

Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Lecture 1

Getting the lecture material

From your browser, go to courses.kersten.org. Under Fall 2006, Click on **Psy 5036W**.

The material for each lecture is in a *Mathematica* Notebook. Find the **Lectures** section.

Download the first lecture:

1.IntroToComputationalVision.nb

Notebooks can contain inserted text and figures, as well as code, numerical results, and graphical outputs.

Sections of Notebooks can be organized by grouping and ungrouping cells and sections.

Practice opening and closing the cells.

Introduction

Understanding visual perception is an important problem in Psychology

One of the great mysteries of psychology is how the human visual system determines what and where objects are just by looking. This is the problem of vision. The perception of what is out there in the world is accomplished continually, instantaneously and usually without conscious thought. The very effortlessness of perception disguises the underlying difficulty of the problem. Vision is important because it is one of the principle routes to our acquisition of knowledge, as well as guide to its utilization.

Understanding vision is an important aspect of brain science

The problem of vision has attracted researchers from many disciplines outside of psychology, including computer science, mathematics, physics, engineering, and neuroscience. Understanding vision is a key problem in brain science, and the life sciences generally. Visual neuroscientists currently estimate that up to 50% of human visual cortex (your gray-matter) may be used for visual processing. The general structure of cortical layers and pattern of inter-connectivity is similar across the neocortex. Thus the hope that if we can understand visual computations in the cortex, this knowledge may generalize to other cognitive domains. With about 10 million retinal receptors, the human retina makes on the order of 10 to 100 million measurements per second. These measurements are processed by about a billion cortical neurons. Vision is a complex process requiring mathematical modeling tools.

Vision is a challenging mathematical problem

The problem of vision is not only important from the point of view of understanding the brain, but it is also formally and mathematically complex. As such vision is an active area of research for computer scientists and mathematicians.

■ *Old Man Picture from Mumford (1995)*



No known algorithm can locate, classify, or determine the shape, of the man in the picture. Or determine the material...skin, hair, cloth...

■ Why is it a hard problem? There are no locally recognizable features (the problem of local ambiguity)



■ From images to actions, objects

What we really want to understand is how to go from images to useful actions:

$$\mathbf{I} \rightarrow \mathbf{A}$$

To get there, vision requires information about objects, surfaces, and scenes. Consider the question: How can one estimate parameters of objects--their colors, shapes, materials, their relationship to other objects, to the viewer, to the viewer's hands, etc.--all from a glance? This is often referred to as the problem of *image understanding*, to emphasize that vision is a problem of perceptual inference. That is, given an image which is just a description of the light intensities at each point in space, and time (e.g. video camera or the sensors at the back of your eye), how can one infer the properties of the scene that caused the image?

Properties of objects and their relationships will be called **scene** attributes, and can be represented by a vector **S**. Let us represent the **image** by the time-varying spatial array of intensities, **I**. (More generally, **I** could be some derived image

measurements, like the location and orientation of local edge segments). Actions **A**, such as grasping an object, or saying "that's a dog", require information about scene attributes **S**. Thus sometimes, we'll think of **A** as a function of **S**: **A(S)**. But many times, we'll assume, **A=S**.

■ **The generative model**


The computer graphics problem is one of "forward optics" -- how to go from a description of the environment to the image:

$$S \rightarrow I$$


In other words, a function that describes what "out there in the world" (object shape, lighting, etc.) causes the image intensities observed.

This is sometimes referred to as the "generative model" because the causal structure leading to the data (image) is well-defined.

The generative model can be simple:



or complex:



In the first case, the signal could be light switch setting that is to "dim" (S1) or "bright" (S2), resulting in average light levels (μ_{S1} and μ_{S2}). But the actual measurements get perturbed by "noise", meaning that there is some fluctuation from one setting to the next, even if the switch is always set to, say, "bright". The second generative model is much more complicated and involves quite different representations for the signals. But later in the course, we'll see how many of the problems of vision share a common structure as shown in these graphs.

