Nonphotorealistic rendering

Computational Photography, 6.882

Bill Freeman

Fredo Durand

May 9, 2006

Drawing from: NPR Siggraph 1999 course, Green et al.
npr_course_Sig99.pdf

Photorealism

- Physically realistic computer graphics rendering
- Images with photographic quality (eg Vermeer, 1632-1675, accused by critics of being cold, inartistic, and displaying ‘spiritual poverty’).

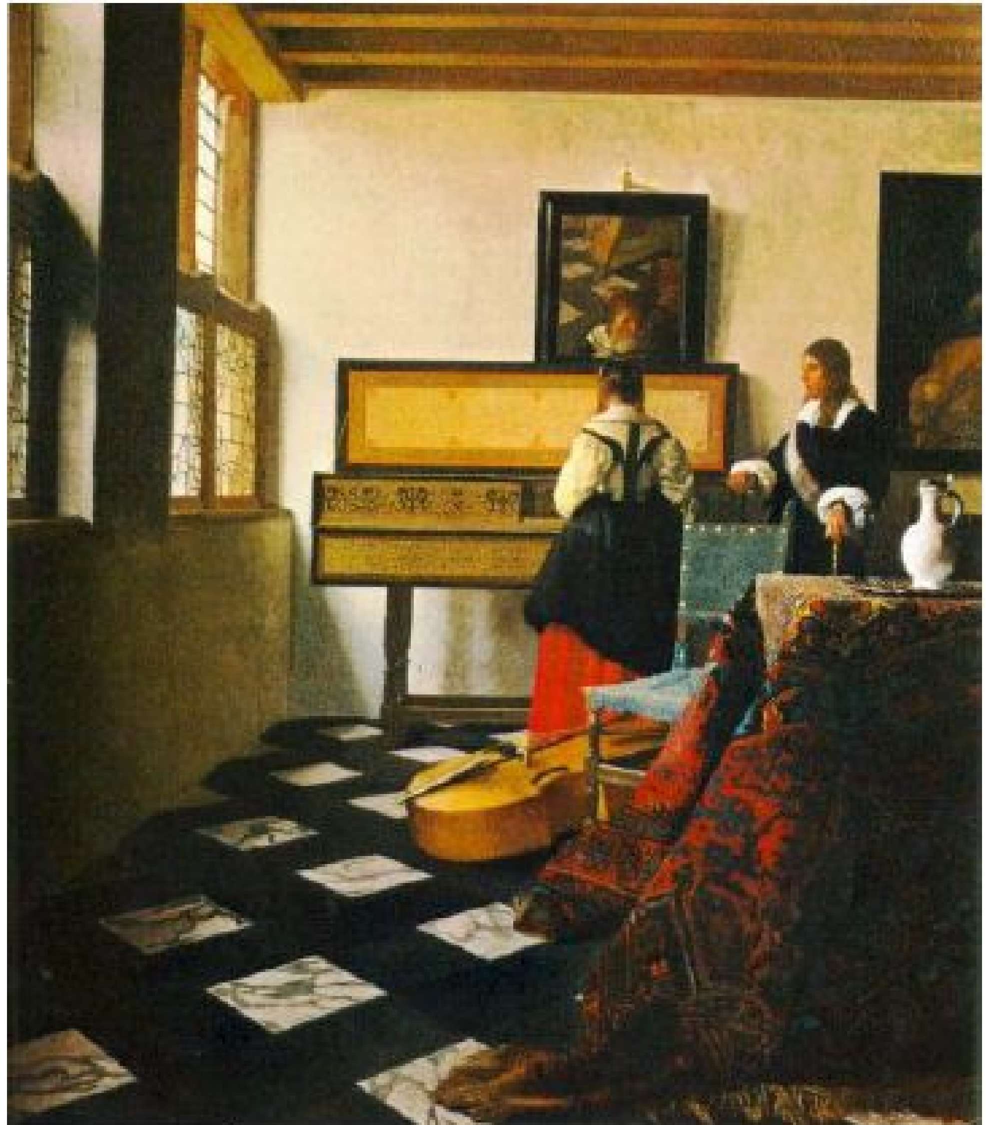

The Music Lesson by Vermeer.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Are these images non-photorealistic renderings?

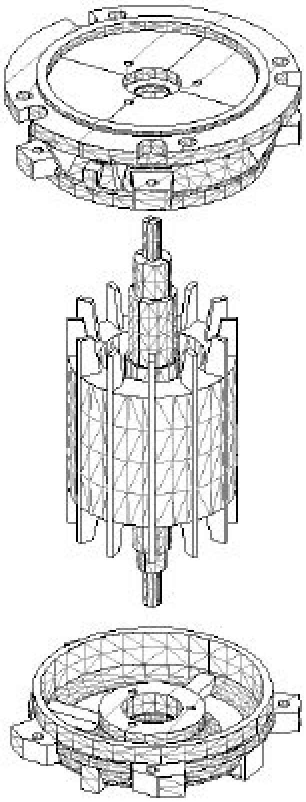

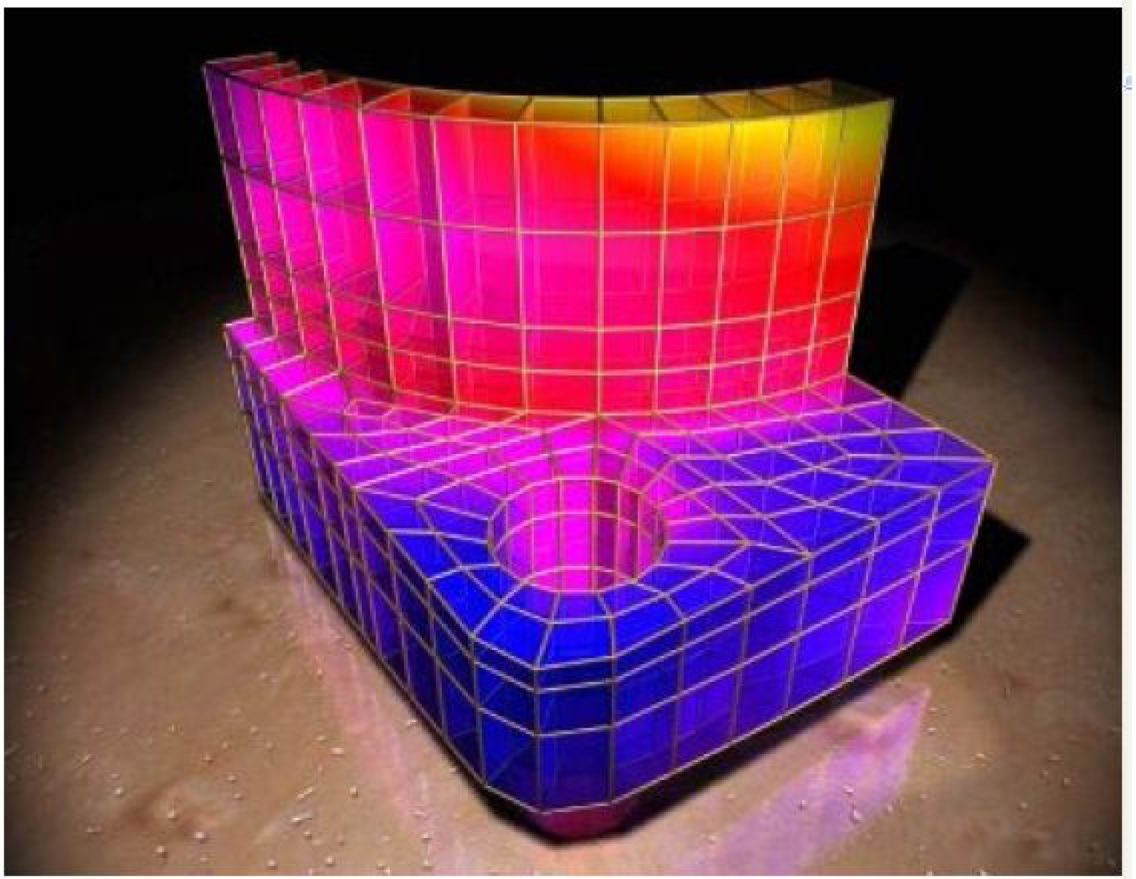

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Non-photorealistic rendering

- Expressive, artistic, painterly, interpretative rendering.
- Not aspiring to realism.
- Early work: natural media emulation
  - Pen and ink
  - Watercolor
  - Oil on canvas
- Attempts to capture the low-level style.
- Simulations of technical illustration.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

An Invitation to Discuss Computer Depiction

Frédo Durand

NPAR 2002

Most authors also agree that the term "non-photorealistic" is not satisfying [NPA00]. The border between photorealism and non-photorealism is fuzzy, and the notion of realism itself is complex [Fer99]. Thomas and Ollie tell an enlightening anecdote about Walt Disney [TJ81], p. 66. Disney would keep asking his animators for more *realism*, which was a cause of misunderstanding, since no one would qualify Disney's animation as realistic. Their interpretation is that he meant *convincing* rather than realistic.

Comparing photorealism and NPR
(Stuart Green)

	Photorealism	NPR
Approach	Simulation	Stylization
Characteristic	Objective	Subjective
Influences	Simulation of physical processes	Sympathies with artistic processes; perceptual-based
Accuracy	Precise	Approximate
Deceptiveness	Can be deceptive or regarded as 'dishonest'; viewers may be misled into believing that an image is 'real'	Honest – the observer sees an image as a *depiction* of a scene
Level of detail	Hard to avoid extraneous detail; too much information; constant level of detail	Can adapt level of detail across an image to focus the viewer's attention
Completeness	Complete	Incomplete
Good for representing	Rigid surfaces	Natural and organic phenomena

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Kurzweil CyberArt Technologies

We create **software** that creates **art.**

Raymond Kurzweil, a name synonymous with the potential of machine intelligence, brings you the latest in CyberArt advancements with **AARON**, the Cybernetic Artist.

AARON is not your ordinary screensaver. Developed by Harold Cohen over a period of nearly thirty years, and productized by Kurzweil CyberArt Technologies, Inc., **AARON** is the first fine art screensaver to utilize artificial intelligence to continuously create original paintings on your PC.

With his continued commitment to human assistive technology, Raymond Kurzweil is proud to introduce this breakthrough screensaver. AARON is offered as shareware! Download a free trial copy of **AARON** and try it for yourself. If you decide to keep the AARON software, we ask you to register it. However, the product will not time out, and the payment requirement is based on trust.

Download
a free trial copy of AARON

Watch
AARON paint right now

Read
a message from Ray Kurzweil

Learn
about AARON's history

Get Help
downloading & using AARON

Poetry
Download Ray Kurzweil's Cybernetic Poet

Ray Kurzweil and Harold Cohen examine AARON's artwork

From the film *The Age of Intelligent Machines*

Non-Photorealistic Rendering - The Artist's Perspective

Simon Schofield
Slade School of Fine Art, University College London

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Statistical techniques to simulate expression

Statistical techniques used to simulate expression - Bad

For us problems arise when statistical techniques are used to simulate a medium's usage by human hand and the decisions of the artist. Many extent systems exhibit such simulation. As described, we too were guilty of attempting to simulate the artist's decisions.

Aside from AI die-hards, its has become an established principal that computers are not inventive or expressive in themselves but are tools through which expression is articulated by users [12]. Hence a word processor handles individual characters but does not construct sentences, and a MIDI system processes musical note, but does not construct tunes. It is interesting that NPR developers often seem to overlook this principal and delegate the task of constructing a completed painting to the computer. The consequences of doing so produce results that are aesthetically similar to allowing computers to generate sentences or tunes - they tend to be either chaotic and unintelligible or flat and predictable.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

"Paintings are not solutions to well-posed problems…"

Statistical techniques are highly inappropriate for imitating the decisions of the artist. The most-often seen form of NPR is the pseudo-impressionist painting - a technique which often relies on randomness to determine the placement, shape and color of marks across the picture plane. It is a constant surprise to see how effective this technique is. However, this technique's similarity to real impressionist painting is extremely trivial – it implicitly reduces the efforts of the most important artistic movement of the last century to experiments in controlled noise. The marks in any impressionist painting are highly specific, even more so than in earlier smoother types of painting, in that the very nature of there visibility made them of heightened concern to the artist.

Neither paintings, nor indeed clusters of marks, decompose into a set of algorithms. Paintings are not solutions to well posed problems; they are determined by whim, a desire to explore and articulate uncertainty. It is probably safe to say, that while we can understand and simulate a medium's physical phenomena, we do not understand and cannot simulate the way people will use that medium.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

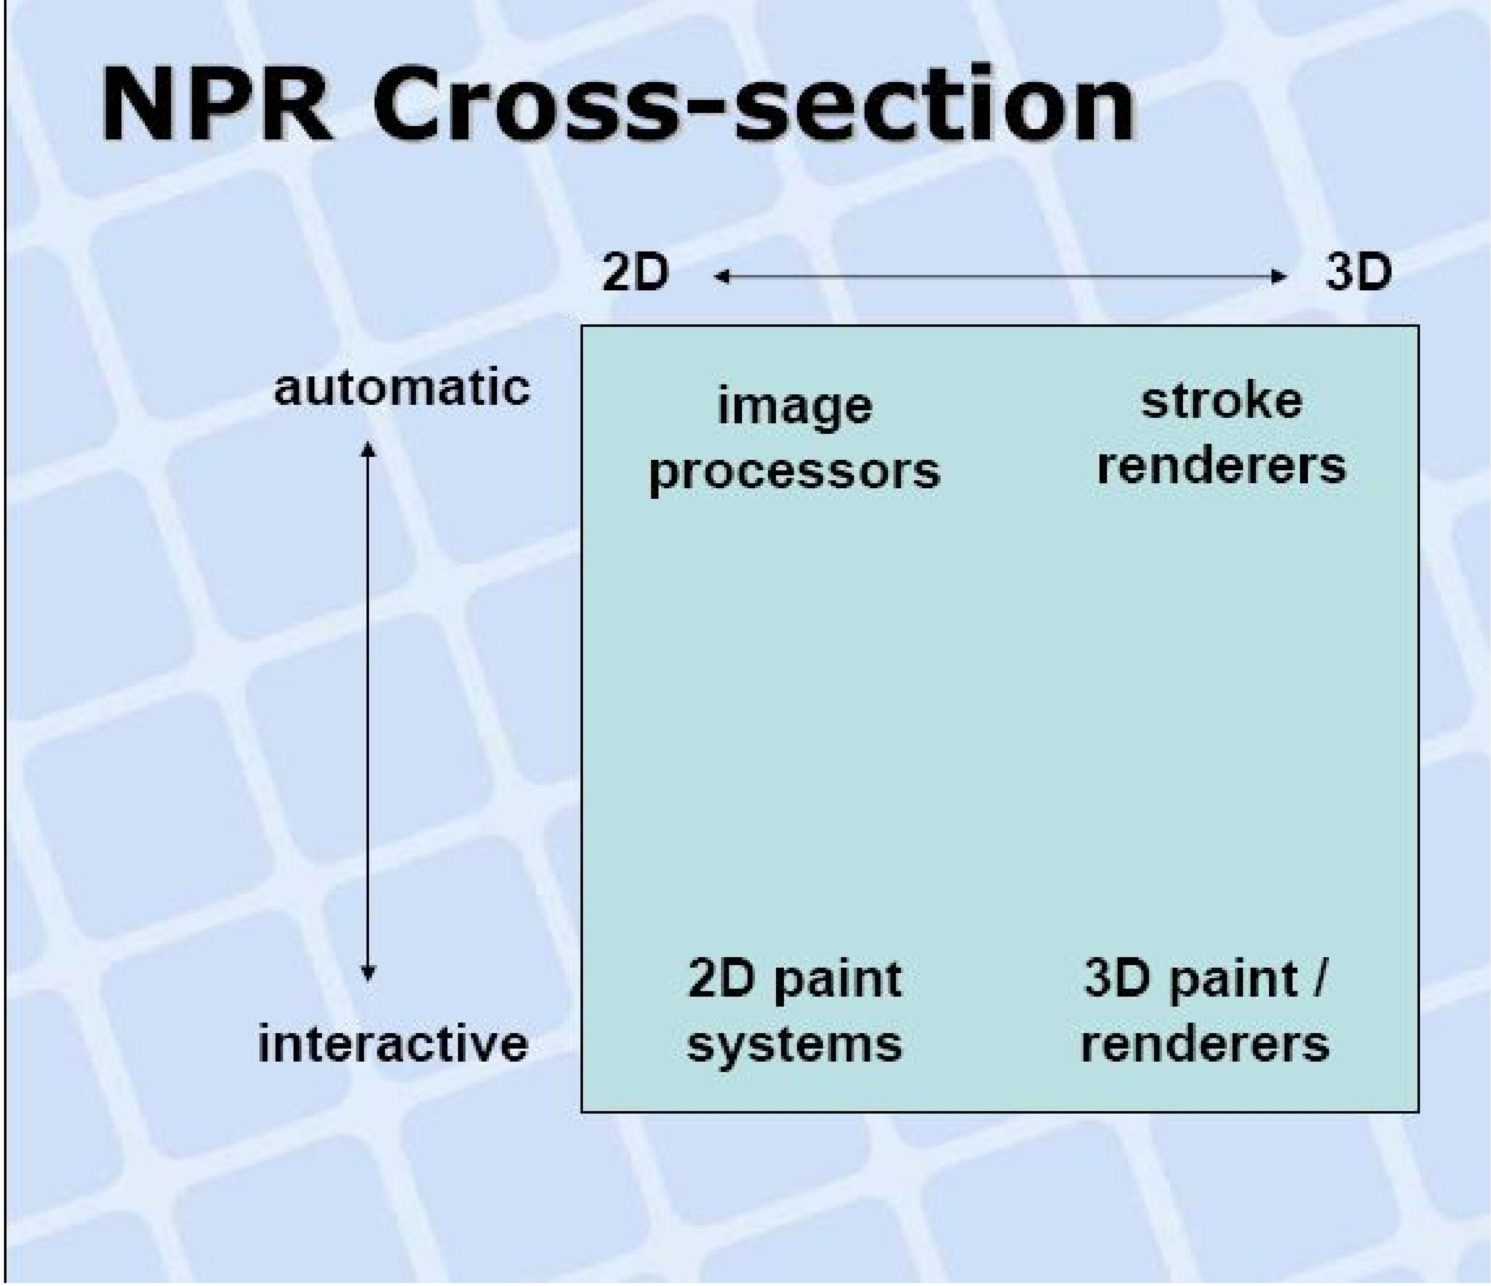

Daniel Teece

http://pages.cpsc.ucalgary.ca/~mario/npr/projects/sigg03/lec8/hand_1.pdf

Organization of NPR methods

- Automated methods
  - 2-d processing
  - 3-d processing
- Interactive methods
  - 2-d processing
  - 3-d processing

Organization of NPR methods

- Automated methods
  - 2-d processing
  - 3-d processing
- Interactive methods
  - 2-d processing
  - 3-d processing

2/2.5 D, no user intervention

Hertzmann has described an approach to hand painting an image using a series of spline brush strokes. A painting is built up as a series of layers of progressively smaller brushes [Hertzmann98].

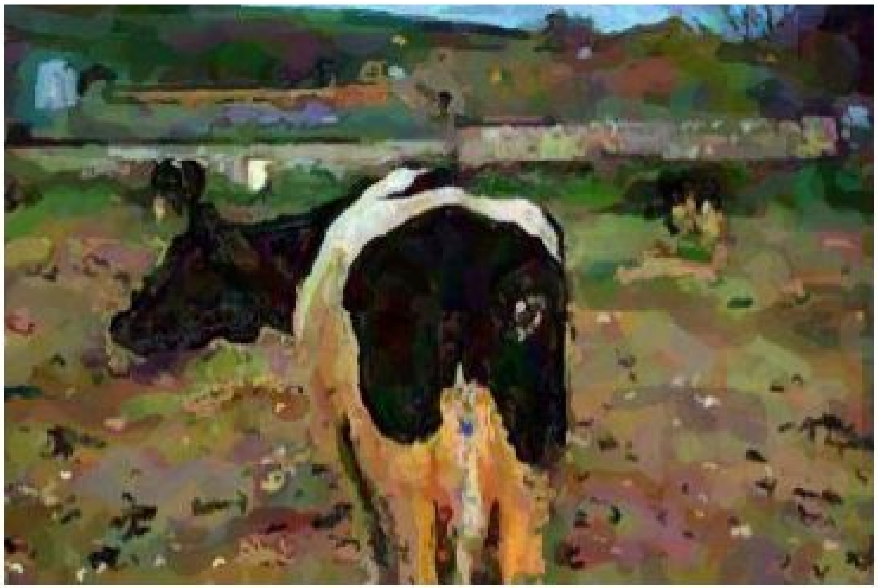

Figure 5: Painterly Rendering with Curved Brush Strokes of Multiple Sizes, from [Hertzmann98]

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Reference images Layers

http://www.mrl.nyu.edu/publications/hertzmann-thesis/hertzmann-thesis-72dpi.pdf

Issues in image style translation

- Fitting
- Translation

Learning Style Translation for the Lines of a Drawing

WILLIAM T. FREEMAN
Mitsubishi Electric Research Labs and MIT Artificial Intelligence Laboratory
JOSHUA B. TENENBAUM
MIT Brain and Cognitive Sciences Department
and
EGON C. PASZTOR
Mitsubishi Electric Research Labs and MIT Media Laboratory

ACM Transactions on Graphics, Vol. 22, No. 1, January 2003, Pages 33–46.

