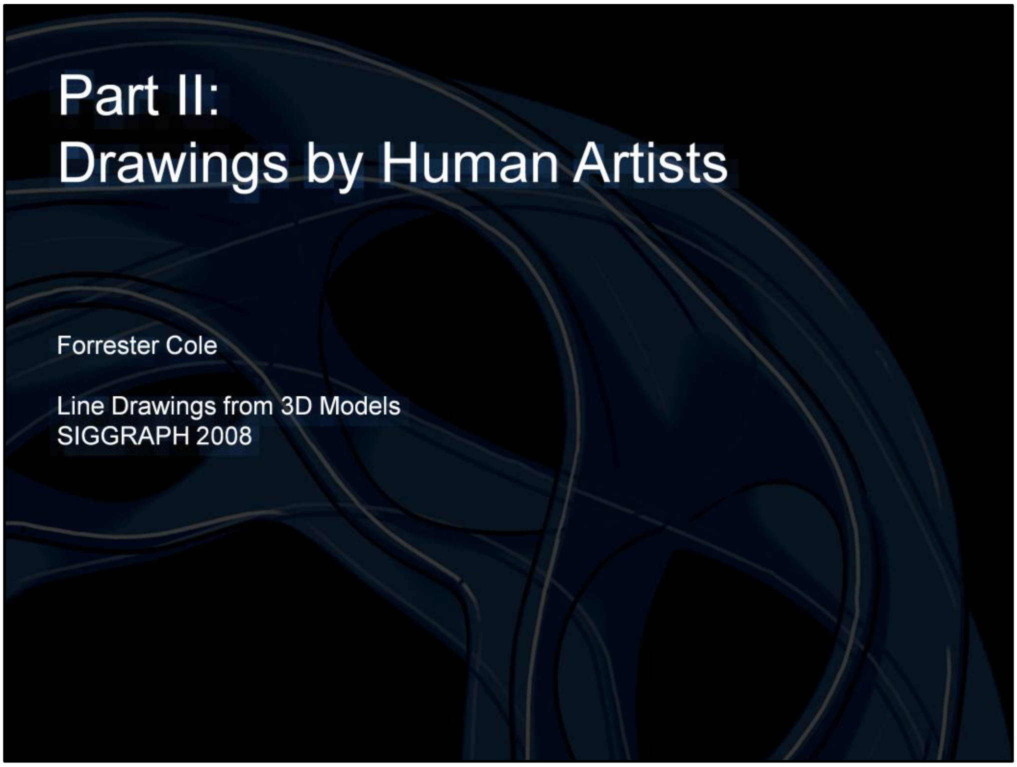
Part II:
Drawings by Human Artists
Forrester Cole
Line Drawings from 3D Models
SIGGRAPH 2008

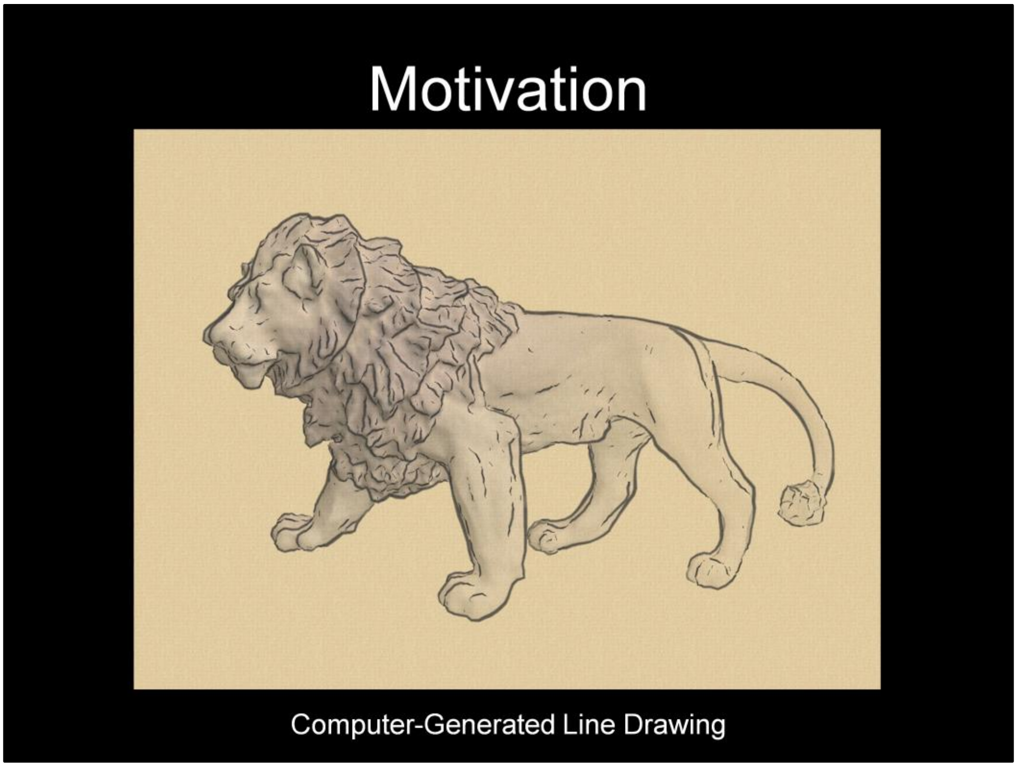

This is an example of a fairly state of the art computer generated line drawing. It uses suggestive contours, slightly stylized strokes, and some visual emphasis effects, and its a pretty nice, effective drawing.

However, its pretty obvious that the work of human artists, at least the best of them, is still in a different class. Somehow this drawing manages to say a lot more than the CG version, even though the amount of shading and lines on the page is roughly equivalent. Of course, this comparison is unfair – the cg model contains much less information than a view of a real lion.

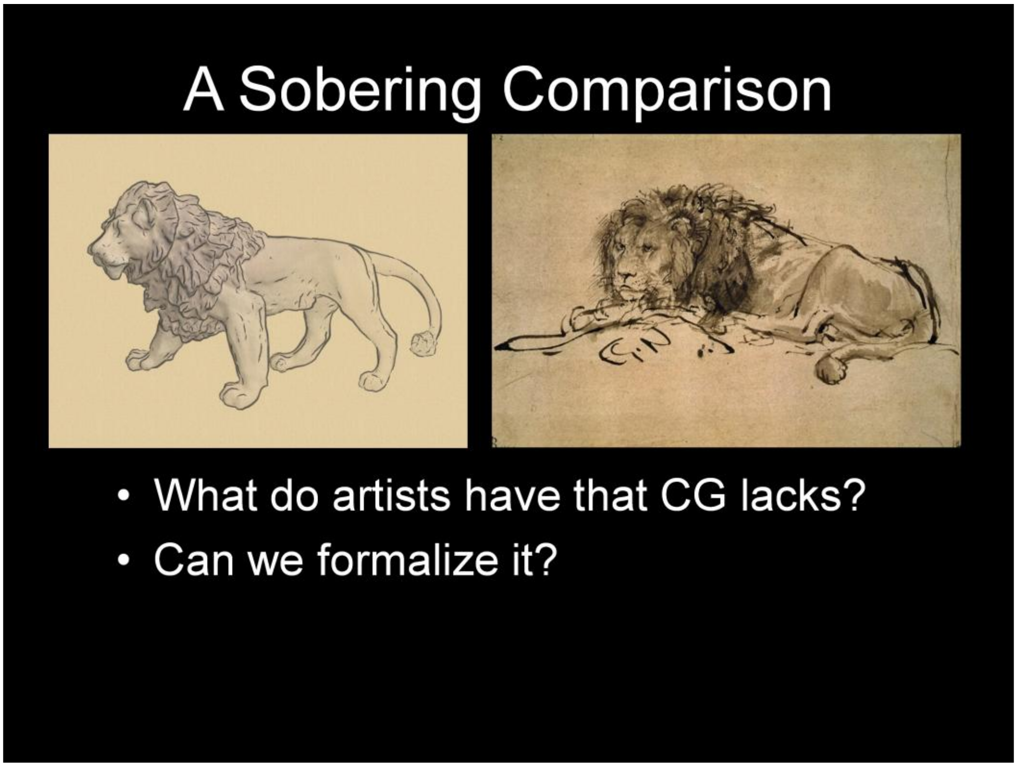

Additionally, there is no shame in not being able to draw as well as rembrandt. There are maybe a few people in history who could. Still, we'd like to learn what the best human artists are doing differently from our current algorithms, and hopefully simulate it. And to simulate it, we need to formalize it.

The Problem

- "Line drawing" is not a single style
- Immense range of approaches
  - Emotional to explanatory
  - Loose to controlled
  - Sparse to developed
  - Veridical to abstract

This course is concerned with line drawings. "Line drawing" is often used as a shorthand for a particular kind of simple, sparse, monochrome drawing, but from the point of view of a computer algorithm, human line drawings include a huge set of variations that need to be examined.