■ The inference or estimation model

The formal problem of vision is to go from \mathbf{I} to \mathbf{S}' , where \mathbf{S}' is an estimate of scene attribute(s) causing the image pattern.-
The mathematical problem of vision is sometimes said to be an inverse problem,

$$\mathbf{I} \rightarrow \mathbf{S}'$$

because it is the inverse of the generative model (a special case is the inverse optics problem, which is the inverse of the computer graphics problem).

Recall the first generative model, where the signal is a light switch setting that is set to "dim" (S_1) or "bright" (S_2), resulting in two different average light levels (μ_{S_1} and μ_{S_2}). How can we infer the value of the switch setting (\mathbf{S}' ="dim" or "bright") from a measurement of \mathbf{I} ? If dim always resulted in a lower value of \mathbf{I} , and a bright setting in a higher value, then the generative model is easy to invert--big \mathbf{I} means the bright switch was set, and low \mathbf{I} means it was the dim switch.

But the signal could be unreliable because the light is "noisy". If the noise is small compared with the differences between the average light levels (μ_{S_1} and μ_{S_2}), then there could still be a clear difference between the result of a dim vs. bright setting--i.e. dim results in little light \mathbf{I} , and bright results in lots of light \mathbf{I} . However, if the light emitted is close to the bright light, and the switch is fairly noisy, sometimes a "dim" setting could cause a lower value of intensity \mathbf{I} than a "bright" setting. Now we have a problem of ambiguity when trying to guess the switch setting from a measurement of \mathbf{I} . We'll study this simple problem of ambiguity in detail later in the course in order to get a better understanding of how to handle more complex generative models.

We will see later that the theory of statistical inference provides a natural framework to model under-constrained problems. Over a century ago, Hermann Helmholtz described perception as "unconscious inference". As we go on, we will justify and amplify on the Helmholtz definition of perception.


■ The problem of ambiguity for objects

■ "Identity Crisis" in Season 4 of Star Trek TNG.






But the problem can be even more challenging. The same object can give rise to different images. Different objects can give rise to the same image.



Current Opinion In Neurobiology

■ An example

The following section has a "mystery" image I , which is just a 2D list (matrix) of light intensity values. Let's make a plot that represents intensity as height:



This is a plot of I as a function of position, (x,y) .

Suppose our goal is to estimate depth, $S'(x,y)$, from I . One naive idea worth trying is to assume that I is proportional to depth, S' at each location. Although naive, it isn't a bad idea--there is a correlation between intensity and depth. Given this assumption, we'd conclude that the high middle ridge is closer than the bottom ridge. But as we will see in a moment, the dominant middle ridge that we see here does not correspond to near depths.

Consider another perceptual inference goal. What if we want to estimate the surface color (pigmentation or "paint") of the image? Let S'' be the surface color, where S'' is big for white surfaces, and low for black surfaces. Now it seems even more plausible that I would correlate very well with S'' .

Let's represent the mystery image in a form where your own visual system can judge:

■ Intensity representation of same data in mystery image:




First, note that the high intensity areas of the teeth are poor predictors of "nearness" in depth. They seem to be better measures of "whiteness". But what about the highlight on the lower lip? This produces a bump in the surface height plot above, but this a highlight due to the glossiness of the lip, and the pigmentation of the actual lip surface is not any lighter than other areas of the lip.

Information about shape and pigment is ambiguous, and the two are confounded in the simple image intensity measurements.

■ The problem of task

Some kinds of information are more important to estimate the other kinds. Face identity vs. illumination direction.



There are HUGE objective changes in the image going from left to right ("variability in illumination") that human perceptual judgments give very little weight to.

■ Where do we stand today?

