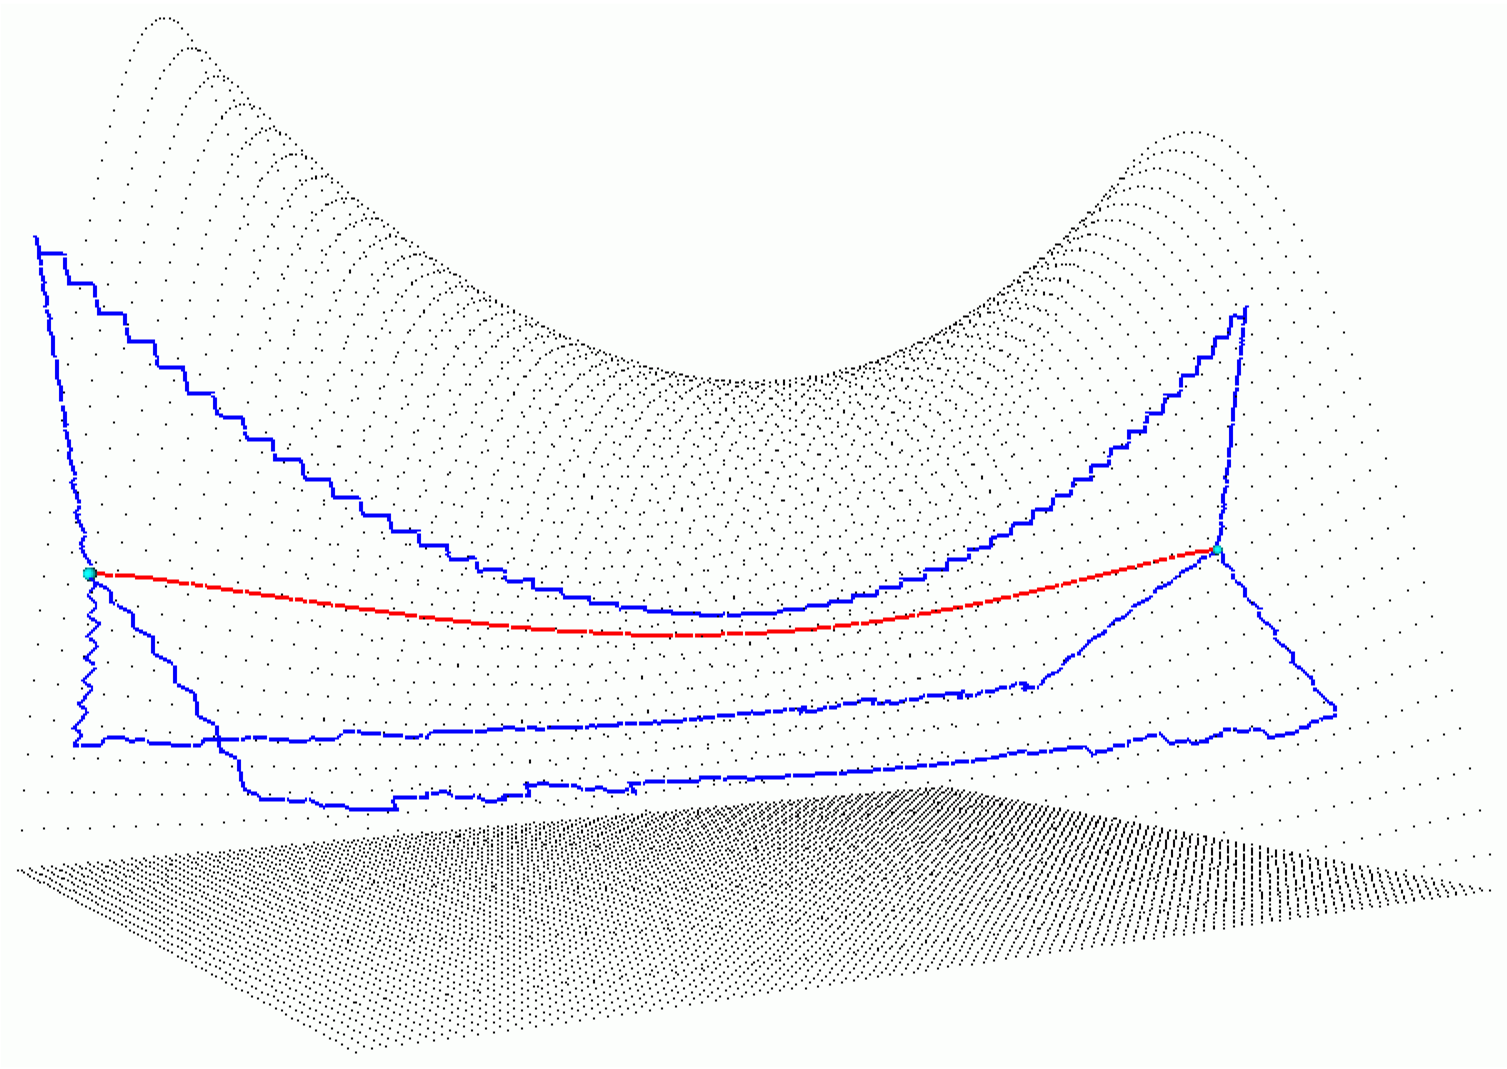

Figures by Frederic Leymarie and Ming-Ching Chang (Brown U.)

One of the transitions is closely connected with ‘salient features of surfaces’ in this case interaction of two ridge curves (crest lines).