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

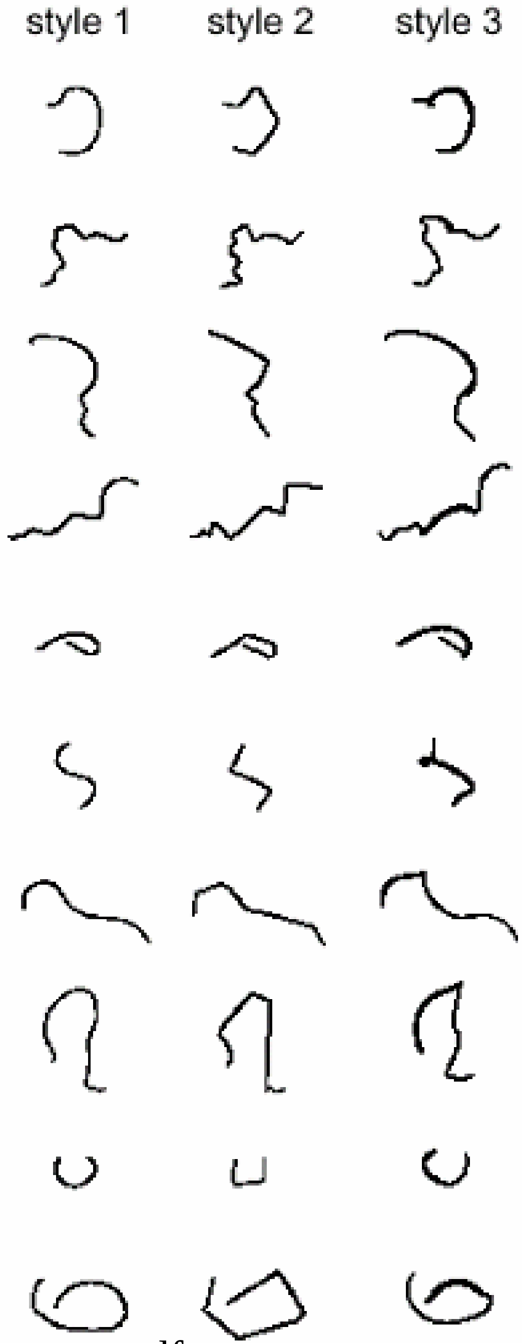

Figure 1: Training data. An artist drew the same lines in several different line styles, which we call generic, jaggy, and brushy. 123 lines were in the training set, in each style. By computer these were scaled in size over 6 scales, ranging from 0.3 to 2.0, over 4 angles, and flipped left/right. The outer product of all such manipulations yielded a training set of 5904 line elements.

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input traced line drawing

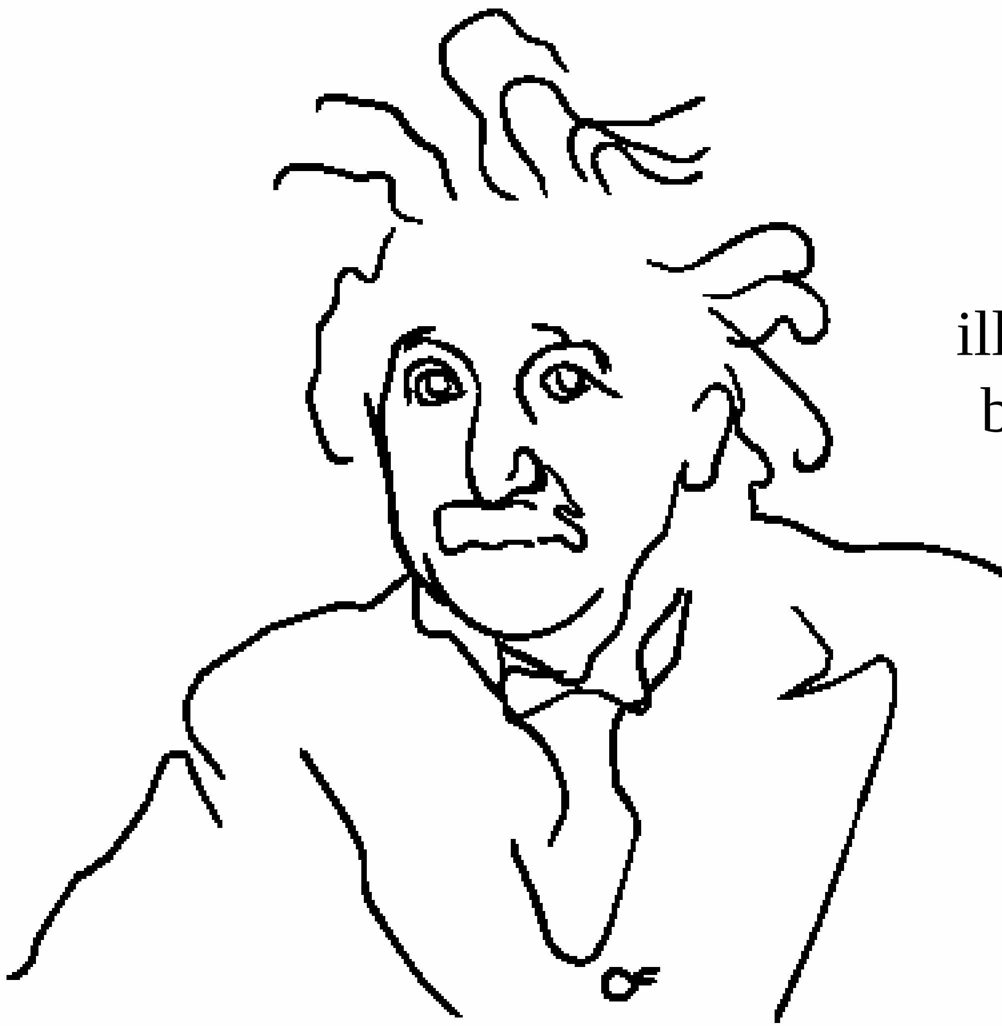

Figure 2: Line drawing traced from photograph of Einstein.

This example will illustrate the tension between fitting and translation

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input drawing

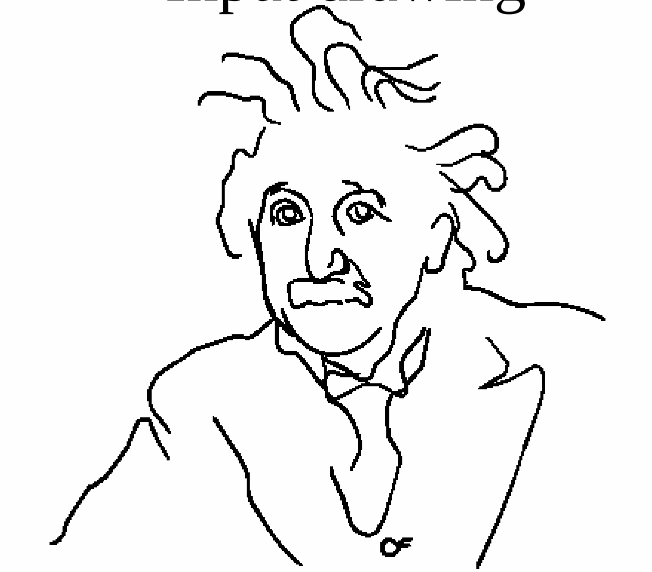

Figure 2: Line drawing traced from photograph of Einstein.

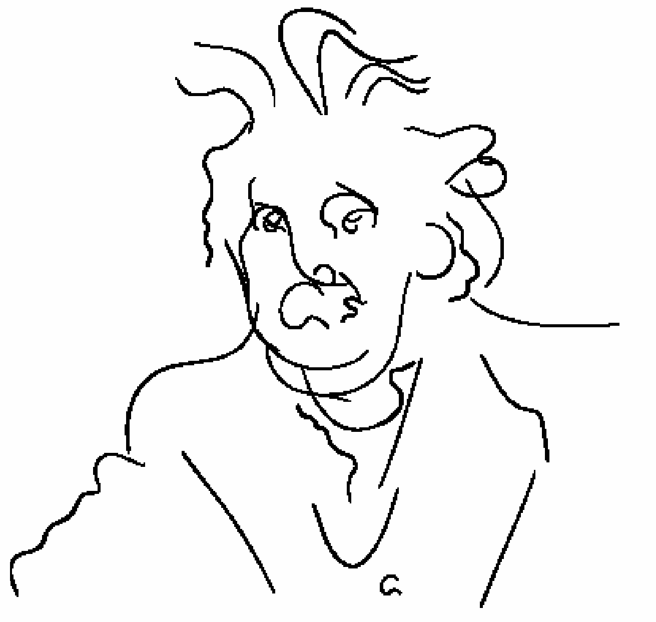

1-NN fit to input, style 1

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input drawing

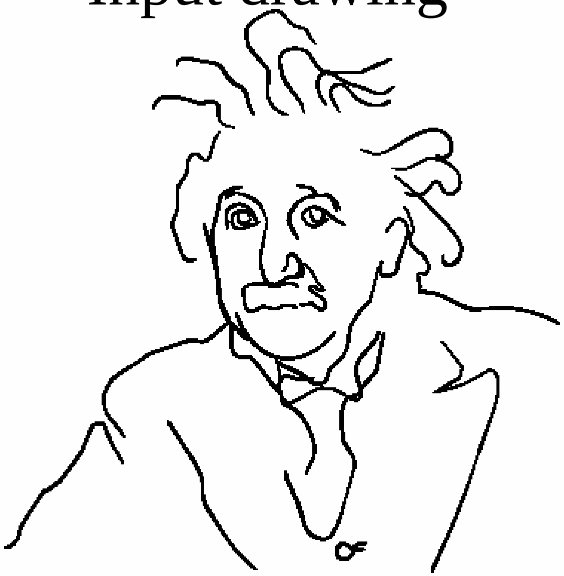

Figure 2: Line drawing traced from photograph of Einstein.

(a)

1-NN fit to input, style 1

(b)

Translation to style 2

Bad fit, good translation

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input drawing

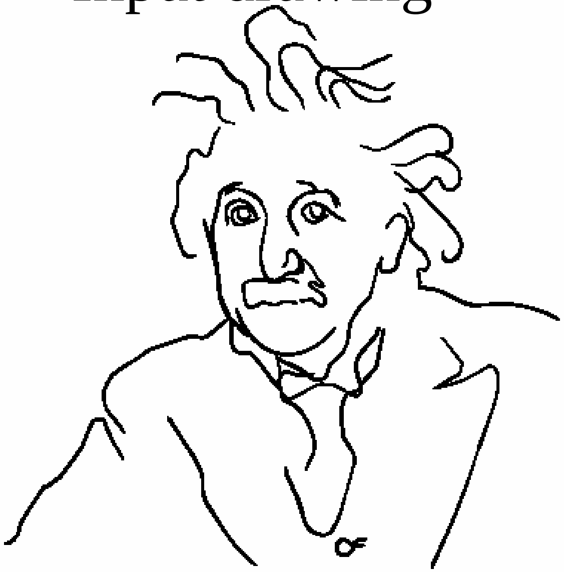

Figure 2: Line drawing traced from photograph of Einstein.

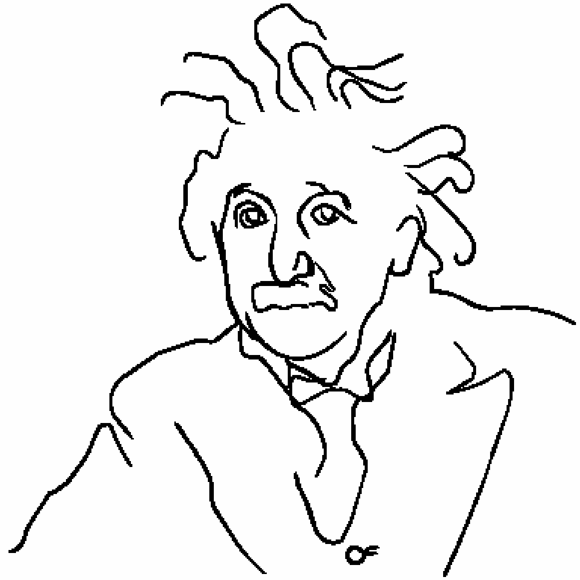

(a)

5904-NN fit to input, style 1

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input drawing

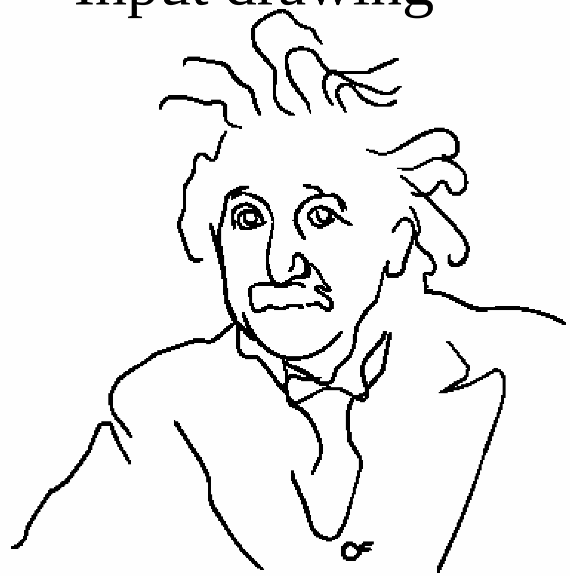

Figure 2: Line drawing traced from photograph of Einstein.

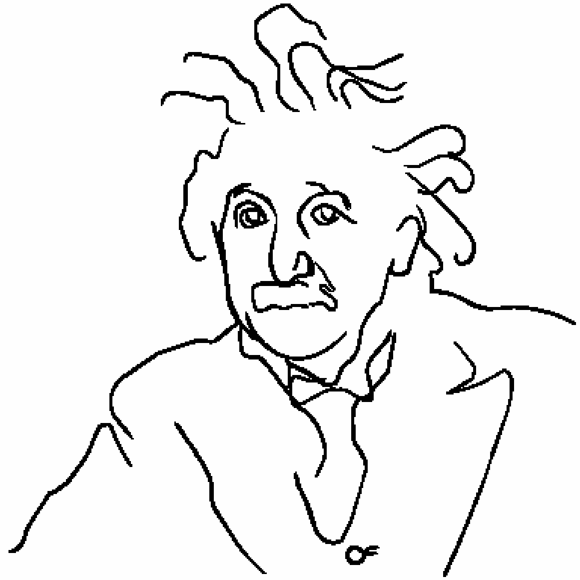

(a)

5904-NN fit to input, style 1

(b)

Translation to style 2.

Good fit, bad translation

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input drawing

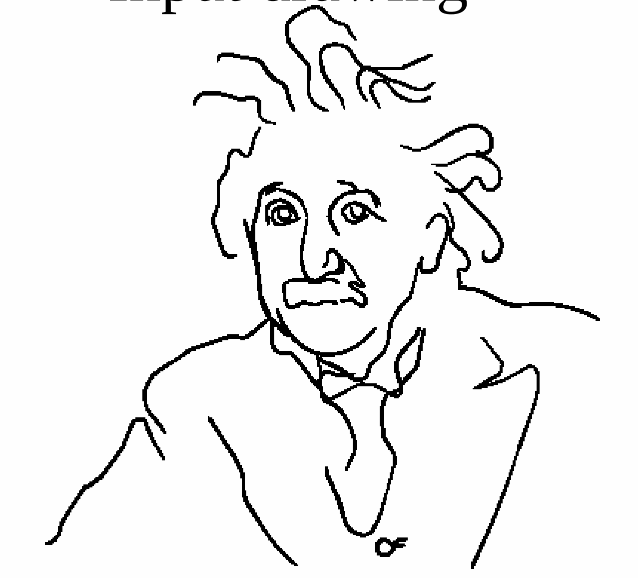

Figure 2: Line drawing traced from photograph of Einstein.

6-NN fit to input, style 1

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

Input drawing

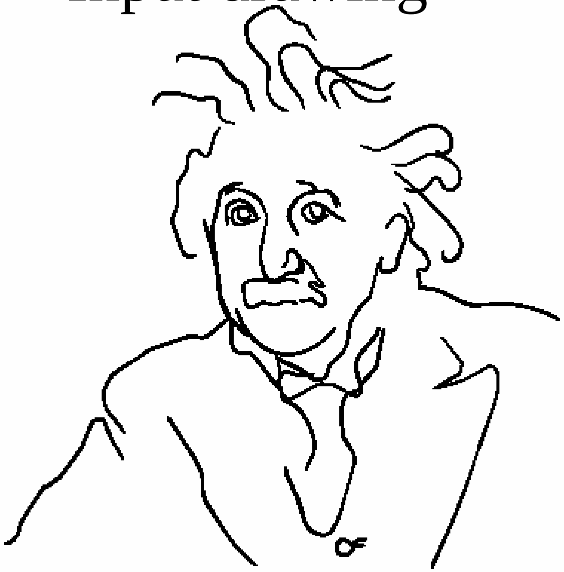

Figure 2: Line drawing traced from photograph of Einstein.