Scientists, despite several decades of research have yet to produce a machine that can solve the general recognition problem: identify objects in natural images from arbitrary viewpoints, illumination conditions and in arbitrary contexts. As pointed out above, a central mathematical problem is the multiplicity of possible scene or object configurations that could have caused a particular set of image measurements. Richard Feynman compared the problem of vision to deciding what jumped into a swimming pool just by measuring the bobbing water height as a function of time using a ruler in the corner of the pool. There are lots of ways of getting a pattern of water heights. Further, as with water height, image measurements are very indirectly related to useful scene information. An understanding of image formation and optics does not sufficiently constrain the number of possible scene descriptions, s that could have given rise to any one image, I . This is one of the defining characteristics of inverse problems in general--the data underconstrain the solution.

One of the major contributions of computer vision has been to define the mathematical problems of vision and to show that these can be quite difficult to solve. Historically, the problem of chess was considered a prototypical problem for Artificial Intelligence. Today we have machines that can beat most of us at chess. A lot of their power comes from high speed brute force manipulations, no doubt quite different from the brain processes of the grand masters. Nevertheless, they can beat us, whereas there is no current system that can pick out the chess pieces from the box (because of variability over viewpoint, lighting, material, and style), and set them up on the board in the right places.

In this course we will study how the visual system deals with variability over illumination, viewpoint and material. In addition to the problem that different illumination conditions (e.g. light source coming from above left or above right), and different viewpoints produce different images of the same object, vision has to cope with occlusion of one piece by another. Like fonts, different chess sets have different styles. To some extent style variation can be modeled in terms of geometric variation. But chess piece styles can be determined by symbols having to do with the formation of concepts.

To further sober (and challenge you), the remarkable limitation of our understanding of visual inference is underscored by the fact that there are no machine systems that can solve the patently simple problem of deciding whether a surface has a light or dark pigment under general illumination. I.e., given a white or black chess piece in isolation, what color is it? The problem is that a black piece in bright light can have the same average intensity as a white piece in dim light. This is a problem that we will return to later in the context of human material and lightness perception.

Understanding vision requires combining approaches from psychology, neuroscience, and computation

Vision is a part of cognitive science -- an interdisciplinary effort to understand the nature of knowledge, its acquisition, storage and utilization. It is also part of Cognitive Neuroscience—the study of the relationship between brain and cognitive behavior. I'd like to spend some time motivating the importance of an interdisciplinary study of vision.

I think some motivation may be required here because of the nature of the course. In this course, we will study vision from a computational point of view. The topics should be exciting because the course involves integrating knowledge across disciplines. But it can be frustrating because although it involves some math and computation, it isn't like most quantitative disciplines that have a structured sequence, and you have to restore, revive, and learn new concepts to make the interdisciplinary picture come into focus.

Philosophy of the Computational Approach to Vision

What is Computational Vision? It is the study of how useful scene information and action parameters can be computed from image measurements. As we will see below, computational vision requires study at several *levels of abstraction*, as well as several disciplines.

We will study human visual behavior to understand what information is and is not used for visual inference.


We will study neural systems because we would like to understand how the machinery of the eye and brain enables actions and the inference of properties of the world from images.

What does it mean "to understand visual perception"?

To truly understand human or biological vision means we should be able to build or simulate a machine "which sees like we do". But what does this mean? There are several levels of abstraction that have to be considered, and the answer to this question requires careful thought. Robot vision, even if excellent, would not necessarily work the same way as human vision even at a very general level. For example, a vision system could rely on the reception of natural image patterns, or could actively send out signals to see how the environment modulates them. An example of the latter case would be to project alternating stripes of light and dark on surfaces, and then use the systematic distortion of the stripes in the image which to decode the shape. In fact, this latter principle is used (laser painted stripes) in commercial applications to measure shapes (e.g. Cyberware scans of the human face). But biological vision (in contrast to echo location) processes nature's patterns without active modulation of its images.

The study of cognitive science, and in particular vision, is necessarily interdisciplinary involving behavioral, biological and mathematical approaches. Let's take a closer look at the methods of three specific areas within each of these general approaches: Psychophysics, Neuroscience and Computer Science. We will gain some familiarity with all three of these disciplines in this course, but let us first have a preview of their respective contributions and see how they relate to different levels of analysis.