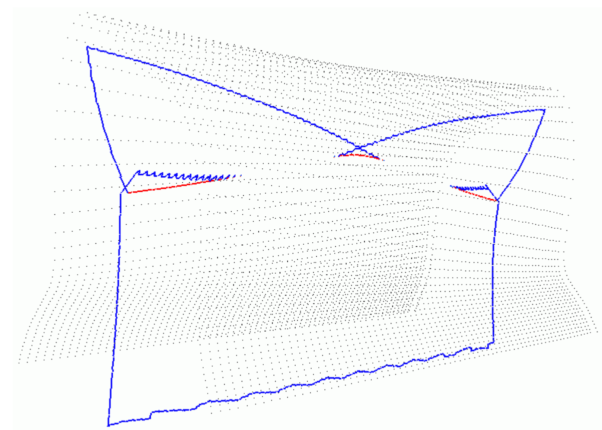

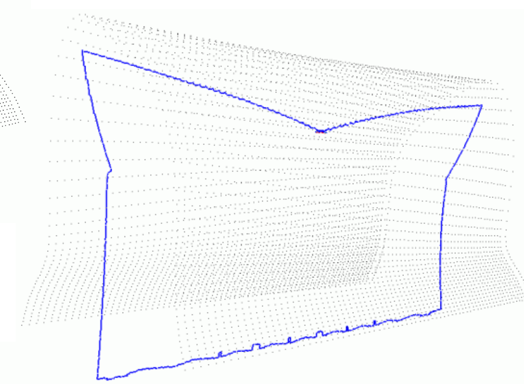

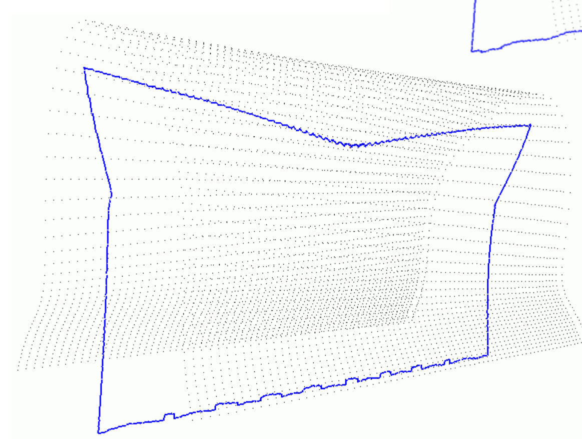

A lot of work has also been done on symmetry measures (medial axes etc.) which are invariant to *affine transformations* (translations + nonsingular linear transformations such as shears, linear magnification etc.)

Some of these constructions are based on *area* rather than *distance*. Some are based on *parallelness*.

Just one example......

Here, for each curve, we take the envelope of chords joining pairs of points where the tangents are parallel (Centre Symmetry Set)