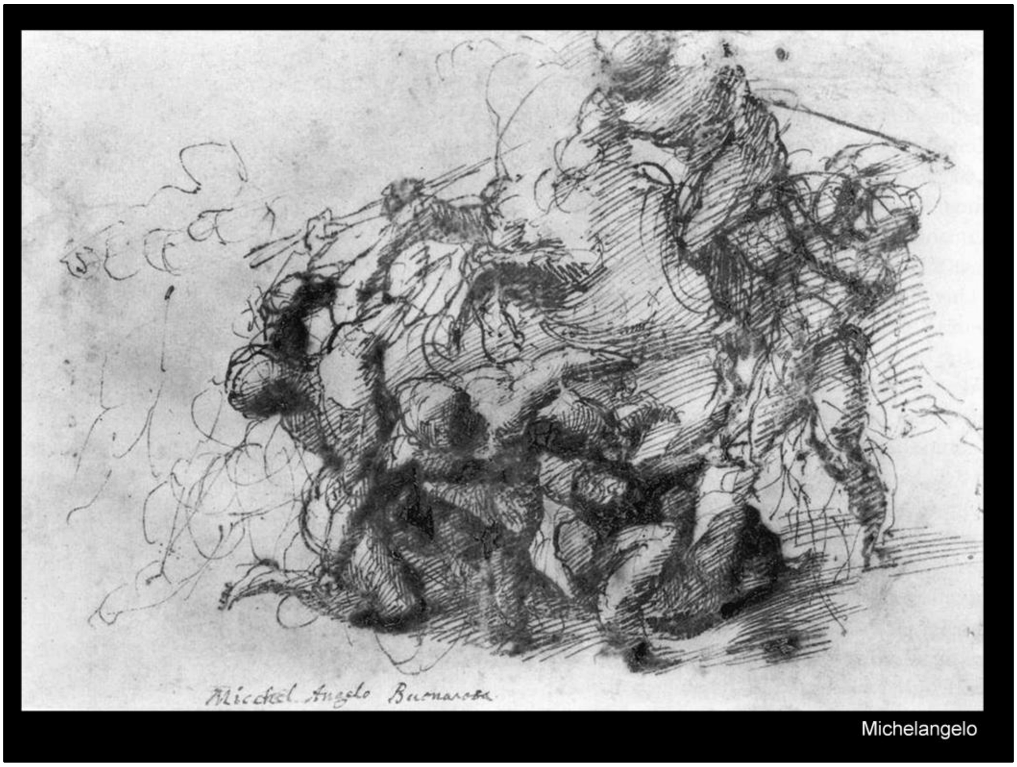

Here is a drawing that would likely be quite difficult to recreate using current computer algorithms. It's a sketch for a later painting, but Michelangelo signed it, so he must have thought that it was reasonably complete. Its hard even for a human to see what exactly the drawing is depicting – something about men struggling, maybe in battle. The point of the drawing is in the way it shows the motion, the emotional content.

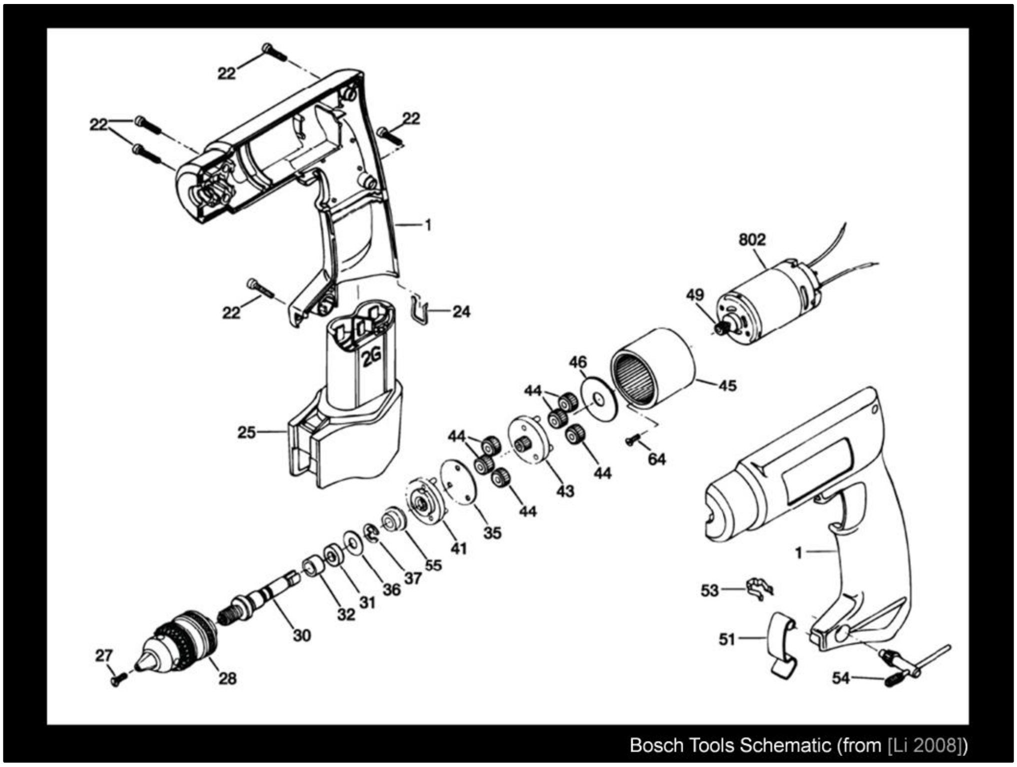

Bosch Tools Schematic (from [Li 2008])

By contrast, here is a drawing that is entirely explanatory. It exists to show the parts exactly, and part of the purpose of the drawing is that the viewer can see completely unambiguously which parts are depicted. The contrast between the two drawings is almost painful; you can feel your brain changing gears when you compare this drawing and the last.

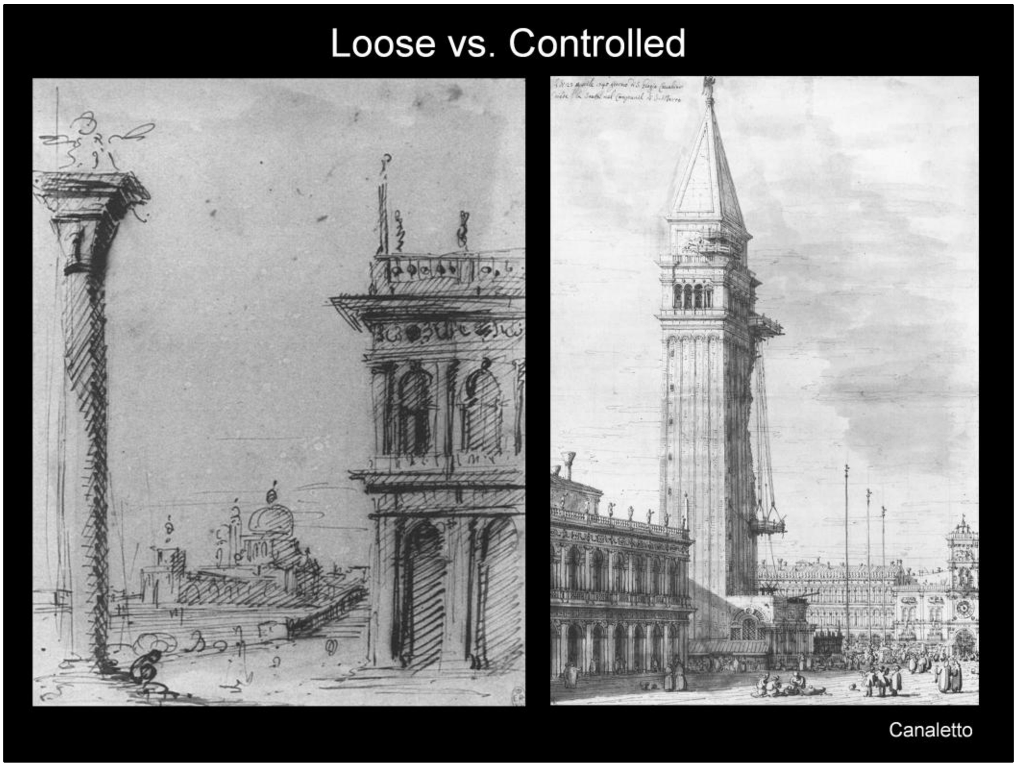

Beyond the basic purpose of a drawing, a major distinction is in how tightly controlled and precise the drawing is.

This isn't just the difference between drawing quickly and taking your time, though the drawing on the left must have been much quicker than the one on the right. A precise, controlled style is best for unambiguously showing detail (you can count the number of people in the courtyard in the drawing on the right). However, a loose style also conveys information; it says that there is more there than you are showing right now, the artistic equivalent of waving hands and saying "et cetera."

All line drawings ask the viewer to use their imagination, and a skilled artist uses this as a great point of strength. The artist can choose, however, how far to guide the viewer's imagination. The artist can leave almost everything unsaid, as in the left drawing (where only the basic outline of a classical interior can be made out), or detail almost every inch of the paper, as in the right (where you can look for a long time and still not catch all the tiny human figures).