6-NN fit to input, style 1

Translation to style 2

Good fit, good translation

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

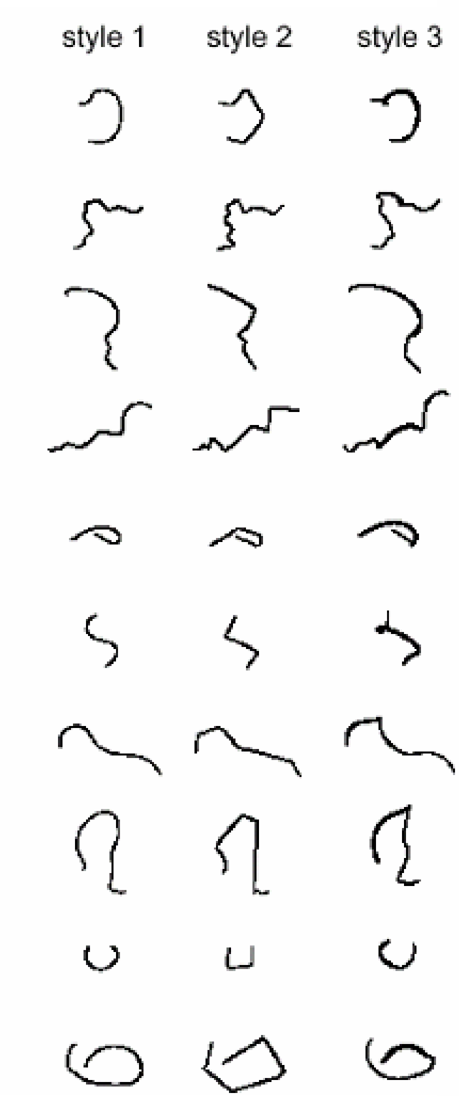

style 1

style 2

style 3

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

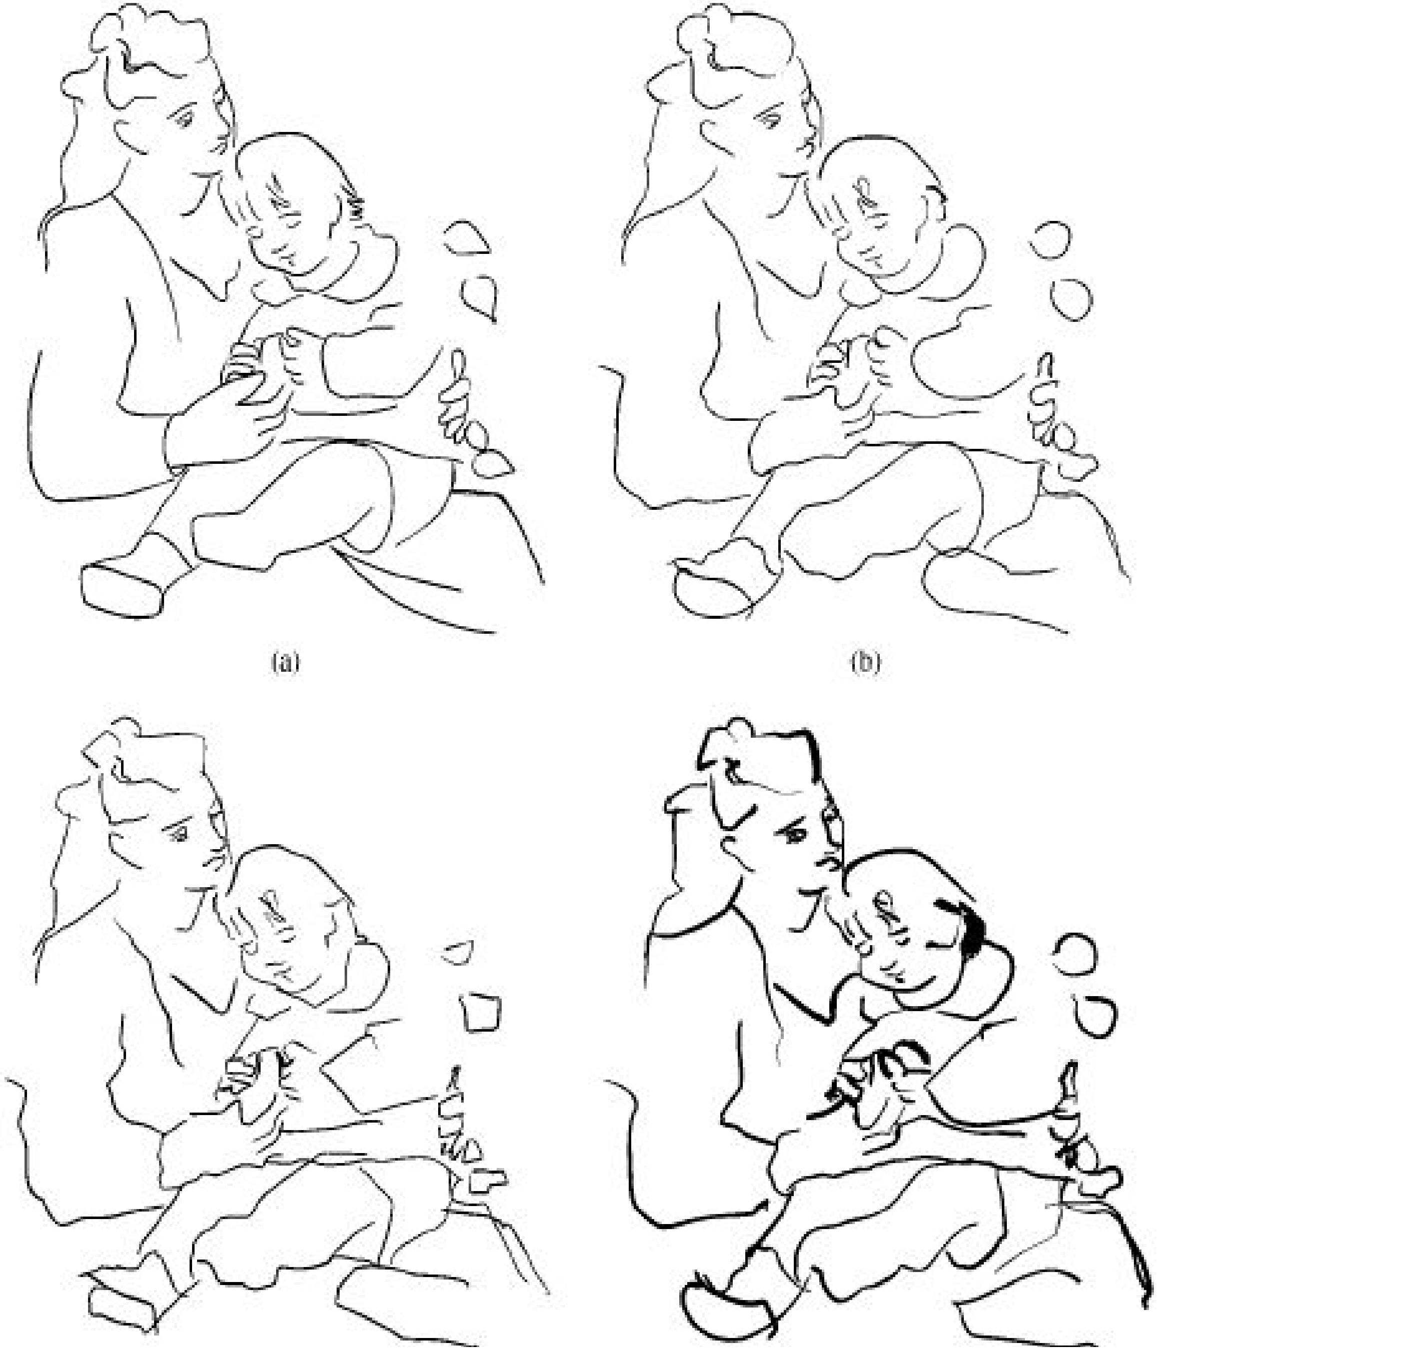

Fig. 9. 6-*NN* style translation example: (a) a tracing of the lines of Picasso's "Mother and Child"; (b) fit to style 1; (c) translation to style 2; and (d) to style 3, both using the 6 nearest neighbor algorithm, and fitting assuming the original was in style 1.

(a)

6-NN fit to input,
style 1

(b)

Translation to style 3

http://people.csail.mit.edu/billf/papers/p33-t_freeman.pdf

http://mrl.nyu.edu/projects/image-analogies/

image analogies

We present a new framework for processing images by example, called "image analogies." Rather than attempting to program individual filters by hand, we attempt to automatically learn filters from training data. For example, the following figure demonstrates an image analogy used to learn a painting style:

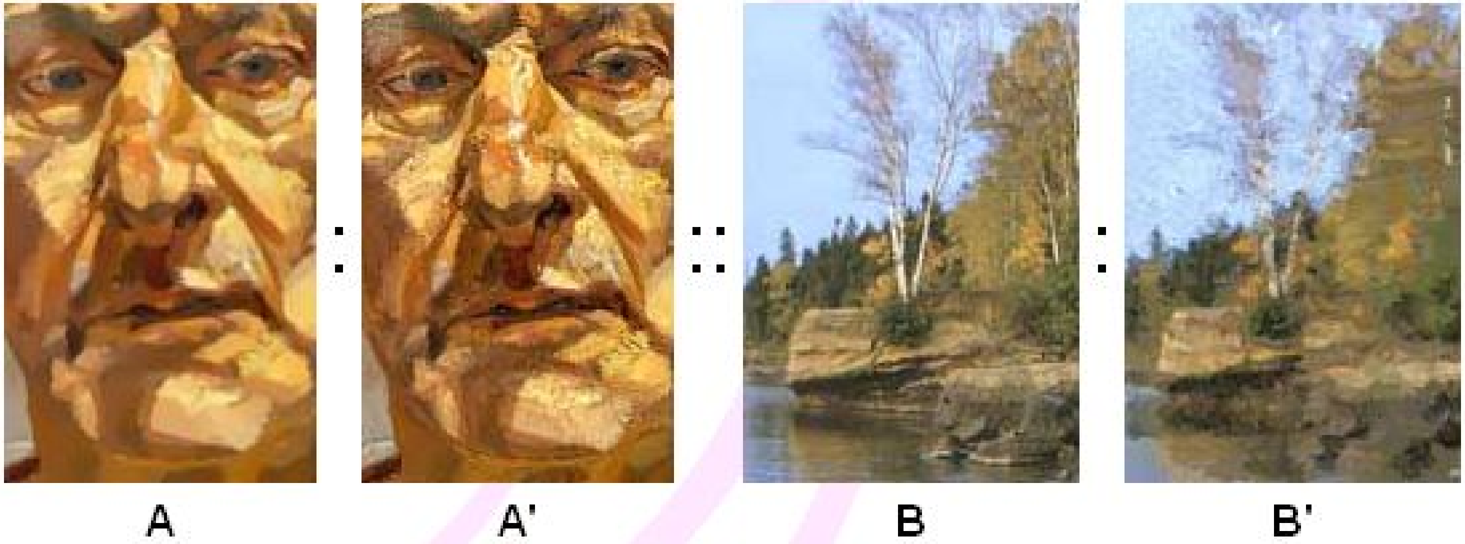

The images on the left are training data; our system "learns" the transformation from **A** to **A'**, and then applies that transformation to **B** to get **B'**. In other words, we compute **B'** to complete the analogy. (Only partial images are shown above; here are the full images).

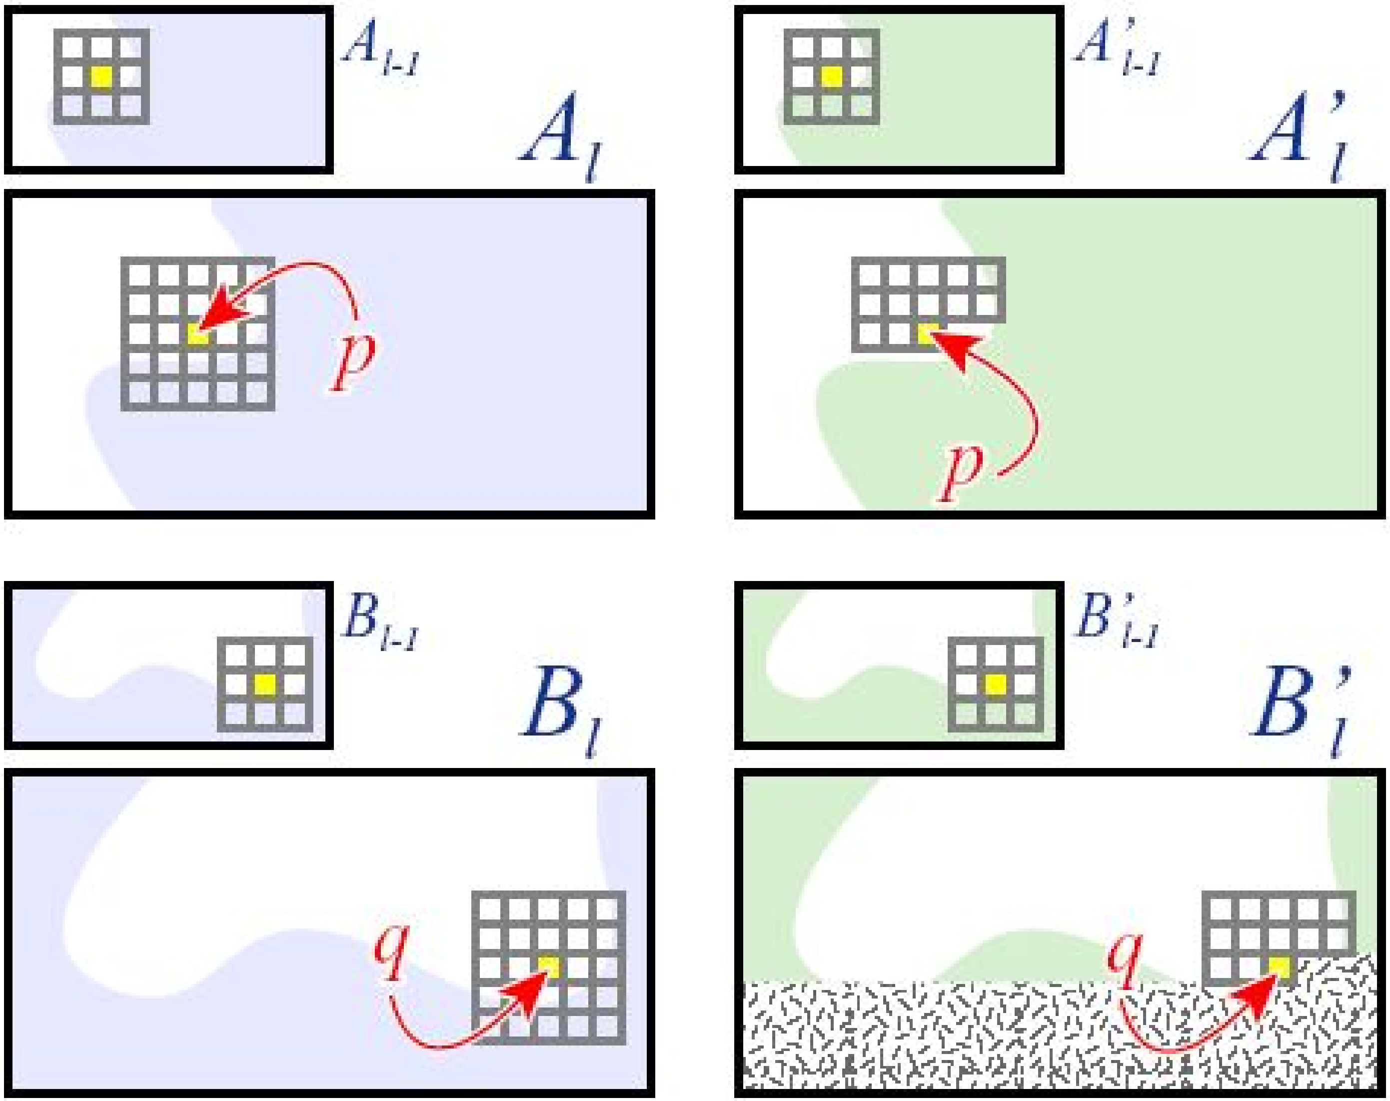

Figure 2 Neighborhood matching. In order to synthesize the pixel value at q in the filtered image B'_ℓ, we consider the set of pixels in B'_ℓ, B_ℓ, $B'_{\ell-1}$, and $B_{\ell-1}$ around q in the four images. We search for the pixel p in the A images that give the closest match. The synthesis proceeds in scan-line ordering in B'_ℓ.

Image analogies applications

- ***traditional image filters***, such as blurring or "embossing" (Section 4.1);
- ***improved texture synthesis***, in which some textures are synthesized with higher quality than previous approaches (Section 4.2);
- ***super-resolution***, in which a higher-resolution image is inferred from a low-resolution source (Section 4.3);
- ***texture transfer***, in which images are "texturized" with some arbitrary source texture (Section 4.4);
- ***artistic filters***, in which various drawing and painting styles, including oil, watercolor, and line art rendering, are synthesized based on either digitally filtered or scanned real-world examples (Section 4.5); and
- ***texture-by-numbers***, in which realistic scenes, composed of a variety of textures, are created using a simple "painting" interface (Section 4.6).

For painterly style translation, how get the A, A’ image pairs?

From a detail of *Reflection (Self-Portrait)*, 1985:

http://mrl.nyu.edu/projects/image-analogies/

Filtered source (*A'*)

http://mrl.nyu.edu/projects/image-analogies/

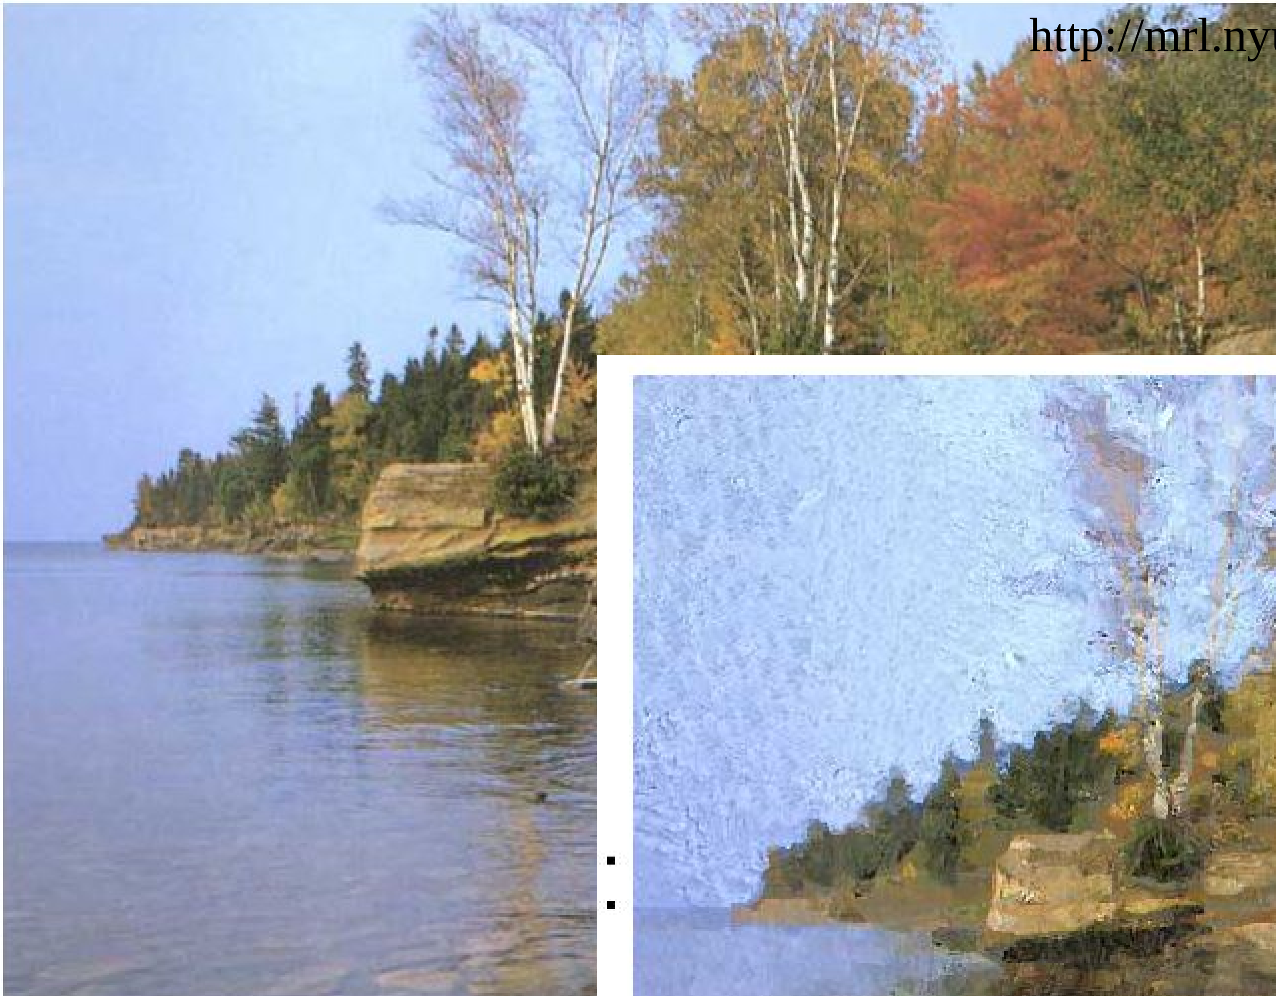

Unfilte

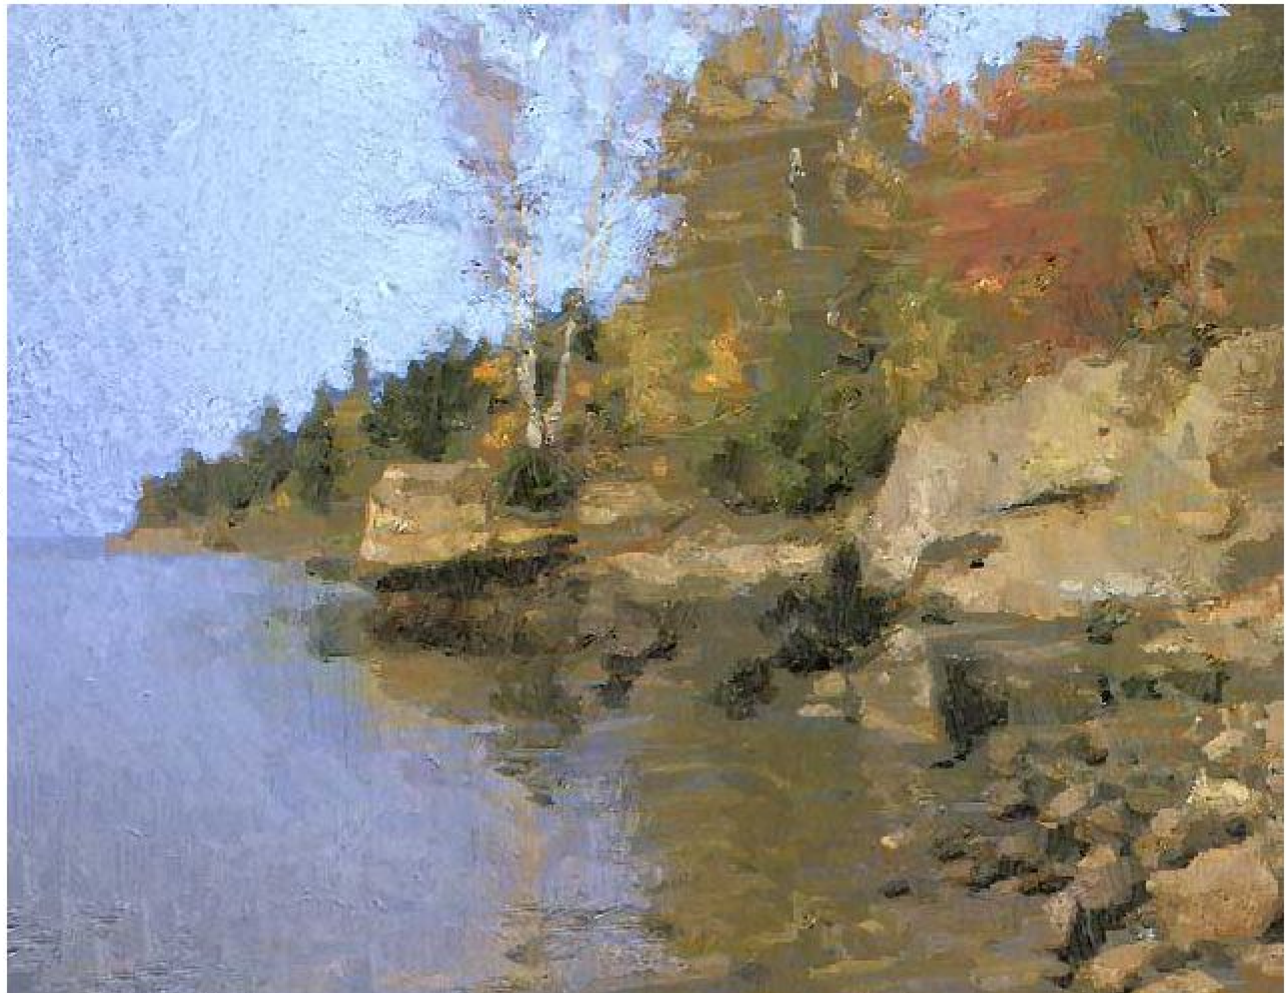

Filtered target (*B'*)