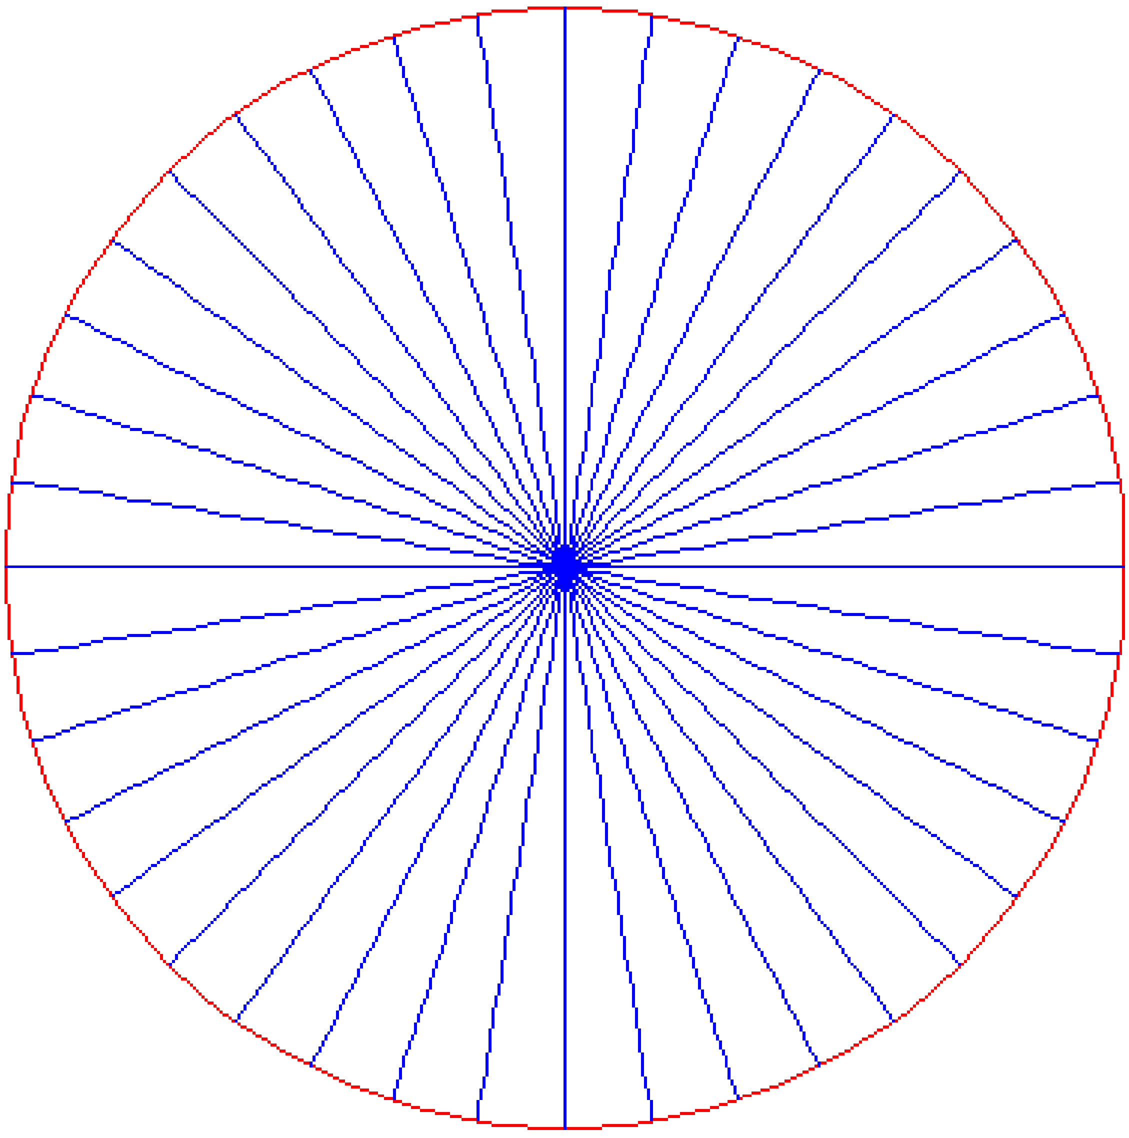

3. Moving views of illuminated surfaces

work in progress with Jim Damon

First paper Int. J. Computer Vision 2008

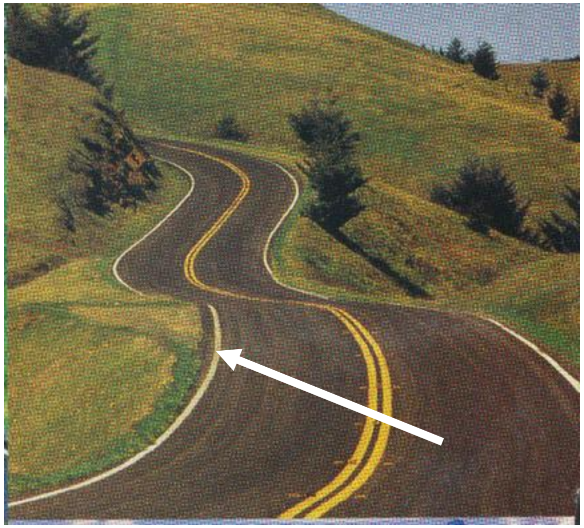
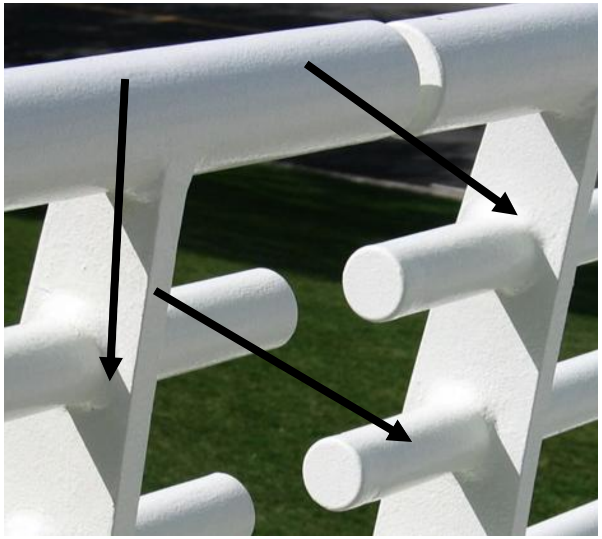
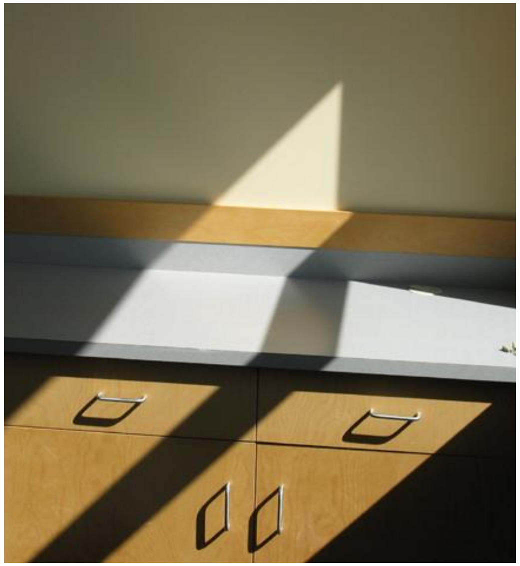

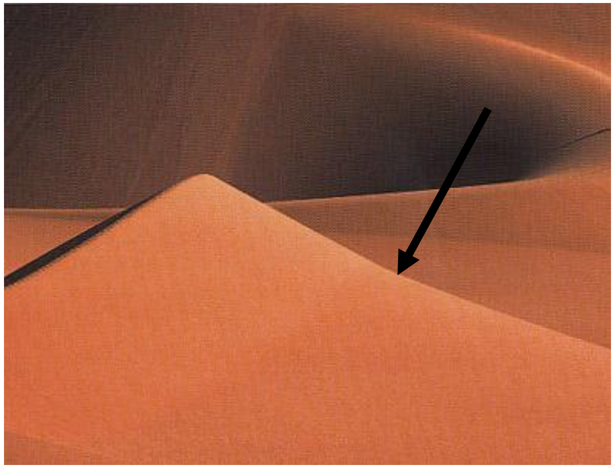


Surface markings, creases (2 surfaces meeting), corners (3 surfaces), shade boundary (terminator), cast shadow, contours (profiles)

One rather discouraging fact is the difficulty in extracting edges and especially junctions in an image.

So the kind of subtleties which we are classifying may not as yet be very practical.