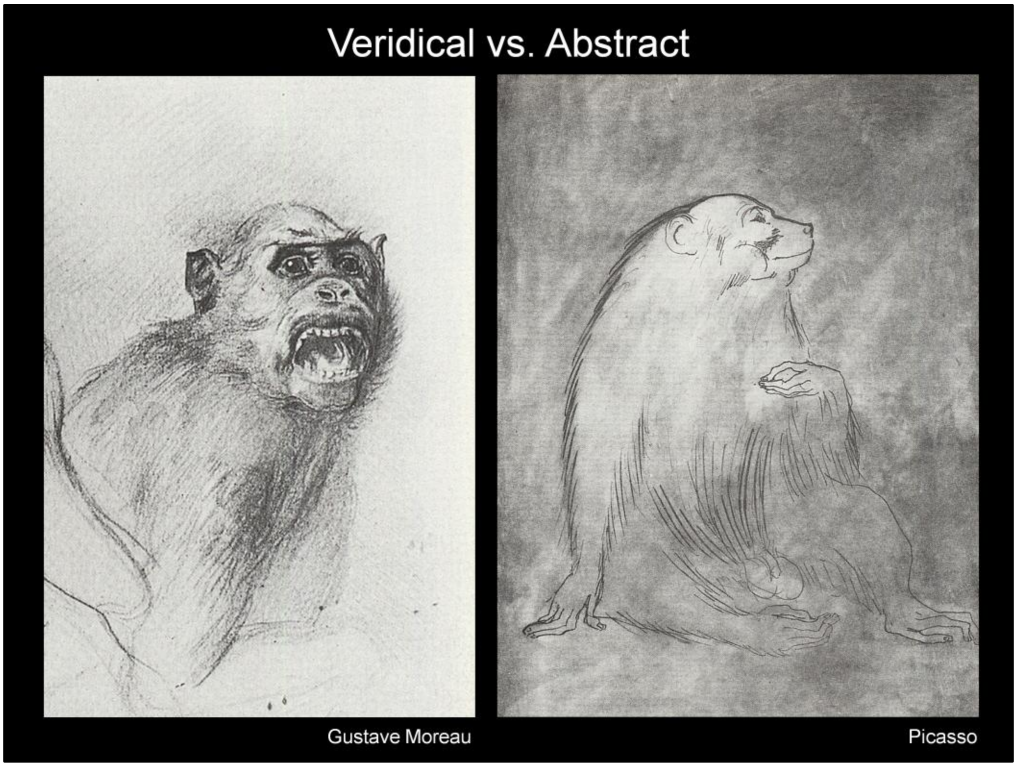

This is a different type of abstraction from using a loose or sparse style – it is a careful depiction of an abstraction of the subject. An extreme example would be a cartoon character. Computer algorithms have so far dealt with this type of abstraction only in very specialized cases.

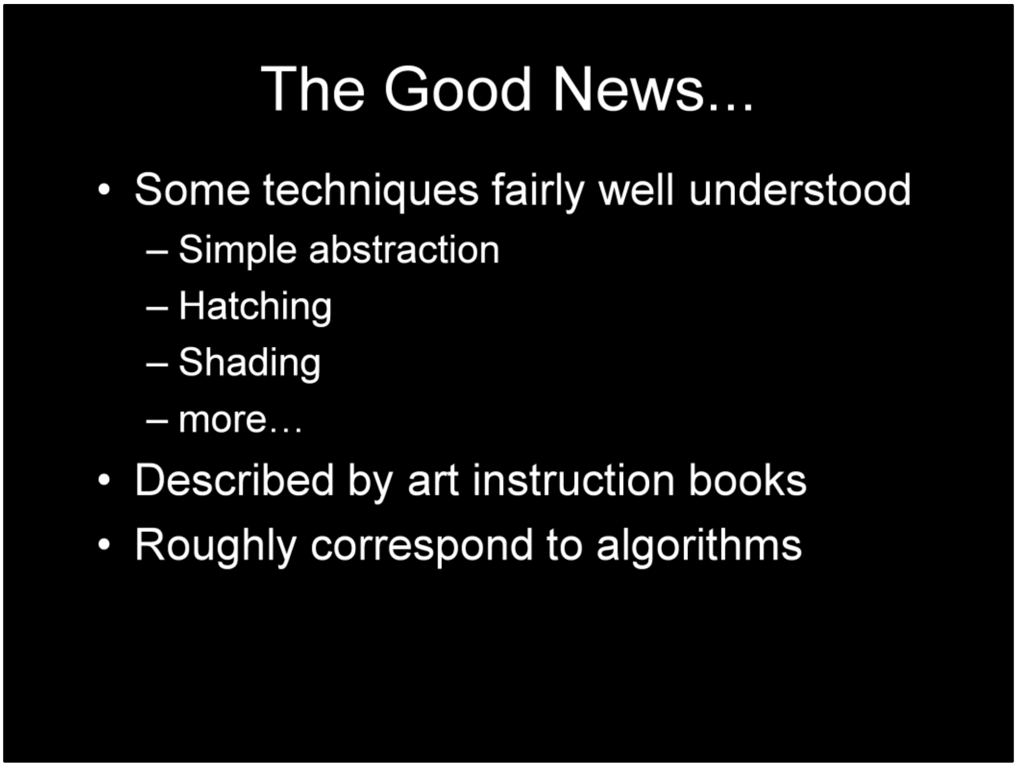

There are a number of techniques used by human artists that are fairly well understood, such as simple abstraction, hatching, and shading. These techniques are described in detail by books on art instruction, and these descriptions can be translated more or less directly into algorithms.

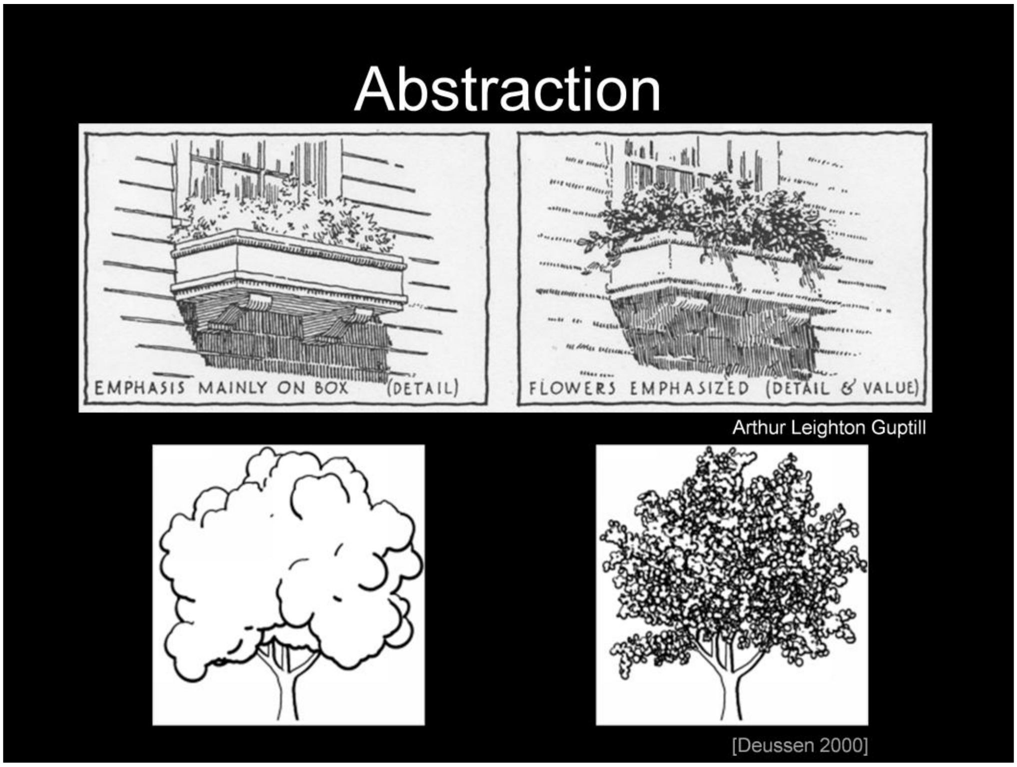

For example, this paper by Deussen et al. simulates quite closely the changes in style laid out by Guptill. The plants are made abstract by rounding the shapes, making the lines more sparse, and making the overall tone lighter. The plants are made more detailed by doing the opposite: adding detail leaf shapes and increasing the overall density of lines.