http://mrl.nyu.edu/projects/image-analogies/

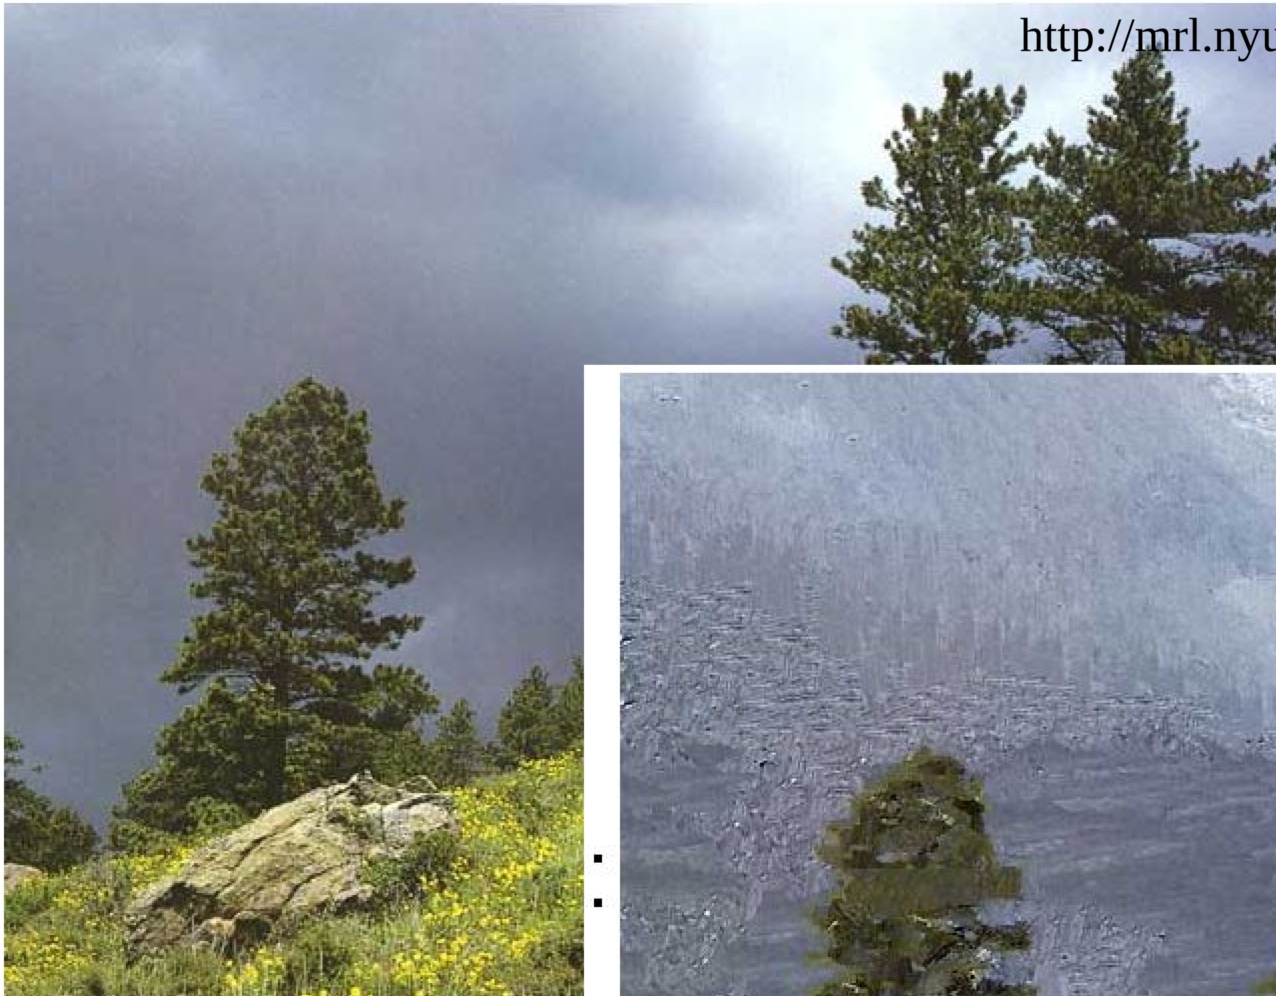

Unfilte

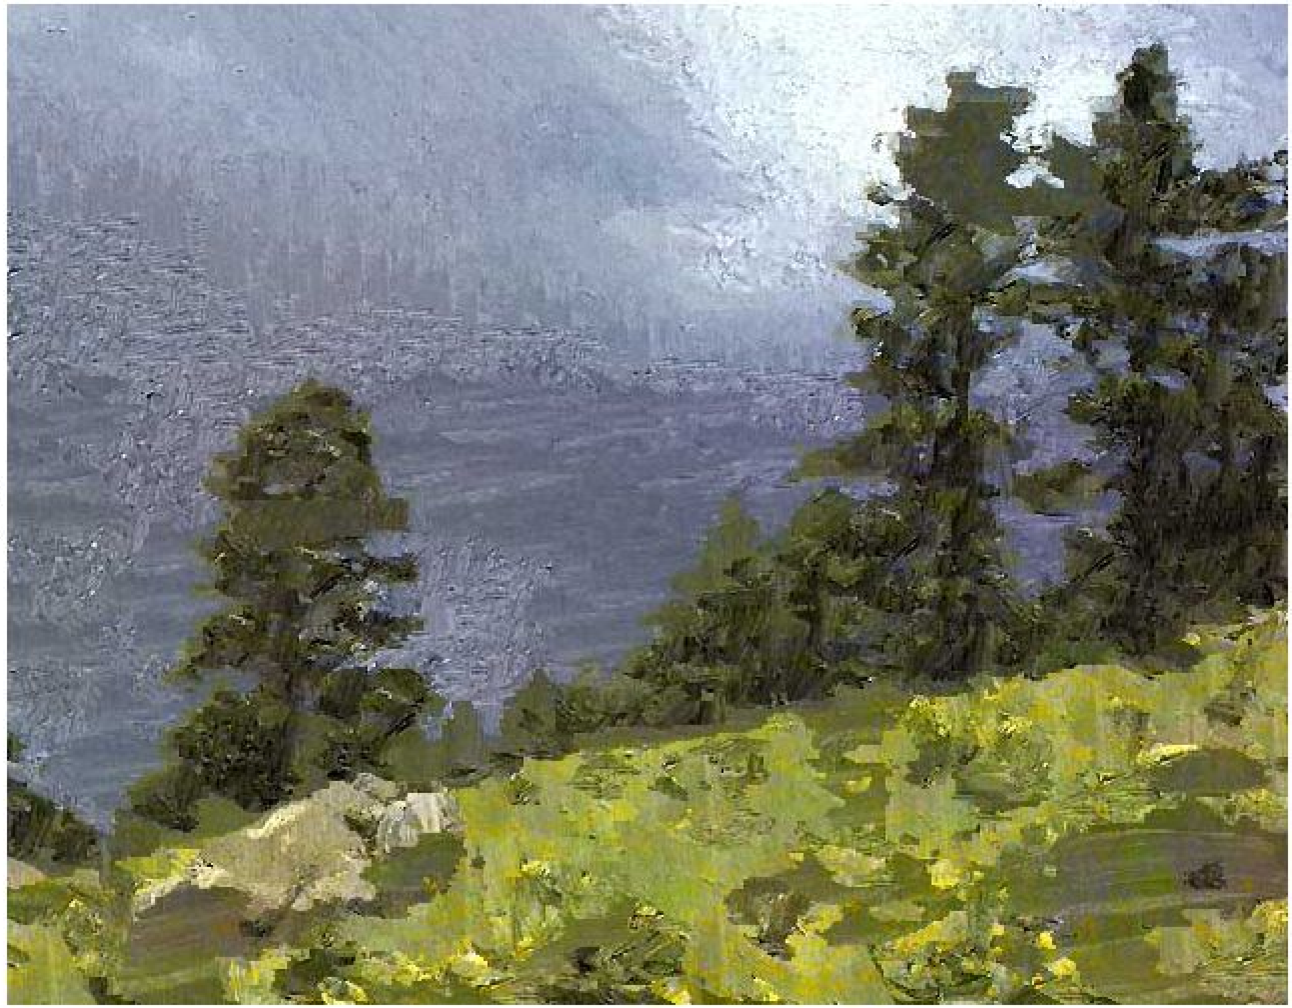

Filtered target (*B'*)

Texture Transfer

- Take the texture from one object and “paint” it onto another object
  - This requires separating texture and shape
  - That’s HARD, but we can cheat
  - Assume we can capture shape by boundary and rough shading
-

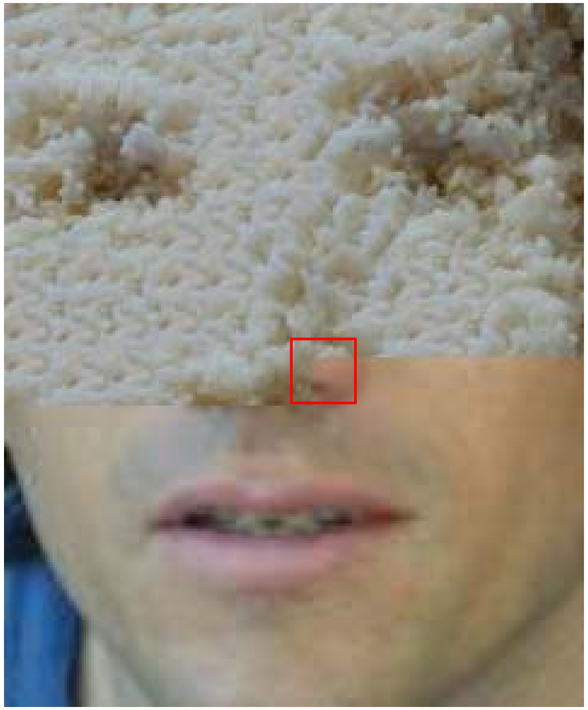

Then, just add another constraint when sampling: similarity to underlying image at that spot

http://people.csail.mit.edu/billf/papers/efrosFreeman.pdf

Source texture

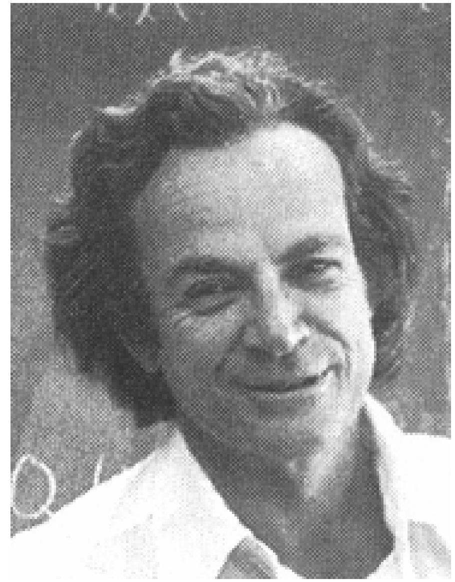

Target image

Source correspondence image

Target correspondence image

http://people.csail.mit.edu/billf/papers/efrosFreeman.pdf

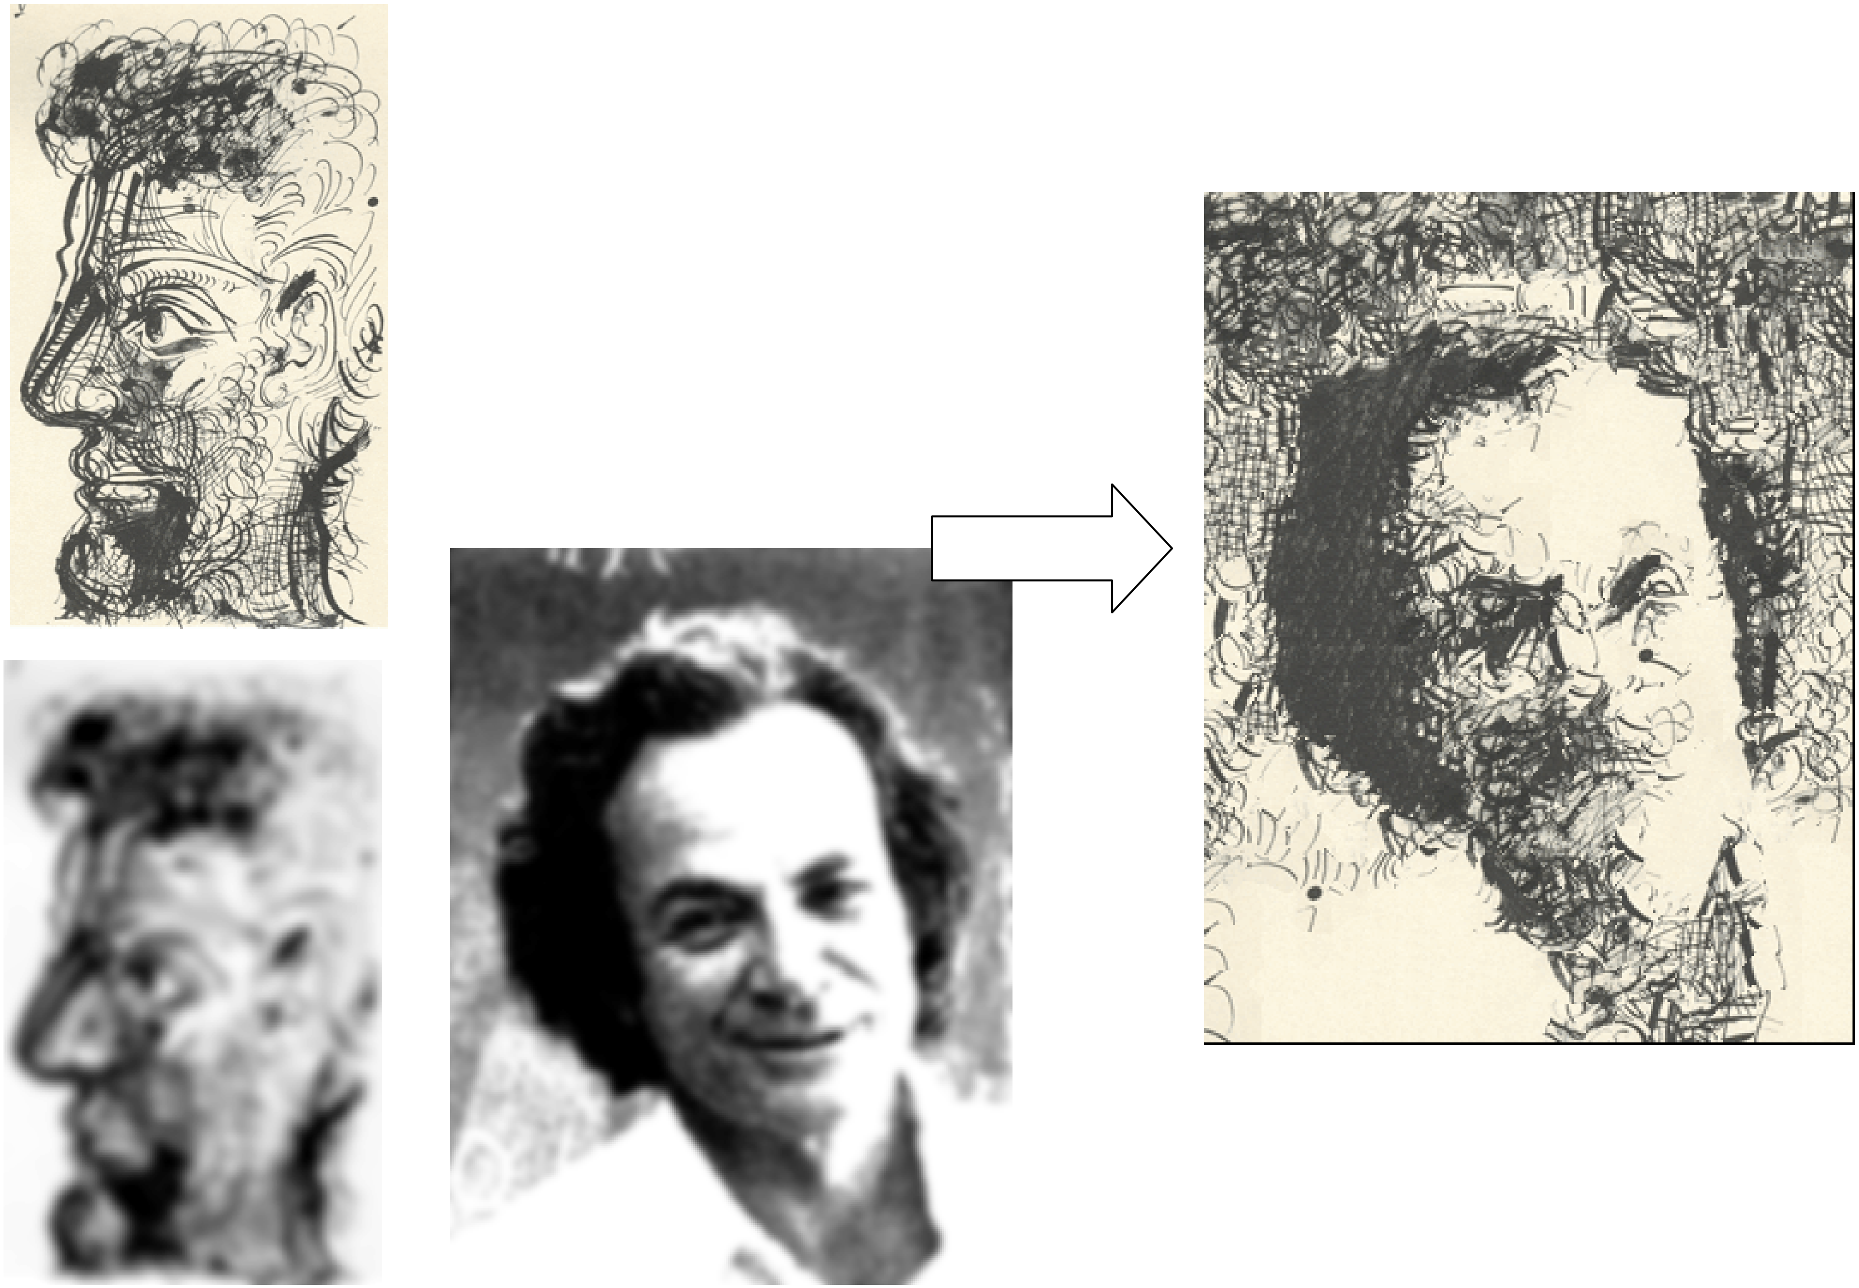

http://people.csail.mit.edu/billf/papers/efrosFreeman.pdf

A

A’

Unfiltered examples (A) **Filtered examples (A')**

Figure 10 Training pairs for the color NPR filters used in this paper. The upper A' image is a detail of *Starry Night above the Rhône* by Vincent Van Gogh; the "unfiltered" source image was generated by processing the painting with Photoshop's "Smart Blur" filter. The lower image pair is a photograph and a watercolor created from it with a semi-automatic digital filter [10].

