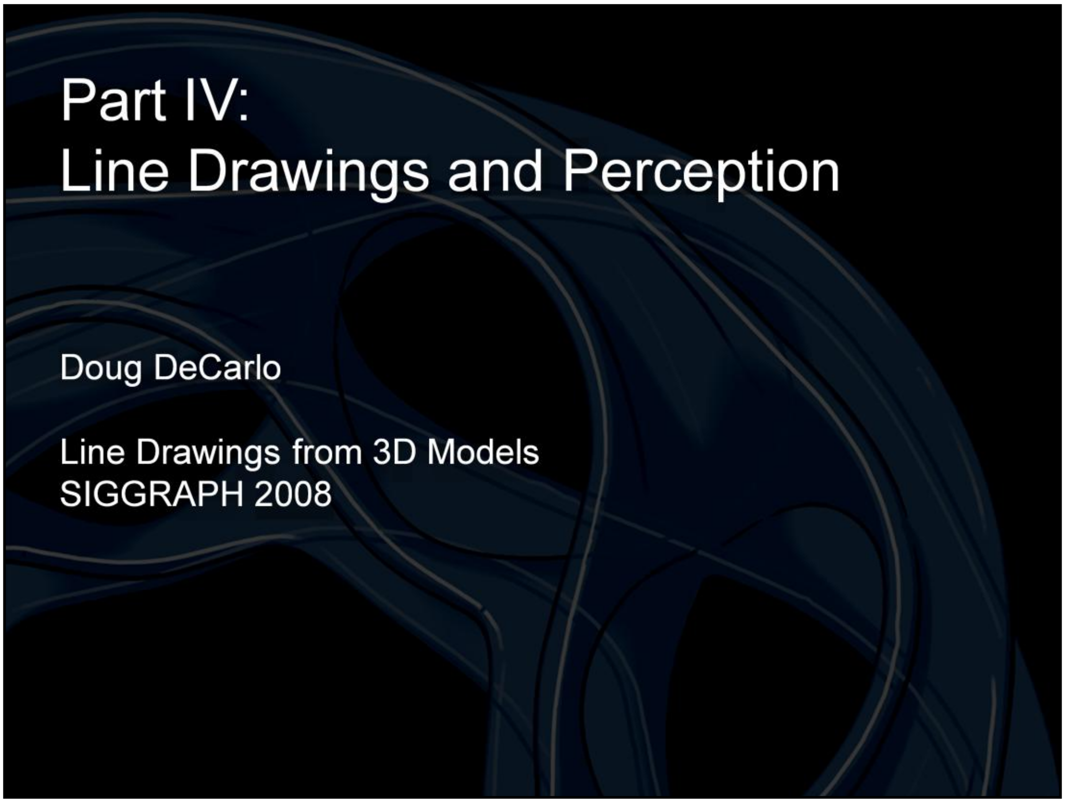

You've heard about the different types of lines that can appear in line drawings.

Now we're ready to talk about how people perceive line drawings.

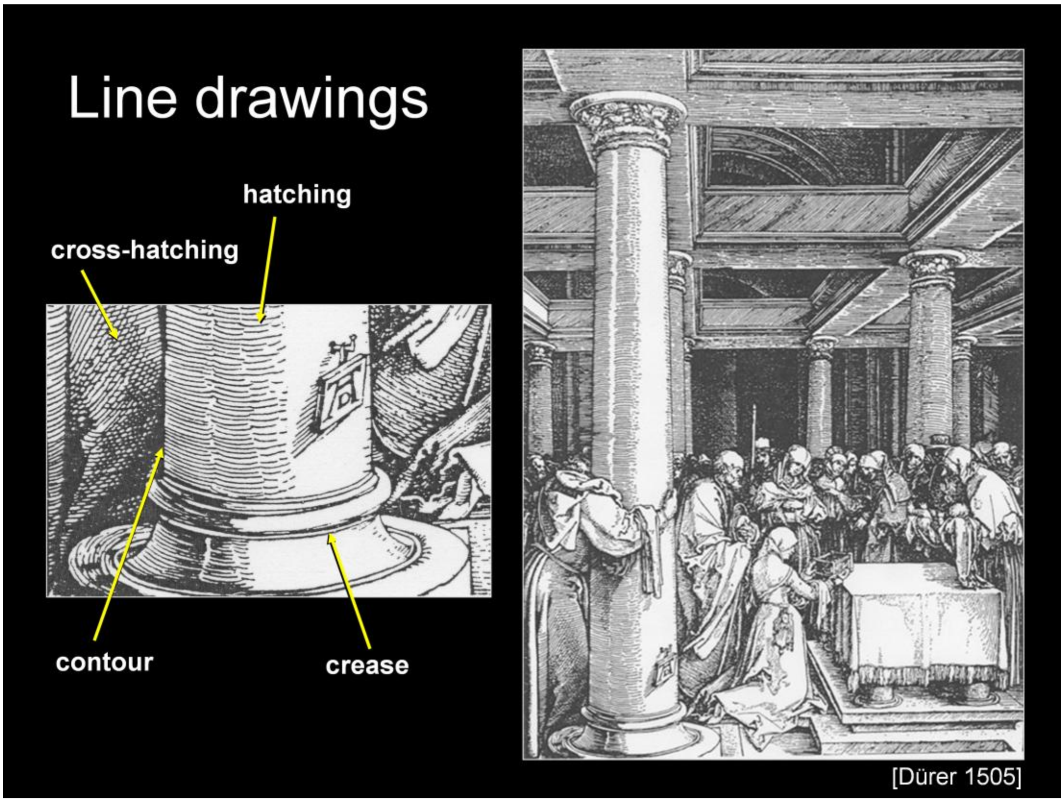

Line drawings bring together an abundance of lines
to yield a depiction of a scene.

Take a look at this print by Dürer.

It uses different types of lines that convey geometry and shading in a way that's compatible with our visual perception.

We seem to interpret this scene easily and accurately.

***

Some of the lines in this drawing only convey geometry.

But the fullness of this drawing comes from Dürer's use of hatching and cross-hatching.

These patterns of lines
convey shading through their local density and
convey geometry through their direction.

Other drawings rely on little or no shading.

In this drawing by Flaxman,
shading is limited to the cast shadows on the floor.

The detail in the cloth here is conveyed with lines
such as contours, creases,
and maybe other lines such as suggestive contours,
or ridges and valleys.

While artists can make drawings like this, they can’t really explain what they’re doing.

They rely on their training,
and use their own perception
to judge the effects of their decisions.

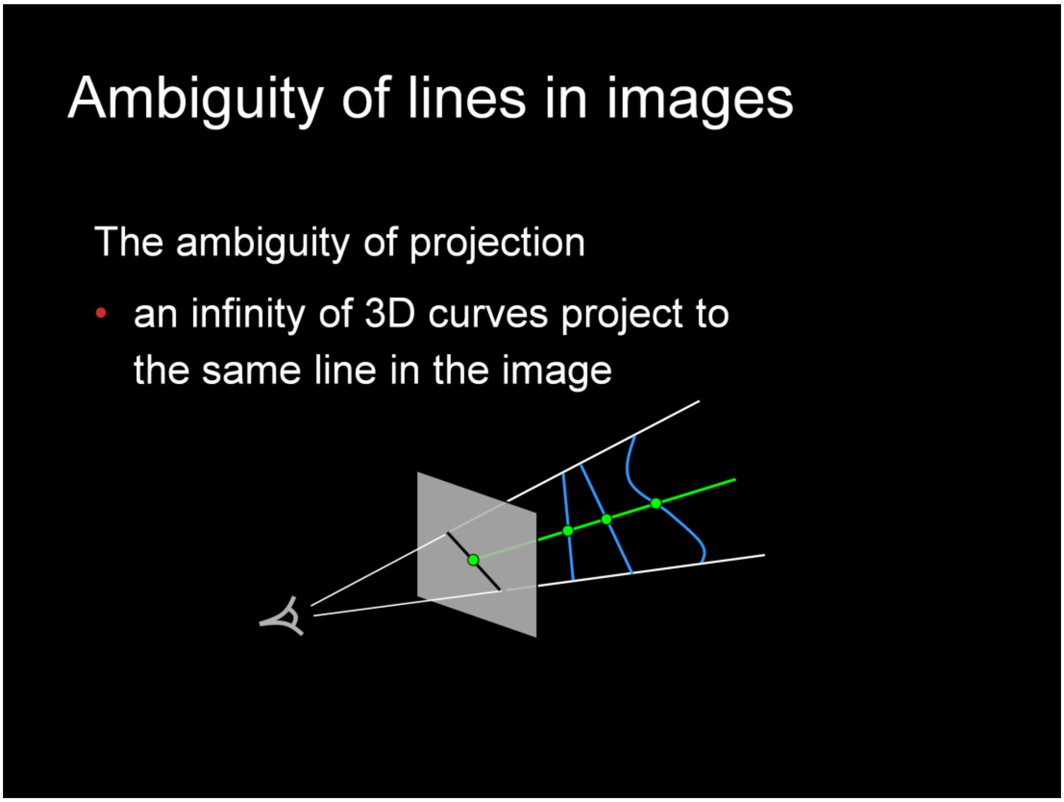

It’s actually a little surprising that line drawings are effective at all.

At first, line drawings just seem to be too ambiguous.

An infinite number of 3D curves
can project to the same line in the image.

All images have this ambiguity, but in photographs,
there are many other visual cues, such as shading and texture,
that help to indicate shape.

Here, we’re just looking at individual lines.

But it turns out that individual lines contain
a wealth of information about shape.

This information is typically local in nature.

But our perception is somehow able to bring all of this together
into a coherent whole.

Well, sort of.

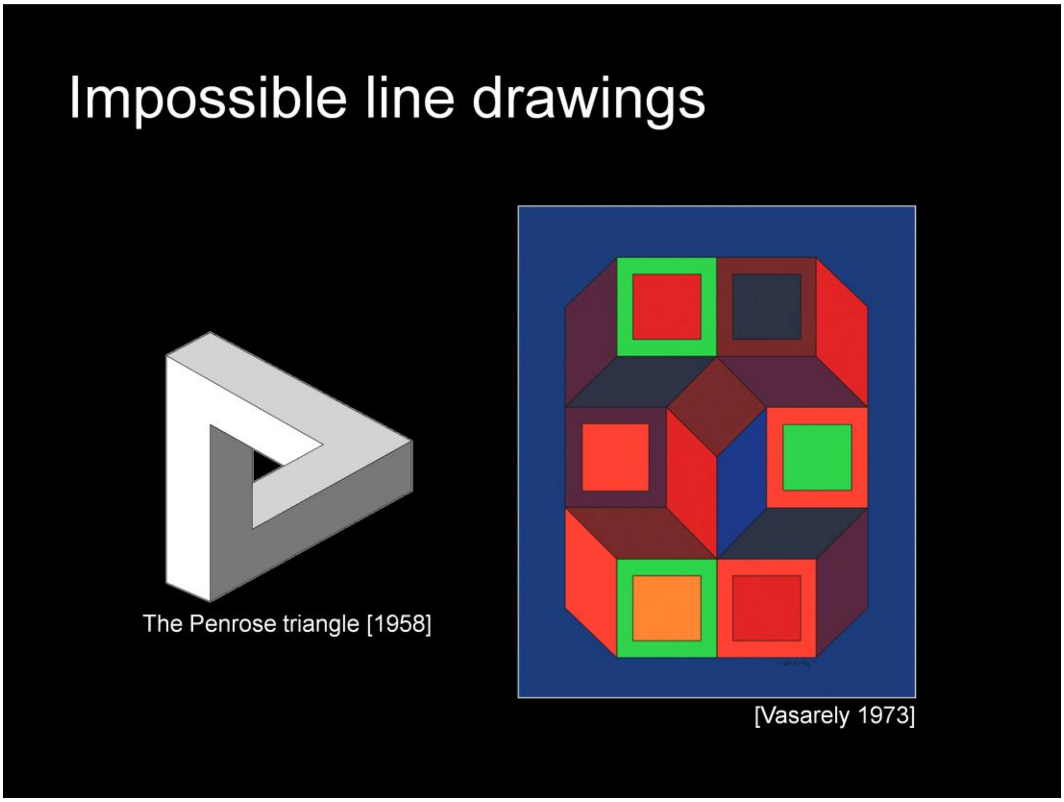

Line drawings of impossible 3D objects show us
that this coherence is NOT global.

The Penrose triangle, which was inspired by the work of Escher,
is perhaps the simplest of the impossible figures.

When you first look at it, it seems to be an ordinary object.

Closer inspection is a little unsettling,
and its inconsistencies are easily revealed.

Vasarely pushed this idea even further,
and made pictures such as this one
that encourage us to explore
several different inconsistent interpretations
at the same time.

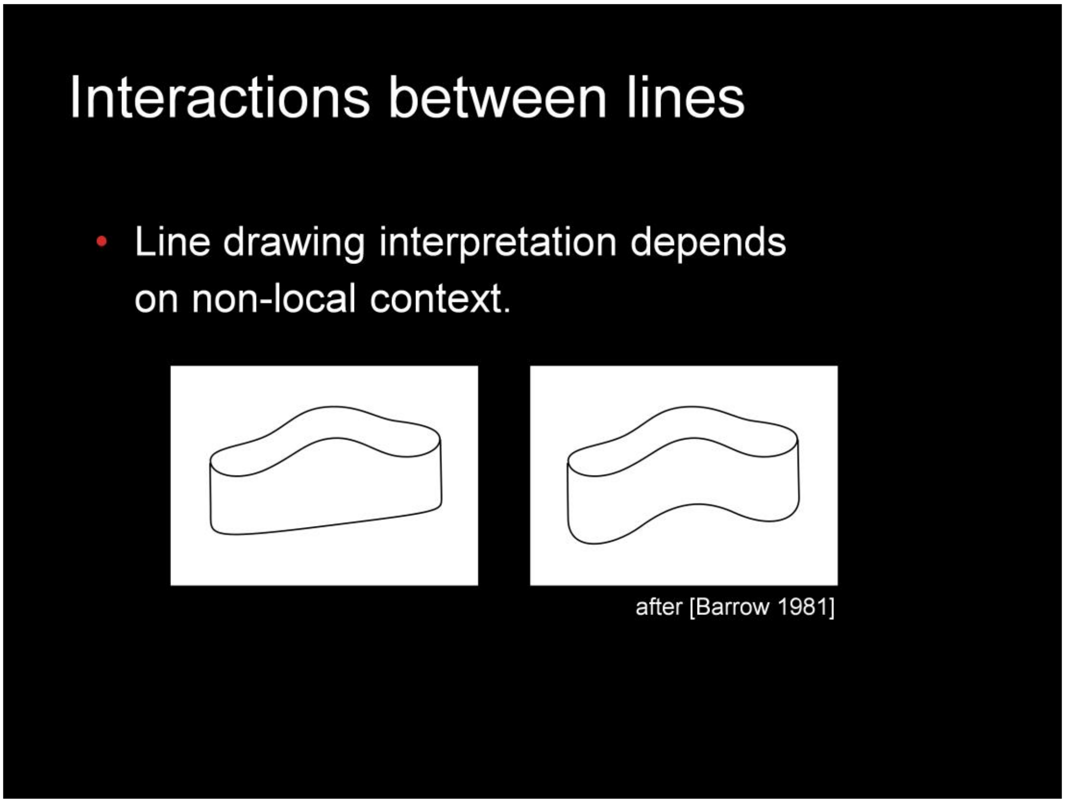

Although you might think the Penrose triangle
shows that there are no global effects for visual inference,
it’s not that easy.

Take a look at these two drawings.

The figure on the left appears to be raised in the center,
while the figure on the right appears to have a flat top,
and bends along its length.

If we compare these two drawings line by line,
the only difference is the line along the bottom.