Behavior and Psychophysics: Black-box approach to construct a model that "sees like we do"



This approach has a *major goal to qualitatively and quantitatively describe visual behavior*. Psychologists, ethologists and "behavioral" neuroscientists all study behavior...but as we will see, even physicists and mathematicians get in the act. Careful description is an essential first step in any science (e.g. Mendel, and genetics).

The study of behavior can be of at least two types.

First, we need to know and understand the visual functions of an organism. Human vision is used to identify objects, to read, to walk, to drive, to steady oneself, to reach and grasp, to throw, to plan, to judge beauty, and the list goes on. Different tasks require different kinds of image processing.

This brings us to the second type of behavioral study, *psychophysical analysis*. Psychophysics measures the behavioral consequence of physical (or informational) variations in the image stimulus with a goal towards understanding underlying neural mechanisms. Examples are: the measurement of apparent brightness as a function of physical light intensity; just discriminable differences in light intensity; changes in sensitivity as a function of adaptation; changes in recognition performance as a function of viewpoint. Clever psychophysical experiments can reveal not only the diverse visual processing requirements, but also test hypotheses about alternative accounts of a given process. Psychophysics goes beyond mere description and historically has made some striking predictions about the nature of the underlying biology. For example, the psychophysics of color matching in the 19th century anticipated three physiological cone receptor types and their relative spectral sensitivities as a function of wavelength. This so-called "trichromacy" theory of was not physiologically established at the cell level until the 1960's. We will see later how psychophysics in the 1940's showed that photoreceptors (rods) in the eye could transduce single photons into an electrical signal. Certain brightness illusions discovered in the 19th century suggested patterns of neural connectivity and spatial interaction (called "lateral inhibition") that were not put on a firm physiological and neural foundation until the 1950's.

A psychophysical approach has clear limits in its ability to give an account of how we see. One reaches a point where too many theories of what is inside the box give the same input/output relation in the psychophysical data. For example, computing y as: $y = x(x - 1)$ gives the same mathematical relationship as a different computation in which x is subtracted from its square: $y = x^2 - x$. The fact that different combinations of wavelengths of light appear the same could have many neural explanations. Researchers eventually "go inside the box", or a animal "model" of the box (like a frog, cat or monkey) to find out what was going on at a finer level of analysis. This brings us to the methods of neuroscience, such as anatomical tracing, electrophysiological recording from single neurons, and brain imaging.

Neuroscience: Going inside the box

What happened when physiologists and anatomists looked at the biological basis of the psychophysical descriptions? Indeed, as discovered using microspectrophotometry in the 1960's, there are 3 distinct types of cone photoreceptors in the retina at the back of the eye. Electrophysiological recordings showed that their spectral sensitivities were remarkably similar, but not identical, to those inferred from psychophysics. And yes, as was shown in the 1970's, photoreceptors can transduce single photons. And there are neural circuits (lateral inhibition) in the retina that behave like Ernst Mach predicted to account for certain brightness illusions. Later, we will see examples of more recent neural accounts of psychophysical observations that go beyond retinal processing to other parts of the brain, accounts that are being tested using both electrophysiological and brain imaging techniques.

In the 1950's and 1960's there was a tremendous excitement that we could understand the brain's function and in particular visual perception in terms of single neurons....but neurobiologists had probably been particularly lucky...at least they were more fortunate than if the brain had been designed like a modern digital computer. In the 1970's, it became increasingly apparent that understanding how biological systems detected light was only scratching the surface of the problem of vision. Computer vision was beginning to show that competent vision was truly a problem of sophisticated inference and estimation. The number of visual areas discovered in the cortex of the brain grew.

Vision was becoming more complicated and harder than expected.