B

Figure 9 Unfiltered target images (B) for the NPR filters and texture transfer. (Leftmost image courtesy John Shaw [38].)

B’

Figure 11 Boat paintings by example. The left image is painted in a style learned from a Van Gogh painting (Figure 10, top row); the right image is in the style of a watercolor filter (Figure 10, bottom row).

B’

I think this one fails

Figure 1 An image analogy. Our problem is to compute a new "analogous" image B' that relates to B in "the same way" as A' relates to A. Here, A, A', and B are inputs to our algorithm, and B' is the output. The full-size images are shown in Figures 10 and 11.

Organization of NPR methods

- Automated methods
  - 2-d processing
  - 3-d processing
- Interactive methods
  - 2-d processing
  - 3-d processing

8 Using Non-Photorealistic Rendering to Communicate Shape

Amy Ashurst Gooch and Bruce Gooch
Department of Computer Science
University of Utah
http://www.cs.utah.edu/

Photograph

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Illustration

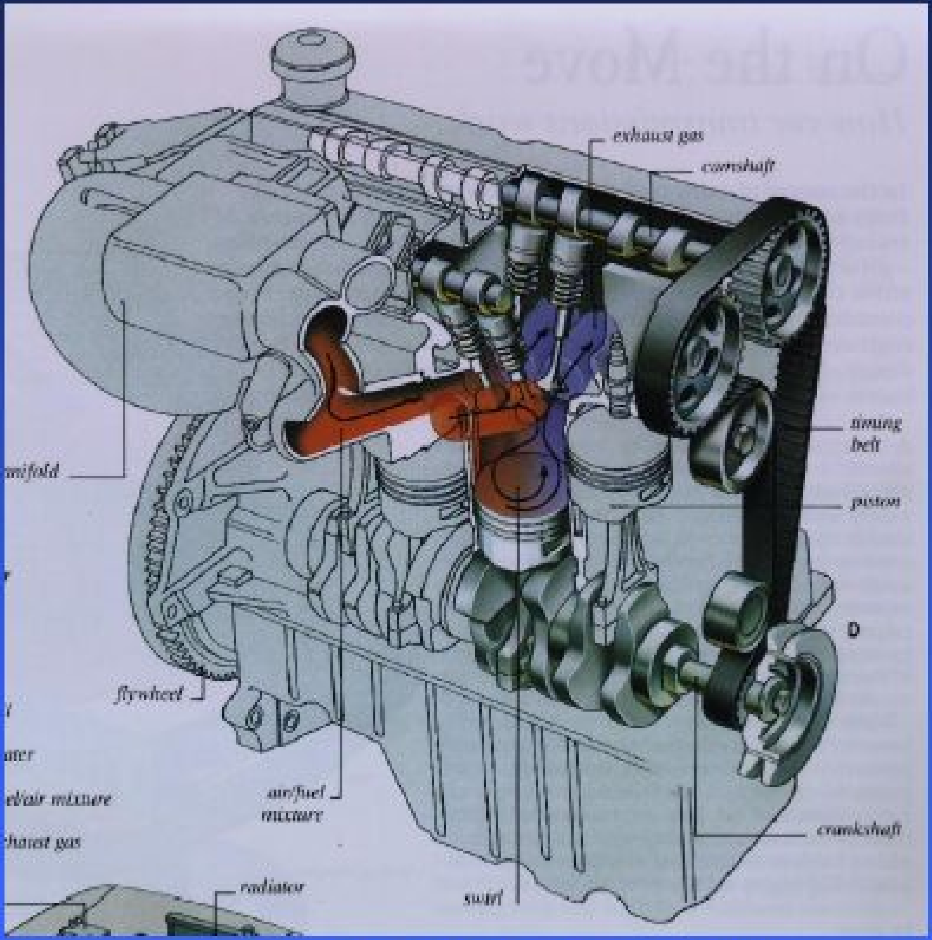

Image from "The Way Science Works".
Courtesy of Macmillan Reference USA, a division of Ahsuog, Inc.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Gooch and Gooch

- Concentrate on the material property and shading aspects of technical illustration.

Some characteristics of technical illustrations

8.2 Technical Illustration

Human-drawn technical illustrations are usually stand-alone images from a single viewpoint presented on a non-stereo medium such as pen on paper. In this section we discuss the components of such illustrations that we use in a computer graphics context: line character, shading, and shadowing.

Examining technical manuals, illustrated textbooks, and encyclopedias reveals shading and line illustration conventions which are quite different than traditional computer graphics conventions. The use of these artistic conventions produces *technical illustrations*, a subset of non-photorealistic rendering. The illustrations in several books, e.g., [20, 23], imply that illustrators use fairly algorithmic principles. Although there are a wide variety of styles and techniques found in technical illustration, there are some common themes. This is particularly true when examining color illustrations done with air-brush and pen. The following characteristics are present in many illustrations:

- edge lines are drawn with black curves.
- matte objects are shaded with intensities far from black or white with warmth or coolness of color indicative of surface normal.
- a single light source provides white highlights.
- shadows are rarely included, but if they are used, they are placed where they do not occlude details or important features.
- metal objects are shaded as if very anisotropic.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Technical Illustration

Shape information is most valued

- Edge lines
- Shading
- Shadows rarely included
- One light

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Technical illustrations

Lines

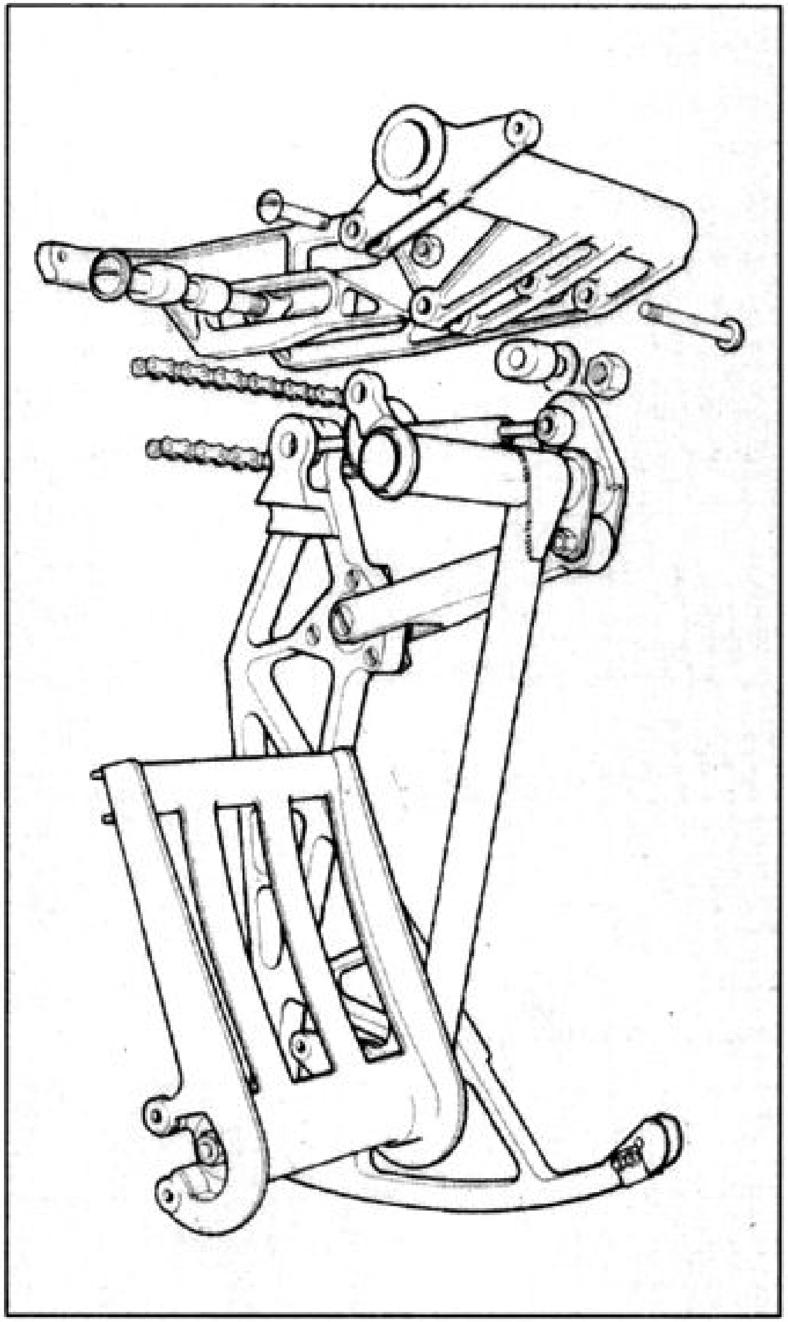

Figure 1: An example of the lines an illustrator would use to convey the shape of this airplane foot pedal. Copyright 1989 Macdonald & Co. (Publishers) Ltd. [20].

Some parameterization dependent lines

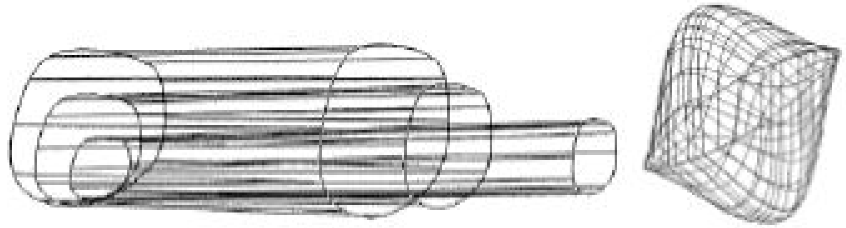

Figure 2: A few examples of a NURBS-based model displayed in wireframe. The large number of isolines makes distinguishing key features difficult.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

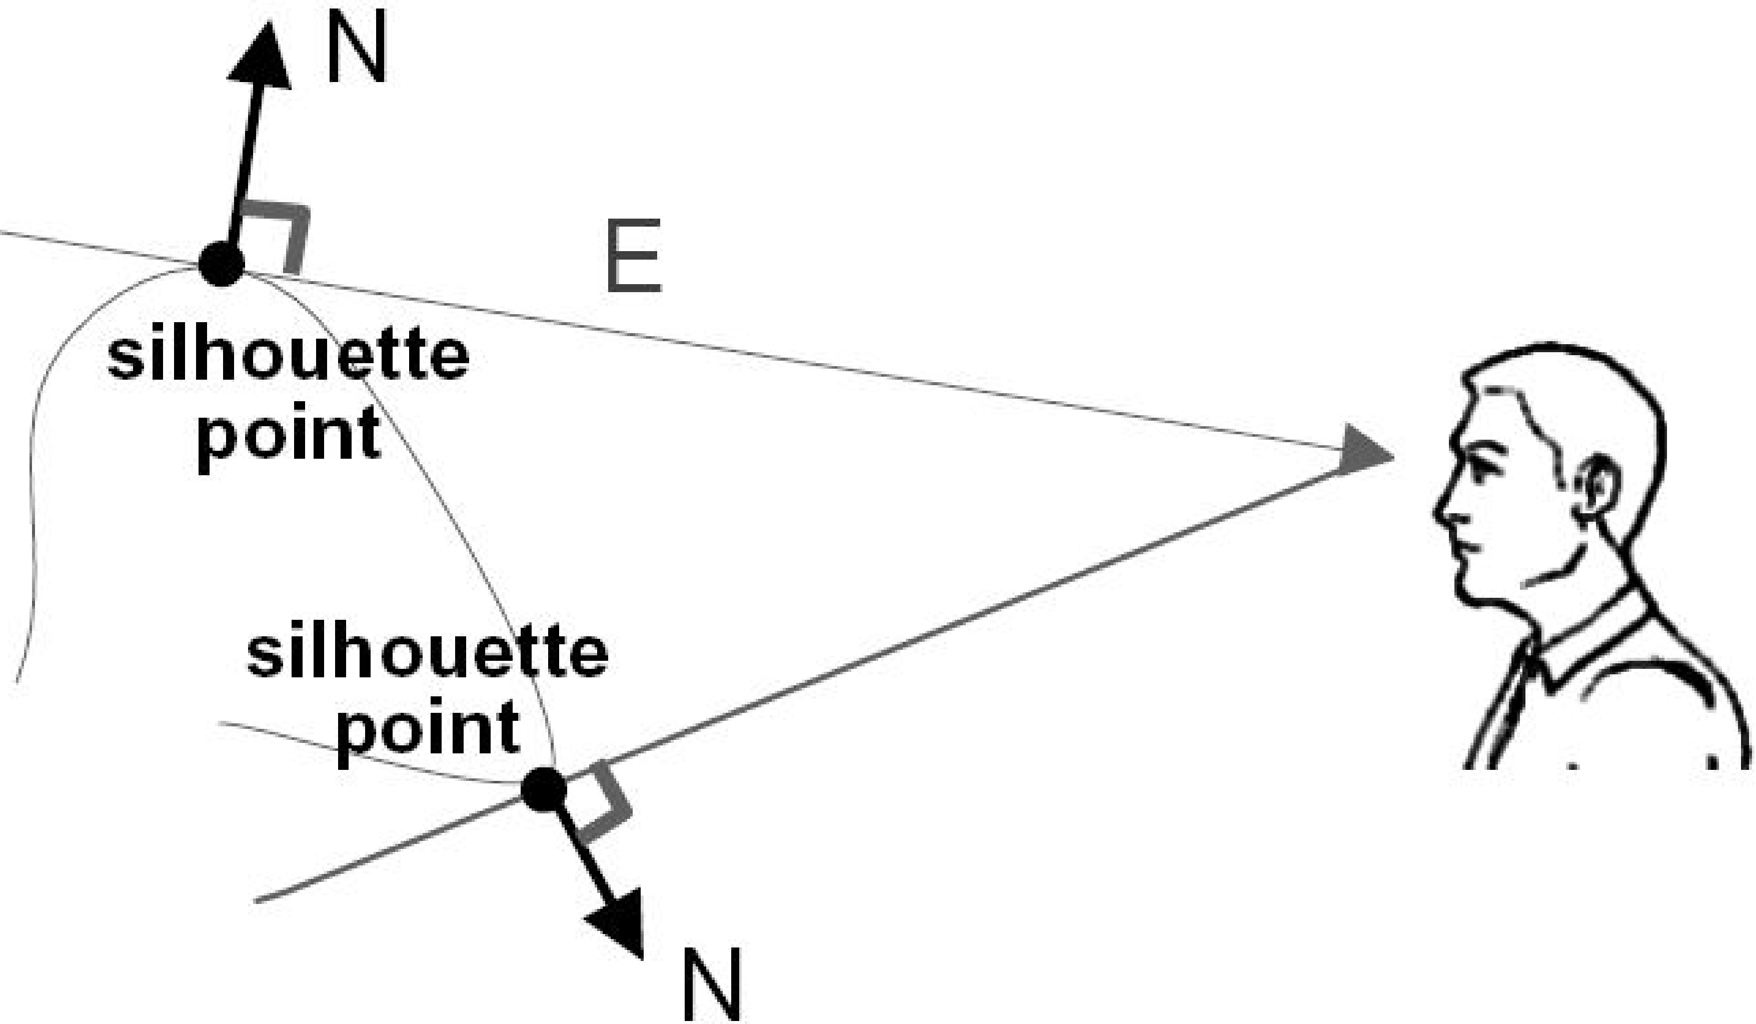

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

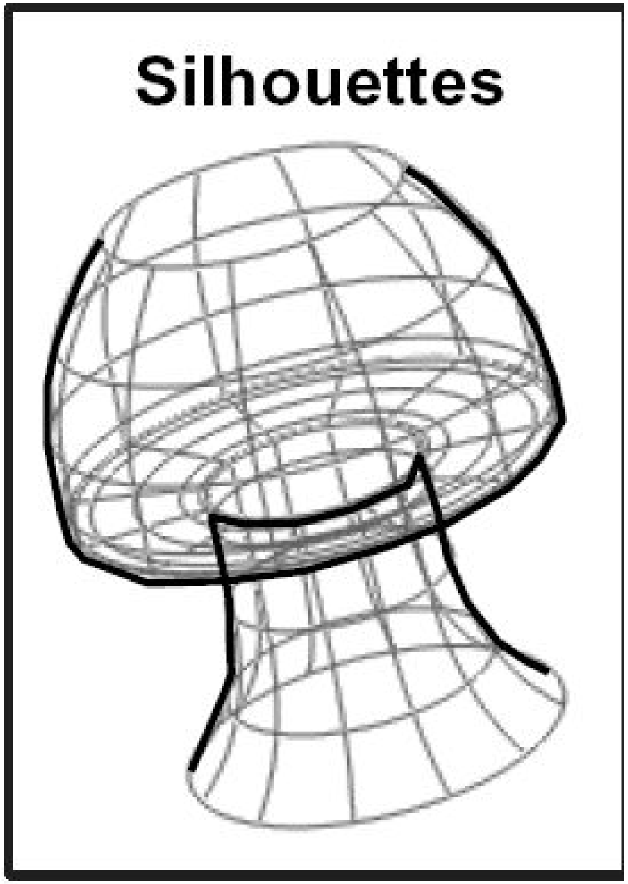

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Line weight variations

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Equal weight

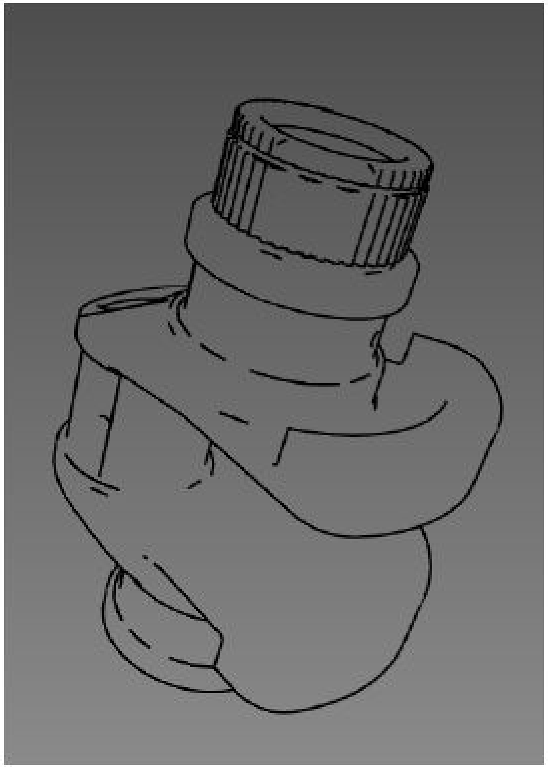

Figure 7: Three line conventions suggested by Martin [20]. Left: single weight used throughout the image. Middle: heavy line weight used for outer edges, other lines are thinner. Right: vary line weight to emphasize perspective.

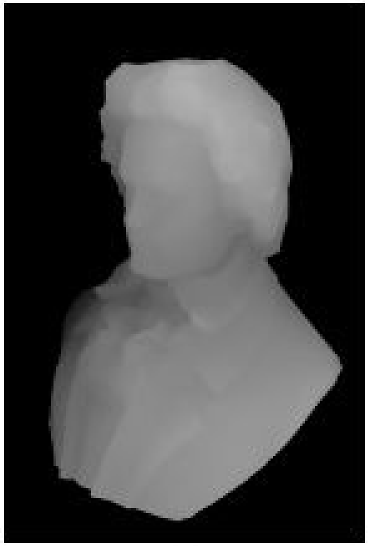

(a)

Figure 3: Outline detection of a more complex model. (a) Depth map. (b) Depth map edges. (c) Normal map. (d) Normal map edges. (e) Combined depth and normal map edges. (See also the Color Plates section of the course notes.)

