

Computational Vision

U. Minn. Psy 5036

Daniel Kersten

Lecture 3: The Ideal Observer

Goals

Last time

Limits to light discrimination

In the famous experiment of Hecht, Schlaer, and Pirenne, measurements were made of the proportion of times observers saw a dim small briefly flashed spot of light. Under the right conditions, a person can detect an amazingly small number of photons. Perhaps even more remarkable was the discovery that single photoreceptors are capable of transducing single photons. We singled out the problem of variability in the photon emissions and absorptions as a key computational problem to understand.

Why was the subject response curve sigmoidal? If there was a specific light intensity (or number of photons) at transition from invisible to visible, we should see a "step" function--but we don't. The answer is *variability or "noise"*. But where does it come from? The variability may be due to several causes: the statistics of photon emission and absorption we've mentioned, "noisiness" in human responses, or in neural transmission. The physical limits to visual detection and discrimination can be traced to the stochastic nature of photon absorption and emission. This stochastic nature or variability is a consequence of the particle nature of light. A careful understanding of these limits in conjunction with psychophysical measurements is an essential component to understanding human sensitivity.


We will see how to understand how to model variability in the physical stimulus and thereby understand the ideal observer's limitations. Eventually we will return to the question of what are the primary sources of variability in the human discrimination of light intensity. The theory will provide a springboard to analyze the limits of human ability to do other more complex tasks.

Preview of visual decisions as inference, hypothesis testing

Inference is choosing an hypothesis based on data. In statistical inference, the decision is based on a model of the probabilities of the hypotheses, H and the data, k (i.e., x above). What are our hypotheses?

Hecht, Schlaer and Pirenne did a simple detection experiment in which the light was always flashed by the experimenter. There were no "catch trials", the observer didn't have to decide whether the light was flashed or not--just whether they thought they saw a flash or not.

The distinction becomes clearer if we generalize the task to discrimination.



Think of the task in terms of signal transmission. The "sender" wishes to set the light switch to one of two positions, corresponding to a high and low average light intensities. The signal (or hypothesis space) is simple, the states of the switches: $H = S_L$ (low), or $H = S_H$ (high). The idea is that we have two hypotheses that each determine a conditional probability distribution.

Thus, given a measurement of x photons, there is an inherent ambiguity in determining the cause or signal-- S_L or S_H ? In the this and the next lecture we will develop signal detection theory (SDT), and then see how to extend the fundamental principles of SDT to general perceptual inference for image patterns and visual tasks. We effectively turn the generative problem on its head, and derive the best guesser for deciding light switch position based on a measure of light intensity.

The computational theory of signal detection and estimation was developed in the 1940's and 1950's. More information about **Signal Detection Theory (SDT)** can be found in a now classic book by Green, D. M., & Swets, J. A. (1974). Treatments from an engineering perspective can be found in Van Trees, H. L. (1968) and Poor (1994).

Although, Hecht et al. measured thresholds for absolute detection; now we will study discrimination because it makes the key points clearer and is more general.

Introduced basic strategy of our approach to vision:

From the starting point of signal detection theory, we'll develop the theory of ideal observers, and as the course goes on, a computational framework to understand perception generally as a process of statistical inference. The general strategy is:

A: Understand the information processing problem

1) Generative modeling--Understand the image (data) formation process: $S \rightarrow I$

e.g. for light detection, S is the switch value or average number of photons for that value, which can take on two values corresponding to "on" or "off". I is the number of photons which varies from trial to trial for a fixed switch position.

The generative model was a "Poisson process". Of particular importance is that *there is variability in the generative process itself*.

2) Inference/estimation modeling--Understand how to make the best guess of causes of data: $I \rightarrow S'$

e.g. I is a measurement of the number of photons, and S' is the guess of the switch value. The strategy for the best guess is limited by the variability inherent to the generative process. We'll call the best guesser an "ideal observer" for the task.

B: Understand how well human vision handles the inference ($I \rightarrow S'$), and from there how it works.

In general, the observer introduces additional variability when making an inference, so its performance will be sub-optimal. Call the observer guesses S'' . Ideal observer analysis compares S'' to S' .

Generative model terminology: Some conceptually related terms:

S -> I

(Causes, scene properties, signals, hypotheses, states of the world) -> (data, image intensities, image measurements, features)

Today

We study some of the key concepts in Signal Detection Theory (SDT) and ideal observer analysis in the context of light intensity discrimination. The essence of the ideal observer approach is to ask: What is the optimal or ideal strategy for a task? What are the theoretical limits to performance on a task? In particular for our example, what are the limitations on the discrimination of light levels that any observer must face?