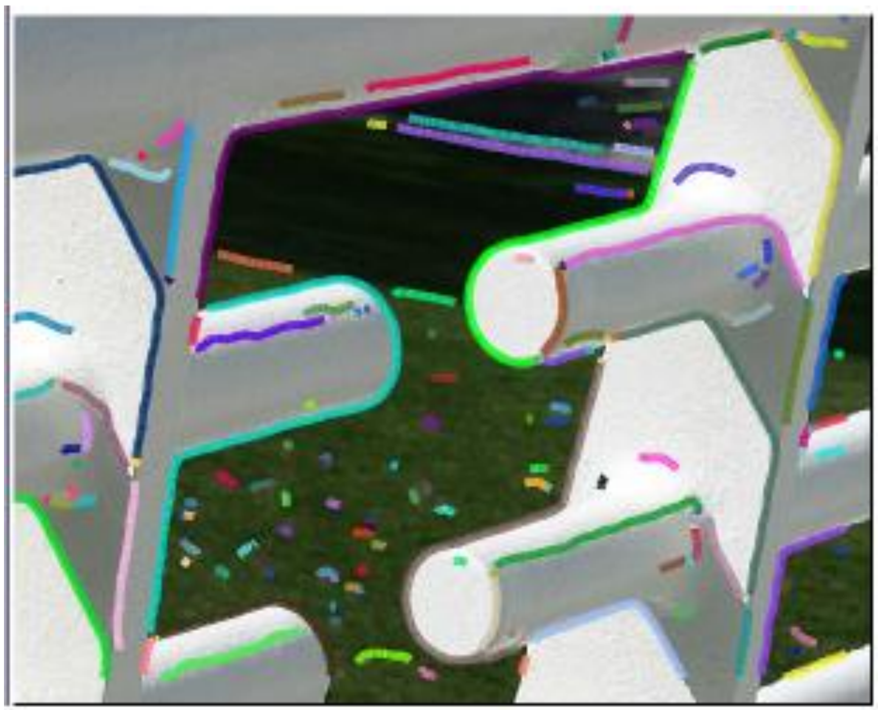

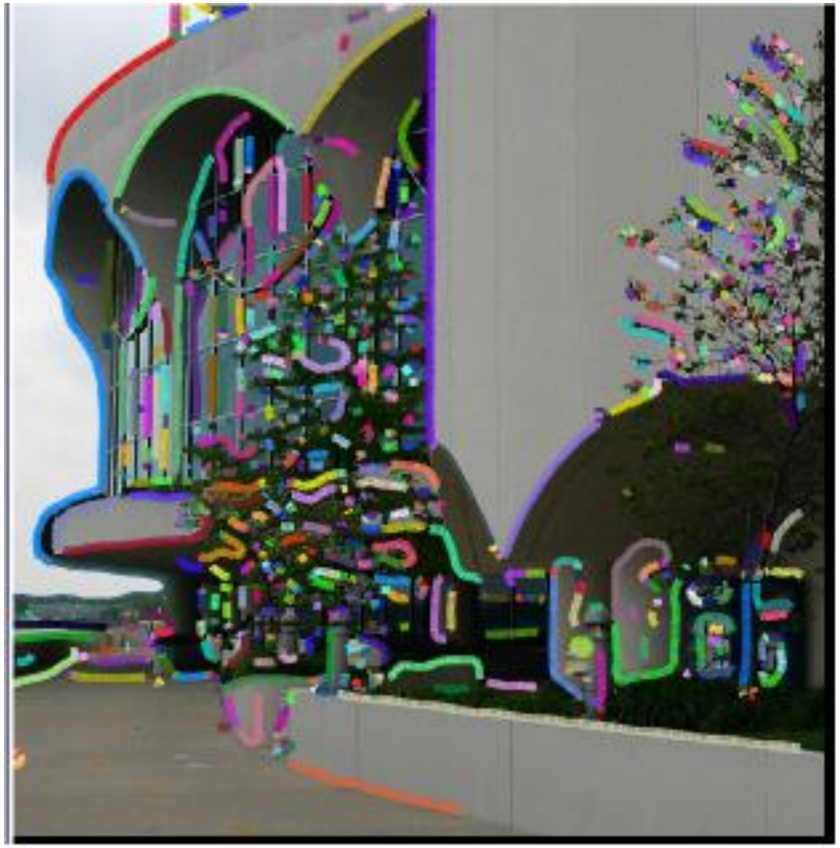

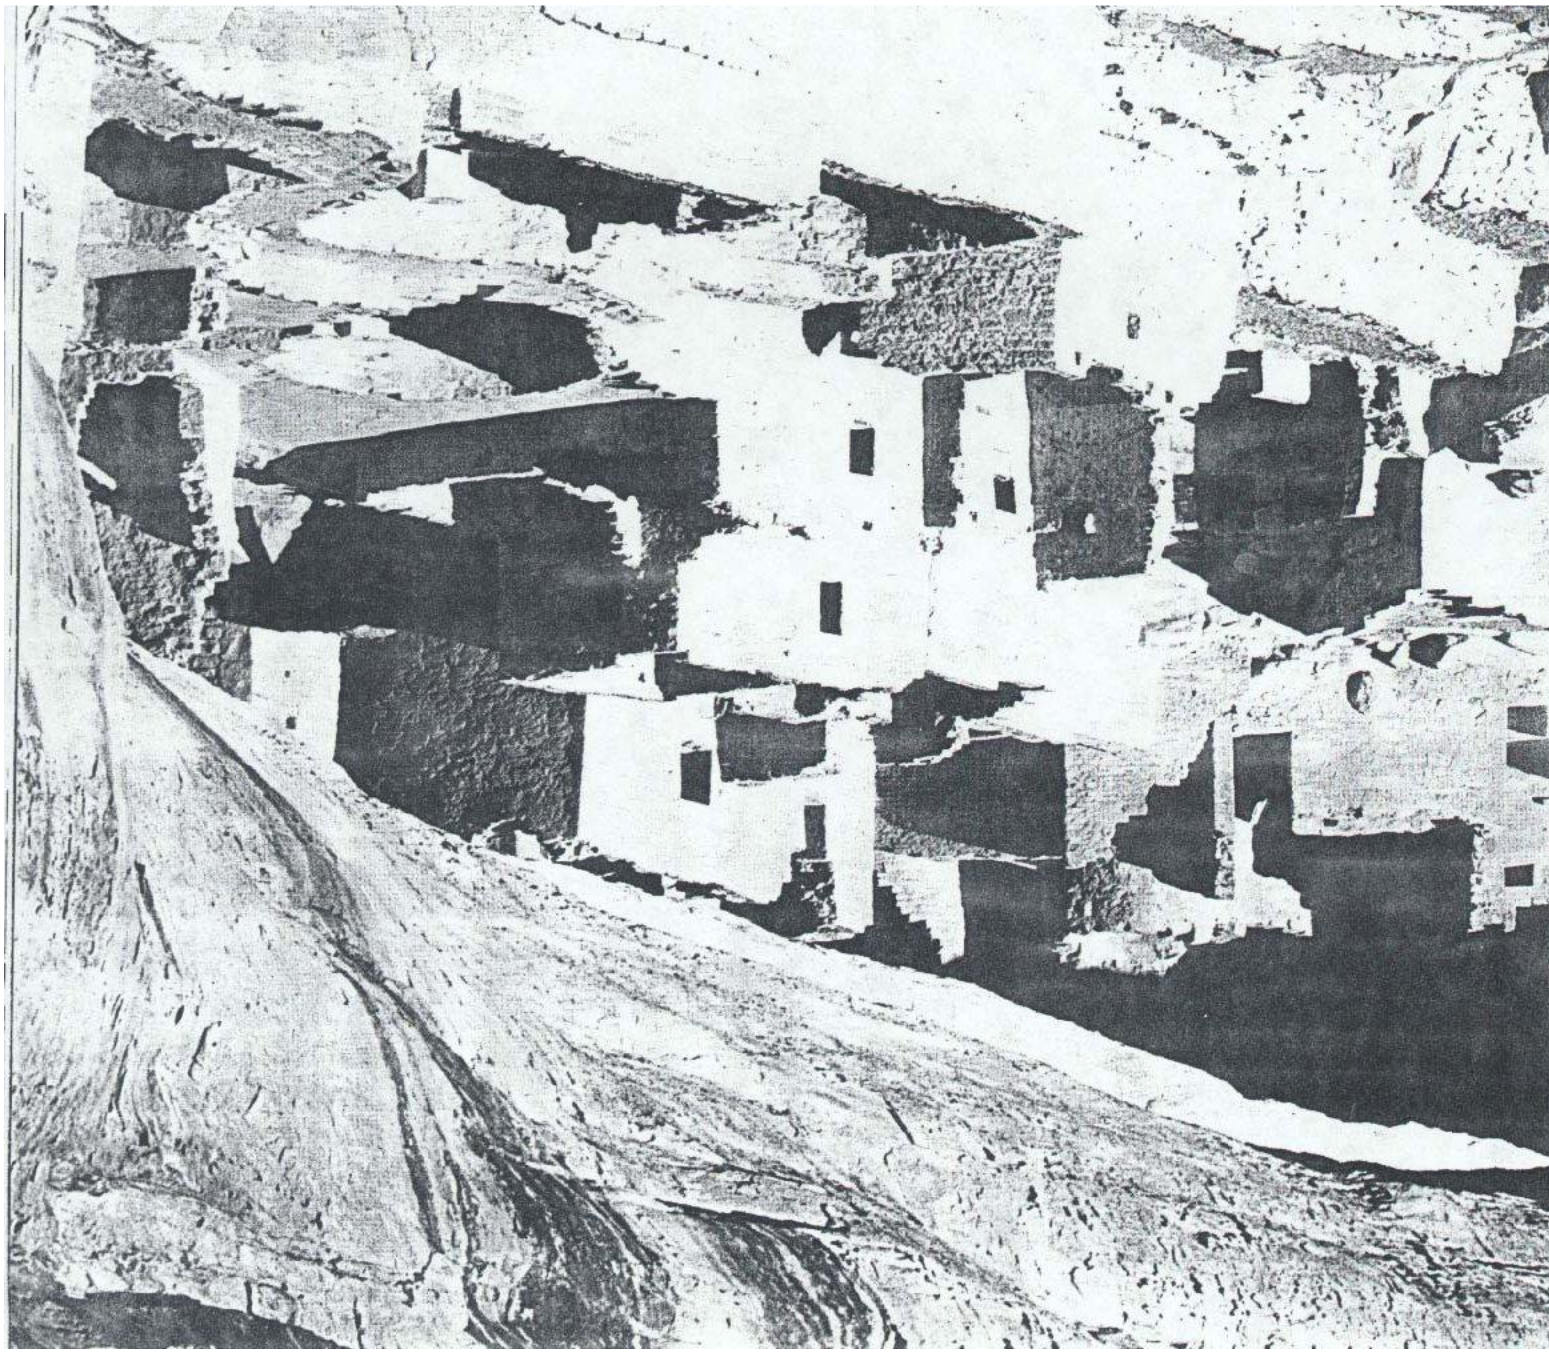

Problem here is to classify all the ways in which

- surface **markings, creases, corners, boundary edges**
- **shade lines** (‘terminators’) and cast shadows
- **contours** (profiles)

can interact (locally) in an image, as the observer ‘flies past’ the scene

The objective is to classify the **visual clues to surface shape** which are available in an image from a ‘1-parameter family’ of viewpoints, that is a fly-past, **just from the above 1-dimensional visible features of the image.**

It looks as if there will be a zillion cases so **singularity theory** is used to sort them out in a sensible way, and to make sure nothing is missing. The singularities here arise again from projecting surfaces into the viewplane.

The setup: One (major) source of illumination, providing the shade curves (gradual 'shading off') and cast shadows (sharp)

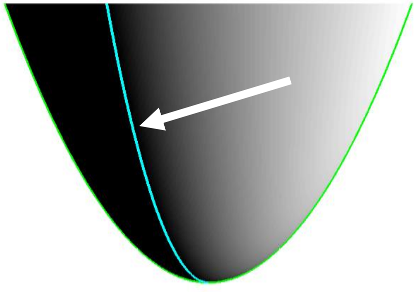

Illumination is *stable*, from a point source far away