Nobody knows whether we see this difference
because we consistently integrate local information,
or perform certain types of non-local inference.

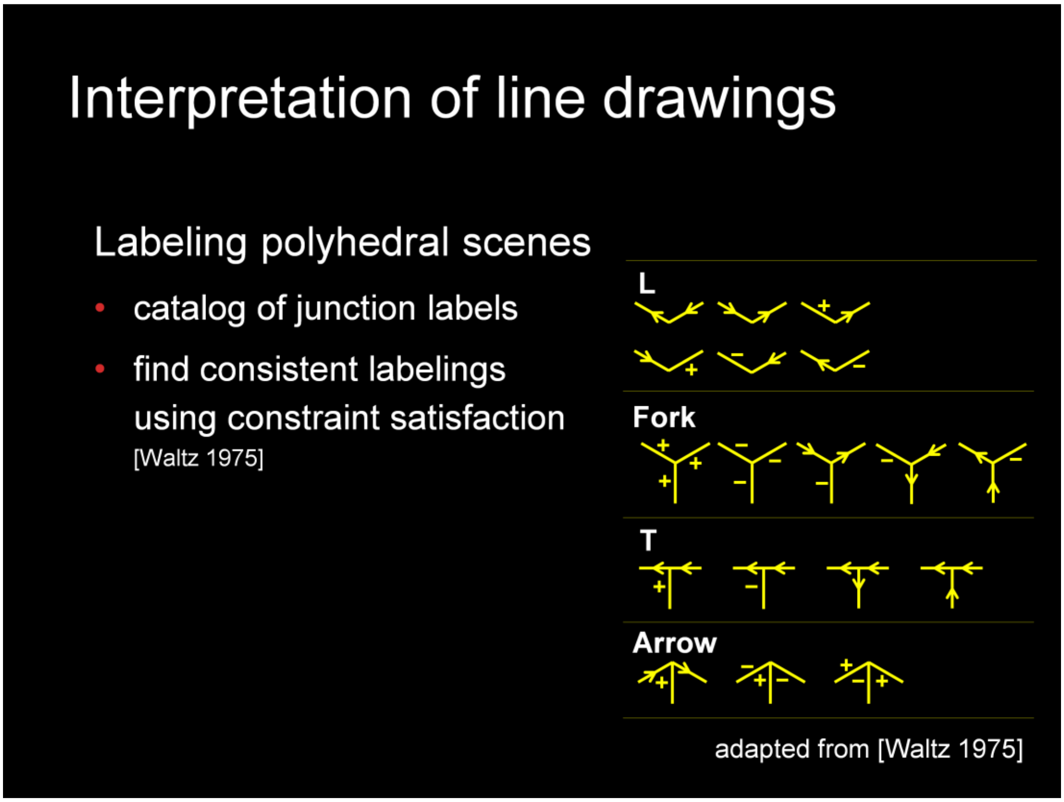

Use of non-local inference is plausible.

Algorithms exist for searching among the space of possibilities.

Waltz's method for line-labeling starts with catalogs of all possible line junctions, which are places where two or more lines meet.

Here's the catalog of 18 junctions that lets you classify any trihedral vertex in a polyhedral scene.

CONVEX lines are labeled with a PLUS, CONCAVE lines with a MINUS.

Arrows mark visual occlusions, where the closer surface is to the RIGHT of the arrow.

Algorithms for constraint satisfaction compute all possible configurations of junctions for a particular picture.

For an impossible figure, this set is empty.

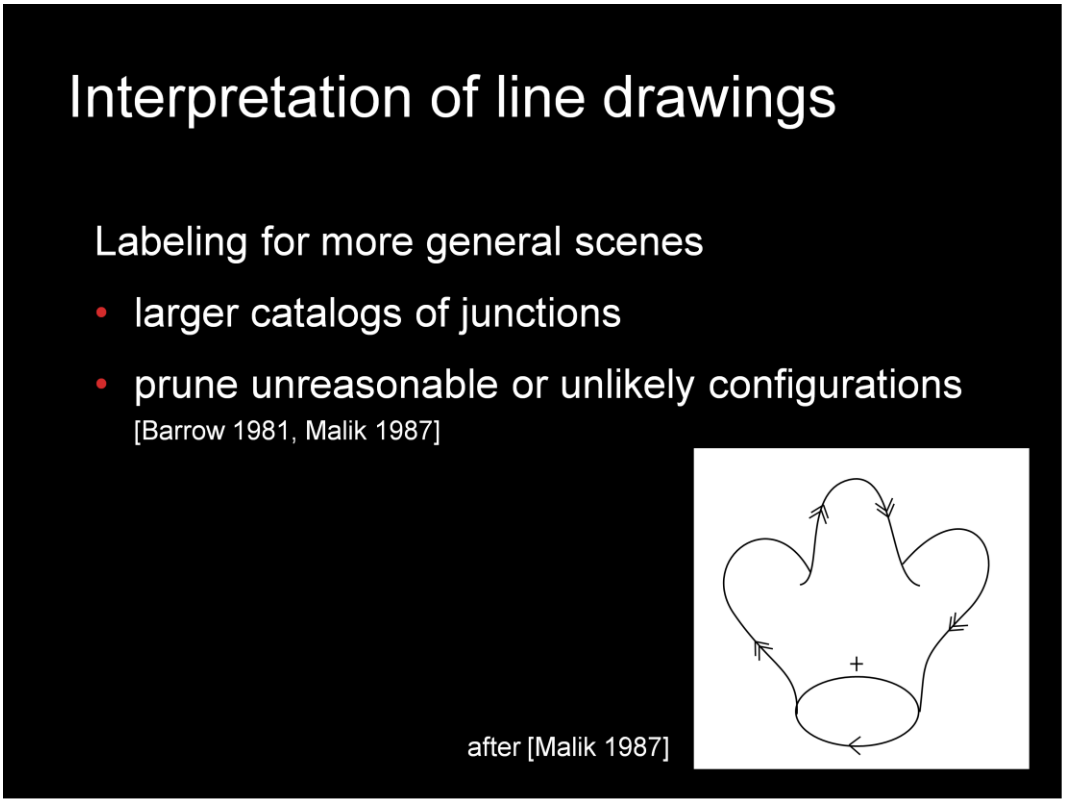

This idea can be extended for line drawings
that contain smooth surfaces.

First, you need a more comprehensive junction catalog.

Then, you need methods that can prune away large numbers
of unreasonable interpretations,
to prevent a combinatorial explosion.

These algorithms only label lines with a type.

They don’t infer geometry.

Furthermore, existing algorithms are restricted to lines
from contours and creases,
and sometimes lines from shadows.

Interpretation of line drawings

Unfortunately, we do not know how **humans** process line drawings.

Nevertheless, much is known about the **information** humans might be using.

While these algorithms suggest that exhaustive search might be a viable method for scene interpretation, they don’t say anything directly about how PEOPLE interpret line drawings.

In fact, not very much is known about that.

Even so, we can still be very specific about what INFORMATION is available in a line drawing.

This is the information that our perceptual systems are probably using.

Interpretation of line drawings

Each line **constrains** the depicted shape

- depending on the type of line

The type can sometimes be inferred from context (within the drawing)

Ambiguity always remains

- some interpretations are more likely than others

Essentially, each line in a drawing
places a constraint on the depicted shape.

In the end, the geometry that results is never unique.

But our perceptual systems excel
at uncovering the most reasonable
and most likely interpretations.

So now let's go through the kinds of information
that different types of lines provide.

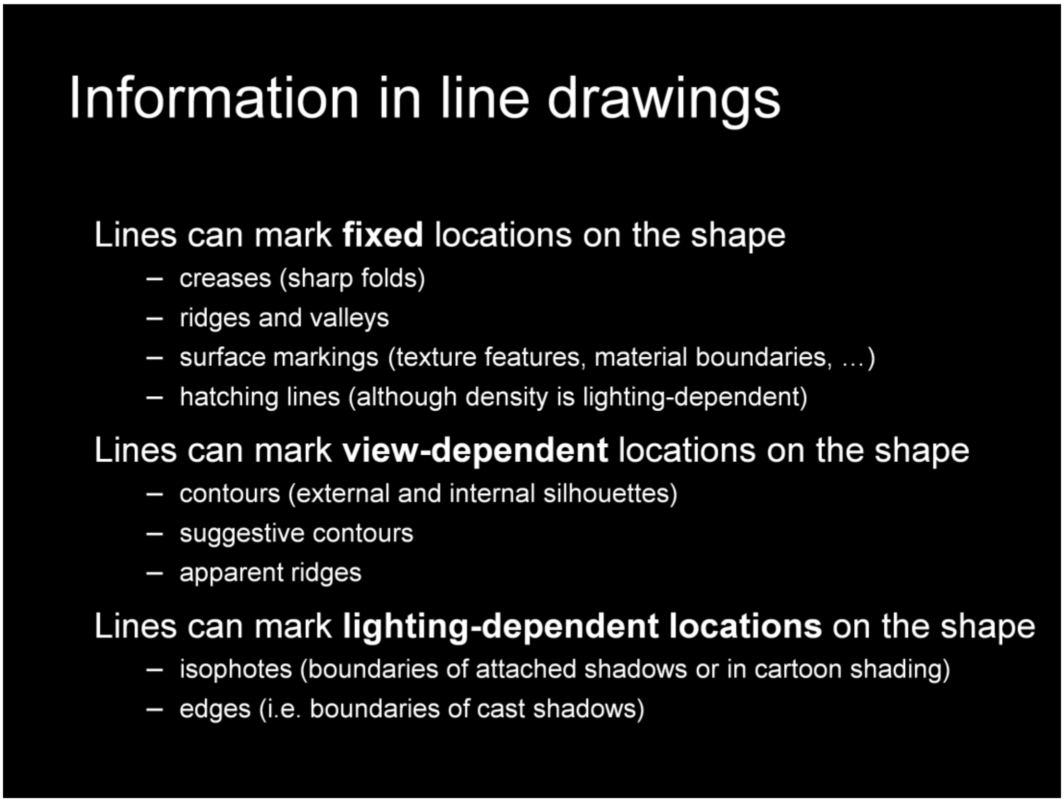

First, we'll consider lines that mark FIXED locations on a shape, such as creases, ridges and valleys,

and surface markings.

Then, we'll consider VIEW DEPENDENT lines.
The most important is the CONTOUR,
which lets us infer
surprisingly rich information about the shape.

There are also lines whose locations are lighting-dependent,
such as edges of shadows, but I'm not going to be discussing those.

Of all these lines, creases and contours are well understood.

Research on the information other types of lines provide is ongoing.

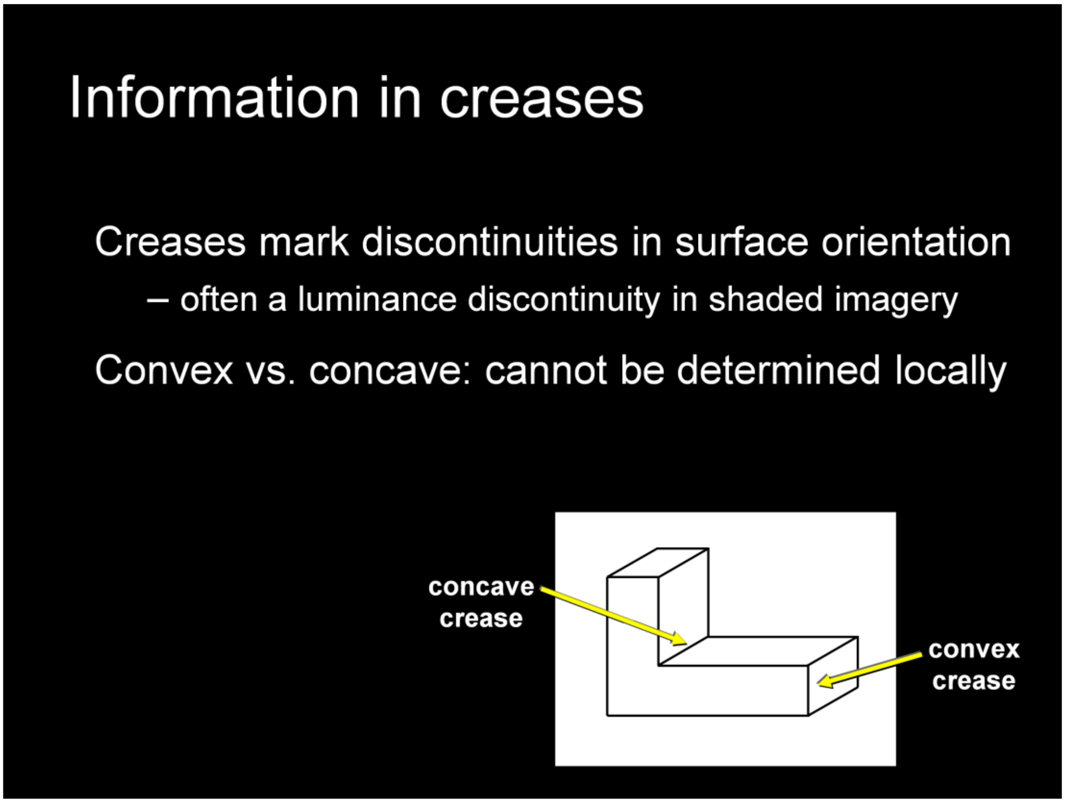

Creases mark discontinuities in surface orientation,
and are typically visible in a REAL image
as a discontinuity in tone.

The crease can be concave or convex.

But local information doesn’t let us determine which.

The algorithms for line labeling I mentioned earlier
proceeded by considering every possibility,
and then enforced consistency across the whole drawing.

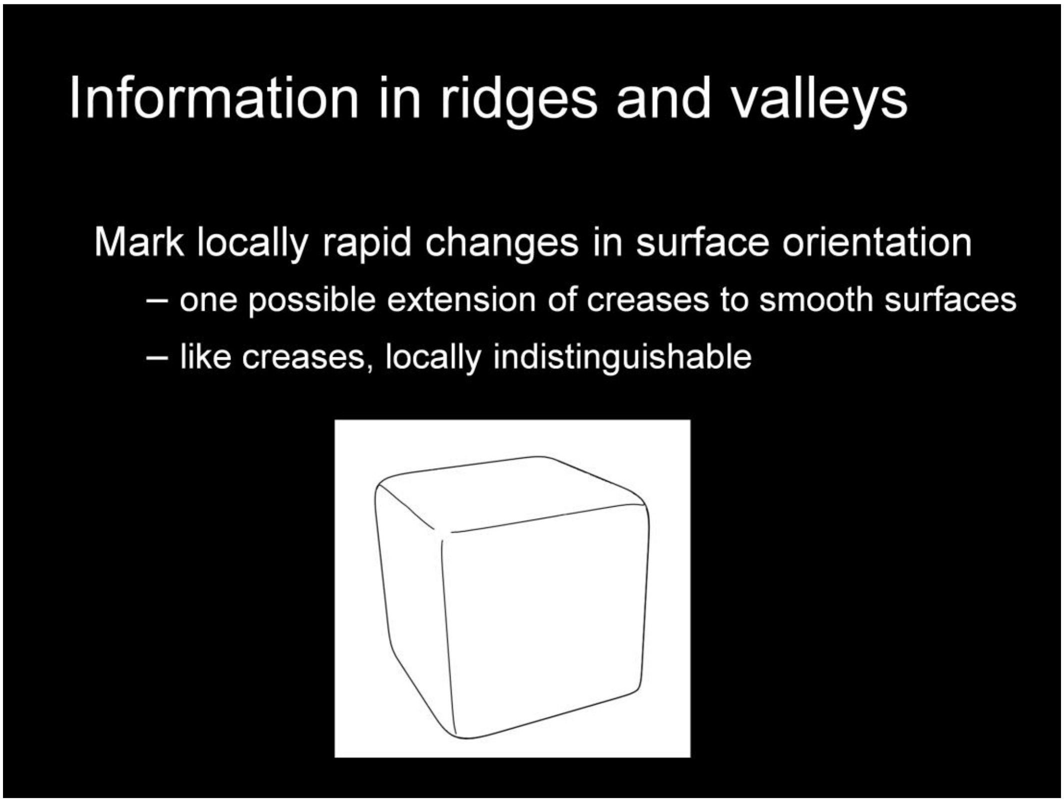

Ridges and valleys mark locally maximal changes in surface orientation.