Computational Vision: The need for a new discipline to handle the complexities of perceptual inference

Imagine the following example. Sometime in the distant future, Martian scientists have acquired a Terran computer device that plays an ancient video game, say Mario Brothers 1. Now consider the various ways these scientists might go about trying to understand this device.

First, they could adopt technique adapted from a neuroscience, "anesthetize" the computer (i.e. just take away the screen, so there is no output), and begin using a volt meter or a logic probe to figure out what the box is doing. This is like doing neuroscience without psychology or psychophysics. The scientists might learn about logic gates, shift registers, and RAM, etc.. But, what are the chances of figuring out that the machine was even designed to play a game? Pretty slim.

But now give the scientists a working system complete with screen and the controls, but minus the logic probe. This is like doing behavioral science--psychophysics. With some careful experiments, they could begin to figure out the rules of the game. (Although, they might be left with questions forever unanswerable, like "why was this machine built in the first place"!). But if asked to "build a machine" that does the same thing, the Martian scientists might still have a hard time. They might be able to build a copy that mimics the behavior of the game, but even if the scientists had a solid body of results using the logic probe and observing the functioning system as a whole with the screen on, there is still something missing. They would have missed the point that the essential structure of the game is not the hardware, nor the input-output relations, but rather a highly complex computer program. They need an understanding of the intermediate level -- the software (firmware), the algorithms, and how that these are related to the hardware that supports it.

In short, what is missing is an understanding of how the pieces fit together to solve a specific information processing task-- a task that involves getting magic mushrooms, escaping turtles, smashing brick ceilings to get coins, jumping up flag poles, etc. Without this knowledge, they would be unable build new video games, like Mario Brothers 2. True understanding of vision should result in the generalization, e.g. the capability of building a machine that sees like us, but which may

differ from the original in ways that we can understand.

This example illustrates the need for a **computational approach to vision**. Although this approach grew out of the early communications theory, cybernetic, and artificial intelligence studies of the late 1940's, 1950's and 1960's (e.g. Turing, von Neumann, Wiener, Shannon), one of the chief protagonists of this approach for the study of visual perception was the late David Marr from MIT in the early 1980's.

The computational level involves understanding how image patterns are formed (generative models), and how scene inferences can be drawn from image patterns. To handle the complexity of natural patterns requires the tools of computer programming, and has been increasingly emphasized in recent years, the mathematics of statistical inference.


Levels of analysis

Although Marr made many specific contributions to understanding human vision, he is well-known for his elucidation of a computational approach as consisting of multiple levels of explanation for puzzles of perception.

■ Functional ("computational") Theory. What is the goal of a computation?

Why is it appropriate? What strategy can carry it out? Both psychology and theoretical analysis help to answer these questions. For example, the Necker cube (below) perceptually flips because the goal of the visual computation is to represent the 3D structure of the objects causing the 2D retinal image. But, there are two equally plausible 3D interpretations.

■ Necker cube




■ **Another example: Transparency and rigidity demo.**

(<http://gandalf.psych.umn.edu/~kersten/kersten-lab/demos.html>)


■ **Representation and algorithm. How to represent input and output? How to get from input to output?**

Psychology and computer program models help answer these questions. For example, mental rotation experiments done with human observers give clues as to how visual information is represented and processed. The time to decide whether a sample figure on the right is a rotated version of the one on the left or not, increases monotonically with the actual angle required to check the match (over a certain range and conditions).

The right figure is an image of the same object as the one on the left, but is rotated by 30 degrees about the vertical axis:



In the next figure below, the right figure is rotated by 80 degrees:



■ Implementation or hardware. How to build it with actual components?

Pencil and paper, tubes, silicon chips, hydraulics, billiard balls, or neurons? Neurobiology and neural network computer simulation help. For example, after-images are a well-known perceptual phenomenon. The after-image of a flash from a camera is an effect of visual perception that has to do with how human vision implements transduction in the retina, and the receptors in particular.