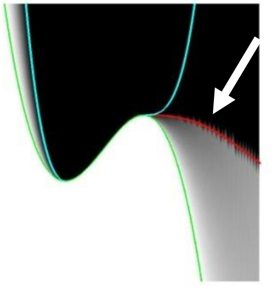

Surfaces are 'piecewise smooth', that is two or three surfaces can meet in a 'crease' or 'corner'

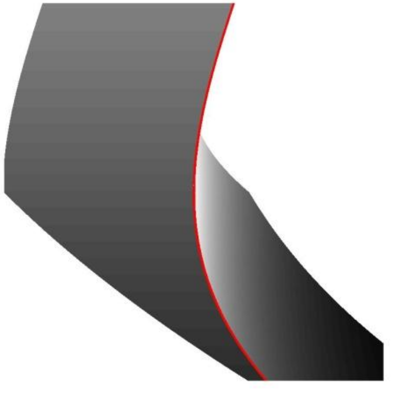

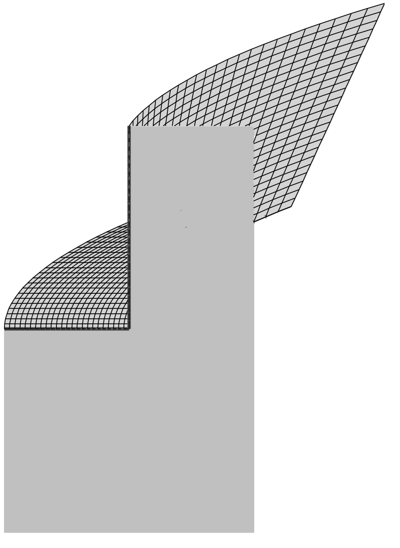

Just to put this in a mathematical context, for a single smooth illuminated surface.

We have a surface M, locally given by say $z = F(x, y)$, and two projection maps

$$\mathbf{R}^2 \overset{\psi}{\leftarrow} M \overset{\phi}{\rightarrow} \mathbf{R}^2$$

where the left-pointing projection is a viewing projection and the right-pointing one looks in the direction of the illumination.

We assume ϕ is stable (submersion, fold or cusp) and classify projections ψ *up to an equivalence which preserves the features created by the illumination, together with any boundary edges and surface markings.*

There is an 'alphabet' of stable interactions of features, shade/shadow and contour. It is divided into 'hard' curves like creases, markings and contours and 'soft' curves like shade boundaries and cast shadows.