Today's goals are:

- o Practice some Mathematica tools useful for modeling variability and ideal decision making.
- o Intuitive introduction to the basics of probability distributions, cumulative distributions.
- o Introduce the idea of an ideal observer that models "external variability" inherent to the generative process. The ideal observer makes the best guesses with respect to some performance goal, such as minimizing error rate.

Our generalization of the ideal observer assumes that to minimize error, a decision should choose the most probable signal (e.g. switch setting) given a measurement. The jargon is that the ideal makes a decision that maximizes the "posterior probability".

- o Introduce the yes/no task, "hits" and "false alarms" (also called "false positives"). Hit and false alarm rates completely summarizes the ideal's performance for this task. (Recall Hecht et al.'s experiment only reported hit rate, the "proportions of yes's" when the light was flashed. They didn't measure what observers said when no light was flashed.). We'll also define "correct rejections" and "misses".

Preview of needed tools & concepts

Basic Mathematica

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 through the Help Browser under the Help menu item. You can also get help on-line, e.g. for the exponential of a function, or for plotting graphs:

?Exp

Exp[z] is the exponential function >>

- 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. >>

- 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 **RandomReal** function do?

Defining functions. You can define your own functions. Here is a gaussian function:

```
gauss[x_, mean_, std_] := Exp[-((x - mean)^2)/(2 std^2)]/(std Sqrt[2 Pi]);
```

And you can view the definition in standard form by evaluating:

```
gauss[x, μ, σ]
```

$$\frac{e^{-\frac{(x-\mu)^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma}$$

And you can check a your own definition:

```
? gauss
```

The underscore, **x_ is important** because it tells Mathematica that x represents a slot, not an expression. Note that we've used a colon followed by equals (:=) instead of just an equals sign (=). When you use an equals sign, the value is calculated immediately. When there is a colon in front of the equals, the value is calculated only when called on later. So here we use := because we need to define the function for later use. A double equals, ==, is a relational operator used for testing whether an equation is true or not.


Also note that we could define our function with **N[]**. *Mathematica* tries to keep everything exact as long as possible and thus will try to do symbol manipulation if we don't explicitly tell it that we want numbers. Try it:

```
ngauss[x_, mean_, std_] := N[Exp[-((x - mean)^2)/(2 std^2)]/(std Sqrt[2 Pi])];
ngauss[x, u, σ]
```

$$\frac{0.398942 \times 2.71828^{-\frac{0.5(x-1.0)^2}{\sigma^2}}}{\sigma}$$

Graphics. Let's plot the graph of the gaussian function:

```
Plot[gauss[x, 0, 1], {x, -5, 5}]
```



- As with the Poisson distribution, *Mathematica* has a built-in function for the gaussian distribution. Can you find it and plot make a plot similar to the above? Hint: It isn't called gaussian, and when you find it, you can just cut and paste from the documentation example.

Lists

We are going to do a lot of work with lists, in particular with vectors and matrices to represent images, as well as transformations on images. Sometimes we'll use vectors to represent 2D images and use matrices to do transformations on those images. Other times, we'll use matrices to represent images. Here is a four-dimensional vector which we'll call x . x could represent the pixel intensities in an image, or the input signal strengths at 4 receptors:

```
x = {2,3,0,1};
```

By ending a line with a semi-colon, you suppress the output after hitting the return key. Now let's make another vector, this one will be a list of "weights", say, representing the efficiency with which the inputs are weighted before being summed in the visual receptive field:

```
w = {2,1,-2,3};
```

Now the output of a model neuron that simply takes a weighted sum of the inputs is just the dot product of the input with the weights:

```
y = w.x
```

```
10
```

This kind of operation is sometimes referred to as a "cross-correlator". It takes a signal x , and cross-correlates it with a template, w . In your homework you calculate the performance of a cross-correlator. It is one of the simplest measures of similarity between two patterns. A cross-correlator is also the simplest model of a neuron, such as a retinal ganglion cell. But more on this later.