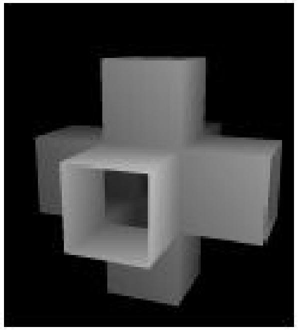

(a)

Figure 1: Outline drawing with image processing. (a) Depth map. (b) Edges of the depth map. (c) Normal map. (d) Edges of the normal map. (e) The combined edge images. (f) A difficult case: folded piece of paper (g) Depth edges. (See also the Color Plates section of the course notes.)

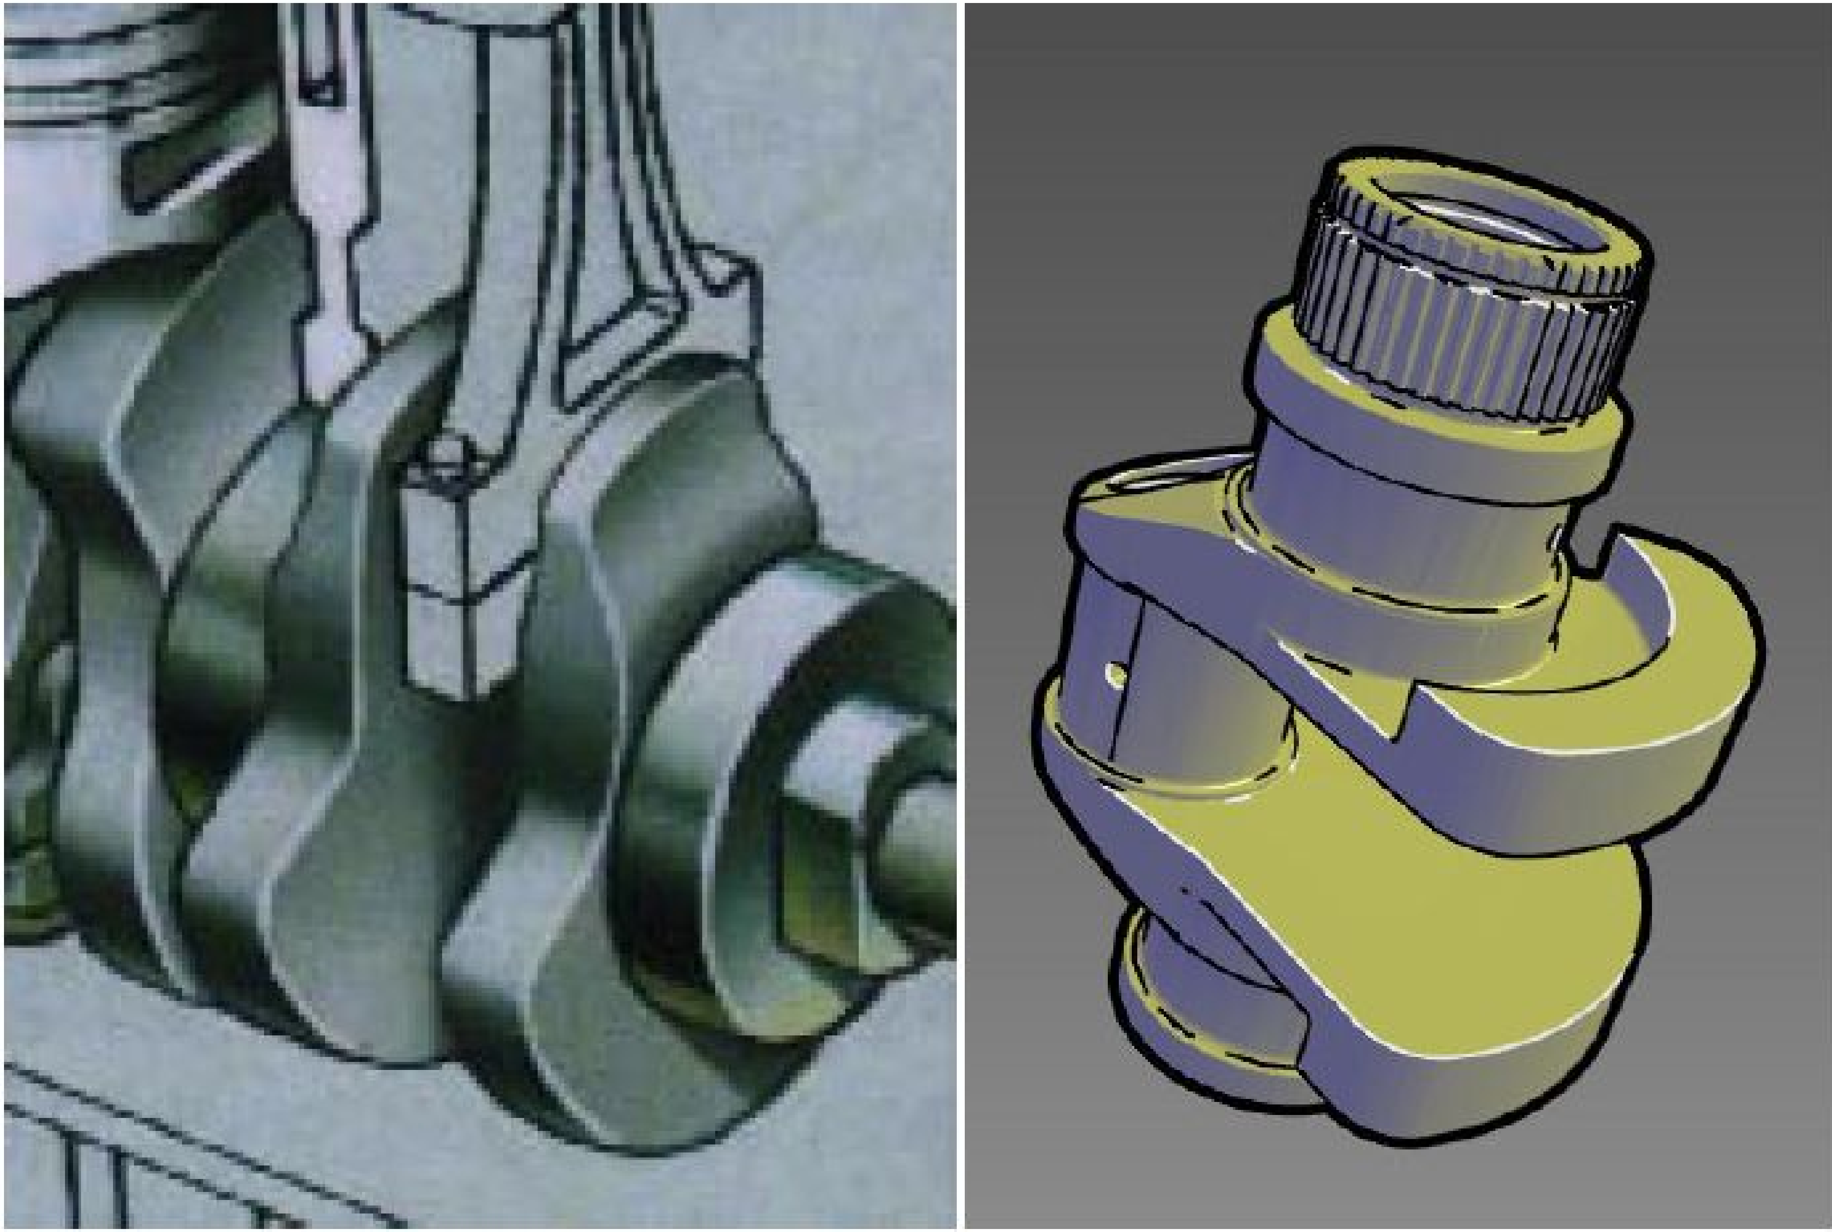

Figure 9: Left: Illustrators sometimes use the convention of white interior edge lines to produce a highlight. Courtesy of Macmillan Reference USA, a division of Ahsuog, Inc. [23]. Right: An image produced by our system, including shading, silhouettes, and white crease lines.

Technical illustrations

Shading

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

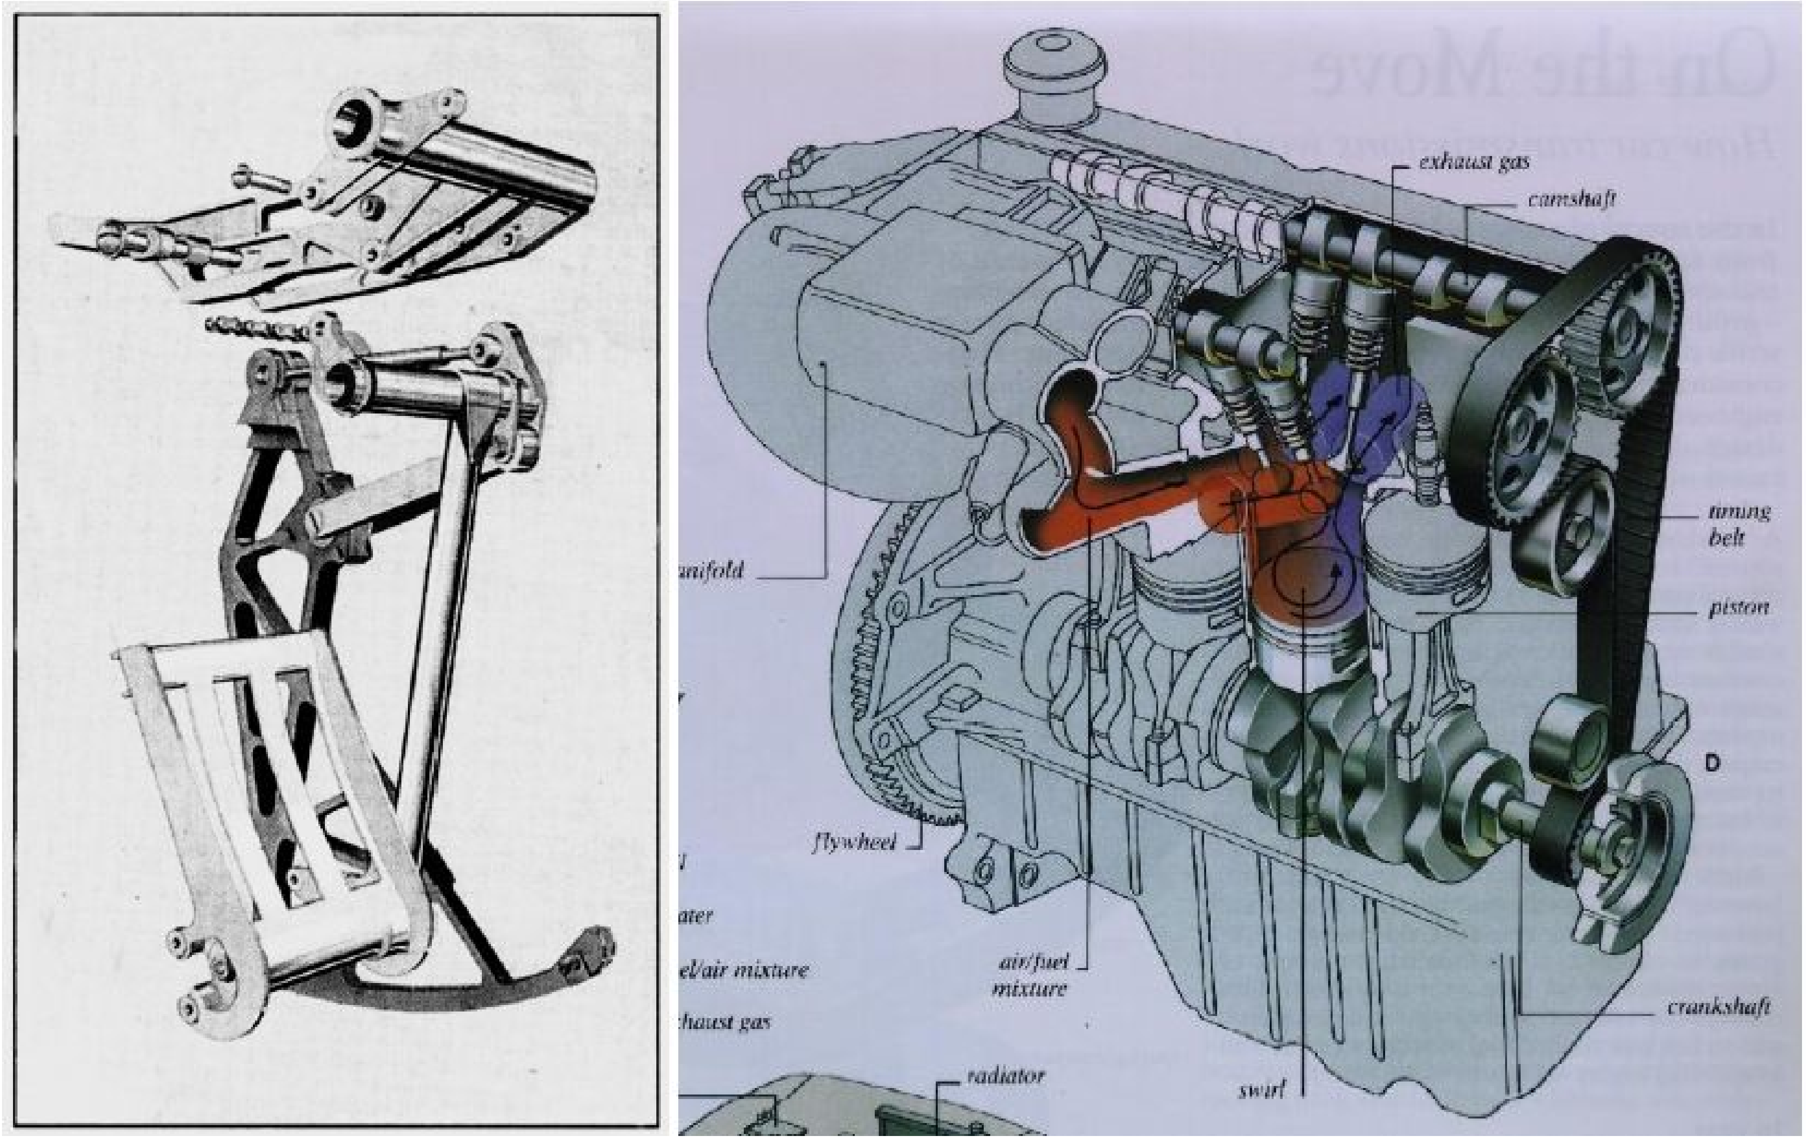

Figure 10: Illustrators combine edge lines with a specific type of shading. Shading in technical illustration brings out subtle shape attributes and provides information about material properties. Left: Compare this shaded image of airplane pedal to the line drawing in Figure 1. Copyright 1989 Macdonald & Co. (Publishers) Ltd. [20]. Right: Engine. Courtesy of Macmillan Reference USA, a division of Ahsuog, Inc. [23].

Diffuse shaded model

$I = k_d k_a + k_d \max(0, L \cdot n)$ with k_d=1 and k_a = 0.

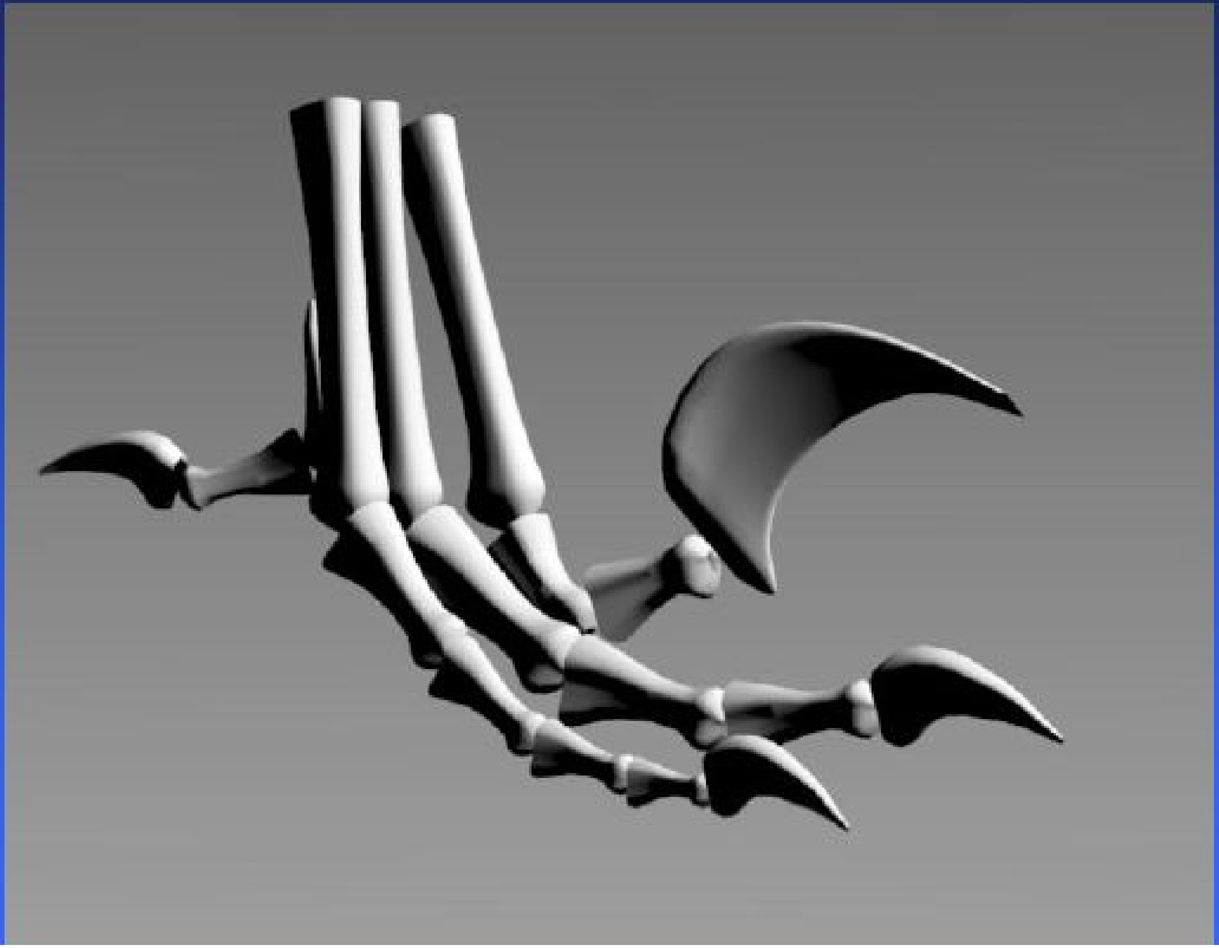

Figure 11: Diffuse shaded image using Equation 1 with $k_d = 1$ and $k_a = 0$. Black shaded regions hide details, especially in the small edge lines could not be seen if added. Highlights and fine details are lost in the white shaded regions.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

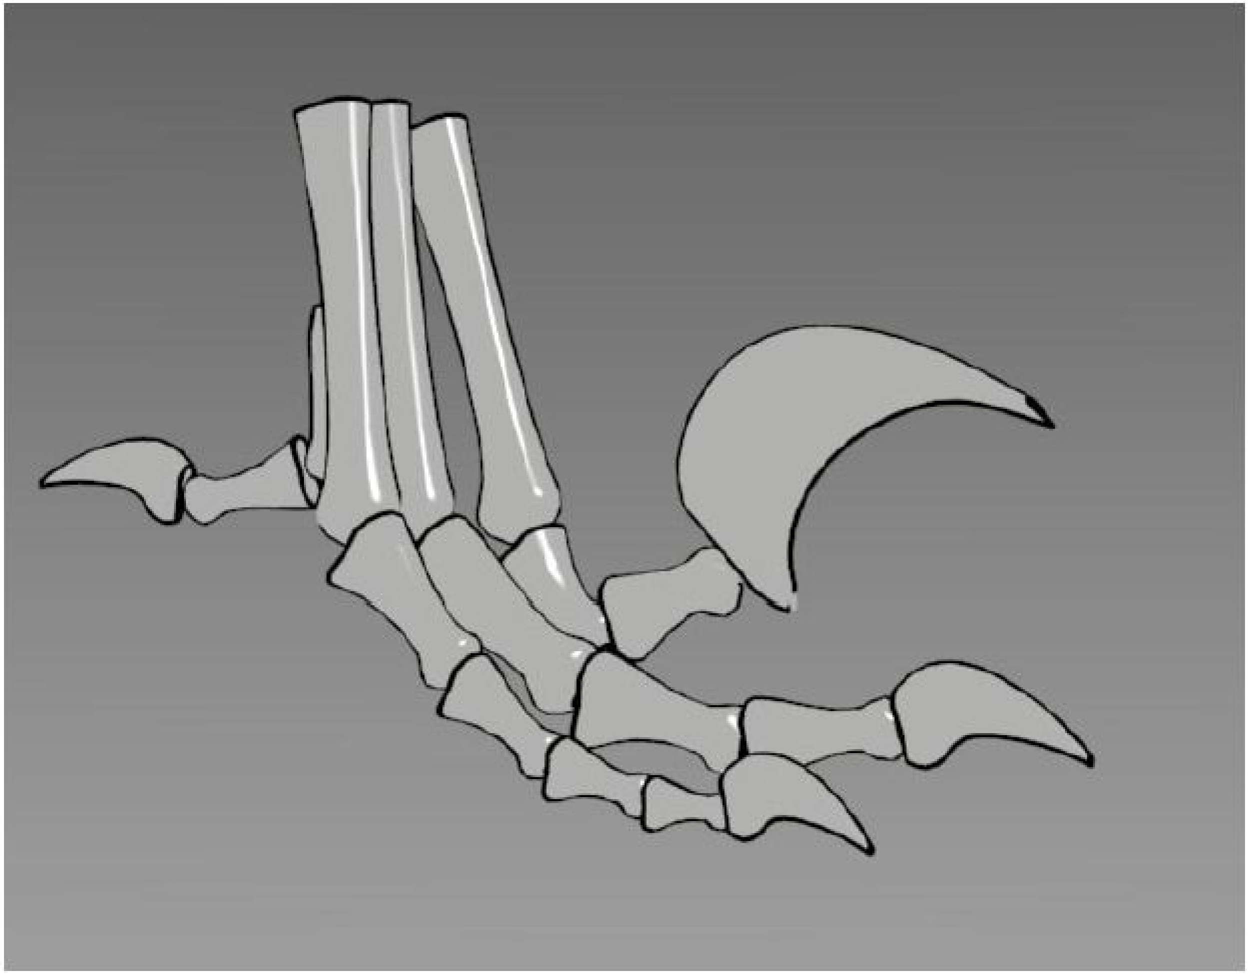

Figure 12: Image with only highlights and edges. The edge lines provide divisions between object pieces and the highlights convey the direction of the light. Some shape information is lost, especially in the regions of high curvature of the object pieces. However, these highlights and edges could not be added to Figure 11 because the highlights would be invisible in the light regions and the silhouettes would be invisible in the dark regions.