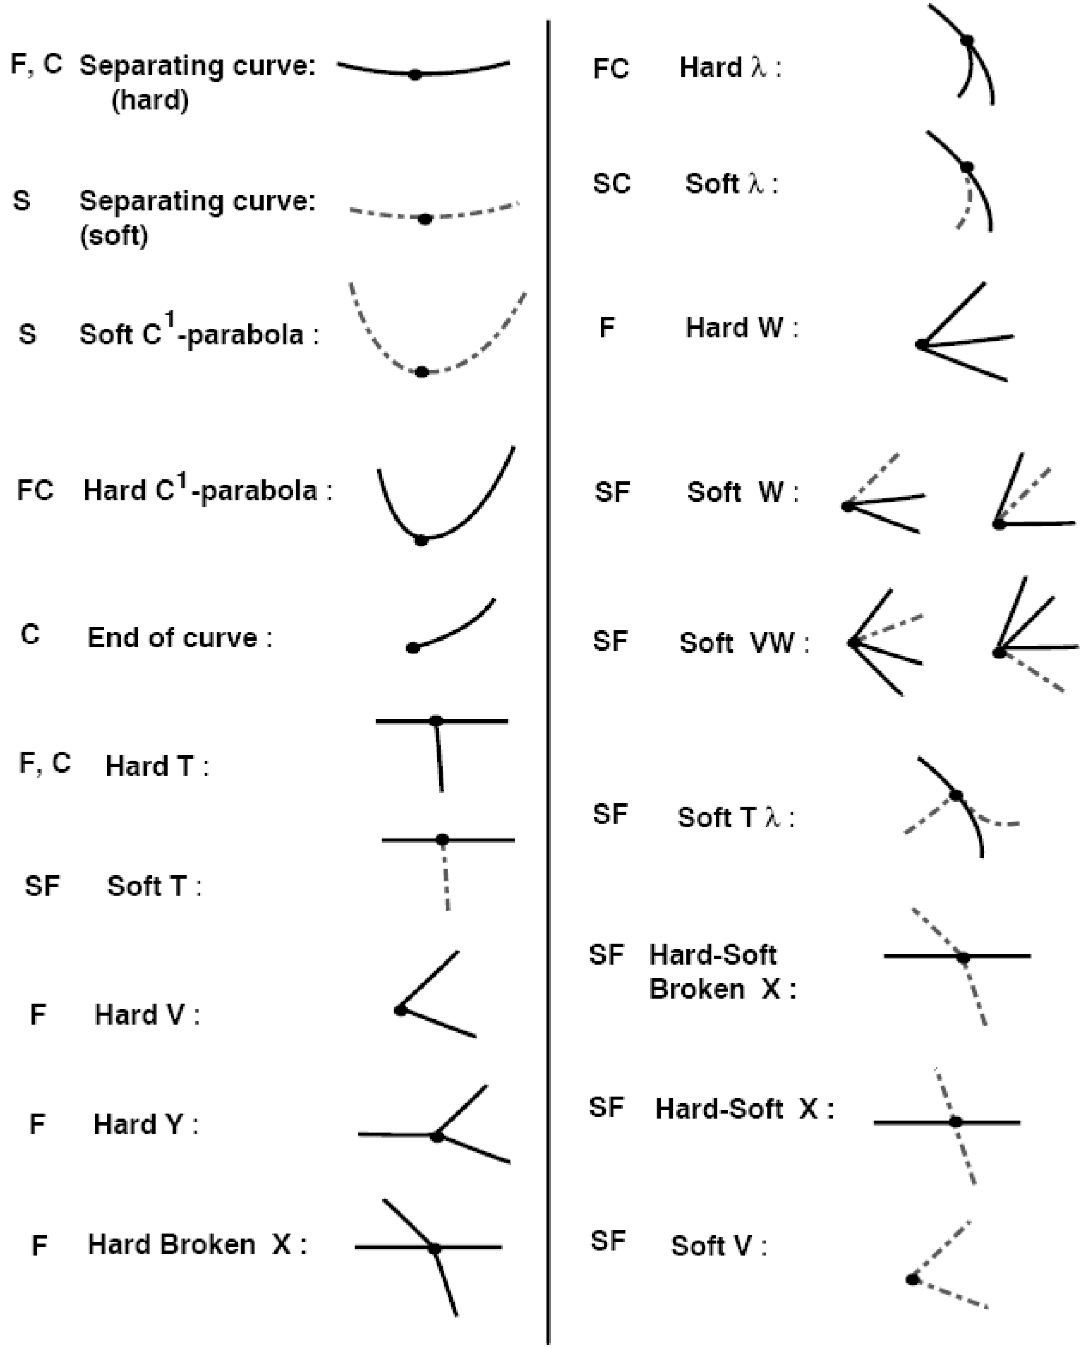

Here are a couple of examples of these.

FC Hard λ :

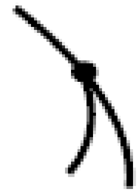

This would be for a surface marking

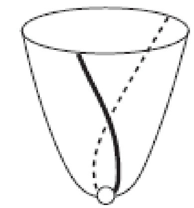

SC Soft λ :

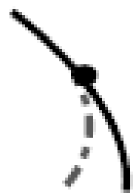

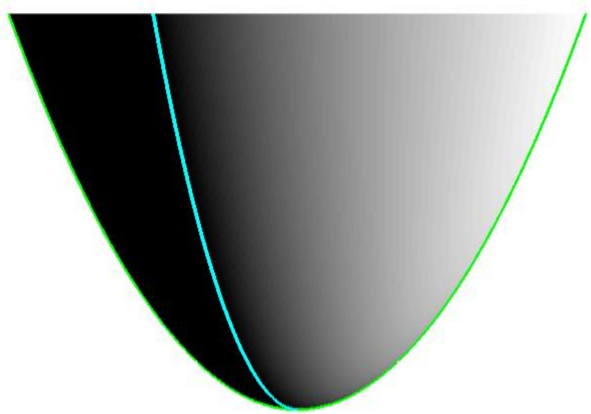

F, C Hard T :

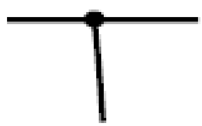

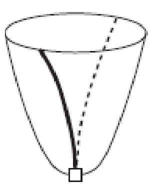

SF Soft T :

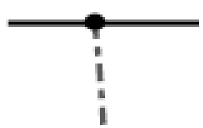

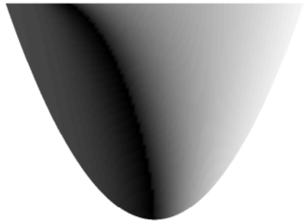

shade curve (terminator) and apparent contour

Another one of the alphabet:

This is two curves meeting with a common tangent line but different curvatures; hence rather harder to detect!