Using `Table[]` to make lists and images

We will need to convenient ways of making vectors and matrices (e.g. in order to synthesize an image or a filter). The `Table[]` function does this for us. For example, make a list whose elements are the squares of the element location.

```
s = Table[x^2,{x, 1, 8}]
```


```
{1, 4, 9, 16, 25, 36, 49, 64}
```

We will use `Table[]` quite a bit, because it provides a compact way to iterate function values into an array. To see the result, we use `ListPlot[]` which takes a list of numerical values, rather than a function as the argument for plotting.

```

Clear[x];
poisson[x_,a_] := Exp[-a] a^x / Factorial[x]
poisson[x,a]
p1 = Table[N[poisson[x, 4]], {x, 0, 20}];
ListPlot[p1]

```




- Try this one-liner that does the same thing, and with a prettier graph:
`DiscretePlot[PDF[PoissonDistribution[4],k],{k,0,20}]`

You can also use **Table[]** to make a list of lists, e.g. to define a 2D image. Here is a 64x64 matrix that we can plot up using **ArrayPlot[]**:

```

A = Table[Cos[ $\frac{2 \pi x}{24}$ ], {x, 1, 64}, {y, 1, 64}];
ArrayPlot[A, Mesh -> False, ColorFunction -> "GrayTones", PlotRange -> {-2, 2}]

```



- What happens if you change the PlotRange values for ArrayPlot?

You can also use **Image[]**.

Image[] has several “high-level” advantages in that the format is used in *Mathematica*’s built-in image processing functions. It also provides quick summaries of the image properties. **Image[]** assumes that values of 0 and 1 map to the maximum values of light and dark--compare the left and right images:

```
{Image[A], Image[A] // ImageAdjust}
```



You can also use: **ListDensityPlot[A, ColorFunction->GrayLevel[]]**, but it will interpolate values, hiding the internal quantization into discrete graylevels. Try it.

Below we are going to use the following functions:

Table, Range, Map, Count, ArrayPlot, Show, PoissonDistribution

We will also use Function[], but in short-hand form.

Probability concepts


In the next lecture, we will systematically go over the elements of probability and inference in detail. But today, we will take an intuitive approach, letting the light discrimination problem motivate the need for the concepts we need. We will motivate the later use and development of the concepts of: **Random variable**, **Probability distribution**, **Conditional probability**, **Prior probability**, **Posterior probability**.

Modeling external variability: The Ideal Observer

Generative model for stimulus: Given the switch setting, what measurements result?

Schematic of set-up for ideal photon counter

Our goal is a quantitative description of distribution of photons measured under the two switch settings. We will show how to simulate the process on a trial-by-trial basis, and how to compare the simulation with theory.



Demonstration of the idea: two switch settings produce various intensities

As at the end of the last lecture, define a Poisson distribution with a mean of **mean**, with a function to draw a sample from this distribution:


```
numberofphotons[mean_] := RandomInteger[PoissonDistribution[mean]];
```

Let highmean = 7 for the "high" setting, and lowmean = 5 for the "low" setting. Now, turn the switch to "high", draw a sample, and then to "low" and draw a sample. Let's plot up the two:

```

highmean = 7; numhighsample = numberofphotons[highmean];
lowmean = 5; numlowsample = numberofphotons[lowmean];
highsample = Table[RandomReal[{0, 1}, 2], {numhighsample}];
lowsample = Table[RandomReal[{0, 1}, 2], {numlowsample}];
highg = Graphics[{PointSize[0.04], White, Point /@ highsample},
  AspectRatio → 1, Frame → False, FrameTicks → None, Background → GrayLevel[0.],
  PlotRange → {{-0.2, 1.2}, {-0.2, 1.2}}, PlotLabel →
  "High Flash: " <> ToString[numhighsample], LabelStyle → Directive[White]];
lowg = Graphics[{PointSize[0.04], White, Point /@ lowsample}, AspectRatio → 1,
  Frame → False, FrameTicks → None, Background → GrayLevel[0.],
  PlotRange → {{-0.2, 1.2}, {-0.2, 1.2}}, PlotLabel →
  "Low Flash: " <> ToString[numlowsample], LabelStyle → Directive[White]];
Show[GraphicsRow[{highg, lowg}, Spacings → Scaled[0.3]], Frame → False]

```




Execute the above cell several times. If you have patience or luck, you will see that sometimes the "Low Flash" setting produces more dots than the "High Flash" setting.

(Note: The two flash demo at the end of Lecture 2 put the dot generating code into a Module[]. More about modules later.)