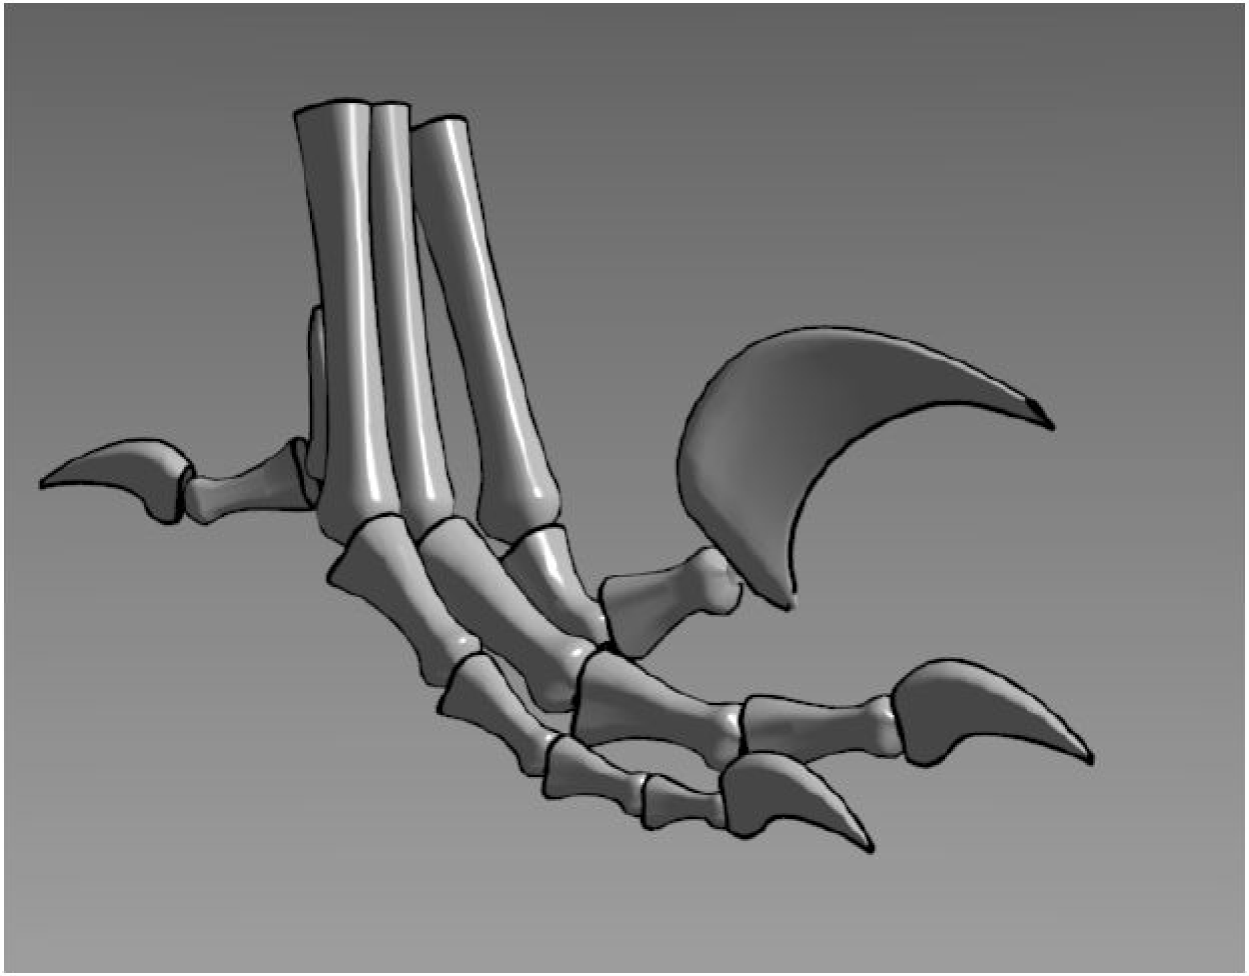

Figure 13: Phong-shaded image with edge lines and $k_d = 0.5$ and $k_a = 0.1$. Like Figure 11, details are lost in the dark gray regions, especially in the small claws, where they are colored the constant shade of $k_d k_a$ regardless of surface orientation. However, edge lines and highlights provide shape information that was gained in Figure 12, but could not be added to Figure 11.

Encoding surface orientation by color temperature

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

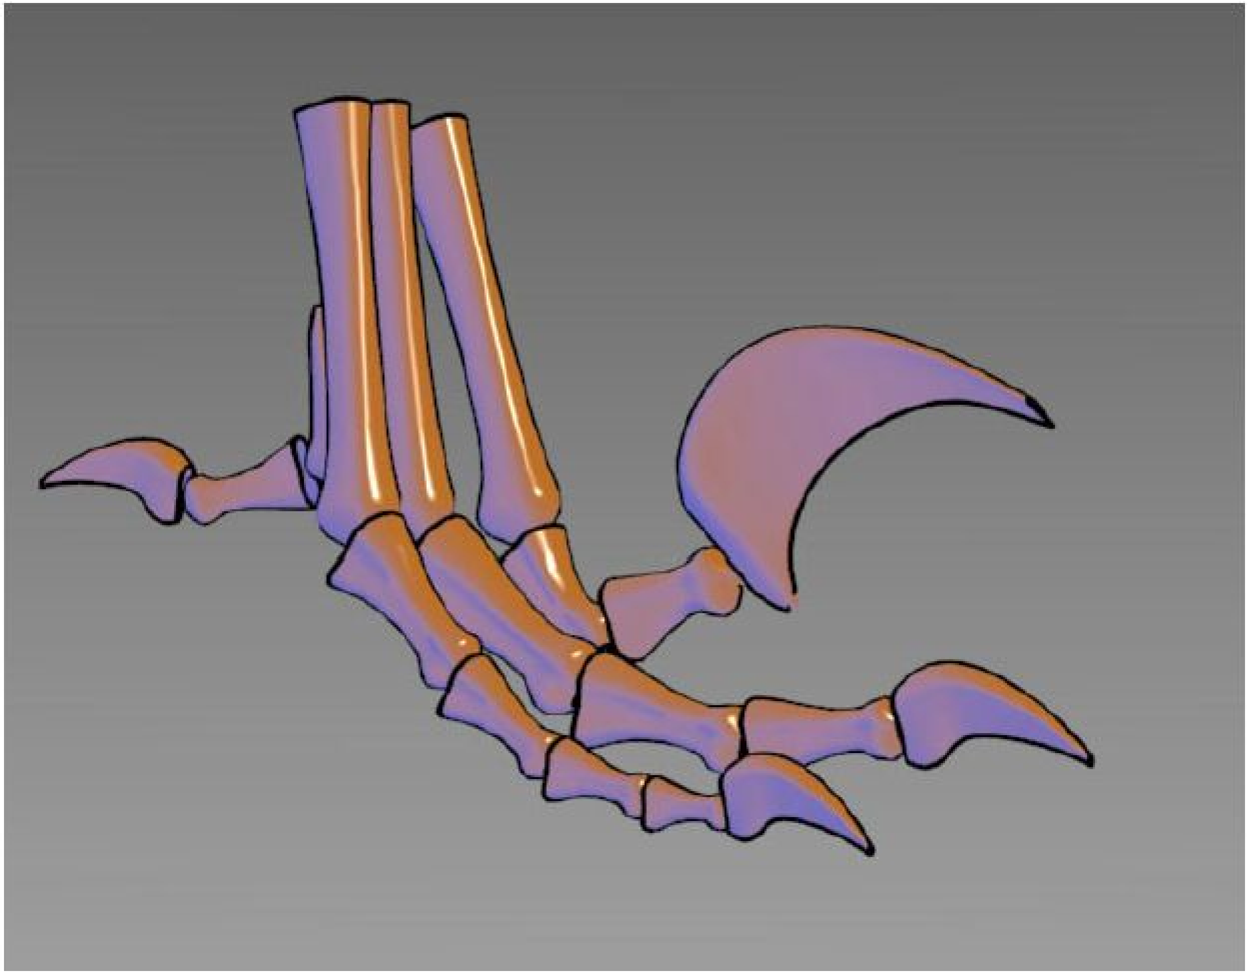

Figure 14: Approximately constant luminance tone rendering. Edge lines and highlights are clearly noticeable. Unlike Figures 11 and 13 some details in shaded regions, like the small claws, are visible. The lack of luminance shift makes these changes subtle. (See Color Plate).

Direction dependent illumination color

Combining color-temp surface orientation coding with some tonal variations in object color

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

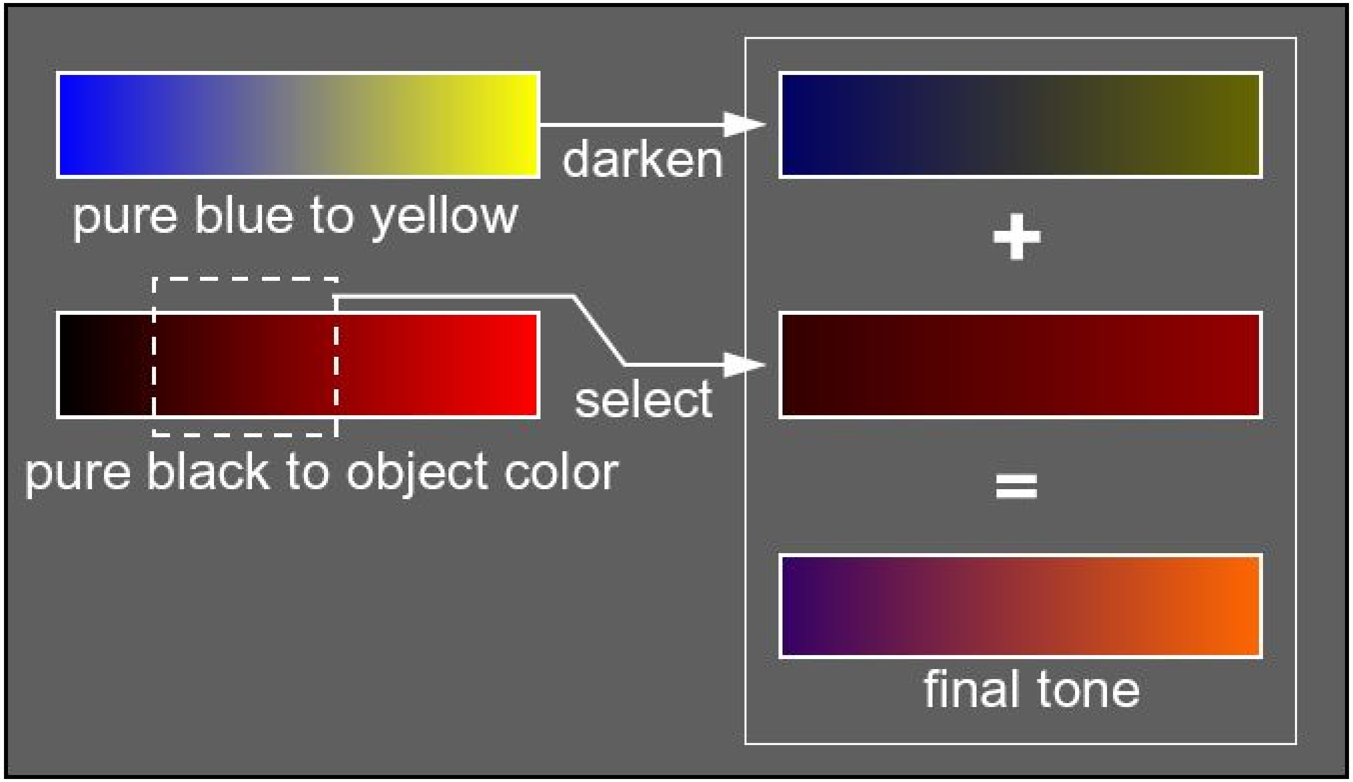

How the tone is created for a pure red object by summing a blue-to-yellow and a dark-red-to-red tone

Parameter setting # 1

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

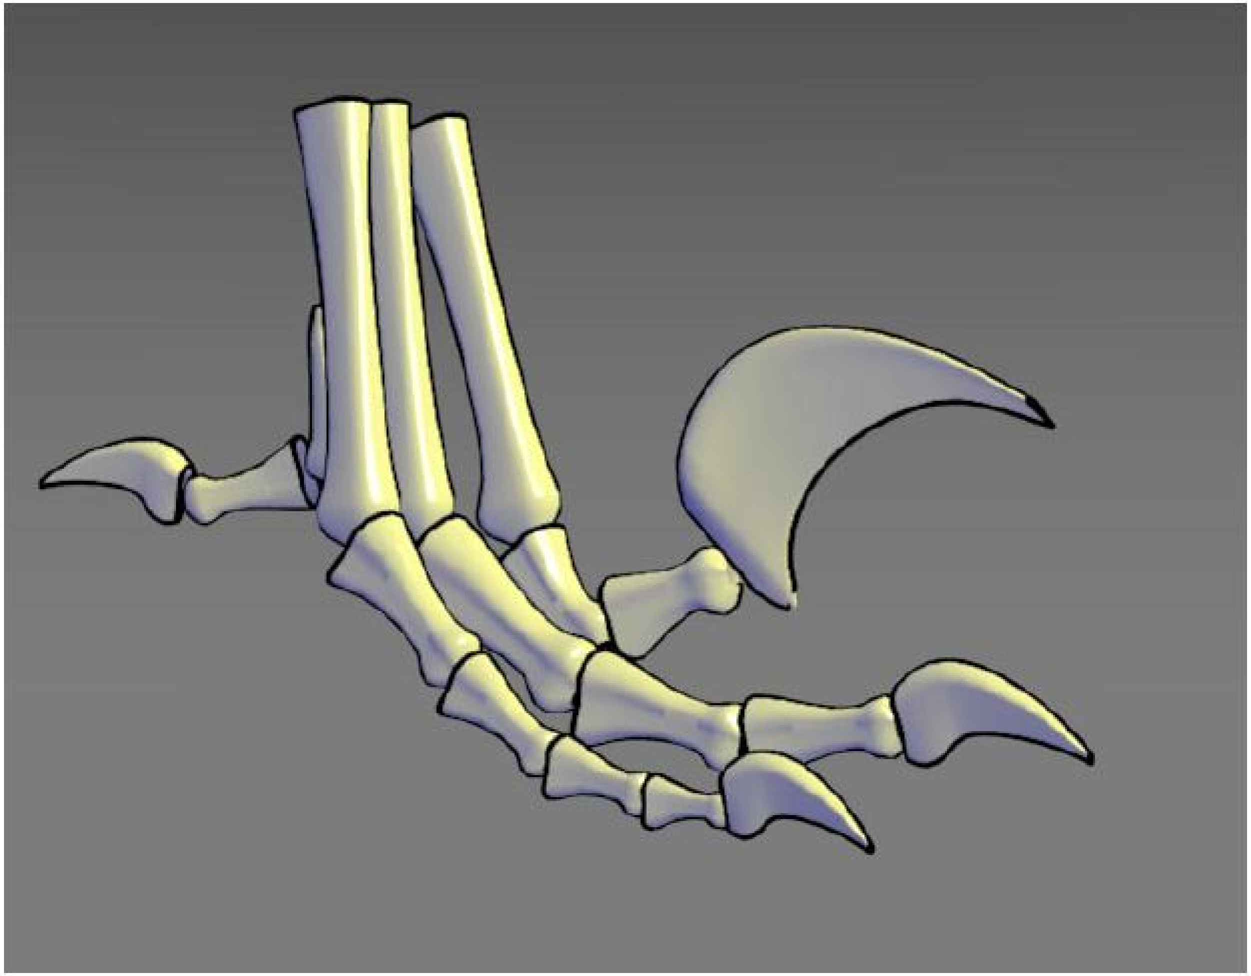

Figure 16: Luminance/hue tone rendering. This image combines the luminance shift of Figure 11 and the hue shift of Figure 14. Edge lines, highlights, fine details in the dark shaded regions such as the small claws, as well as details in the high luminance regions are all visible. In addition, shape details are apparent unlike Figure 12 where the object appears flat. In this figure, the variables of Equation 1 and Equation 1 are: $b = 0.4$, $y = 0.4$, $\alpha = 0.2$, $\beta = 0.6$. (See Color Plate).