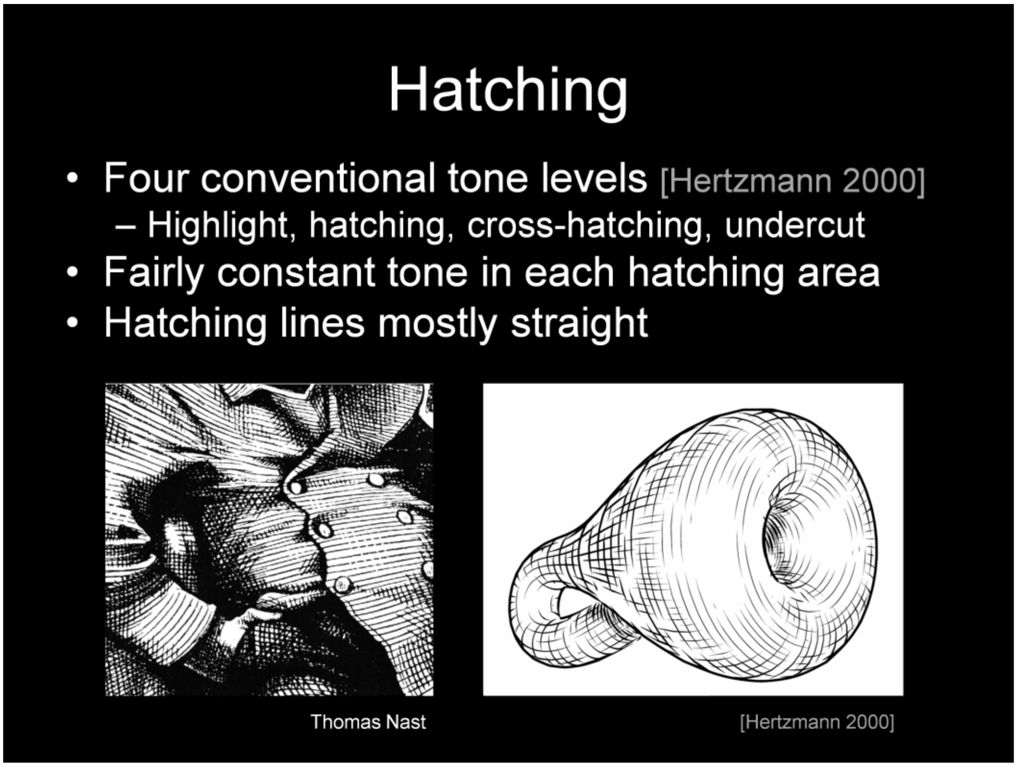

Hatching is a technique for creating shading and suggesting shape using strokes of roughly similar width and length. Rules for effective shading are well laid out in the art literature. Four more or less constant levels of tone are commonly used: highlights, or essentially blank spaces, such as on the hand or the folds in cloth; single hatching, as on the body of the jacket; cross-hatching, as in the shadows of the folds; and undercuts, essentially blacked out areas in deep shadow. An important point is that hatching lines are usually rather straight, though tending to bend somewhat with the underlying surface. Hatching lines that follow the surface too closely for too long tend to appear as markings on the surface itself. Hertzmann's 2000 paper does a nice job of working these observations into an effective algorithm, with an example result on the right.

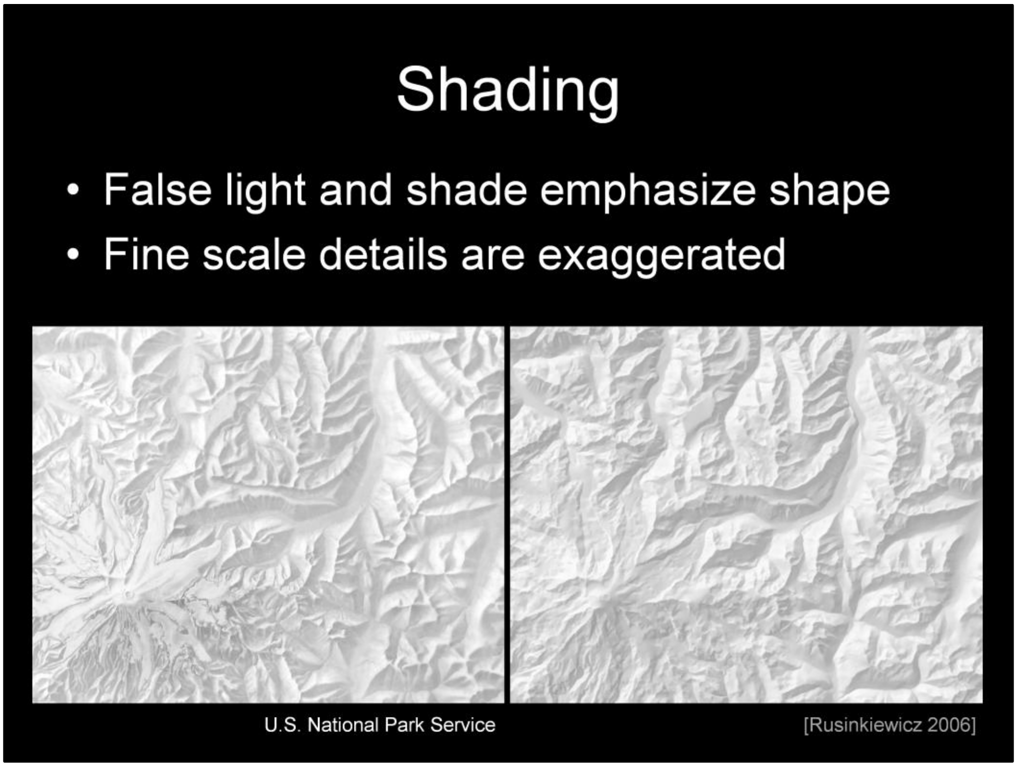

Shading style varies widely, but it is common for artists to exaggerate shading and use false light to bring out the shape of their subject. In the case of terrain relief, rules for shading are codifed well enough that an artist's shading (left) can be reproduced quite closely (right). This algorithm uses a multi-resolution approach to locally manipulate the light direction much as an artist might, giving small details extra prominence and the curving sides of valleys roughly similar levels of shade.

... And The Bad News

- Not everything so well explained
- Even “simple” drawings are quite complex

Unfortunately, not all the interesting effects are so well explained in the art literature. Even the simplest of line drawings by human artists contain effects that are difficult to effectively describe with our current knowledge.

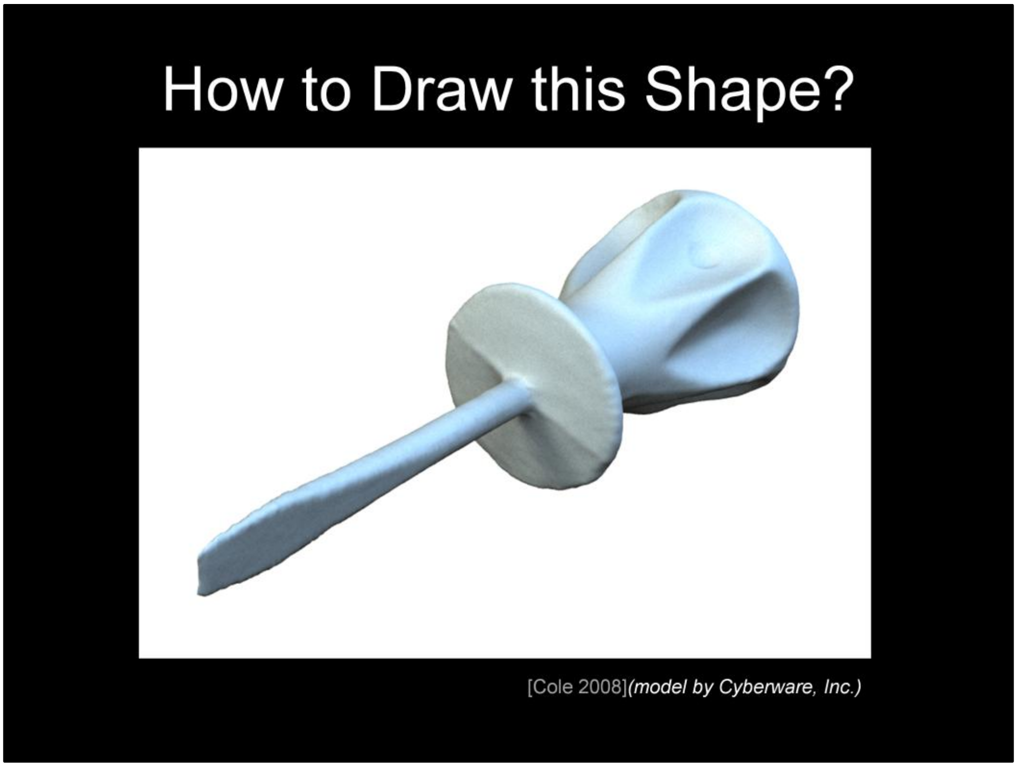

For example, consider making a simple, pure line drawing of this shape, with the object of convey the shape as well as possible. Ignore abstraction, hatching, or shading entirely.

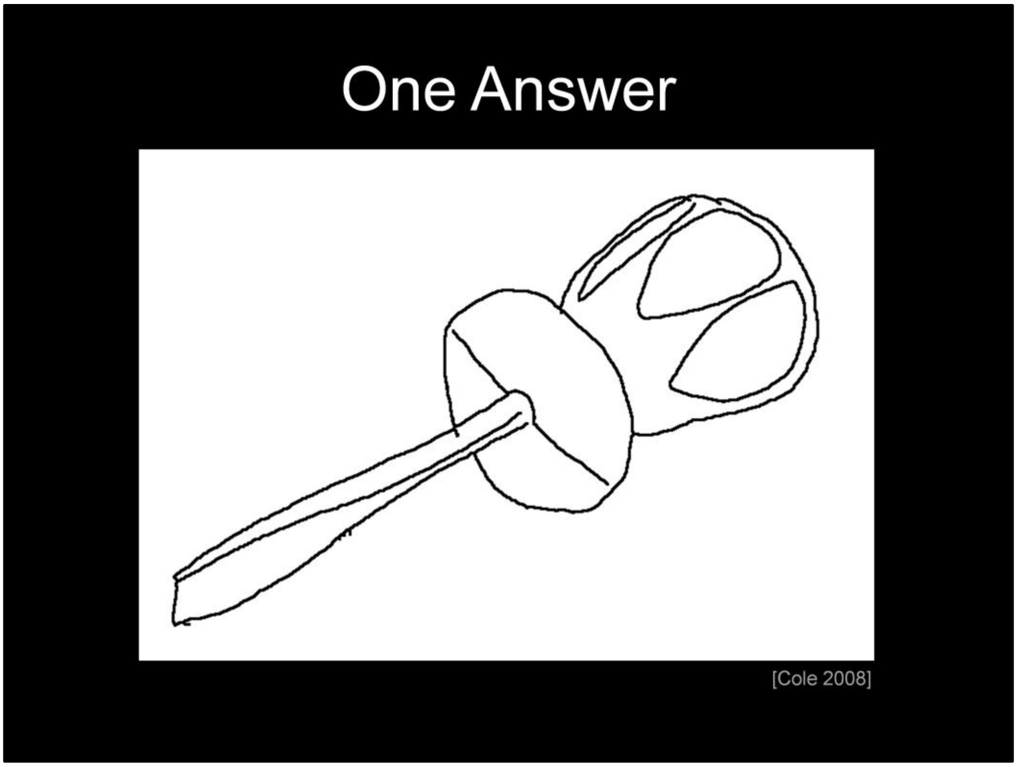

This is one answer to that prompt from an actual artist. Note that the drawing has been scanned and filtered so that each line is constant thickness and strength. This artist made some fairly obvious decisions, such as drawing all the silhouettes, but also some idiosyncratic decisions, such as extending the crease line at the front of the screwdriver far along the shaft.