Histograms and population distributions

Imagine first that we do an experiment as described above. First set the light level to S_L . Now we begin to send flashes by opening and closing the shutter. As emphasized in the previous lecture, no matter how hard we try to make the apparatus perfect, we would discover that the photon counter doesn't always record the same number of photons for each flash. Then we do the same for the switch setting, S_H . So, we compile a histogram of photon counts and find something like this:



The horizontal axes represents a photon count, and the vertical axis represents the proportion of times that a particular count was measured. The histograms provide a quantitative representation of the variability. As the number of samples grows (i.e. test flashes in our case), the histogram becomes a better and better approximation of a "true", underlying Poisson distribution.

Although we don't prove it here, the Poisson distribution arises because of the independence of photon absorption.

The Poisson distribution is characterized by the mean (or expected) number of photons flashed when the switch is set to S_L . The spread in the distribution (measured by the standard deviation) is equal to the square root of the mean. As above, the probability of k photons being detected for a Poisson distribution is given by:

?? PoissonDistribution

PoissonDistribution[] represents a Poisson distribution with mean μ . >>

Attributes[PoissonDistribution] = {Protected, ReadProtected}

PDF[PoissonDistribution[a], k]

$$\begin{cases} \frac{a^k e^{-a}}{k!} & k \geq 0 \\ 0 & \text{True} \end{cases}$$

with a mean level and variance of a . The standard deviation is the square root of the variance. (This is one of many standard statistical distributions that describe various kinds of uncertainty.

Let's simulate the process of photon generation and do some histogram counting using *Mathematica*.

Simulation of the generative process-- Photon counts, and histograms corresponding to high and low switch settings

We compile a histogram of photon counts conditional on the switch setting. Define high and low mean models:

```
highpdist = PoissonDistribution[7];
lowpdist = PoissonDistribution[4];
nsamples = 1000;
```

Draw nsamples (e.g. = 1000 flashes) for each switch setting.

```
sample[ntimes_] :=
  Table[Random[lowpdist], {ntimes}];
```

```
zlow = sample[nsamples];
```

- Note: This kind of simulated random draws occurs so frequently, that you can do it all in one line: **RandomVariate[PoissonDistribution[4],1000]**

How frequently do trials with 1, 2, ...n dots occur? We can use the built-in function: **Histogram[]** to show us,

using: **Histogram[{zlow,zhigh}]** .

- Try it.

Or as an exercise to get exposure to a bit of "functional programming", we can count the frequencies,

lowFreq, ourselves,

```
domain = Range[0, 15];
lowFreq = N[Map[Count[zlow, #] &, domain] / Length[zlow]];
```

- What does Count[zlow,#]& [4] do? Use the Help Browser to look up the definitions of: **Count[]**, **Map[]**, and **Function[]**.


Now let's measure a histogram for the high switch setting:

```
sample[ntimes_] :=
  Table[Random[highpdist], {ntimes}];

zhigh = sample[nsamples];
domain = Range[0, 15];
highFreq = N[Map[Count[zhigh, #] &, domain] / Length[zhigh]];
```

Plot up both of the histograms showing the frequency that the various possible photon counts occur:

```
BarChart[Transpose[{highFreq, lowFreq}], PlotRange -> {0, 0.25}, BarSpacing -> {1, 0}]
```



Theoretical -- Compare Poisson probability distributions to high and low simulated histograms

PDF[PoissonDistribution[x], a]

$$\begin{cases} \frac{e^{-x} x^a}{a!} & a \geq 0 \\ 0 & \text{True} \end{cases}$$

In standard probability notation, if the mean value is a , then the probability of random variable X taking on the value of x photons is:

$$p(X=x \text{ photons}) = \frac{a^x e^{-a}}{x!}$$


The mean or expectation of X is $E(X) = a$, and for a Poisson distribution, the variance, $var(X) = E[(X - a)^2]$ is also equal to a . The mean gives a measure of "central tendency", and the variance is measure of "spread" of the distribution. (Note that we could write our own poisson function simply as: **poisson[x_,a_] := Exp[-a] a^x / Factorial[x]**);

Let's make two lists with the theoretical values, corresponding to means of 4 and 7:

```
theorylowFreq = Table[N[PDF[PoissonDistribution[4], x]], {x, 0, 15}];
theoryhighFreq = Table[N[PDF[PoissonDistribution[7], x]], {x, 0, 15}];
```


Compare simulated photon count probabilities with theory predictions. What does the plot look like if nsamples is much smaller, say nsamples = 50?

```
BarChart[Transpose[{theorylowFreq, lowFreq, theoryhighFreq, highFreq}],
PlotRange -> {0, 0.25}]
```