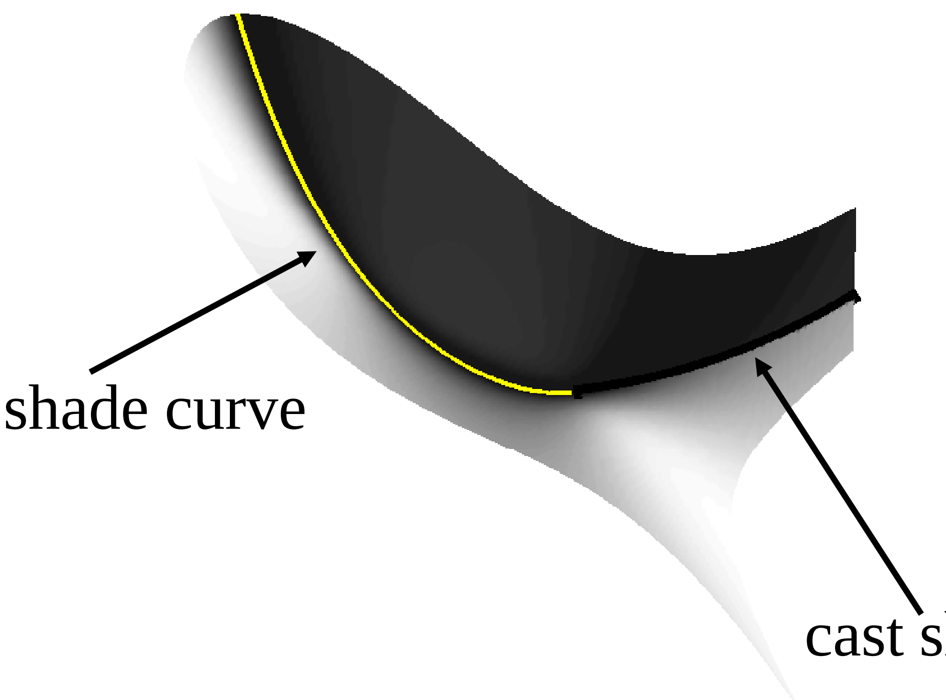

It occurs for when a cast shadow and a shade curve (terminator, emphasized in yellow here) meet locally. So this surface is casting a **local** shadow on itself.

(This happens when the projection in the light direction is a cusp map.)

To turn briefly to **corners:** there are four basic types

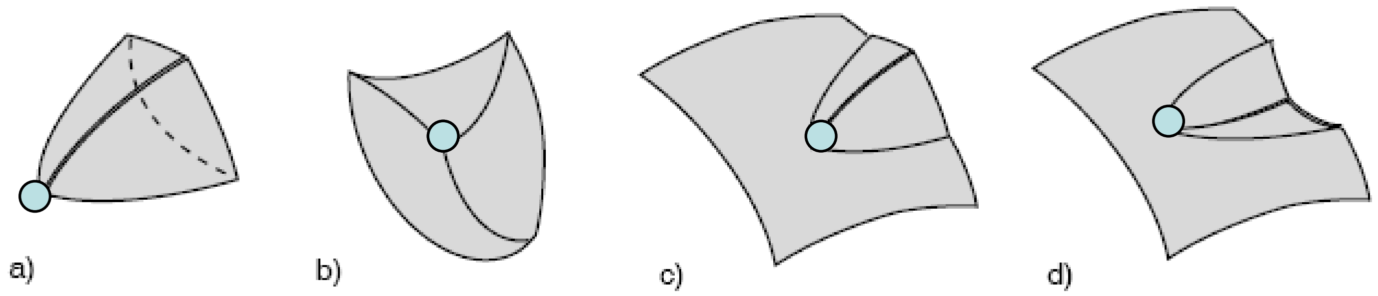

convex concave saddle notch

To turn briefly to **corners** it is not possible to tell if this corner is concave or convex

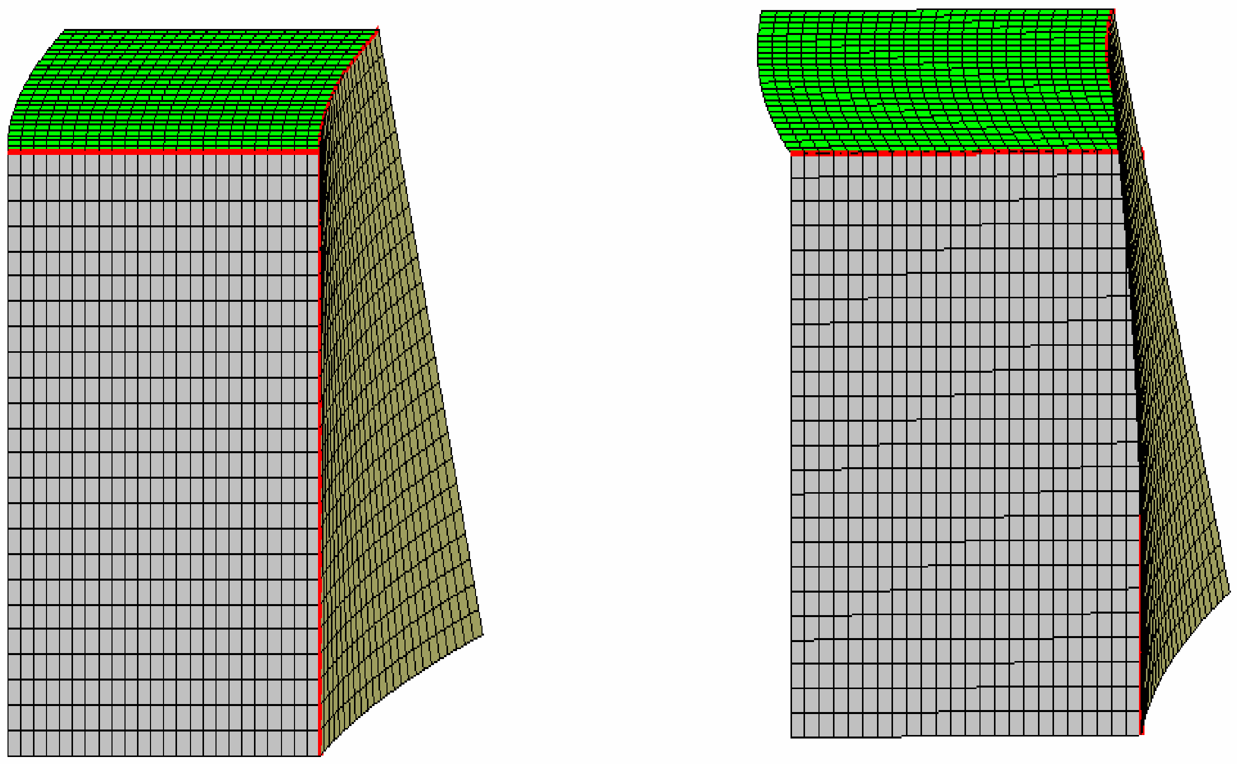

but if you move the viewpoint slightly, the contour tells you that this is a concave corner