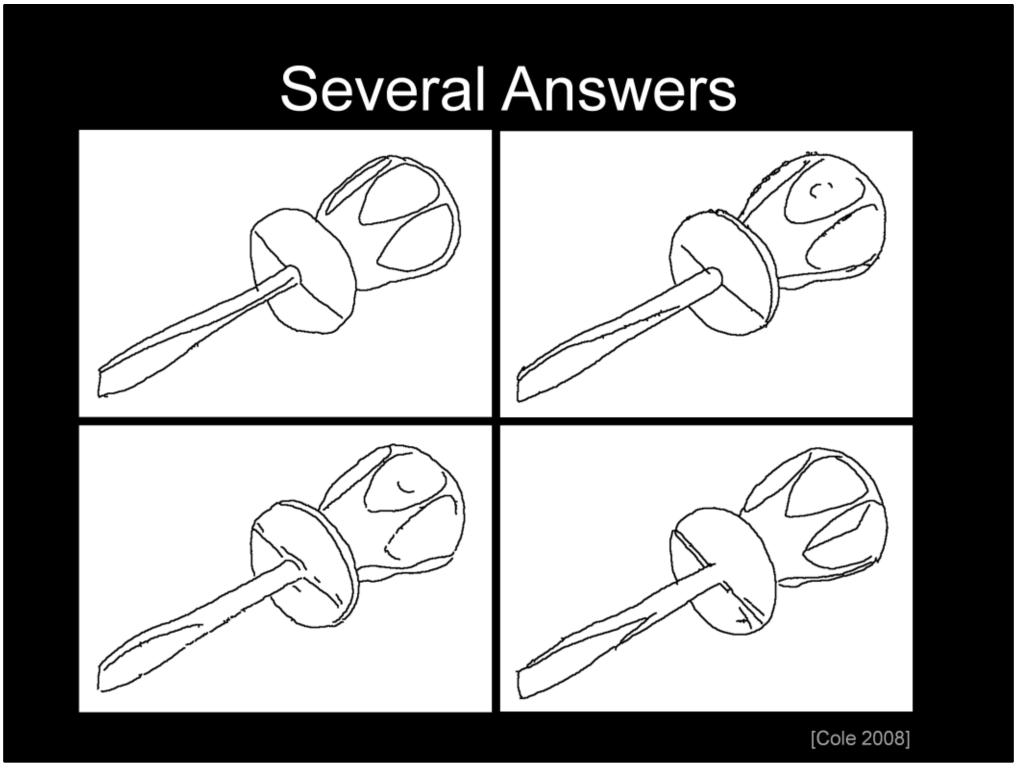

Other artists made similar decisions in many cases, but each drawing is unique. One shared property between all the drawings is the way the rear of the screwdriver is represented by three independent components, though some artists chose to complete each loop and some chose to leave the rest of the loop implied. Each drawing has its strengths and weaknesses, but they all get the important bits pretty much right. Lets look at how we might try to formalize these lines.

Feature Line Models

- Use CG drawing algorithms as models
- Known algorithms
  - Image edges
  - Image ridges and valleys
  - Geometric ridges and valleys
  - Extension and anticipation of contours

There are several line drawing algorithms that can provide models for artists' lines. These algorithms will be described in detail later on in the course, but for now I'll just explain briefly.

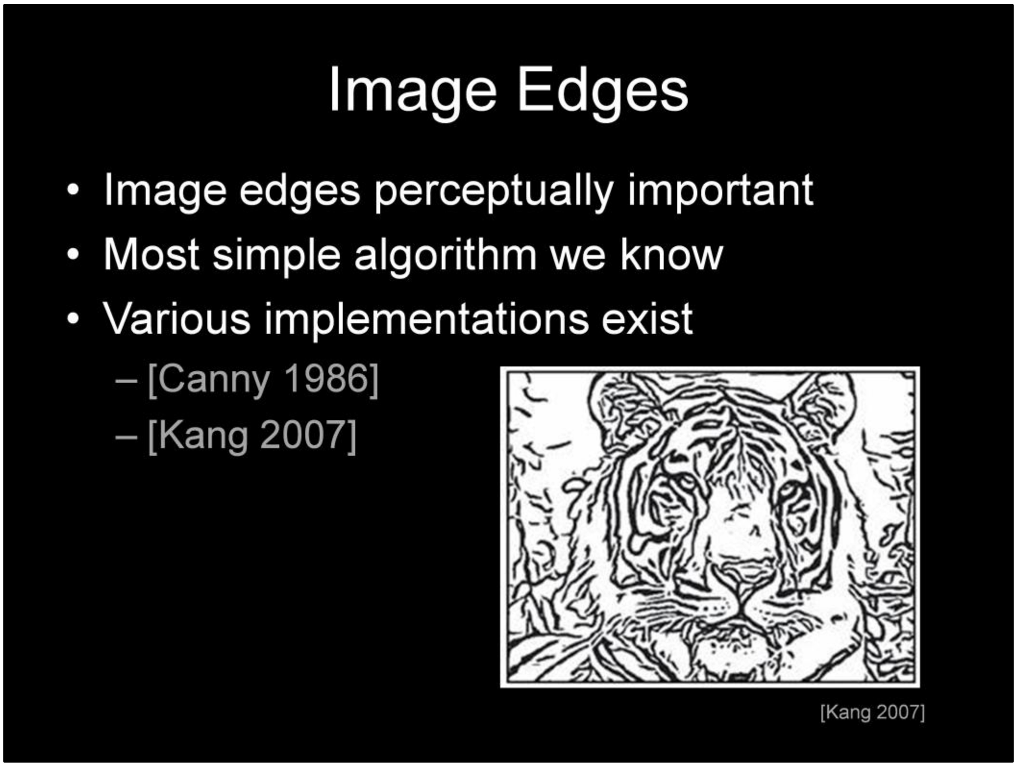

We know that edges in an image are perceptually important. Extracting edges from an image is also simple, in fact, it is probably the simplest line drawing algorithm we know. Variations on the theme include Canny edge detection and more sophisticated methods such as Kang 2007.

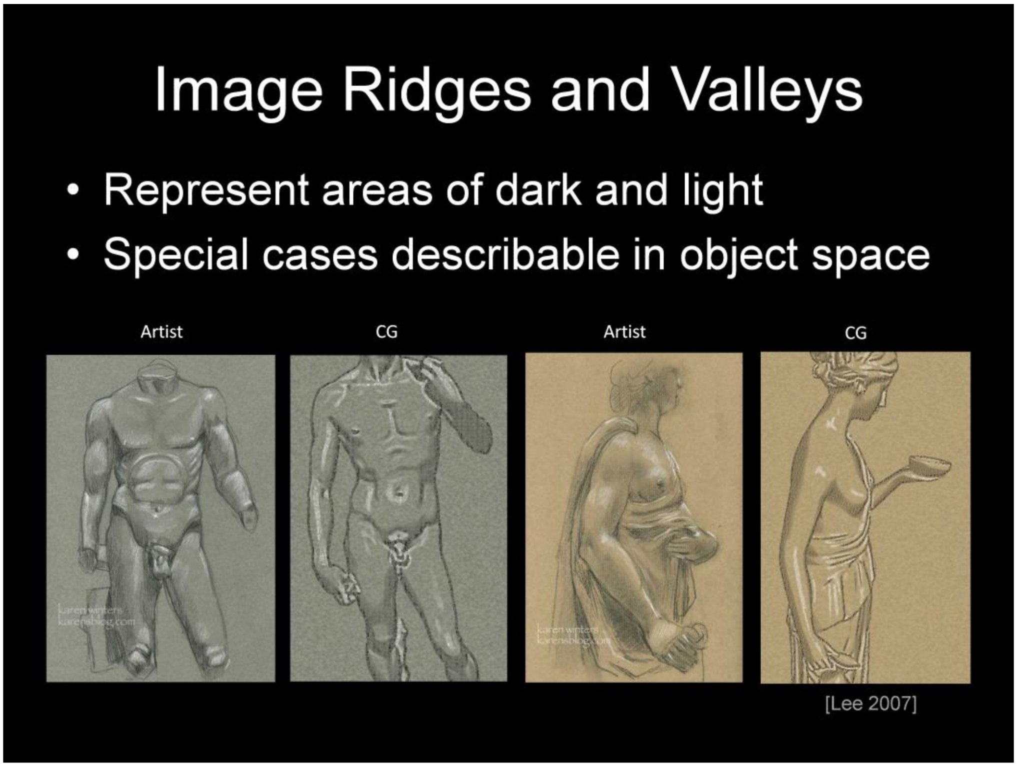

Besides image edges, we can look at ridges and valleys of the image. There is some reason to believe that artists sometime abstract shading with lines, which would correspond to making lines along the image intensity ridges and valleys. Although these lines are most obviously image space features, later on we will see how to describe some special cases in object space.