In real images, they can appear
as smooth but sudden changes in tone.

The ridges on this rounded cube are particularly effective
at conveying its shape, when drawn with the contours.

Information in ridges and valleys

Their use is still unresolved

- many seem to convey shape
- others seem to convey surface markings
- humans can locate them in shaded imagery [Phillips 2003]

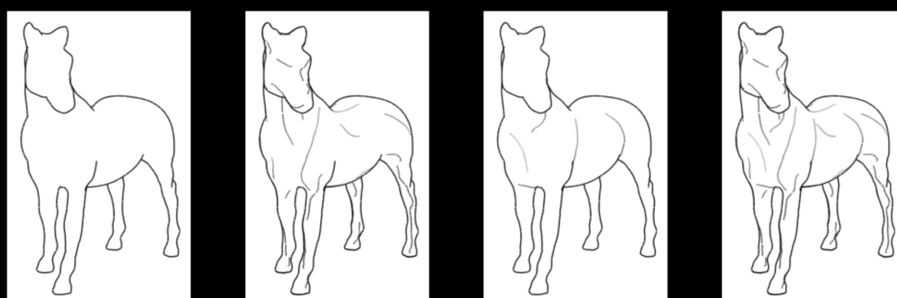

contours | with ridges | with valleys | with both

Research on the use of ridges and valleys in line drawings
is ongoing.

When used alongside contours,
ridges and valleys can produce an effective rendering of a shape.

The valleys on the side of the horse are quite convincing.

In other cases, they look like surface markings,
such as the ridges on its head.

Ridges and valleys are reasonable candidates for line drawings,
as there is psychological evidence that viewers
can reliably locate them in realistic images.

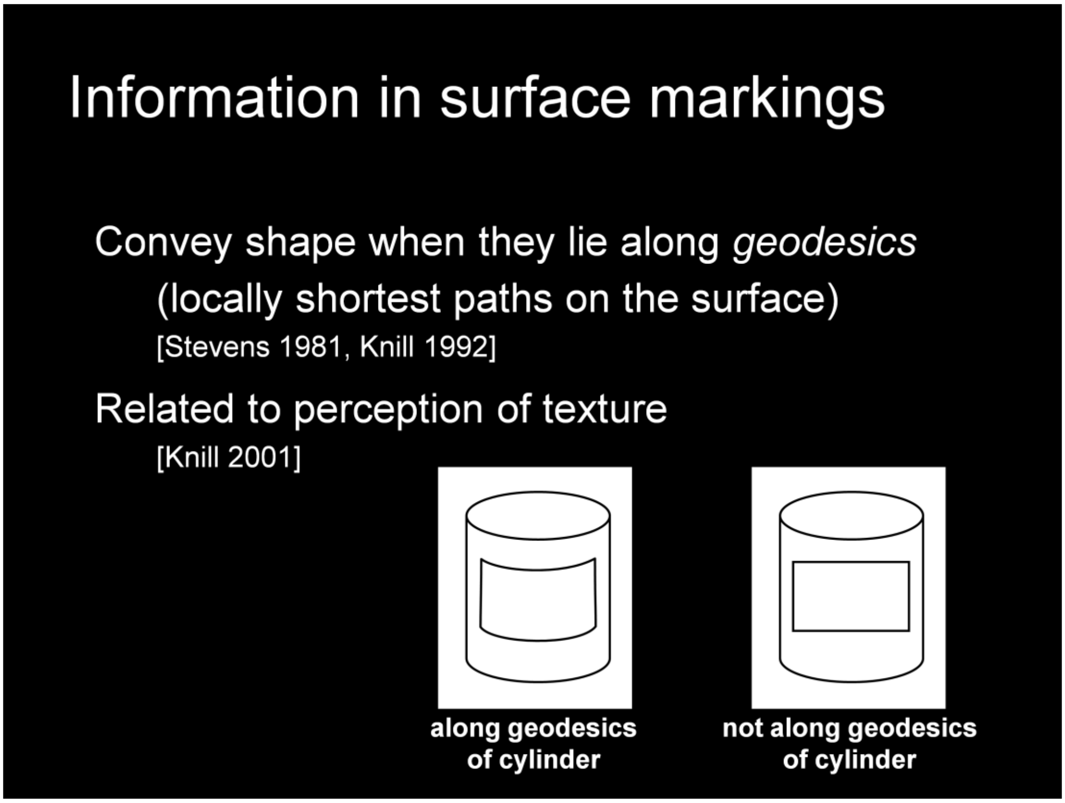

Markings on a surface can appear
as arbitrary lines inside the shape.

However, for a certain type of line known as a geodesic, they can also convey shape.

Geodesics are simply lines on the surface
that are locally shortest paths.

Stevens points out that for many fabricated objects,
surface markings are commonly along geodesics.

Take for instance the label on the cylinder on the left.

For a more general class of surfaces,
Knill draws connections between texture patterns
and sets of parallel geodesics.

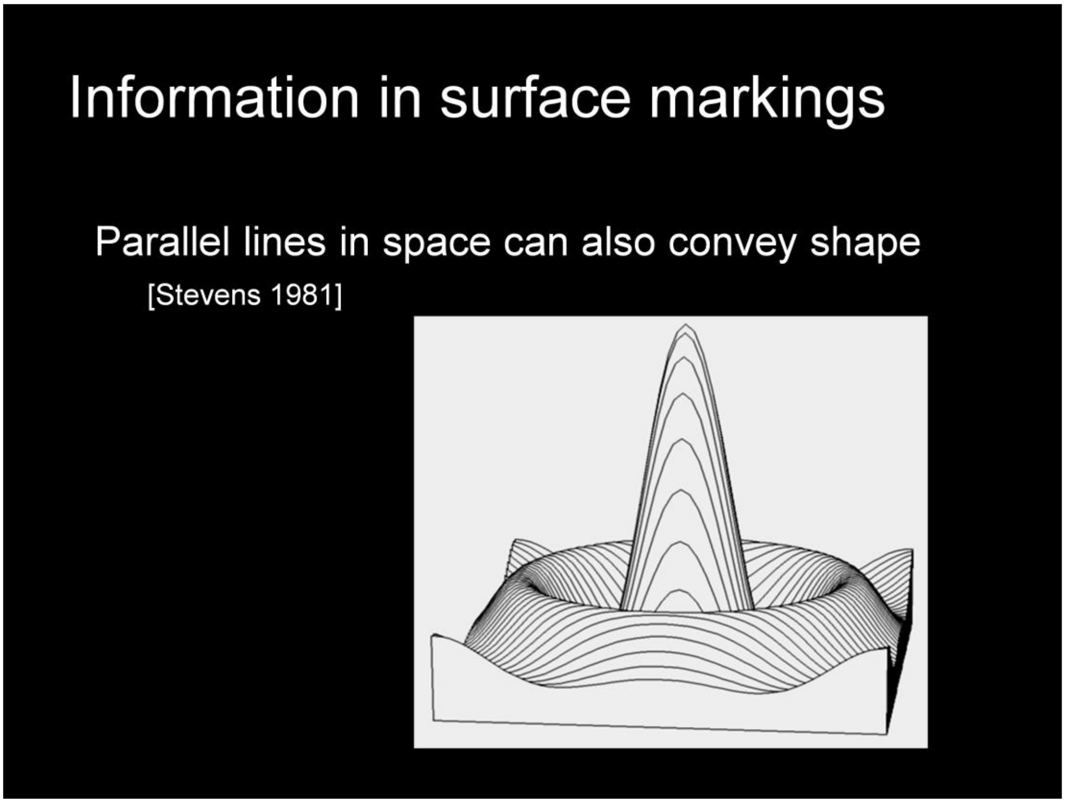

When used in repeating patterns,
other curves can be effective as well.

Sets of parallel lines,
which are often used to construct plots of 3D functions,
are one notable example.

The images that result are analogous to using
a periodic solid texture.

Stevens points out that all one needs to do
to infer the shape is
to build correspondences between adjacent lines,
matching up points with equal tangent vectors.

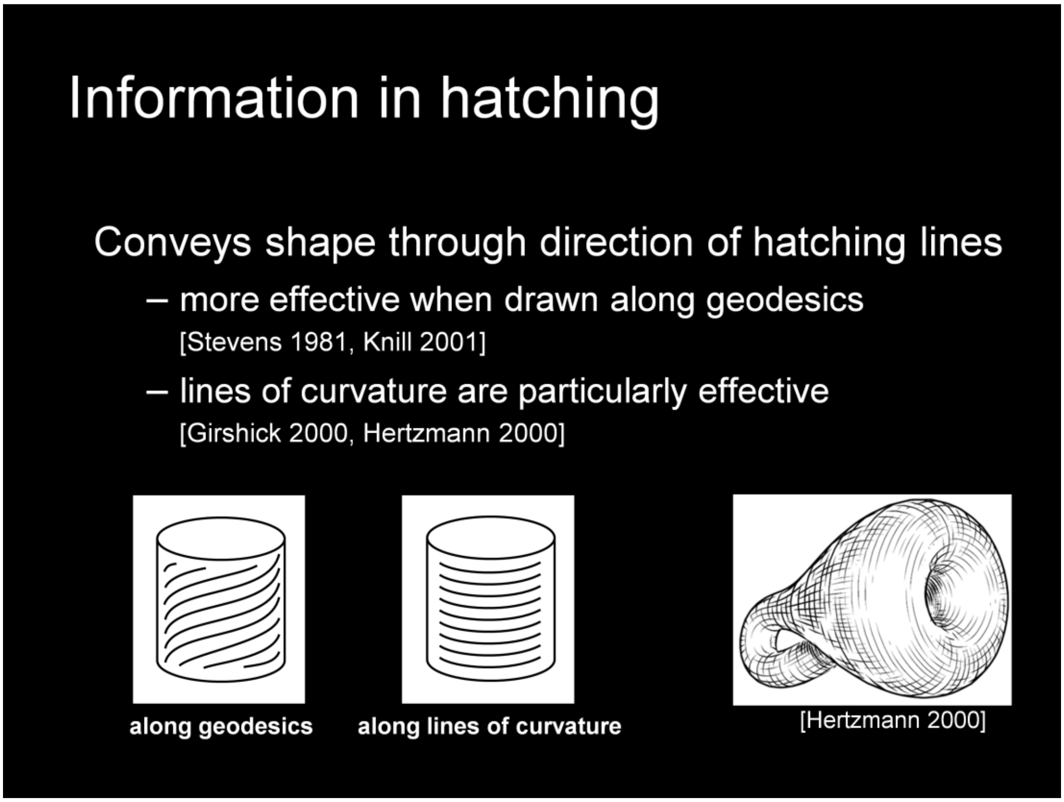

The use of repeating patterns of lines
forms the basis of hatching.

These lines convey shape in two different ways;
they convey shape directly when they are drawn along geodesics.

And they convey shape indirectly
through careful control of their density,
which can be used to produce a gradation of tone across the surface.

Particularly effective renderings are obtained
when lines of curvatures are used,
which align with the principal directions of the surface.
These also happen to be geodesics.

So that’s it for lines
whose locations are FIXED on the shape.

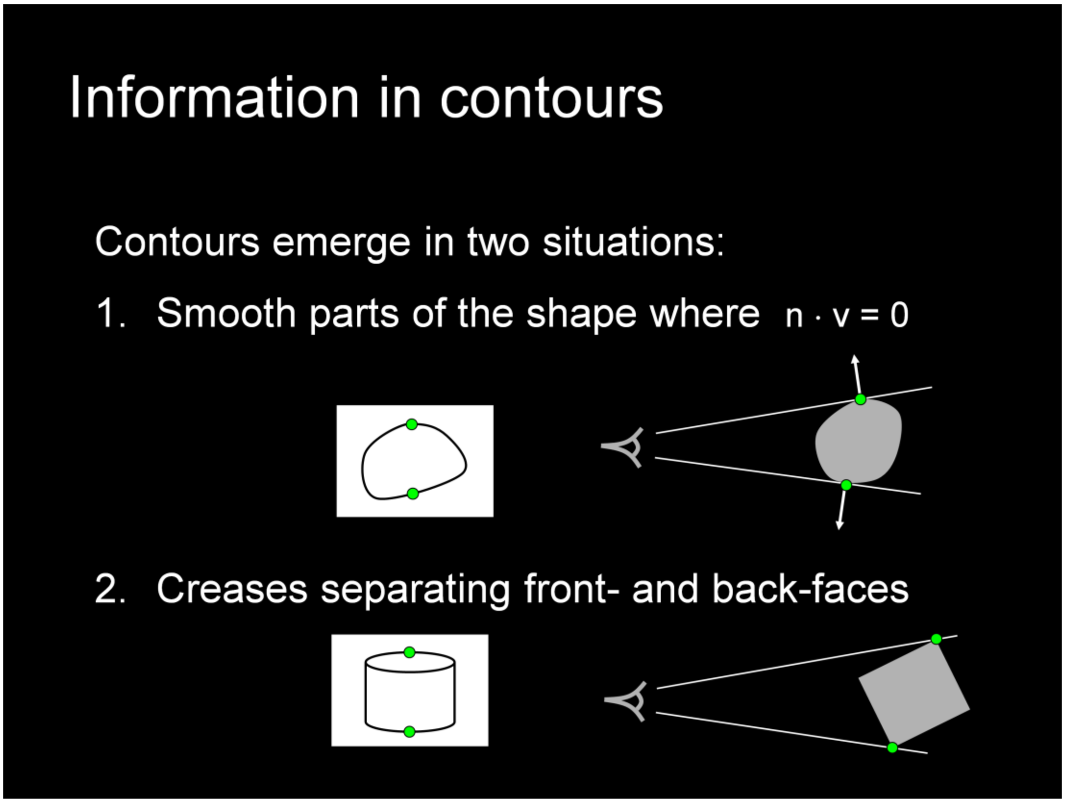

Next are lines whose location
depends on the viewpoint.

The contour is the most notable example.

There are two situations when contours are formed.

On a smooth surface,
contours are produced when the surface is viewed edge-on.

On an arbitrary surface, contours can also appear
along a crease.