For a compact simulation and fancier display, we can do all of the above here:

```
Show[Histogram[Table[RandomVariate[PoissonDistribution[i], nsamples], {i, {4, 7}}],
{1}, "PDF", PlotRange -> All],
Plot[Evaluate[Table[PDF[PoissonDistribution[μ], x], {μ, {4, 7}}]],
{x, 0, 15}, PlotStyle -> Thick, PlotRange -> All]]
```



- You can play with the **nsamples** variable. Note that the more samples we draw, the better the experimental histograms match the theoretical. Try a small value of nsamples, say **nsamples = 100**;

The main point is: photon production is inherently statistical, and because of the independence of photon absorption, the histogram for ideal photon counting can be modeled as a Poisson distribution.

Conditional probabilities: a preview

Our histograms are estimates of the *conditional probabilities* $p(x|S_L)$ and $p(x|S_H)$. I.e. they represent the probability of observing a photon count of x , given (or "conditional on") the switch setting $H = \{S_L \text{ or } S_H\}$. We will treat the switch setting as an hypothesis, which is also a random variable, i.e. H . As we've noted earlier, the hypothesis is about the state of the world, such as a signal state (high or low setting), or an object property.

A note on terminology: $p(x|S_L)$ is the probability of x conditional on S_L . But below when we consider optimal inference, we will talk about likelihood functions, $p(x|H)$, where for example $p(x|H = S_L)$ is the *likelihood* of S_L given a measurement of x .

The knowledge from a generative model gets embodied in the conditional probability--given a cause H (where $H = S_L$ or S_H), the conditional probability tells us the probability of a measurement x . In a concise form, $p(x|H)$, is one component of the generative model, and as we see below, $p(H)$ (the "prior") is

the other part.

Inference model for light intensity discrimination: *Given a measurement, what was the switch setting?*

Our goal now is to model how, from a measurement of the photon count, one can make the best guess as to the switch setting. The best guess is based on choosing the most probable switch setting given a measurement. This is called choosing the setting that "maximizes the posterior distribution". In particular, we'll see how probability theory can be applied to showing that maximizing the posterior can be done quite simply for this problem.

We'll see that the answer is exactly what your intuition would be: you can't do any better than to base your decision on a count of the number of photons.

Yes/no task

In a yes/no task, we designate one switch position as the "signal" and the other as "not the signal", or "noise". We'll have the observer say "yes" when the switch is set to high.

Posterior probability distribution

Given a measurement, what was the switch setting? We've seen that because of variability in the measurement, it is impossible to answer this question correctly every time just from the data. However, we can ask how the decision should be made so that an observer is correct as often as possible (equivalent to minimizing the probability of error). So we ask:

What decision strategy will minimize the probability of error?

The observer that achieves a pre-defined optimality criterion like this is called an "*ideal observer*". We'll see other ideal criteria later.

Intuitively one might guess (correctly) that, given a photon measurement count, the ideal observer should decide "signal", based on the switch setting (i.e. decide the hypothesis) which has the bigger probability of the two. In other words, measure x , and then calculate:

**If $P(S_L | x) > P(S_H | x)$, choose S_L
otherwise, choose S_H .**

$P(S_L | x)$ is "the probability of the switch taking on the value of S_L , conditional on a measurement x ", and similarly for $P(S_H | x)$. This rule is called the "maximum *a posteriori*" or MAP rule. ...but these aren't the conditional probabilities we calculated and graphed above.

$P(S_L | x)$ and $P(S_H | x)$ are called *posterior* probabilities. Why "posterior"? Even without any data measurements, knowledge of $P(S_L)$ and $P(S_H)$ can provide important information for optimal decisions, and these are called *prior* probabilities, i.e. the information the observer has *prior* to the arrival of the data measurements. Posterior probabilities are "informed by the data".

Now it may be hard to directly specify the posterior probability, and it is easier to model the prior probabilities and the likelihoods, and then use Bayes rule to relate them to each other:

$$p(H | x) = \frac{p(x | H) p(H)}{p(x)} \propto p(x | S) p(S)$$

..ah, looks like now we will be able to plug in the conditional probabilities we calculated earlier.

Bayes' rule is a classic rule in probability theory, attributed to the Rev. Thomas Bayes (1702-1761). After you've gone over the basics of probability (see ProbabilityOverview.nb) you'll be able to easily prove it.