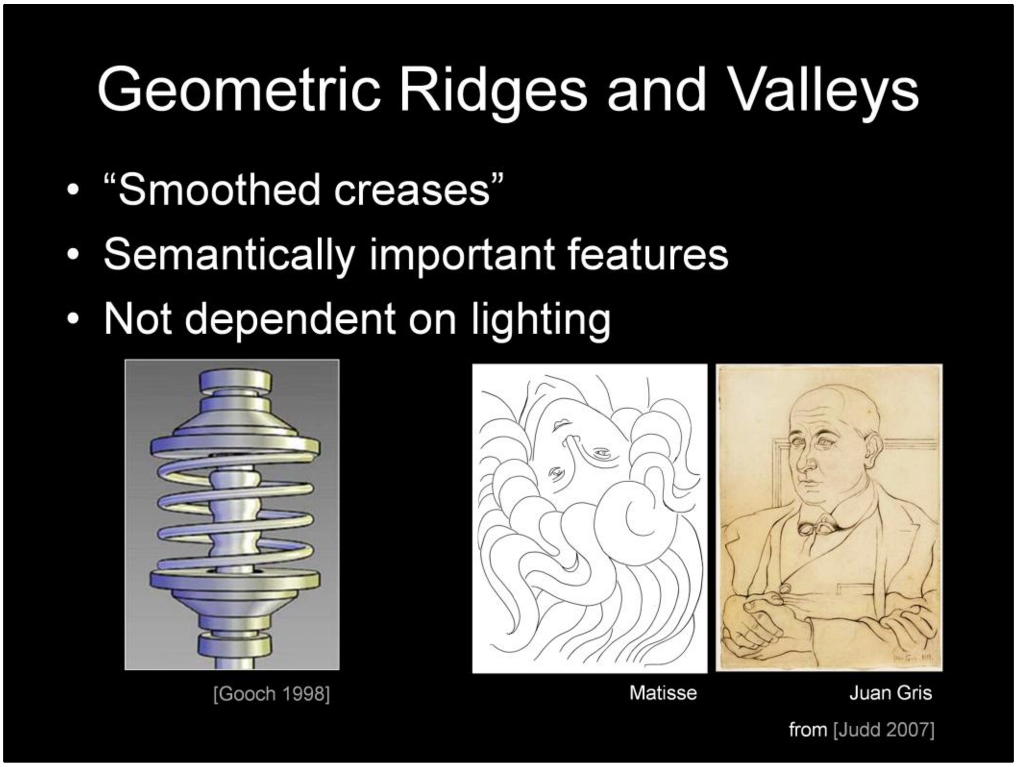

Creases (lines of normal discontinuity) are obviously important lines to draw, as shown by the left image. Geometric ridges and valleys are essentially smoothed creases, i.e., lines where the normal is changing rapidly in one direction. These lines are intrinsic properties of the shape, and are therefore not dependent on lighting. Some examples of artists renderings using ridges are shown at right. The most prominent ridge-like features in both are along the nose.

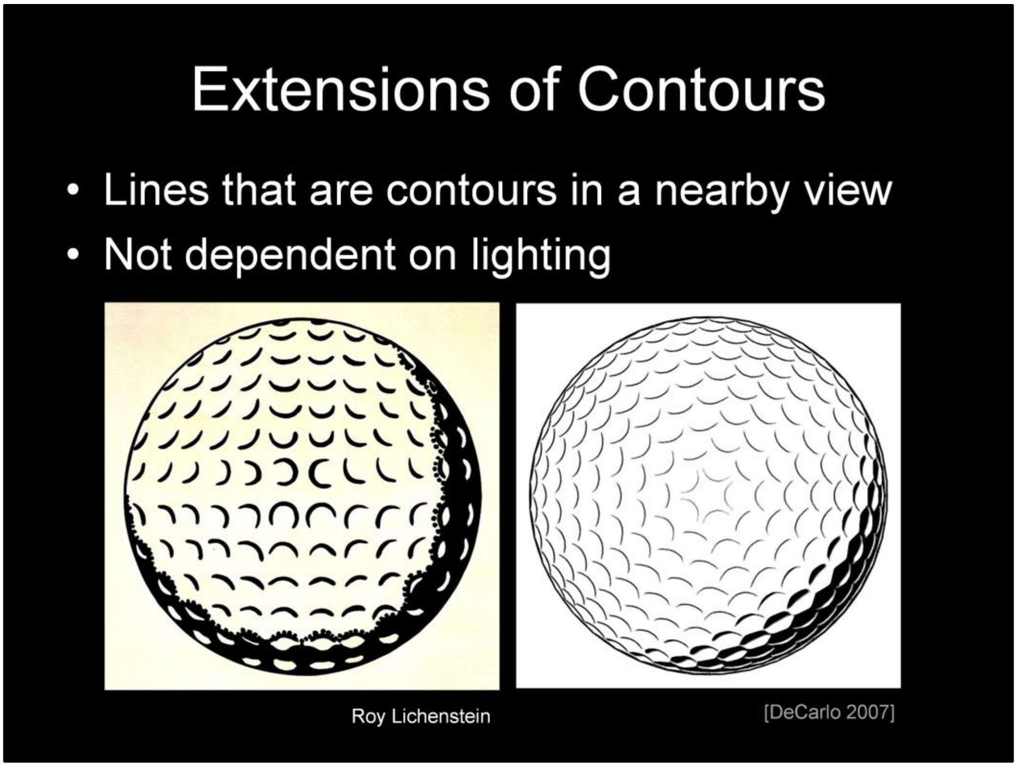

Occluding contours are known to be important lines to draw. Artists also seem to extend these contour lines, and draw lines where no true contour exists, but a contour would exist if the viewpoint were slightly changed. These lines are also not dependent on lighting – note that the lines on the lower right of the golf ball (closer to the shade) are the same as the lines on the upper left (closer to the light). Later in the course you'll see how to formalize these lines to produce images such as the one on the right.

Which Model Is Best?

- All seem to work sometimes
- But which, where, and when?

We are left with the question of which model of lines to use. All the models appear to work in some cases, but it is not clear when and where to use each one.

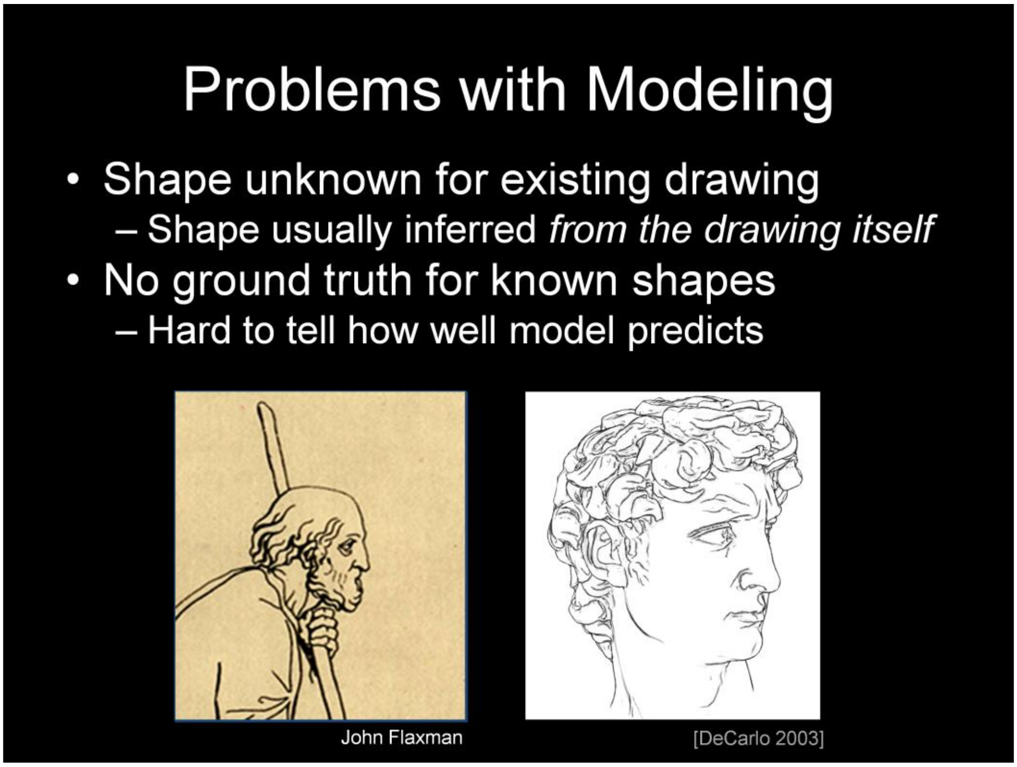

There are greater problems with the modeling approach, however. Even if we come up with a model that produces very similar lines to a given drawing, we have no way to evaluate how well the model truly matches. The shape that the artist considered when making the original drawing is unknown, and in fact, we often beg the question by inferring the shape from the drawing itself. Further, even when we do know the exact shape (such as when we have the 3D model), we cannot make a direct comparison back to the original drawing, since the resulting drawings are completely different. We need a way to directly compare an artist's drawing with our models.