In future lectures we will see the relationships between human behavior (psychophysics), physiological mechanisms, and computational theory. Sometimes the computational theory comes first, but sometimes we will work backward from the experiment or visual phenomenon to the theory:

- o quantum limits to vision-- What are the theoretical limits to light discrimination? -> Computational theory of discrimination.

- o lateral inhibition-- for detecting edges? or to reduce redundancy? -> Computational theory of neural image coding.

In the next lecture we will see how a computational theory of discrimination can be applied to understanding the limits to the human ability to discriminate brightness.

Notes on getting started with *Mathematica*

- You can read *Mathematica* files free with MathReader.

- Go to the Help menu in *Mathematica*. Go through the *Tutorial*.

- If you wish, take a look at the overview from *Mathematica*'s "intro.nb" (*now a bit out of date*).

- **Numerical Calculations.** You can do arithmetic. For example, type $5+7$ as shown in the cell below, and then hit the "enter" key. Try other operations, 5^3 , $4*3$ (note that $4\ 3$, where a space separates the digits is also interpreted as multiplication). Note that if you try division, e.g. $2/3$, you get the exact answer back. To get a decimal approximation, type $N[2/3]$.

```
5+7
```

```
12
```

```
N[2/3]
```

```
0.666667
```

You can go back and select an expression by clicking on the brackets on the far right. These brackets are features of the Macintosh interface and serve to organize text and calculations into a Notebook with outlining features. You can group or ungroup cells for text, graphs, and expressions in various ways to present your calculations. Explore these options under Cell in the menu. You can see the possible cell types under the **Style** menu.

- **Built-in functions.** *Mathematica* has a very large library of built-in functions. They all begin with an uppercase letter. You can get information about a function, e.g. for the exponential of a function, or for plotting graphs:

```
?Exp
```

Exp[z] is the exponential function.

```
?Plot
```

Plot[f, {x, xmin, xmax}] generates a plot of f as a function of x from xmin to xmax. Plot[{f1, f2, ...}, {x, xmin, xmax}] plots several functions fi.

If you type two question marks before a function, **??Plot**, you'll get more information. Try it. What does the **Random** function do?

Making the lecture figures with Mathematica

- The **Necker cube** can be drawn by making a list of the coordinates, and ordering them by doing a mental "Etch-a-Sketch". The coordinates are placed in a list we call **points**. Then we use the **Line[]** graphics primitive to draw lines between the points:

```
points = {{0, 0, 0}, {1, 0, 0}, {1, 1, 0}, {0, 1, 0}, {0, 0, 0}, {0, 1, 0},
          {0, 1, 1}, {1, 1, 1}, {1, 1, 0}, {1, 1, 1}, {1, 0, 1}, {0, 0, 1},
          {0, 1, 1}, {0, 0, 1}, {0, 0, 0}, {0, 0, 1}, {1, 0, 1}, {1, 0, 0}};
Show[Graphics3D[{Thickness[.025], Line[points]}, Boxed → False,
               Shading → False, ViewPoint → {100, 150, 100}]];
```

- The **mental rotation** figures are made by generating a list of 4 3-D coordinates, whose values are randomly assigned between 0 and 1. Then we draw lines between them, and make 3D plots from two different views. To get the different views we use an "add-on" function in a library **Geometry`Rotations`** that we read in:

```
<< Geometry`Rotations`
```

Now generate the list of 4 3-D coordinates:

```
randomwire1 = Table[{Random[], Random[], Random[]}, {5}];
```


The `RotationMatrix3D[theta,0,0]` will rotate a point by $\text{Pi}/2$ radians about the z-axis.

```
degtheta = 80;
theta = degtheta / 57.3;
(*convert degrees to radians*)
```

```
view1 = {1, 1, 1}; view2 = RotationMatrix3D[theta, 0, 0].view1;

glv1 = Show[Graphics3D[{Thickness[.03], Line[randomwire1]},
  Boxed → False, Shading → False, ViewPoint → view1],
  DisplayFunction → Identity];
glv2 = Show[Graphics3D[{Thickness[.03], Line[randomwire1]},
  Boxed → False, Shading → False, ViewPoint → view2],
  DisplayFunction → Identity];

Show[GraphicsArray[{glv1, glv2}]];
```