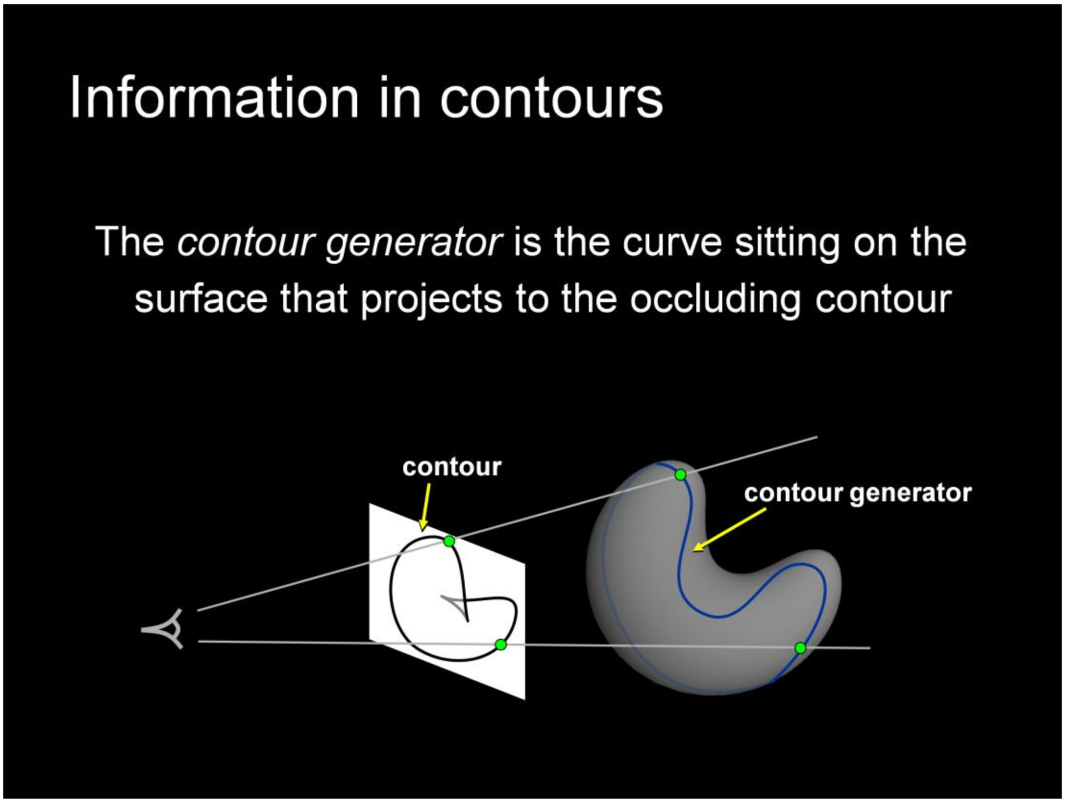

In either case,

sitting on the surface
is a 3D curve known as the CONTOUR GENERATOR.

This curve marks all local changes in visibility across the shape.

For a typical viewpoint, the contour generator
consists of a set of isolated loops.

It projects into the image to become the contour.

So not all parts of the contour are visible.

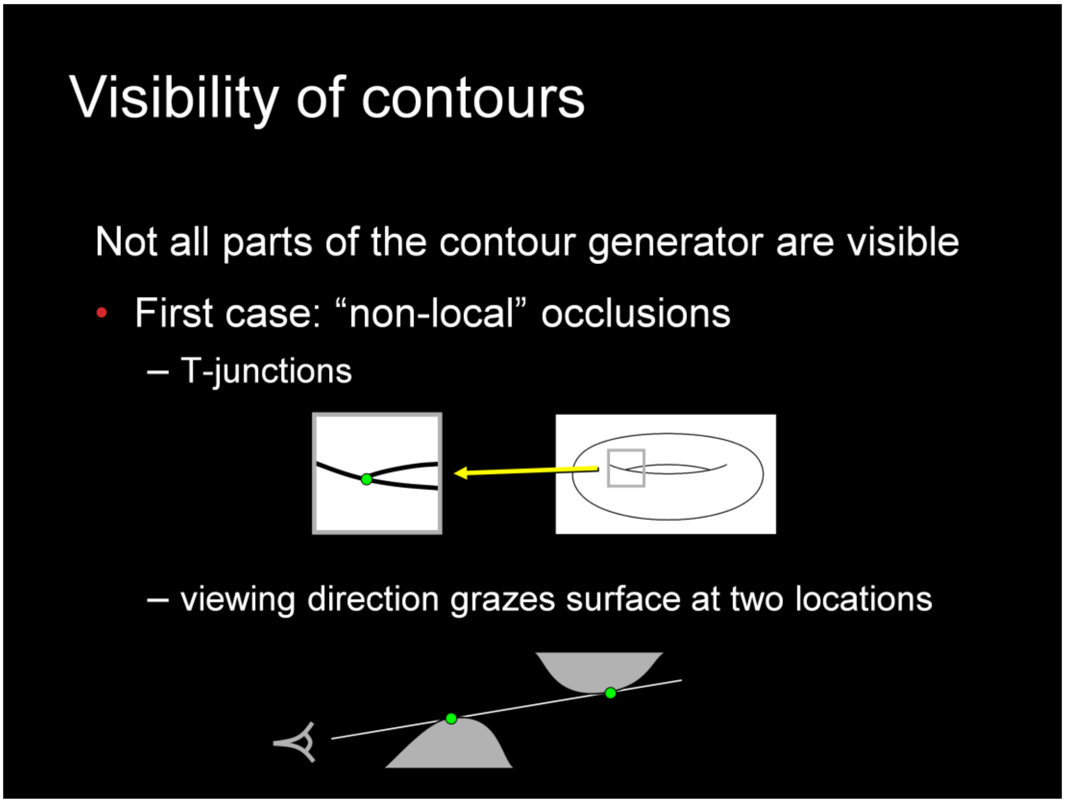

Let’s consider the different cases of visibility for contours.

On a smooth surface,
the first case is when one part of the shape
occludes another more distant part.

This appears in the image as a T-junction,
where the contour goes behind another part of the shape.

****

At the location where the visibility changes,
the visual ray is tangent to the surface in two places.

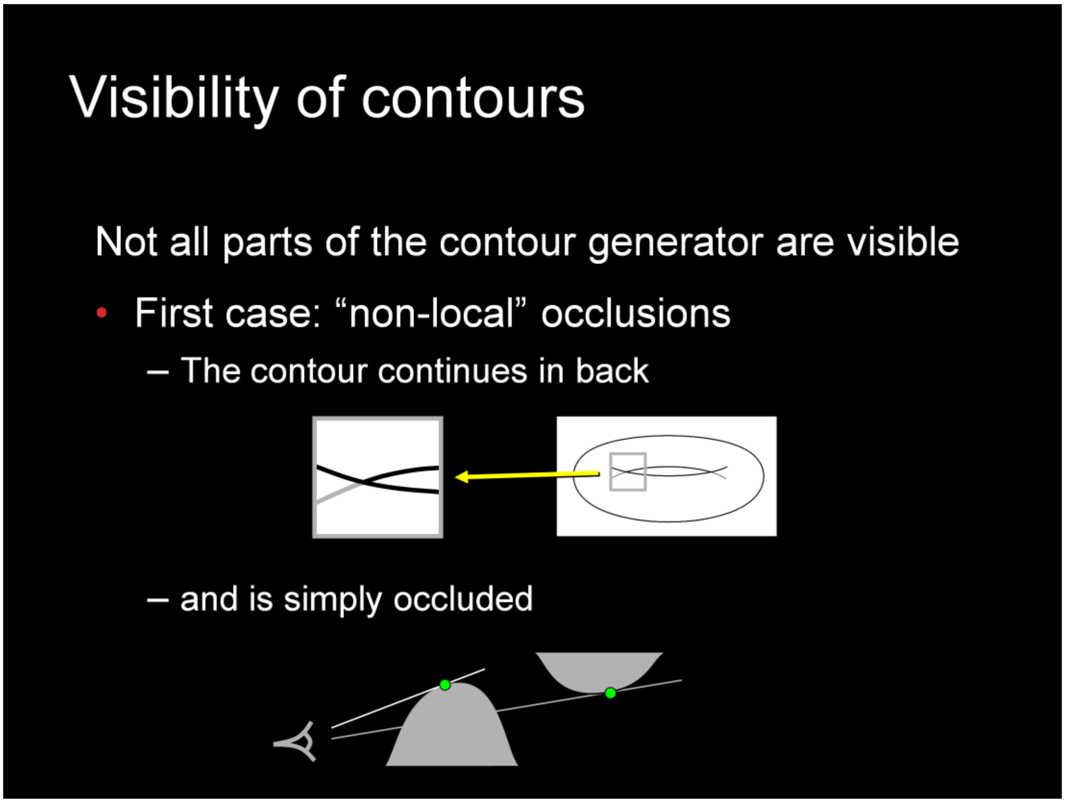

The contour then continues behind the shape,

****

and is occluded.

This can be seen in this transparent line-drawing of a torus.

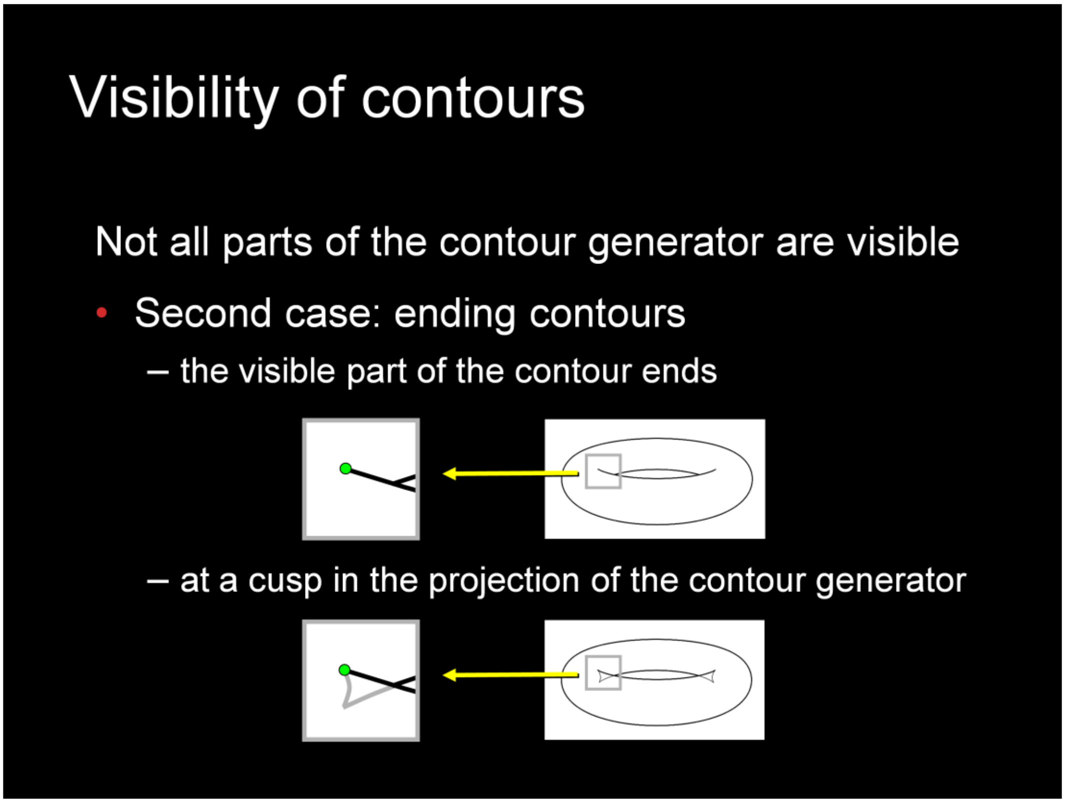

The second case occurs
where the contour comes to an end in the image.

An ending contour.

****

When the occluded part of the contour continues,
it does so at a cusp in the contour.

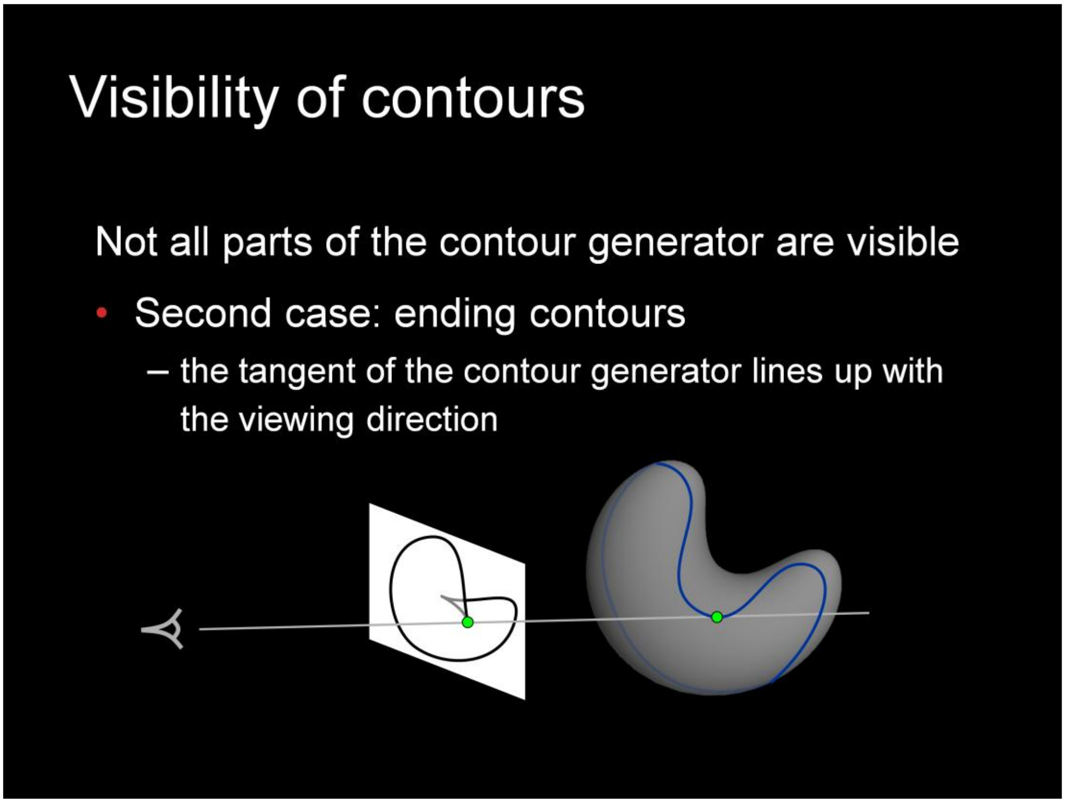

This cusp occurs because the contour generator
lines up with the viewing direction,
so that its tangent projects to a point.

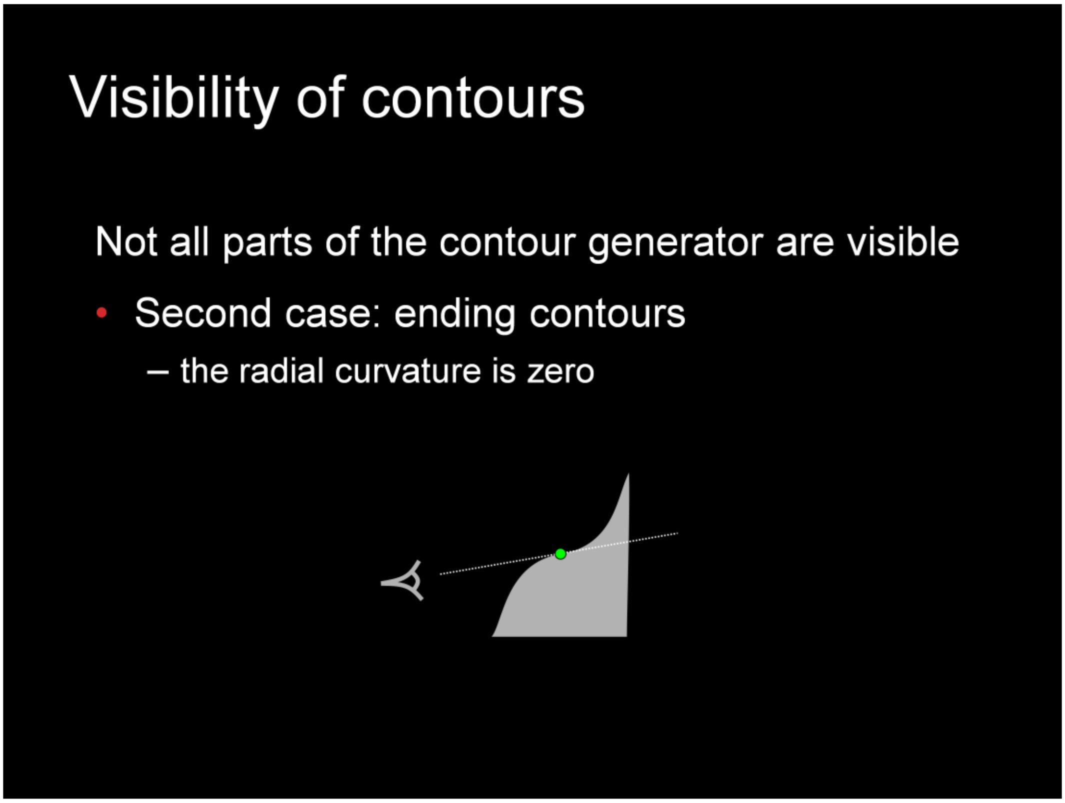

At an ending contour, the radial curvature is zero, which means that we're looking along an inflection; an asymptotic direction of the surface.