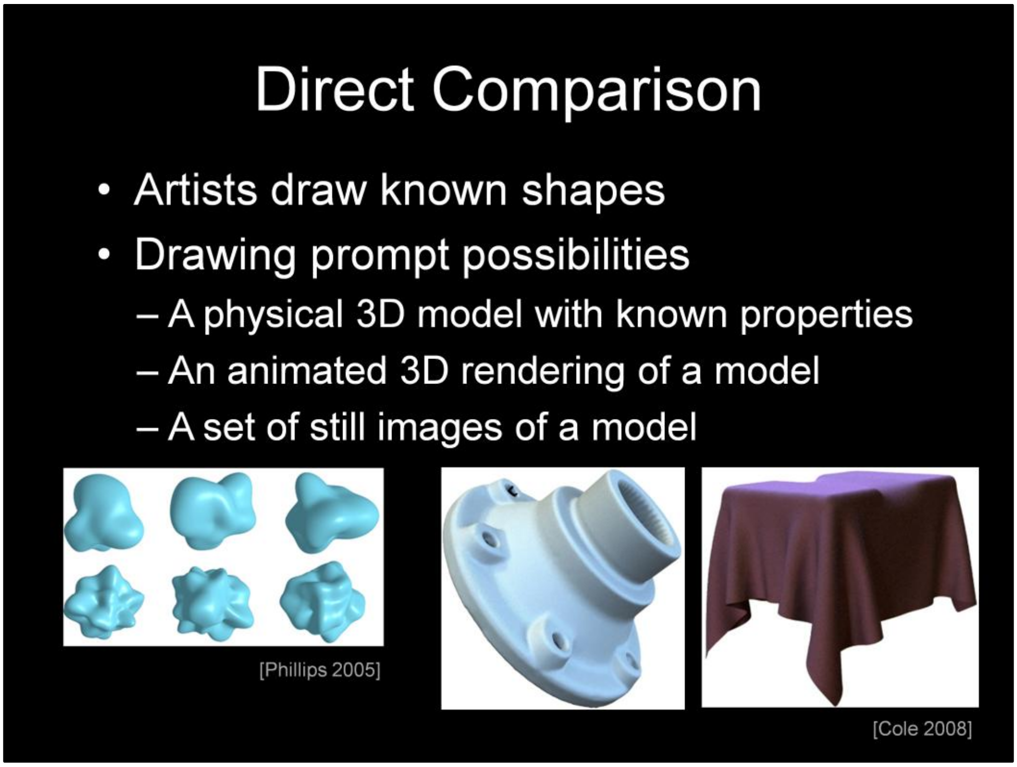

Recently a few studies have attempted to capture drawings of known shapes. The artists are asked to draw under controlled conditions, and their results captured. There are at least three options for drawing prompts: a physical 3D model, and animated rendering, and a set of still images. Phillips 2005 uses all three, while one of our current papers does only the last.

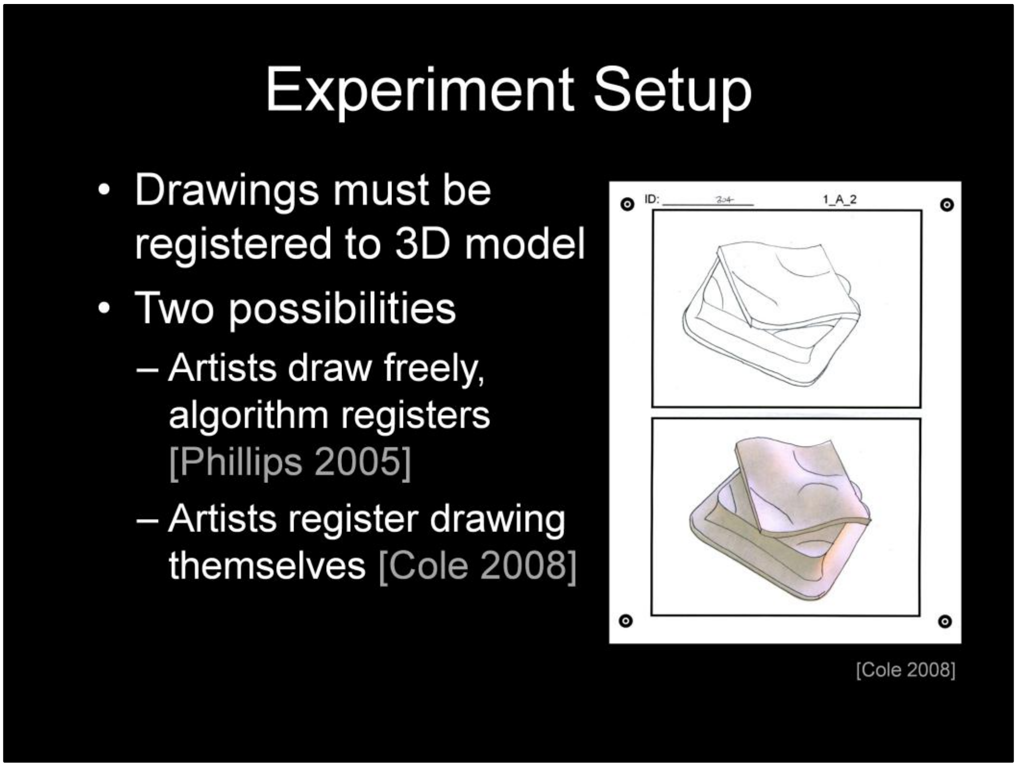

Of course, for the drawings to be useful they must be registered somehow to the 3D model. Phillips allows the artists to draw freely, and then uses a matching algorithm to register the drawings to the model. While this method works for any prompt and provides the artist with maximum freedom, it does not offer good accuracy and the matching problem is hard and ambiguous.

Another option is to ask the artists to register the drawings themselves. The artists first make a freehand drawing, as in the top of the page, and then copy their drawing on to a faint shaded version of the prompt. This method gives good accuracy and works around the matching problem by having the artists solve it for us.

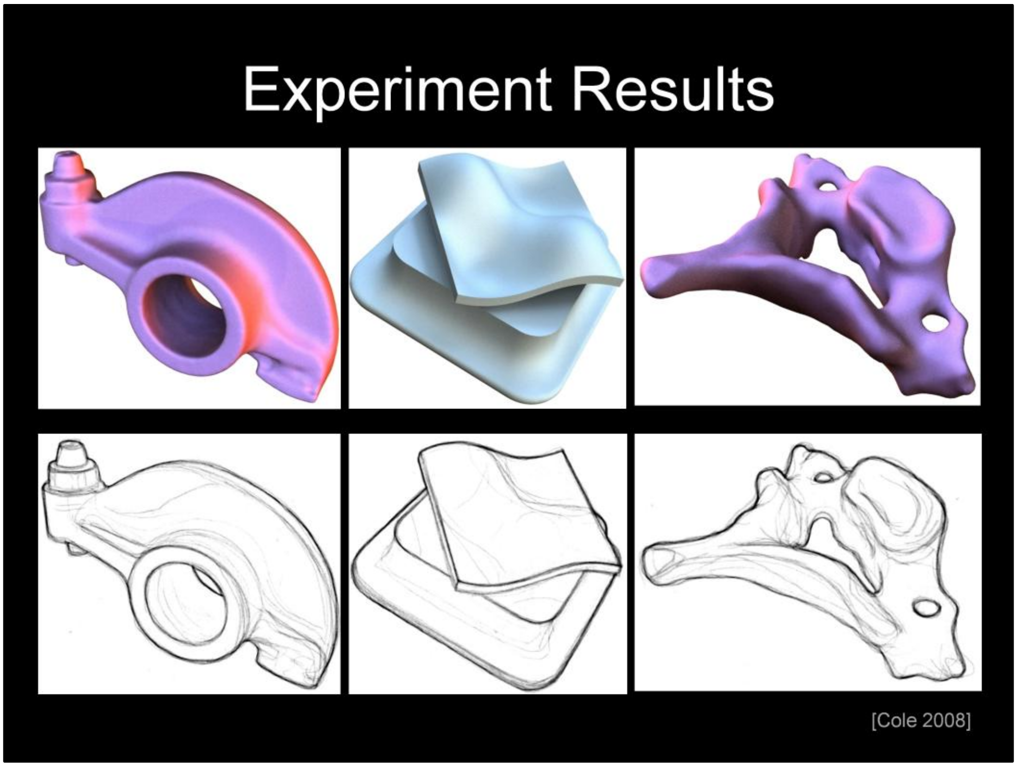

Here are some example results from our recent study. At the top you see the prompt image, and at bottom the average of many different artists drawings of that prompt. The artists seem generally guided by image space shading features, especially edges, which confirms the intuition that artists "draw what they see." Occluding contours are obviously very important, but features such as the circular ridge of the rockerarm are also highly agreed upon by different artists.