- The **mystery** figure was read in from a jpeg file. The first thing is to set the directory to wherever the file is located. This shows where it was on my powerbook the day I wrote this:

```
In[6]:= SetDirectory [
  "/Users/kersten/Sites/kersten-lab/courses/Psy5036W2006/Lectures /
  1_MultidisciplinaryStudy"]
```

```
Out[6]= /Users/kersten/Sites/kersten-lab/courses/Psy5036W2006/Lectures/1_MultidisciplinaryStudy
```

Then use the import function to read in the image:


```
In[7]:= Show[gimage = Import["16x80mystery.jpg"]];
```



The image is input as "Graphics" object, and you can get the image list data out by pulling out the first element:

```
mysteryimage = gimage[[1, 1]];
```

```
ListPlot3D[mysteryimage, Mesh -> False];
```



Below, we've listed the mystery image so that it will be stored in this notebook, and we won't have to find it and read it in next time. If the image was really big, we wouldn't do this.

■ **mysteryimage**

```
mysteryimage =
{{175, 182, 166, 147, 137, 117, 101, 107, 115, 117, 121, 124, 123,
  117, 110, 109, 110, 107, 108, 115, 121, 120, 116, 130, 136, 135,
  149, 157, 156, 160, 163, 155, 152, 161, 173, 173, 185, 189, 189, 189}},
{192, 184, 167, 154, 129, 103, 100, 101, 98, 92, 90, 84, 79, 76, 78,
  90, 91, 86, 88, 99, 105, 113, 101, 90, 87, 88, 98, 116, 136, 154,
  158, 156, 156, 151, 166, 177, 181, 181, 190, 200}},
{197, 178, 157, 147, 119, 102, 106, 90, 74, 61, 60, 62, 74, 86, 91,
  107, 108, 102, 103, 124, 139, 151, 127, 100, 89, 72, 65, 75, 88,
  119, 140, 153, 155, 157, 168, 173, 176, 181, 189, 192}},
{203, 182, 155, 136, 121, 114, 99, 61, 41, 45, 58, 75, 92, 100, 103,
  112, 125, 119, 115, 145, 171, 175, 139, 111, 98, 85, 71, 66, 55,
  73, 101, 140, 157, 165, 170, 172, 179, 181, 183, 192}},
{201, 181, 151, 131, 127, 113, 73, 37, 31, 47, 62, 76, 87, 91, 98,
  107, 106, 101, 92, 114, 141, 146, 123, 103, 87, 82, 72, 68, 63, 58,
  61, 95, 129, 157, 172, 171, 173, 180, 185, 191}},
{199, 173, 147, 130, 120, 88, 45, 33, 42, 54, 61, 62, 64, 67, 77, 91,
  106, 111, 103, 99, 108, 112, 108, 88, 81, 67, 64, 64, 66, 62, 52,
  57, 80, 124, 160, 171, 170, 181, 188, 183}},
{194, 167, 146, 132, 104, 66, 42, 42, 39, 39, 41, 39, 58, 95, 133,
  176, 194, 208, 199, 179, 192, 196, 194, 163, 98, 86, 76, 50, 50,
  54, 51, 51, 56, 79, 115, 160, 178, 182, 186, 191}},
{195, 171, 152, 133, 88, 50, 39, 30, 20, 24, 58, 85, 128, 166, 174,
  199, 214, 222, 196, 172, 214, 229, 225, 190, 142, 165, 156, 117,
  129, 104, 57, 42, 48, 53, 68, 122, 169, 186, 187, 198}},
{186, 166, 154, 110, 64, 48, 25, 1, 21, 64, 109, 116, 134, 124, 98,
  119, 127, 128, 105, 96, 124, 148, 139, 118, 89, 148, 166, 132, 153,
  168, 101, 85, 66, 41, 49, 71, 125, 176, 193, 195}},
{185, 169, 145, 97, 48, 19, 1, 2, 22, 57, 70, 58, 67, 68, 62, 63, 59,
  71, 78, 78, 74, 77, 77, 66, 62, 72, 71, 59, 80, 103, 74, 62, 93, 59,
  41, 26, 59, 127, 174, 189}},
{182, 170, 124, 67, 35, 16, 3, 8, 26, 49, 60, 59, 64, 69, 77, 76, 80,
  97, 109, 112, 106, 98, 94, 90, 87, 77, 73, 70, 62, 59, 48, 41, 39,
  33, 32, 8, 3, 44, 116, 180}},
{180, 146, 86, 49, 47, 46, 50, 65, 79, 80, 86, 81, 79, 91, 112, 133,
  150, 166, 177, 184, 185, 173, 152, 131, 127, 125, 128, 135, 131,
  122, 115, 113, 107, 87, 66, 48, 43, 48, 85, 149}},
{159, 134, 100, 91, 99, 101, 104, 109, 116, 122, 129, 126, 126, 142,
  163, 183, 190, 191, 201, 209, 205, 206, 202, 179, 168, 165, 165,
  165, 167, 160, 149, 149, 147, 140, 124, 111, 109, 107, 118, 154}},
{154, 145, 131, 124, 123, 123, 122, 118, 127, 138, 138, 146, 151,
  161, 179, 188, 199, 191, 191, 200, 203, 207, 208, 197, 181, 173,
  174, 172, 170, 161, 154, 153, 152, 156, 152, 144, 137, 140, 147, 160}},
{156, 150, 147, 147, 142, 127, 117, 123, 121, 127, 131, 147, 153,
  162, 177, 189, 195, 186, 176, 182, 196, 203, 203, 201, 187, 175,
  171, 169, 166, 157, 155, 149, 147, 152, 157, 163, 160, 162, 165, 169}},
{137, 151, 159, 155, 152, 136, 117, 118, 119, 123, 141, 148, 150,
  166, 171, 185, 177, 168, 165, 171, 176, 187, 198, 193, 181, 179,
  172, 166, 160, 148, 152, 147, 151, 154, 152, 164, 173, 176, 173, 178}}
```

```
ListDensityPlot [mysteryimage, Mesh → False, Frame → False,  
AspectRatio → Automatic];
```