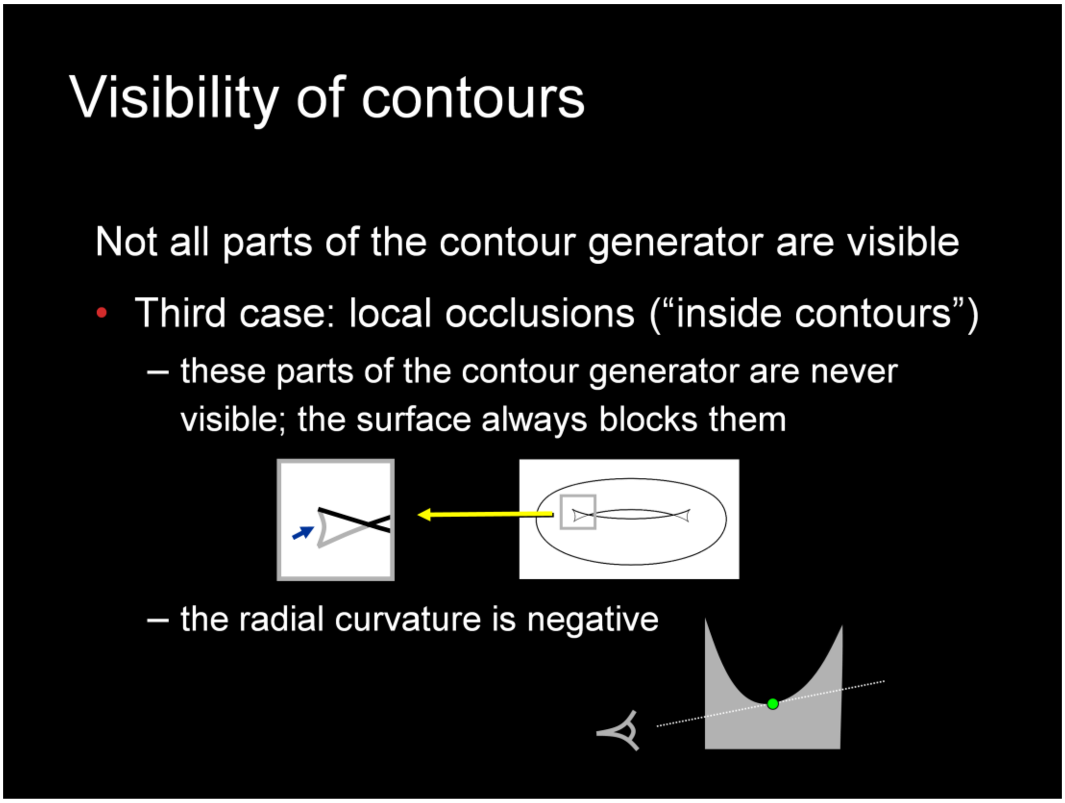

The last case is a local occlusion;
places where the surface has no choice but to occlude itself.

****

These are locations where the radial curvature is negative.

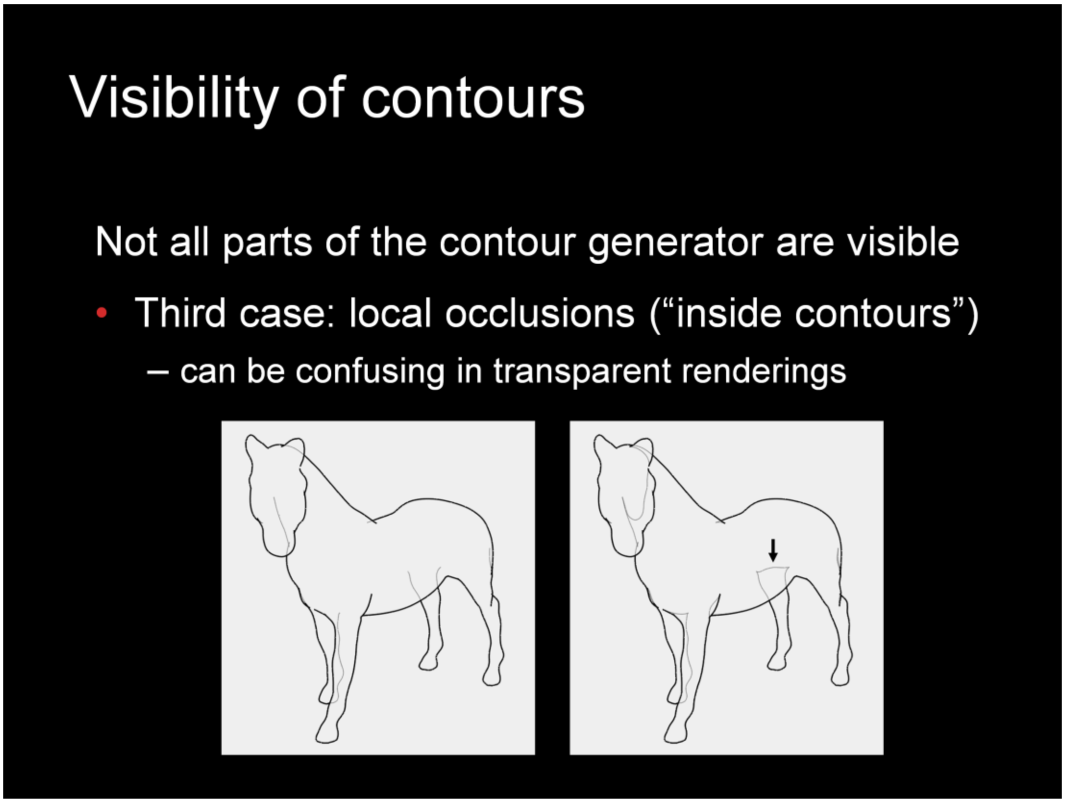

In transparent renderings of contours,
one typically does not draw the local occlusions,
as the results can be confusing.

The image on the right here draws these contours.

One is marked with an arrow.

These curves actually correspond to regular contours
for an inside-out version of the surface.

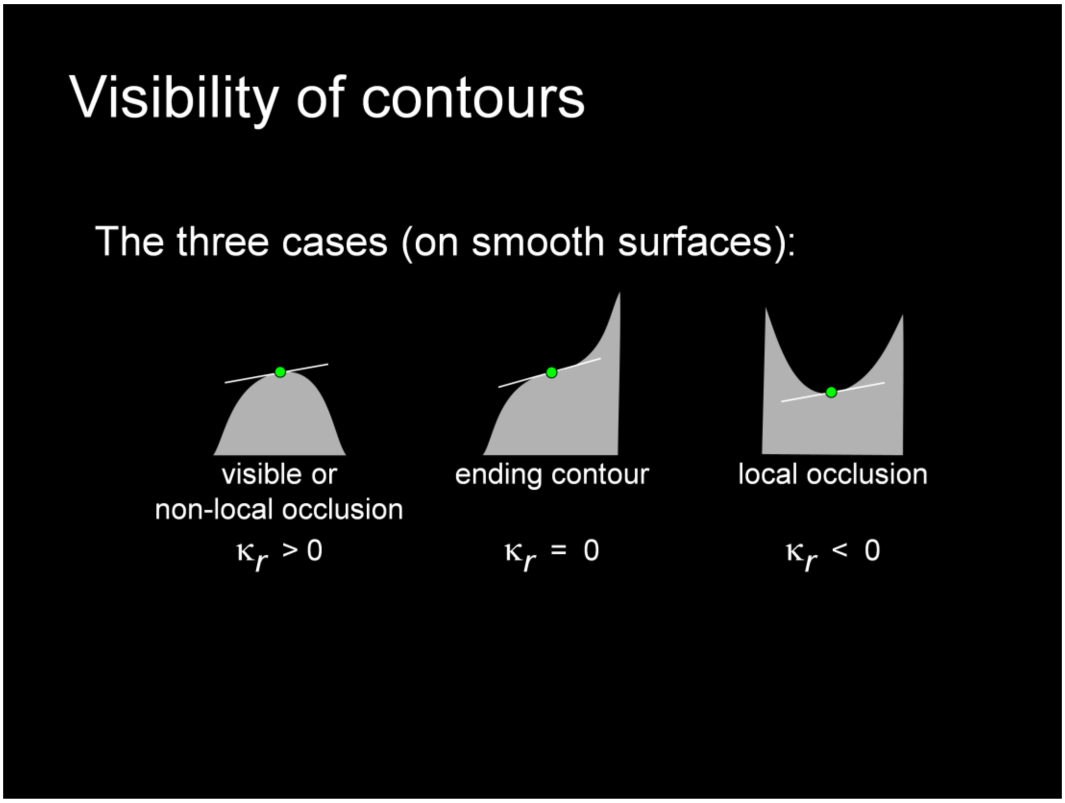

Here are the three cases, all together.

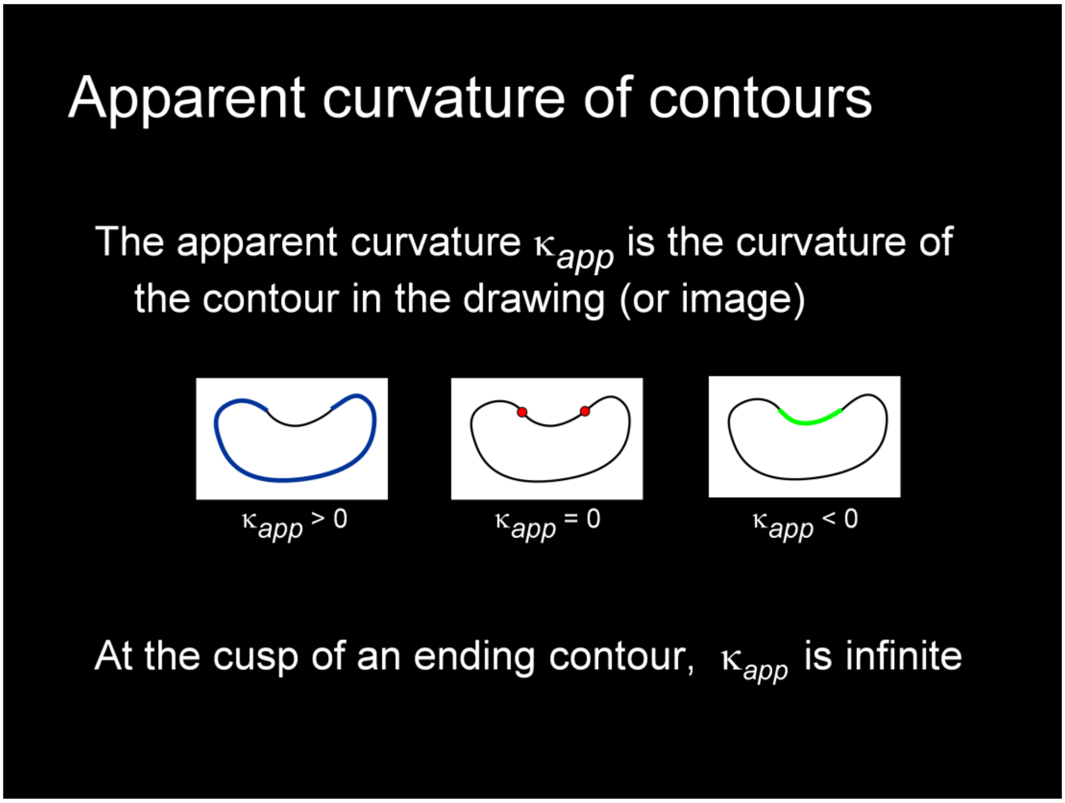

Now, let’s consider what the contours look like in the image.

The apparent curvature is simply the curvature of the contour in the drawing.

The convex parts of the contour
have positive apparent curvature,

the concave parts have negative apparent curvature,
and it’s zero at the inflections.

At the ending contours,
the apparent curvature is infinite due to the cusp.

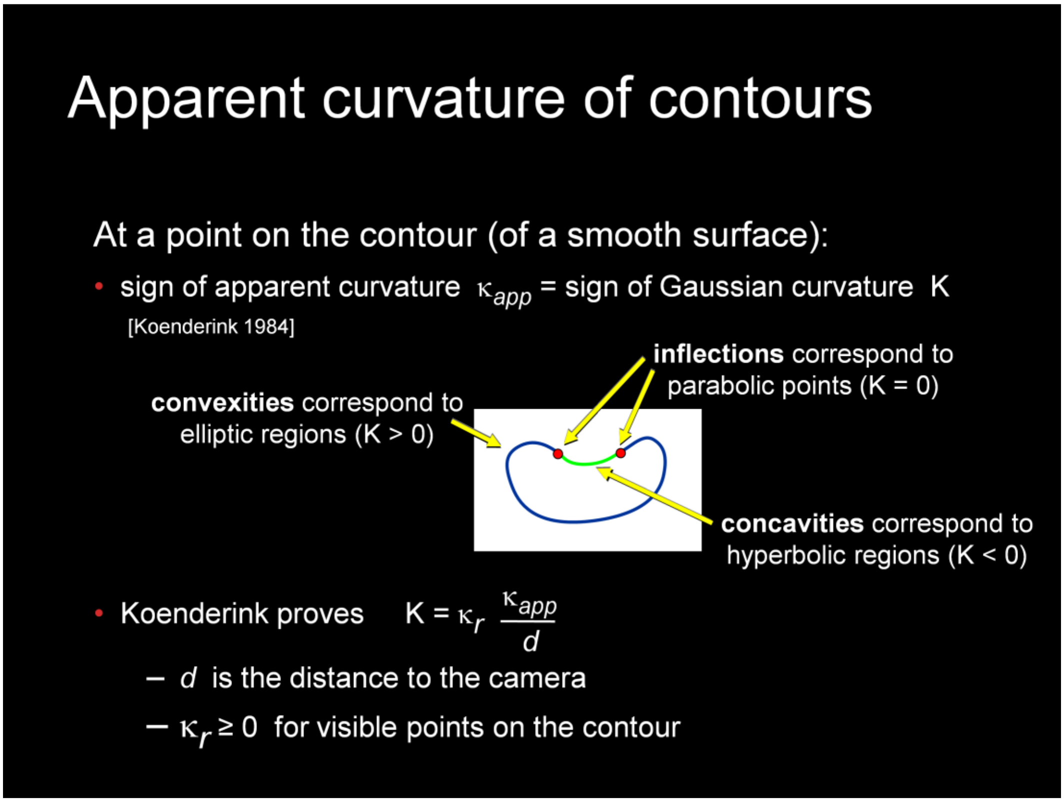

Koenderink proved a surprising and important relationship between the apparent curvature and the Gaussian curvature.

Specifically, for visible parts of the contour on a smooth surface, they have the same sign.
This means we can infer the sign of the Gaussian curvature simply by looking at the contour.

****

CONVEX parts of the contour correspond to locations
where the Gaussian curvature is POSITIVE: elliptic regions.

****

INFLECTIONS on the contour correspond to locations
where the Gaussian curvature is ZERO.

****

CONCAVE parts of the contour correspond to locations where the Gaussian curvature is NEGATIVE: saddle-shaped regions.