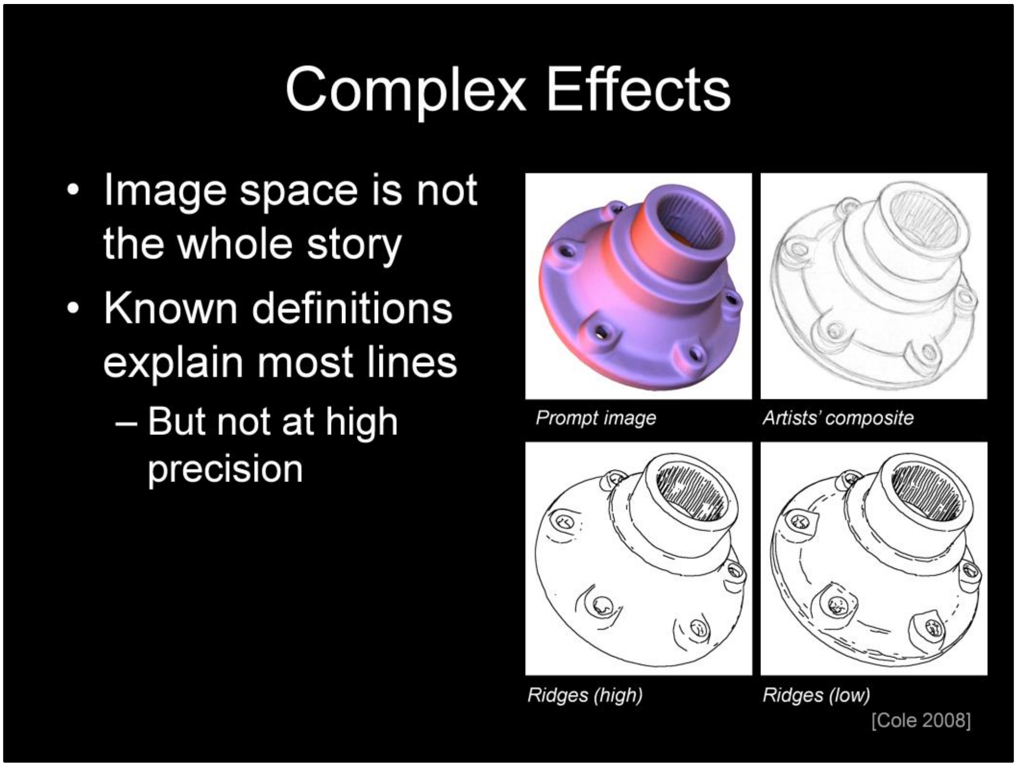

Image space features are not the whole story, however. The average drawing for this shape seems to be best described by the geometric ridges and valleys, because the circular lines remain relatively constant through changes in the image intensity.

However, while all the artists' lines can be classified as ridges or valleys, just selecting ridges and valleys by strength does not give us a matching to the artists' drawings. The artists made some complex editing decisions to leave out the lines along the top of the mounting holes, while leave in the lines around the whole circumference of the shape. As you can see from the two lower images, using a high enough threshold to exclude the short valley lines excludes the long ones as well, while using a low enough threshold to include the long lines includes the short lines as well.

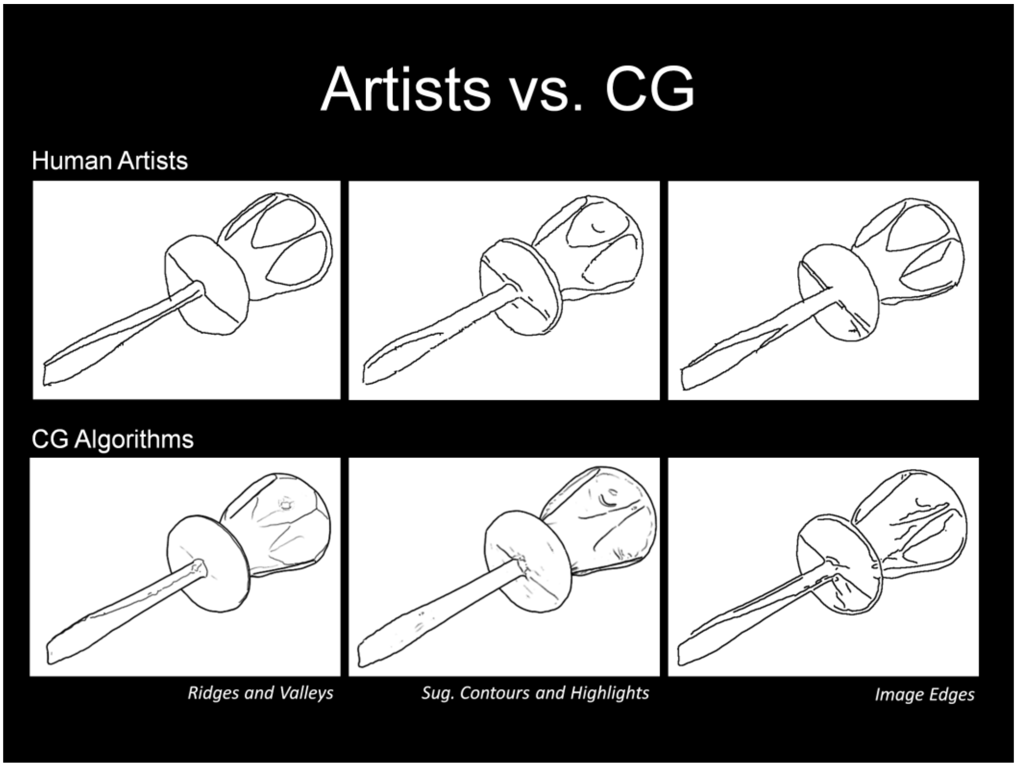

This brings us back to the screwdriver example. While not mentioned before, the screwdriver drawings were made by artists in our study. We therefore have the 3D model and view, and can create CG line drawings for a direct comparison. At top are three of the humans' drawings, and below are three CG drawings, roughly selected to match the humans' drawings.

In the first pair, the shaft line is extended all the way to the hilt, and the loop features at the rear are completed. Note, however, that in the ridge and valley image, lines between loops are shared, unlike in the artist's drawing.

In the second pair, the dimple in the top of the screwdriver is recorded, and the lower section of a loop in left implied as in the artist's drawing. However, the shaft line is missing, as is the back of the top loop.

In the third pair, the crooked line in the middle of the side loop is present, along with the shaft line and some of the messy lines at the front of the screwdriver. However, the top loop is almost omitted in the image edges drawing, while it is completely present in the artist's drawing.

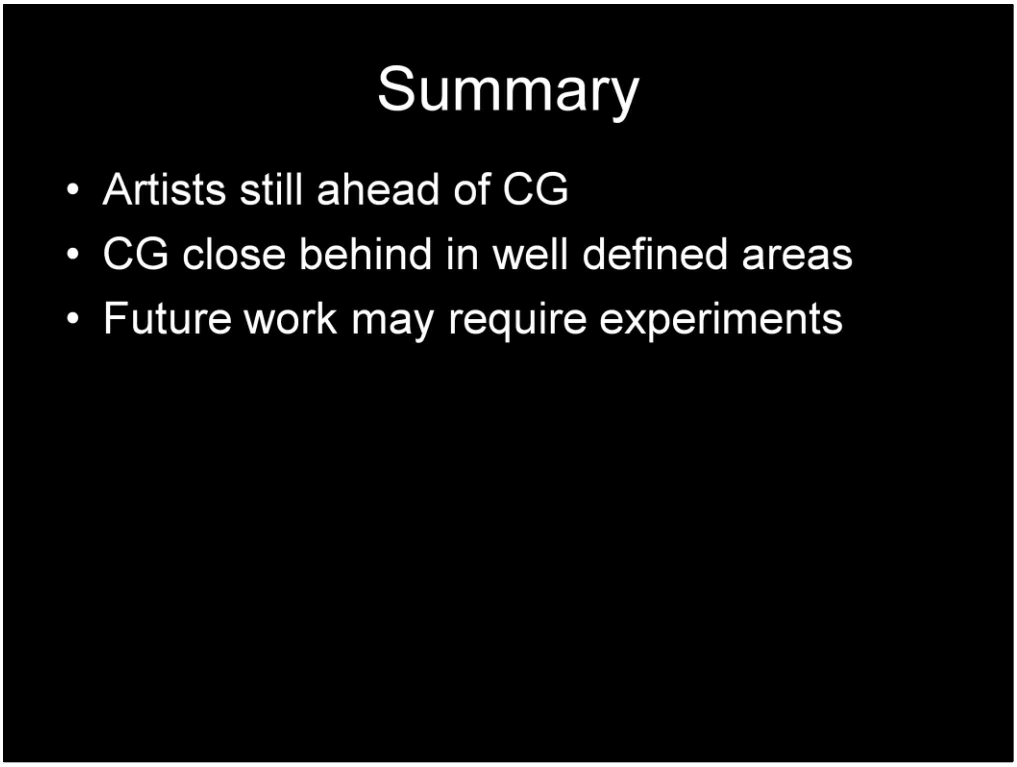

In conclusion, human artists are still ahead of CG algorithms in many respects. However, CG is close behind in the case where the artistic technique is well explained and can be easily formalized. Since no well known formalizations exist for much of artistic technique, however, future study of human line drawings may require experiments with artists in order to make progress.