References

Helmholtz, H. v. (1867). Handbuch der physiologischen optik . Leipzig: L. Voss.

<http://www.psych.upenn.edu/backuslab/helmholtz/>

Hoffman, D. D. (1998). Visual Intelligence . New York: W. W. Norton & Company.

Kersten, D. High-level vision as statistical inference. (1999) *The New Cognitive Neurosciences*, 2nd Edition, Gazzaniga (Ed.). MIT Press.

(Download now: pdf preprint)(http://vision.psych.umn.edu/www/kersten-lab/papers/CogNeuro98-final.pdf)

Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13(2), 1-9.

Knill, D. C., & Richards, W. (1996). Perception as Bayesian Inference . Cambridge: Cambridge University Press.

Marr, D. (1982). Vision: A Computational Investigation into the Human Representation and Processing of Visual Information . San Francisco, CA: W.H. Freeman and Company.

Mumford, D. (1995). Pattern theory: A unifying perspective. In D. C. Knill, & R. W. (Ed.), Perception as Bayesian Inference (Chapter 2). Cambridge: Cambridge University Press.

Poggio, T. (1984). Vision by Man and Machine. *Scientific American*, 250, 106-115.

Zeki, S. (1993). A Vision of the Brain . Oxford: Blackwell Scientific Publications.

© 2006 Daniel Kersten, Computational Vision Lab, Department of Psychology, University of Minnesota.
kersten.org