****

Koenderink gives a formula that connects these two quantities, that involves the distance to the camera and the radial curvature.

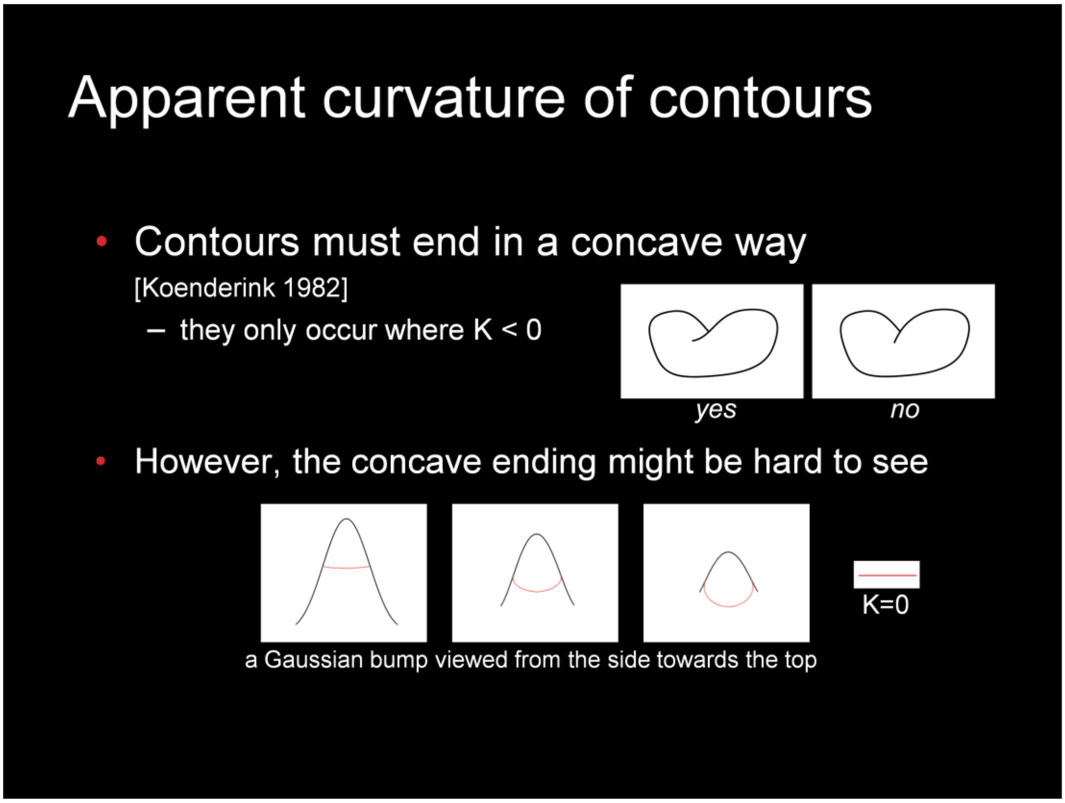

A related result is that since ending contours only occur
where the Gaussian curvature is negative,
the contours must end in a concave way,
approaching their end with negative apparent curvature.

****

But Koenderink and van Doorn also noticed
that artists tend to draw lines
that are missing these concave endings.

It turns out this concave ending can be difficult to discern,
as is the case for this Gaussian bump.

DEMO

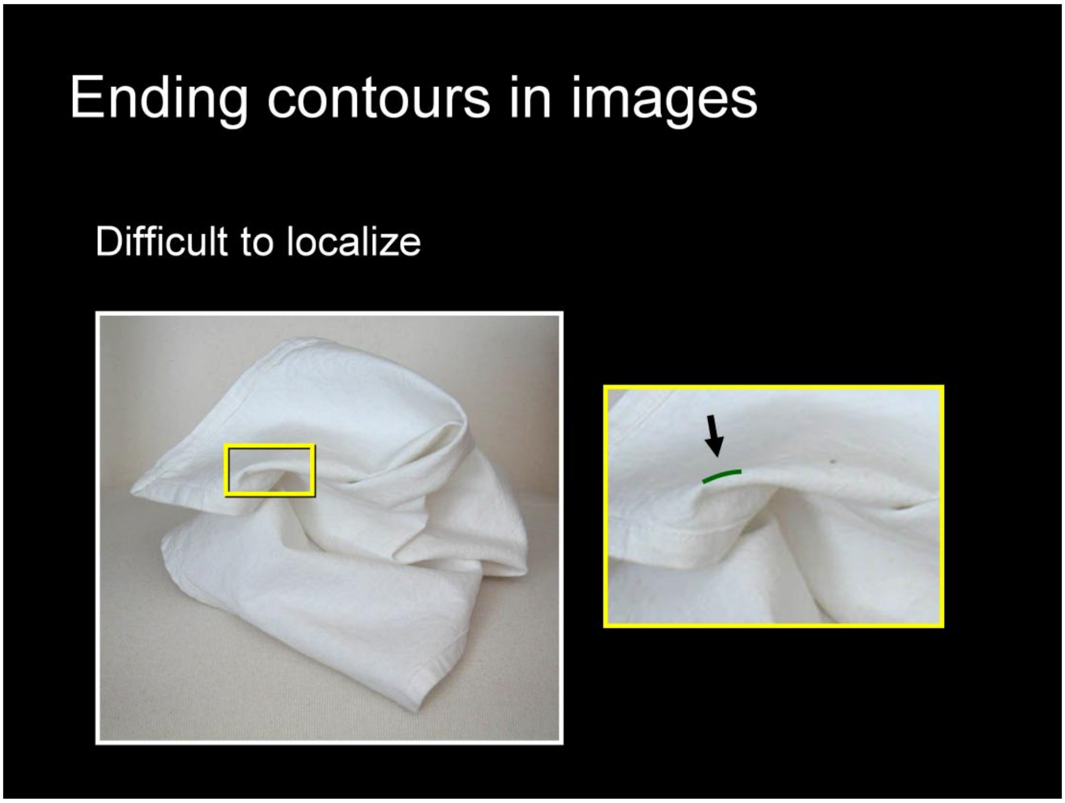

Contours are typically easy to detect in real images, at least when the lighting is right.
And there are many studies that demonstrate how people use them for visual inference.

However, in many cases, it’s not easy to determine where a contour ends.

Here’s an example photograph of napkin.

Even if we zoom in, it’s still not clear
whether the surface occludes itself
or whether it’s simply heavily foreshortened.

Observations like this make sense of line types that extend ending contours: suggestive contours and apparent ridges.

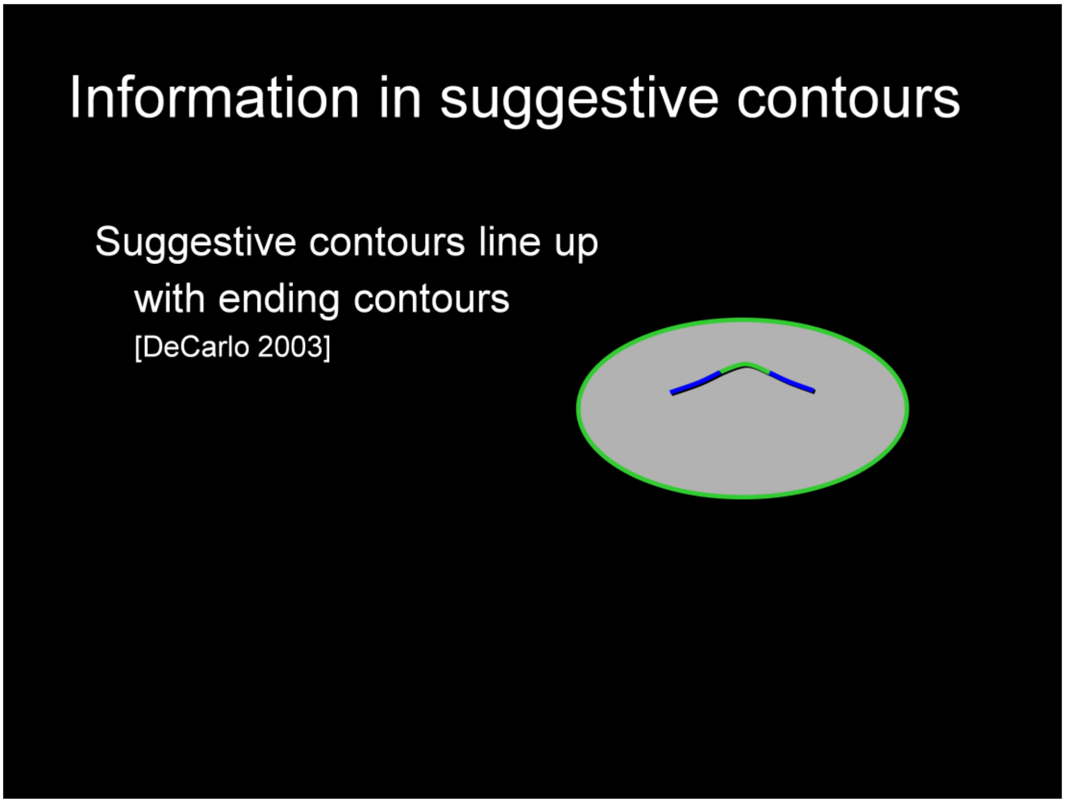

Suggestive contours are another type of line to draw, and whether they are in fact detected and represented by our perceptual processes is still an open question.

They do seem to produce convincing renderings of shape in many cases.

The fact that suggestive contours smoothly line up with contours in the image is encouraging.

****

In fact, if the lines aren't color coded, it's difficult to tell where one starts and the other ends.

Information in suggestive contours

Suggestive contours line up
with ending contours
[DeCarlo 2003]

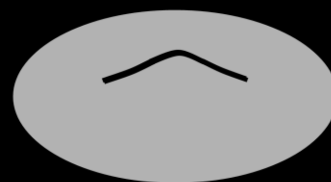

This makes it hard to tell where contour ends

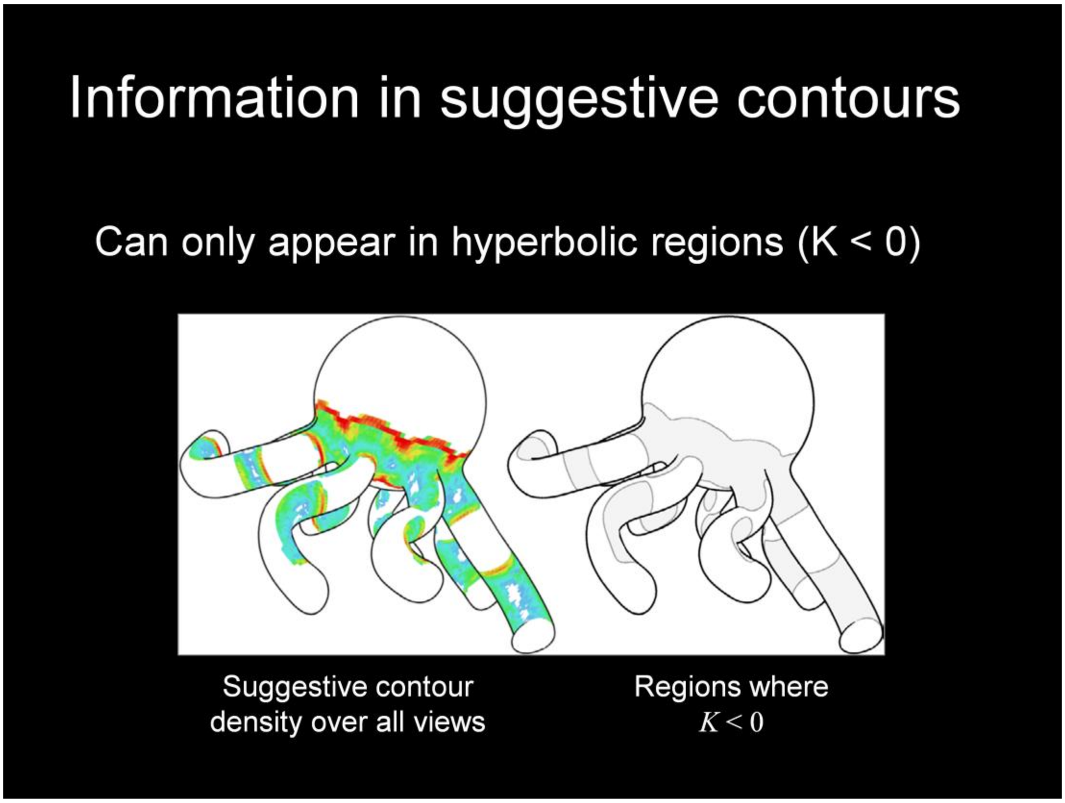

We can say something about what information they provide.

Recall from earlier how suggestive contours
can only appear
where the Gaussian curvature is negative.

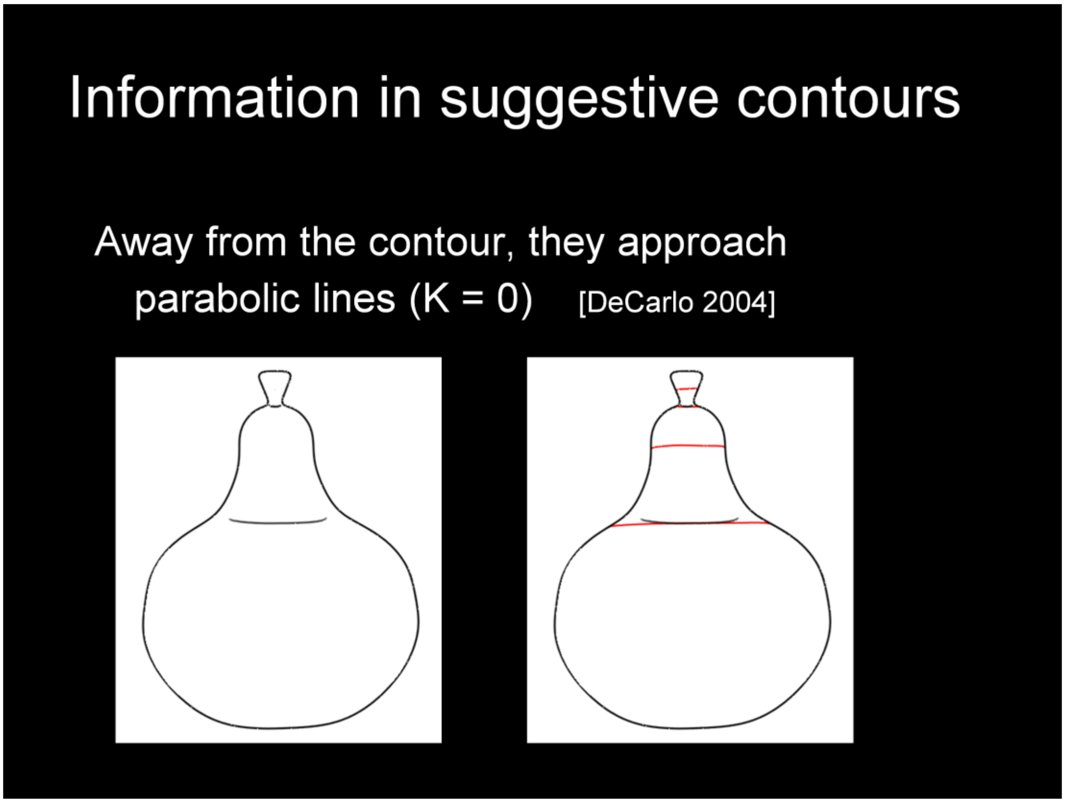

In many cases, the suggestive contours
approach the parabolic lines
away from the contour.

On this pear, we see how the suggestive contour
skims along the parabolic line.

DEMO

We hope to be able to say more about this in the future.