Parameter setting # 2

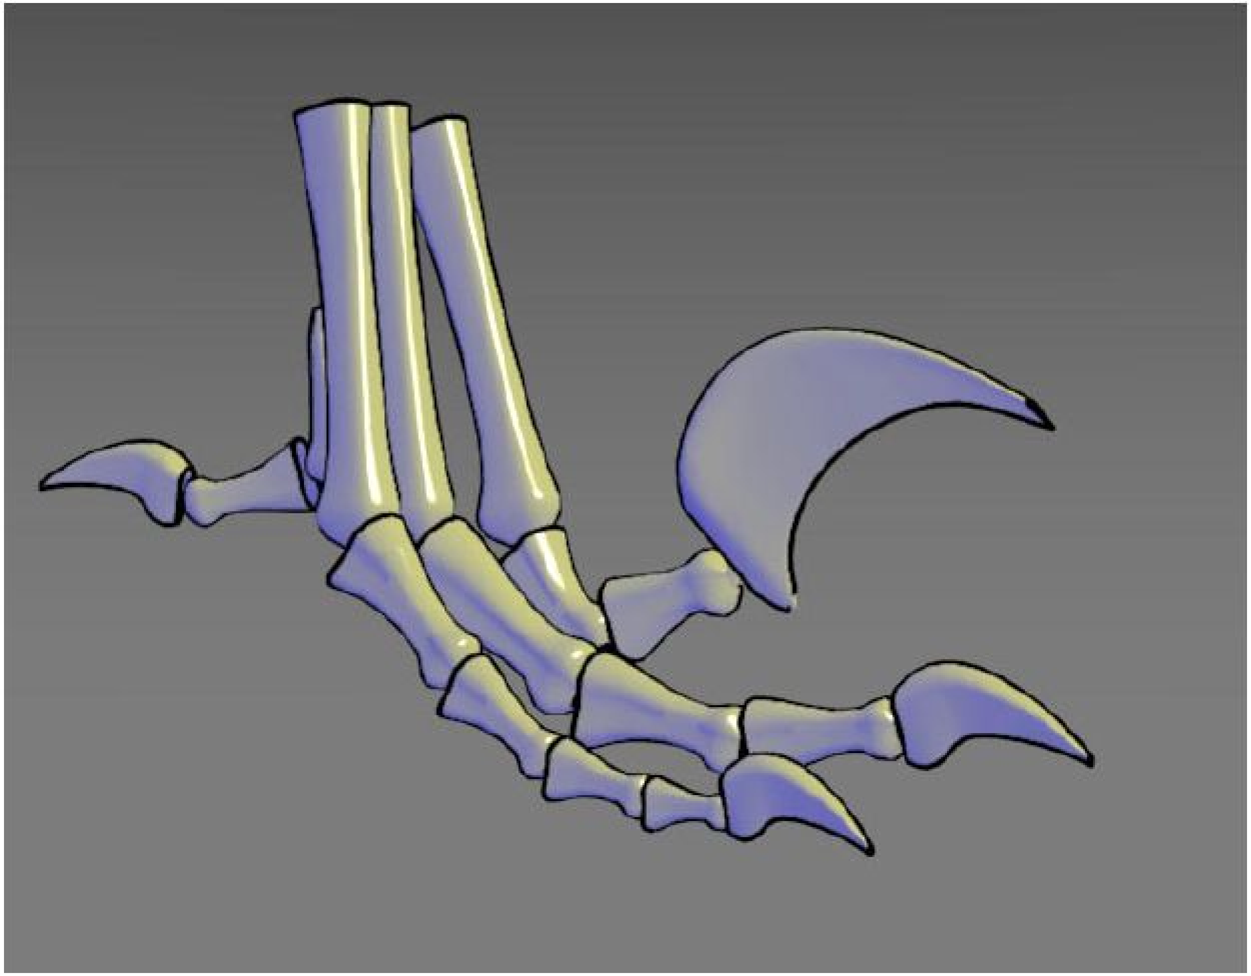

Figure 17: Luminance/hue tone rendering, similar to Figure 16 except $b = 0.55$, $y = 0.3$, $\alpha = 0.25$, $\beta = 0.5$. The different values of b and y determine the strength of the overall temperature shift, where as α and β determine the prominence of the object color, and the strength of the luminance shift. (See Color Plate).

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

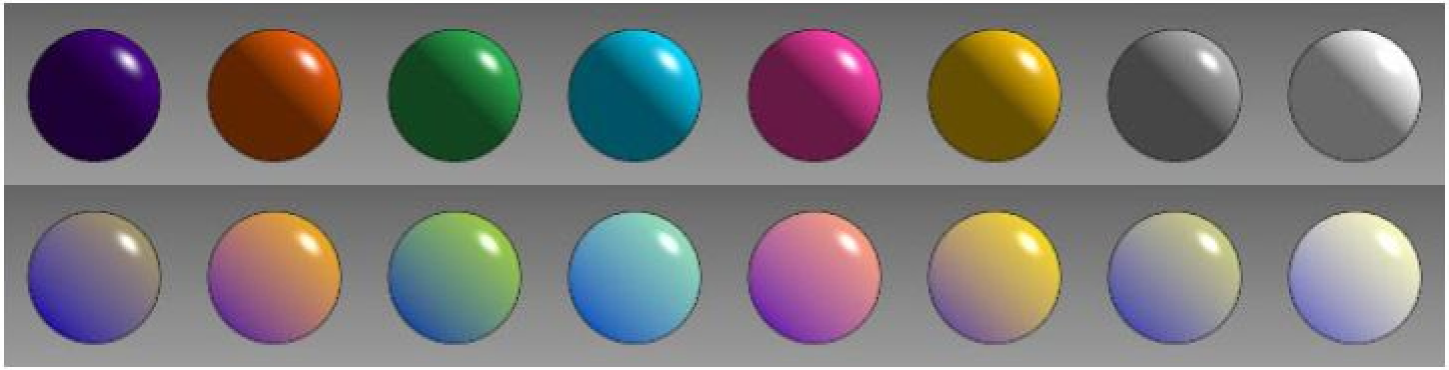

Figure 18: Comparing shaded, colored spheres. Top: Colored Phong-shaded spheres with edge lines and highlights. Bottom: Colored spheres shaded with hue and luminance shift, including edge lines and highlights. Note: In the first Phong-shaded sphere (violet), the edge lines disappear, but are visible in the corresponding hue and luminance shaded violet sphere. In the last Phong-shaded sphere (white), the highlight vanishes, but is noticed in the corresponding hue and luminance shaded white sphere below it. The spheres in the second row also retain their "color name." (See Color Plate).

Metal object with anisotropic reflections

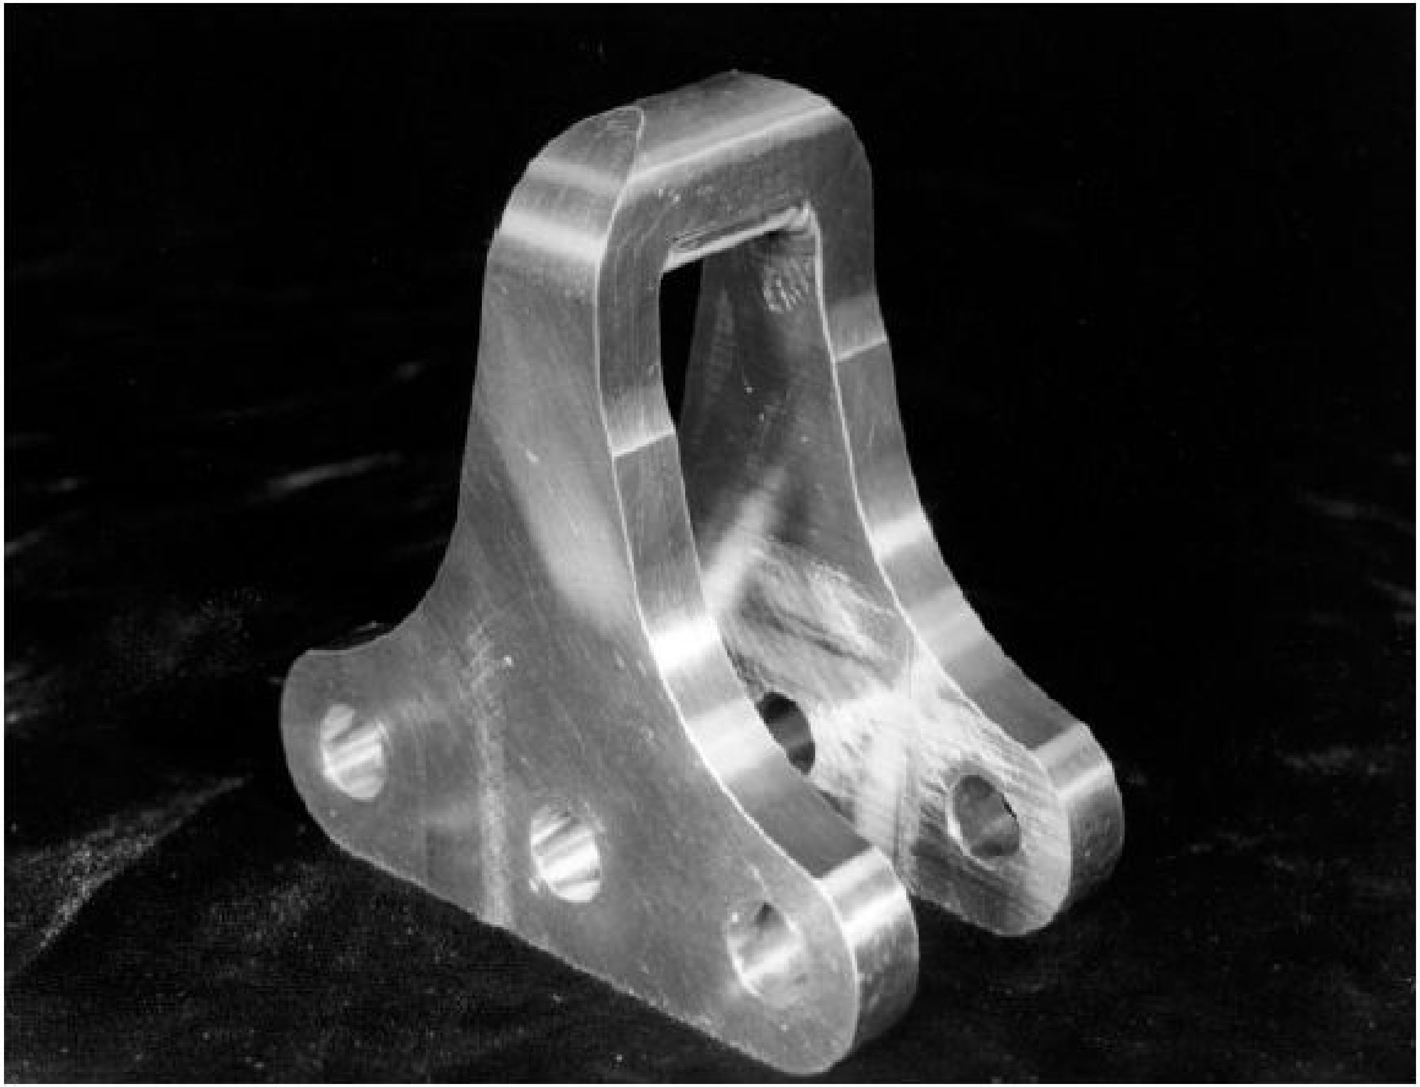

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

Metal object with anisotropic reflections

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

An anisotropic reflection can be seen in the metal objects in this photograph.

“Lines are streaked in the direction of the axis of minimum curvature, parallel to the milling axis.”

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

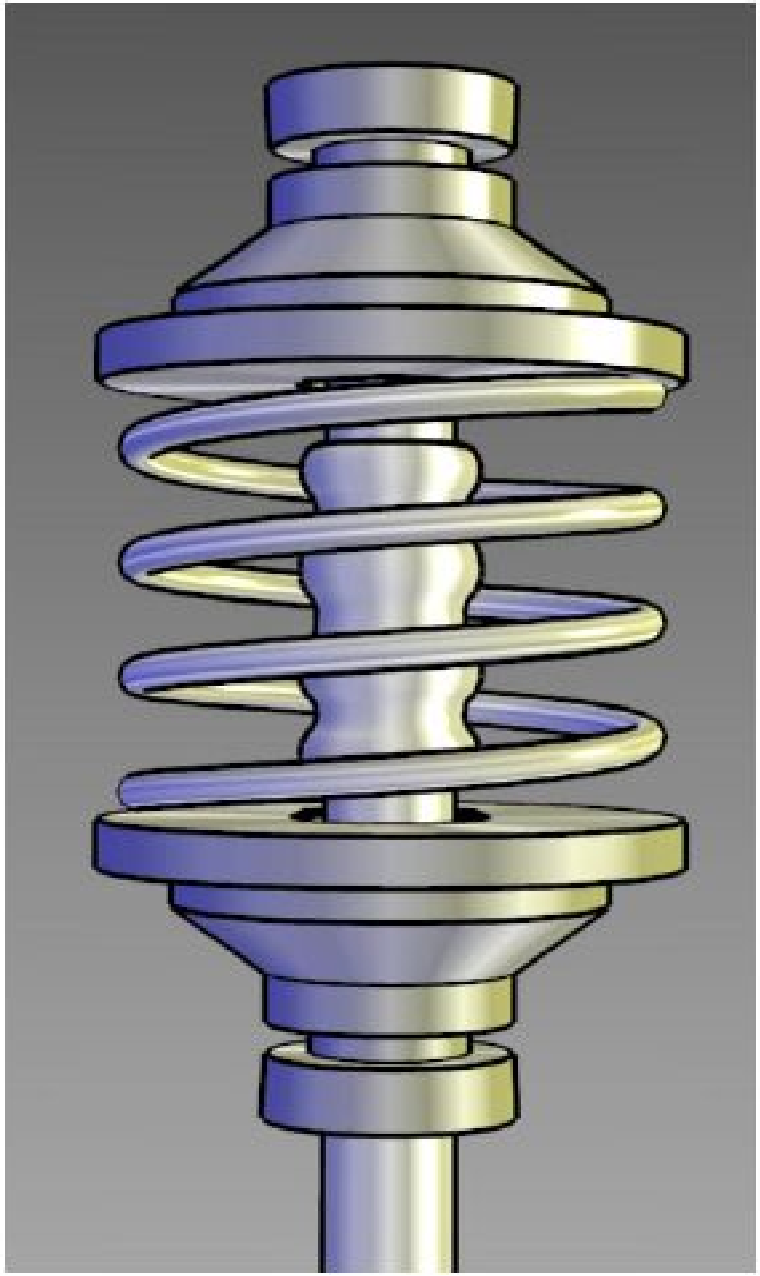

(d) Metal-shaded object with a cool-to-warm shift.

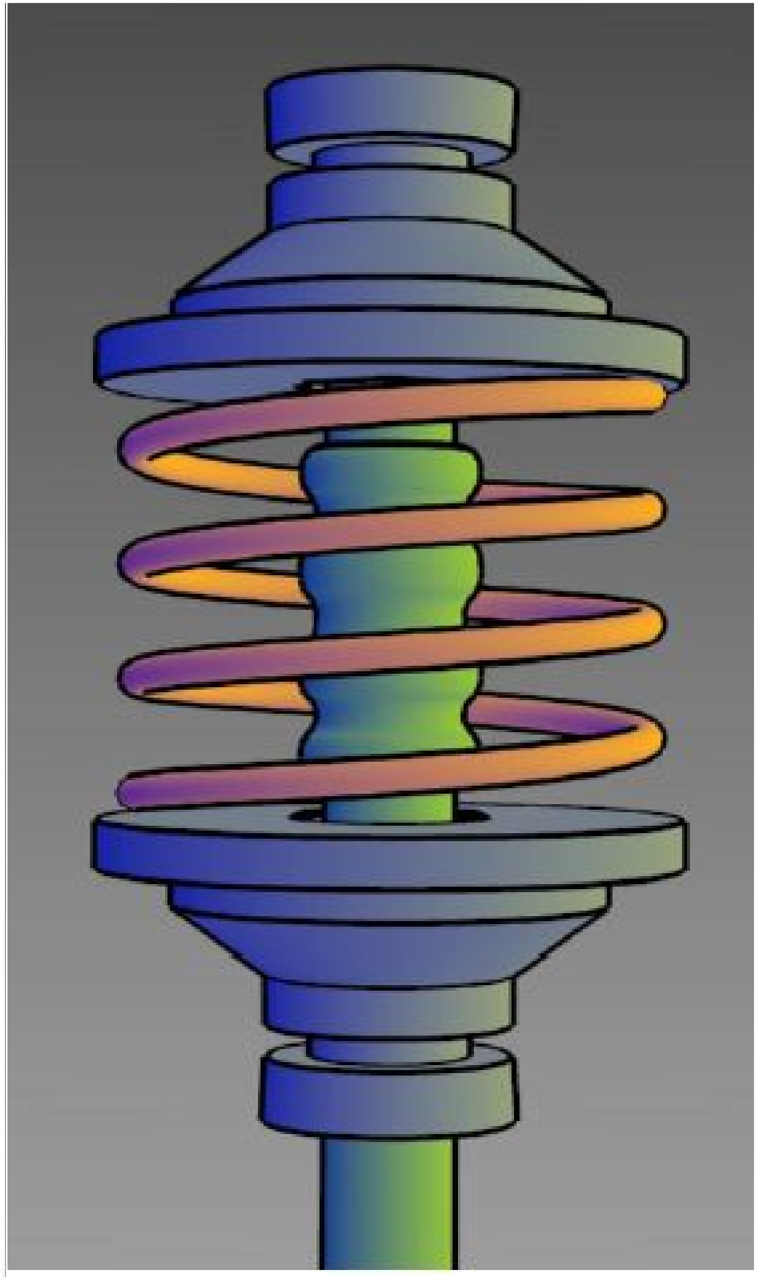

(d) Approximation: Phong shading, two colored lights, and edge lines.

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf

3D, little user intervention

Gooch has developed a non-photorealistic lighting model that attempts to emulate the richness of hand-drawn technical illustration [Gooch98]. The lighting model uses luminance and changes in hue to indicate surface orientation, and gives a clearer picture of shape, structure and material composition than traditional computer graphics methods.

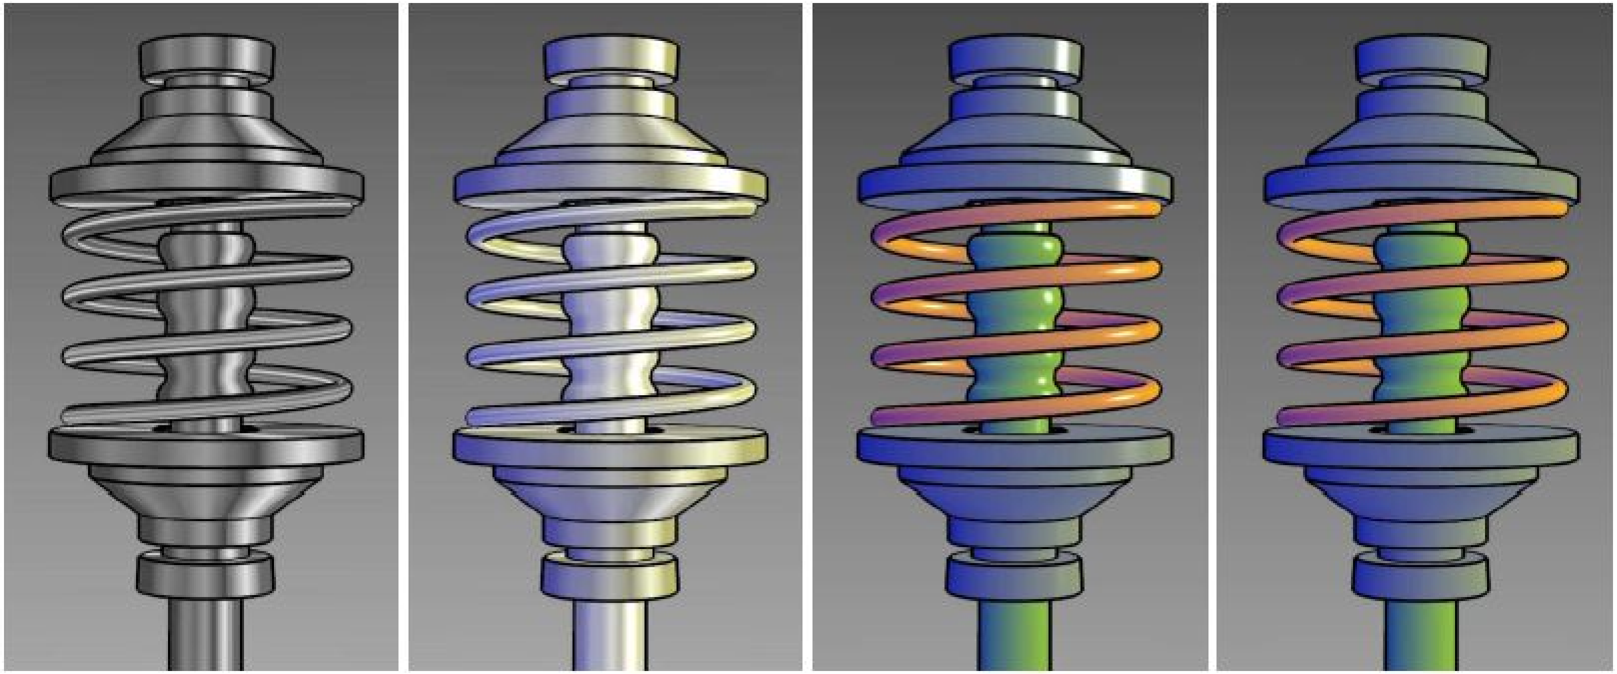

Figure 7: Non-Photorealistic Lighting Model for Automatic Technical Illustration, from [Gooch98]

http://www.cs.utah.edu/npr/papers/npr_course_Sig99.pdf