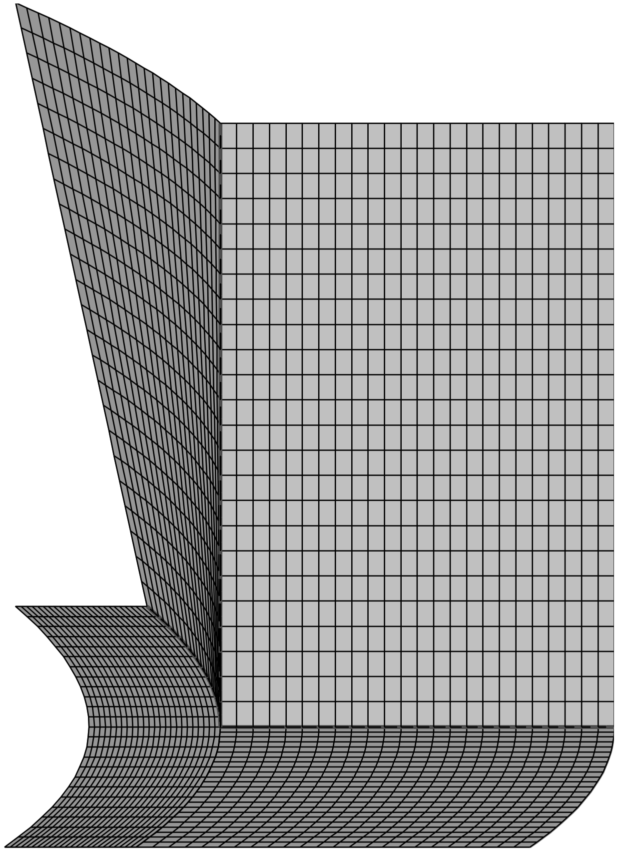

Likewise this could be a saddle or a notch

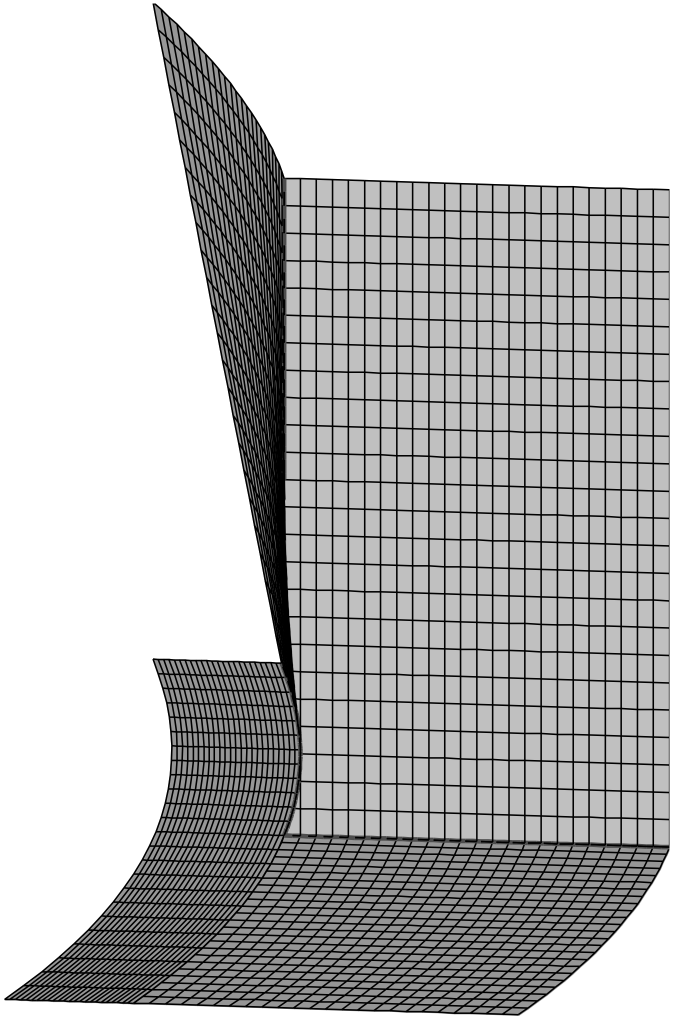

but moving the view reveals it to be a notch

It is also possible for shade boundary *S* (terminator), cast shadow *SS*, countour *C* and feature (crease, *Cr*) all to interact at once during a fly-past

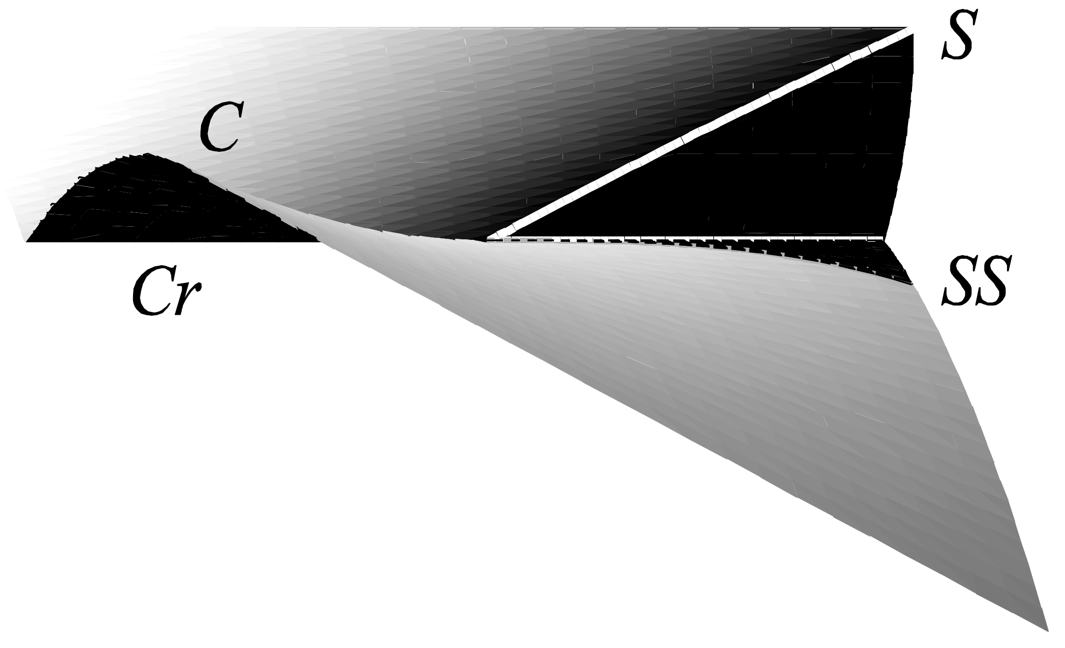

The complete list of cases and their transitions in a ‘fly-past’ can be compiled and it gives a significant amount of information about the local shape of a surface.

Integrating this into global shape descriptors is a different matter and that is not something I personally have been involved in.

The End

Thank you for your attention