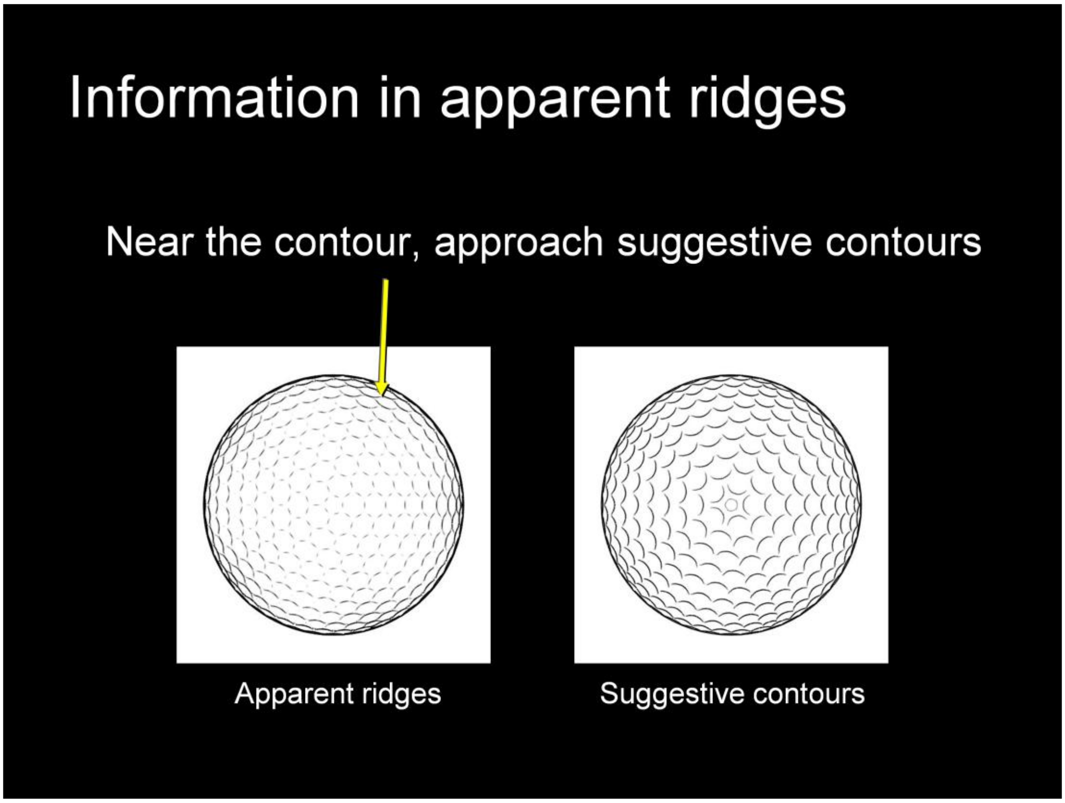

We can make similar statements about apparent ridges.

Near the contour, apparent ridges behave like suggestive contours. They extend ending contours.

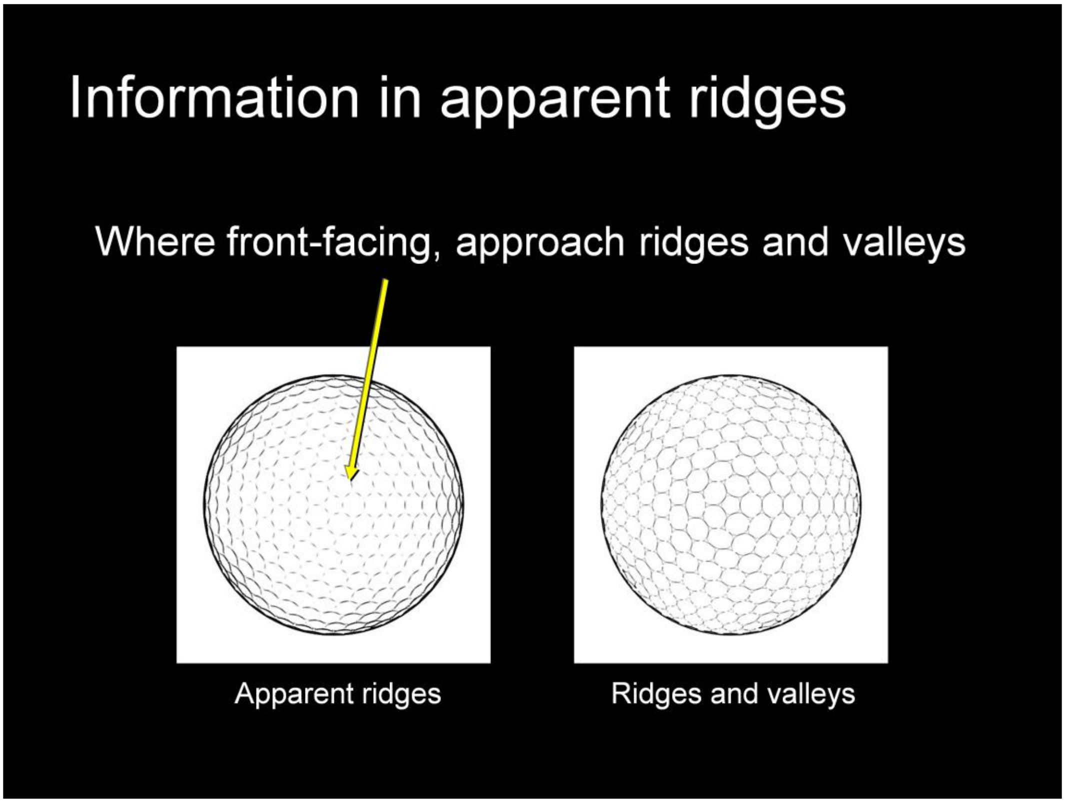

As the surface faces more towards the viewer,
the location of apparent ridges
approaches ordinary ridges and valleys.

And of course, in both of these cases,
apparent ridges are surface locations
where the normal vector is changing maximally.

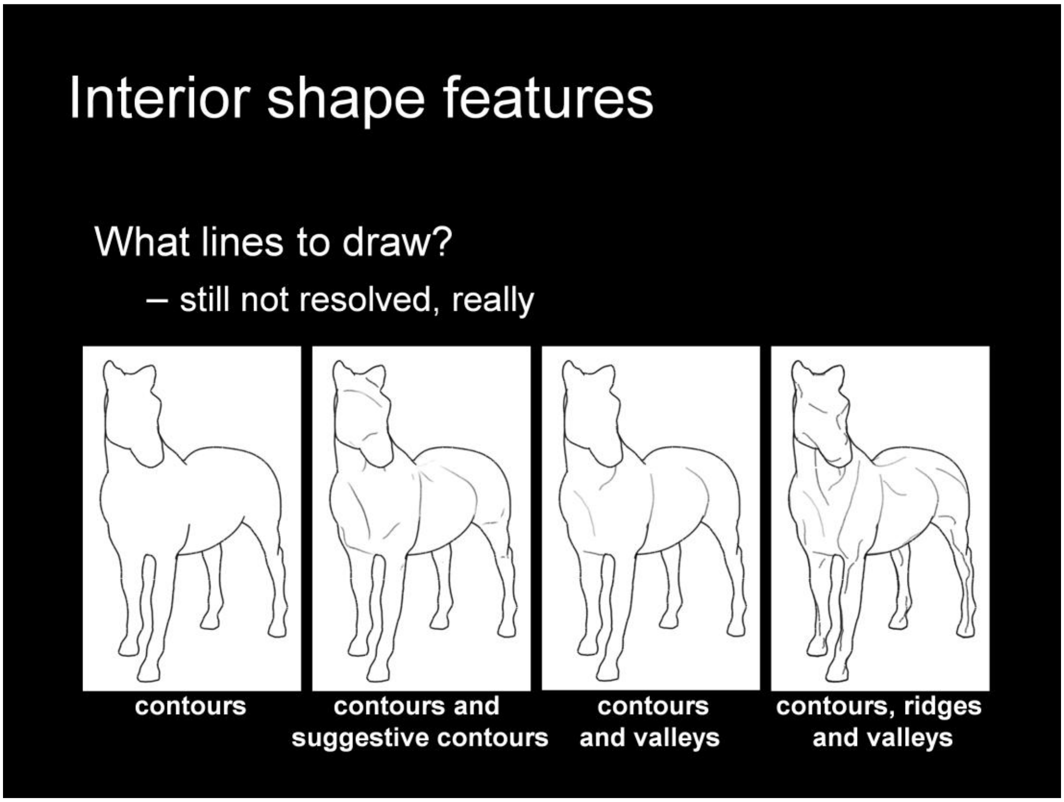

We can compare
renderings with ridges and valleys
to renderings with suggestive contours.

On the horse from this viewpoint,
the rendering with just valleys is actually quite convincing.

As noted earlier, many of the ridges appear as surface markings here.

For the valley rendering, some features are missing,
but the more salient features on the side of the horse are depicted.

Note the slight differences between the lines
from suggestive contours, and from valleys.

The shapes they convey appear to be a little different.

Clearly there is a lot of interesting work to do here.

This concludes our discussion
of what information particular lines provide.

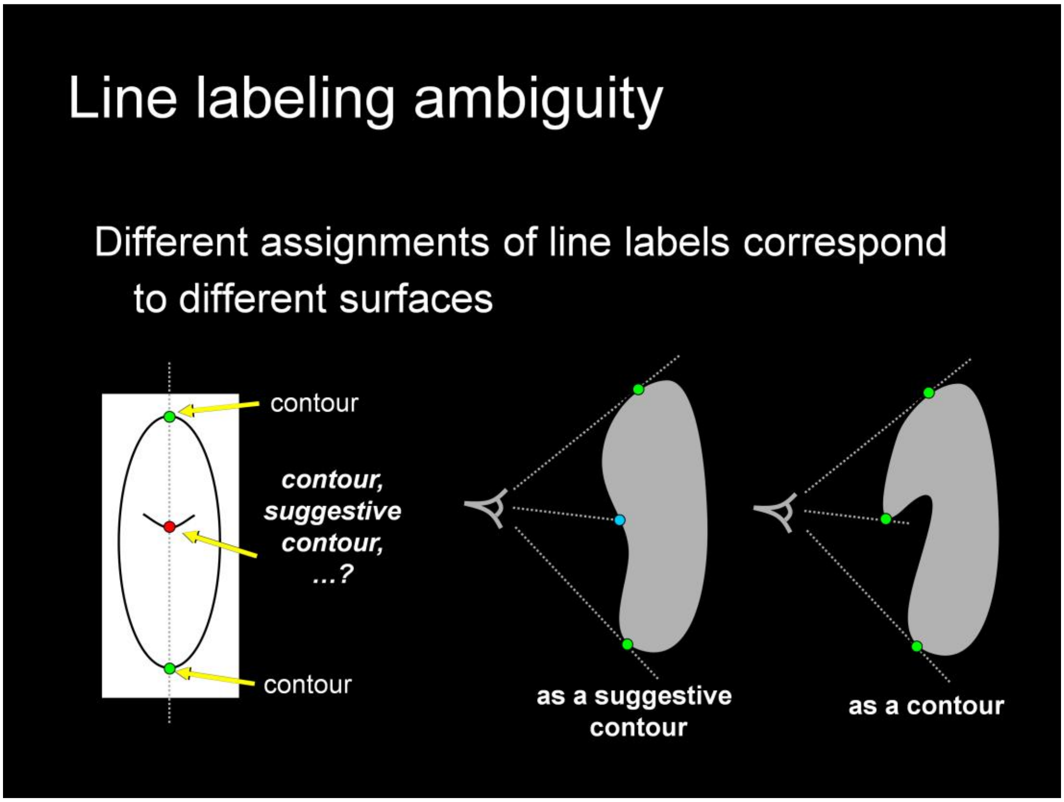

Of course, this information can only be used
if we know the TYPES of the lines
when we're given a drawing.

Earlier we discussed algorithms for line drawing interpretation;
approaches like this are reasonable to consider for this purpose.

But even if we do use these algorithms,
there are often several different labelings that are consistent.

Given the line drawing on the left
which depicts an elliptical shape with a bump,
we can successfully label the green points as contours.

The red point, however, can be either a contour
or suggestive contour.

Two possible shapes that match these labelings
are shown on the right.

Presumably this problem cannot be solved in general.

There will always be ambiguity.

It's possible that when artists make line drawings,
they're careful to shape the remaining ambiguity
so it won't be a distraction.

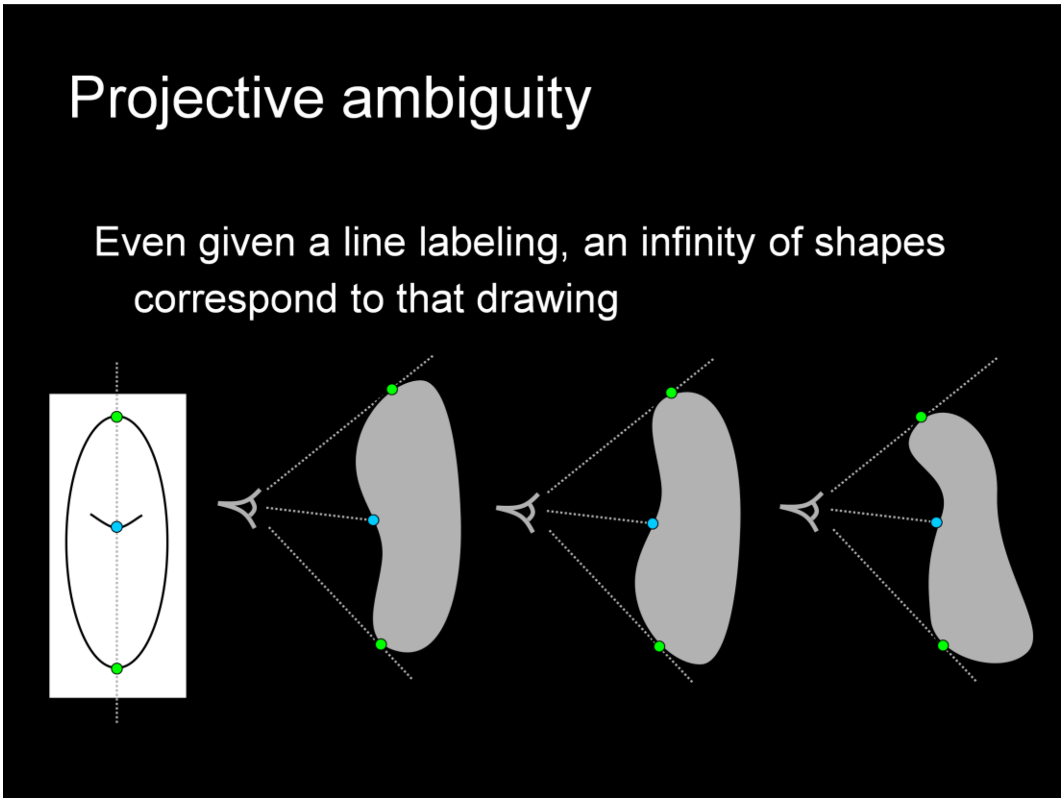

And even with a line labeling,
there is the ambiguity of projection.

These three interpretations have the same line labeling,
but different geometries.

At first, this seems hopeless.

Yet sketching interfaces like Igarashi’s Teddy
seem to be quite successful by using inflation.

How can this be?

Well, there are reasonable constraints on smoothness
that we can expect of the underlying shape.

We also presume
that the artist has drawn all of the important lines,
so that no extra wiggles remain.

These issues are the source of one crucial challenge
for sketch-based shape modeling.

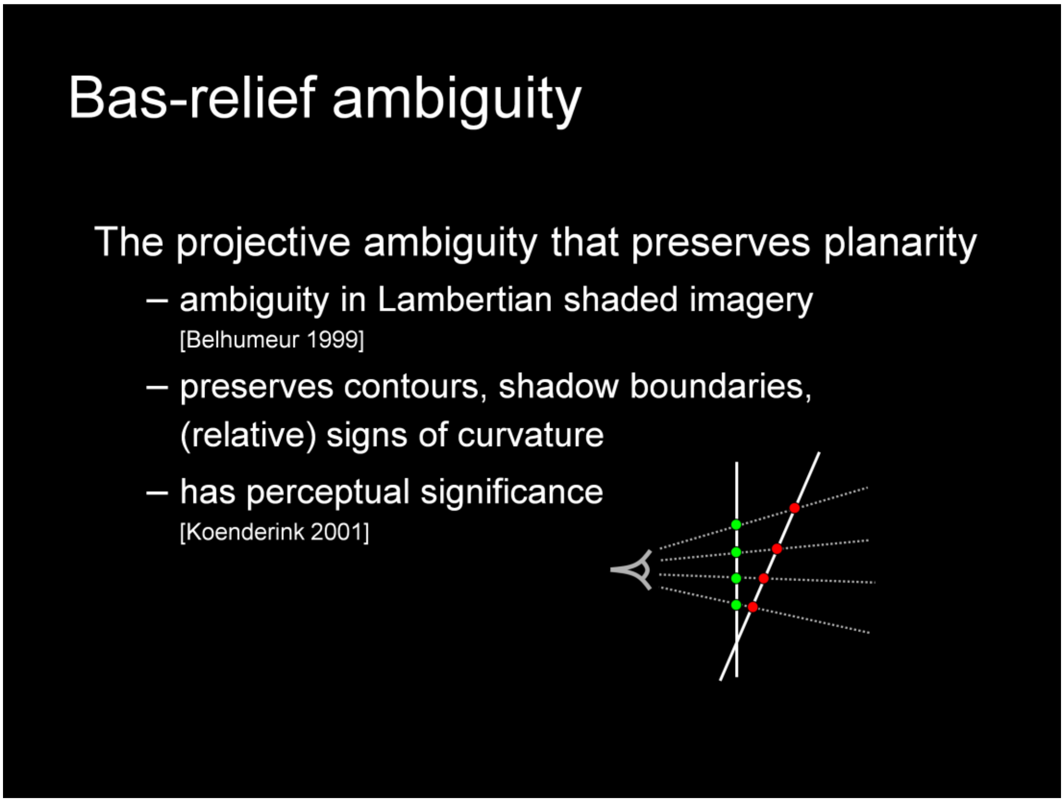

We can be more specific with regard to this ambiguity.

For real images, there are well defined ambiguities for particular types of imagery.

One notable example is the ambiguity that remains when viewing a shape under Lambertian illumination.

There is a group of shape distortions that can be applied to a shape, that with an corresponding transformation of the lighting positions, approximately produce the same image.

This is the three-dimensional projective mapping known as the generalized bas-relief transformation. As shown here, it moves points along visual rays and preserves planes. It also preserves contours, boundaries of shadows, and the relative signs of curvature on the shape.

Perhaps most interestingly is that when you ask people to describe the shapes they see in shaded imagery, they answer consistently modulo this ambiguity transformation.

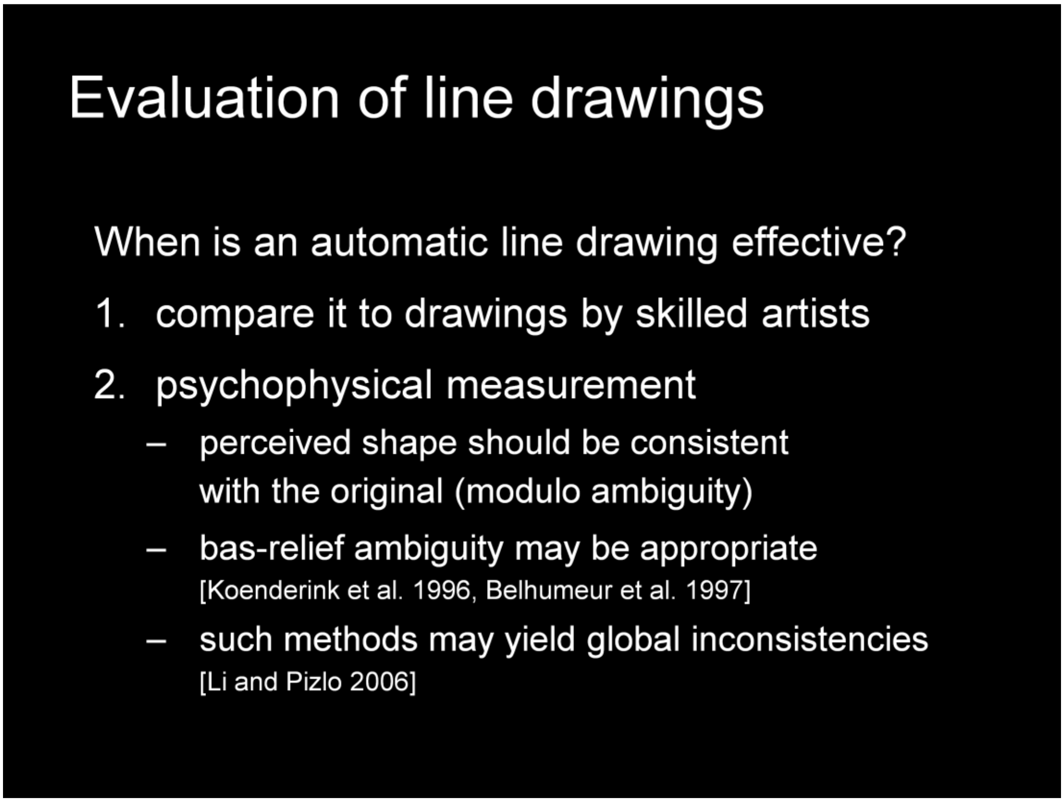

So how can we be sure that a line drawing we make is perceived accurately?

As you saw earlier, one possible path is to compare that line drawing to those made by skilled artists.

Another way, based in psychophysics, is to simply ask the viewer questions about the shape they see.

If this is done right, you can reconstruct their percept and compare it to the original shape, given the appropriate ambiguity transformation.

Koenderink and colleagues already performed a study like this on a single line drawing.

Their results suggest that the bas-relief ambiguity might be the appropriate one to consider here.

However, this ambiguity may only be resolved locally, where different parts of the shape are locally consistent, but not necessarily in a global sense.

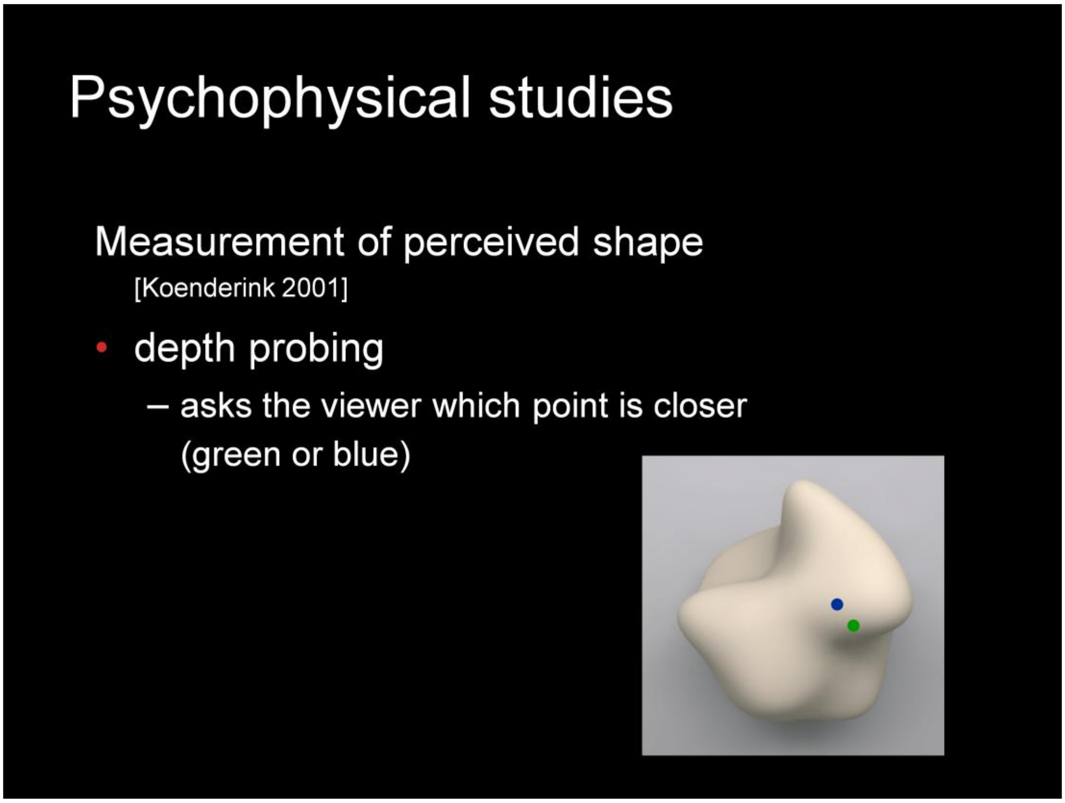

So what kinds of questions can you ask viewers?

In psychophysics, the answer is:
very simple ones, and lots of them.

Koenderink describes a set of psychophysical methods
for obtaining information
about what shape a viewer perceives.

The first they describe is called RELATIVE DEPTH PROBING.

The viewer is shown a display like this one,

and is simply asked which point appears to be closer.

They are asked this question for many pairs of points.

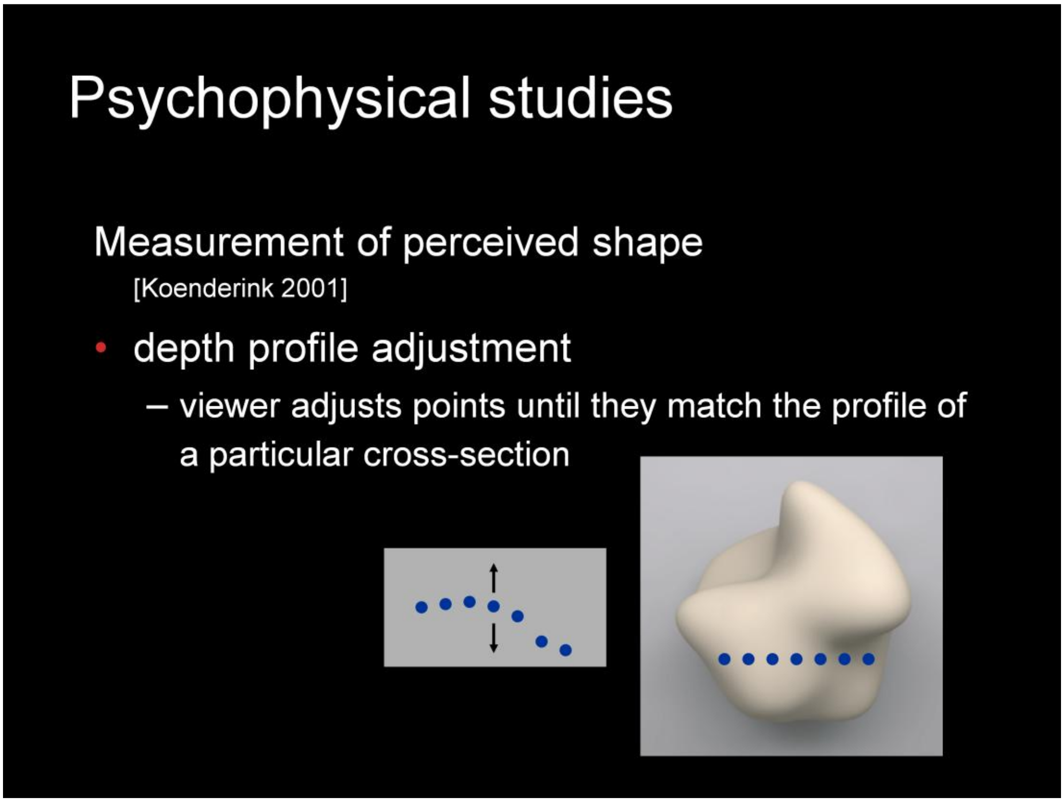

Another method is known as DEPTH PROFILE ADJUSTMENT.

Here, the viewer adjusts points
to match the profile
of a particular marked cross-section on the display.

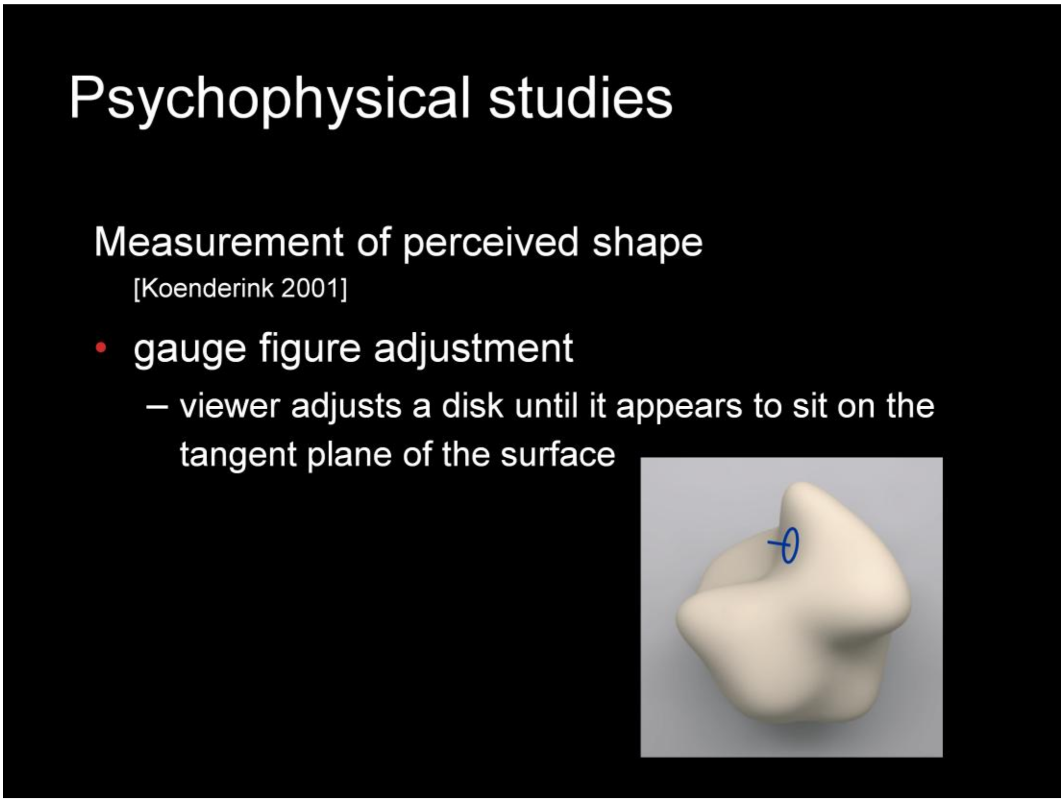

Their third method is known as GAUGE FIGURE ADJUSTMENT.

Here, the viewer uses a trackball
to adjust a small figure that resembles a thumbtack,
so that it looks like its sitting on the surface.

All of these methods are successful.

But gauge figure adjustment seems to give
the best information given a fixed number of questions.

Summary

- Specific information is conveyed by each type of line
- Psychophysical evidence suggests humans use this information
- How exactly humans interpret line drawings is still unknown

So in summary,

Each type of line in a line drawing
conveys specific information about shape.

There is a fair amount of evidence
that people use this information.

But how exactly people use this information is still unknown.

We’re very encouraged by how a combination
of computer graphics and psychophysics
can lead to answers to these questions.