Optimal performance is determined by the posterior probability, and that is why Bayes' rule is of central importance in determining optimal theories of inference. Let's consider the roles the priors and likelihoods play separately.

Prior only

Suppose the light transmission switch is set to S_L 4/5 of the time, and to S_H 1/5 of the time.

```
pH = {4 / 5, 1 / 5};
(*pH = {1 / 2, 1 / 2};*)
```

But the photon counter isn't working at all--so there are no data. What is the observer's optimal strategy to guess whether it was the high or bright flash? First we have to specify the receiver's goal, and then it makes sense to answer this question. If the receiver wants to minimize the average probability of error, then it should always pick S_L .

Like the maximum a posteriori rule, the rule that minimizes the error rate is:

```
choose  $S_H$  if  $p(S_H) > p(S_L)$ 
choose  $S_L$  if  $p(S_L) > p(S_H)$ 
```

So if $p(S_H) > p(S_L)$, the rule says to always choose S_H .

(Interesting enough, people often don't do this. With feedback, they tend to match probabilities. This is a phenomenon which has been studied by cognitive psychologists and economists, cf. Green et al., 2010).

Likelihood

We use *Mathematica* to define $p(x|H) = \text{poisson}[x, a]$, where the hypothesis or switch setting H , determines the average number of samples, x :


```
poisson[x_, a_] := PDF[PoissonDistribution[a], x]
```

Then we make two lists of probabilities corresponding to means of 4 and 9:

```
p1 = Table[N[poisson[x, 4]], {x, 0, 20}];
p2 = Table[N[poisson[x, 9]], {x, 0, 20}];
```

Then use **BarChart[]** to view the distributions.

```
BarChart[Transpose[{p1, p2}], PlotRange -> {0, 0.25}, AxesLabel -> {"x", "p(x|H)"}]
```



As earlier, the plot shows the frequencies of observing x photons under the two possible switch settings, **low** and **high**.

Let's put the two likelihoods together into one 2x21 array:

```
poispxH = {p1, p2};
Dimensions[poispxH]
{2, 21}
```

Choosing the hypothesis with the largest likelihood will minimize the probability of error when the prior probabilities are uniform--i.e. both the same. Then maximum likelihood decisions are the same as maximum a posteriori decisions in this case. But what if the priors are not uniform (i.e. the same constant probability for all outcomes)?

Combining the prior and likelihood: Maximum a posteriori (MAP) rule

Suppose (as we assumed for the light discrimination example) that we know the likelihoods: $p(\mathbf{x}|\mathbf{S}_L)$ and $p(\mathbf{x}|\mathbf{S}_H)$. Bayes' rule tells us how to combine them to obtain the a posteriori probability of the hypotheses conditional on the data:


$$p(H | x) = \frac{p(x | H) p(H)}{p(x)} = \frac{p(x | H) p(H)}{\sum_H p(x | H) p(H)}$$

We can calculate the posterior using the following, somewhat opaque, line of *Mathematica* code:

```
posterior = Transpose[Transpose[{poispxH pH}] / Plus @@ (poispxH pH)];
```

Below is a plot of the posterior probabilities, $p(\mathbf{S}_H|\mathbf{x})$ and $p(\mathbf{S}_L|\mathbf{x})$, as function of the data \mathbf{x} :

```
BarChart[Transpose[{posterior[[1]], posterior[[2]]},
  PlotRange -> {0, 2}, AxesLabel -> {"x", "p(H|x)"}]
```



Note that the plot shows the probability of the switch being **low** or **high**, i.e. $p(H=S_L)$ or $p(H=S_H)$ given that x photons were counted. (Why does the posterior plot look so different from the likelihood plot?)

Let's apply the *maximum a posteriori* rule (MAP)

For *maximum a posteriori* estimation, the rule is:

choose S_H if $p(S_H|x) > p(S_L|x)$
 choose S_L if $p(S_L|x) > p(S_H|x)$

More generally, a function which chooses the value of argument ($H = S_L$ or S_H) that maximizes a function is sometimes written as:

$$\operatorname{argmax}_H p(H|x)$$

- Suppose 8 photons were counted, pick the most probable switch setting by inspecting the above graph.
- Compare this plot, which assumed $p_H = \{4/5, 1/5\}$, with: $p_H = \{1/2, 1/2\}$ and a measurement of 8.

Let's apply the Maximum likelihood rule

In maximum likelihood estimation, the rule is:

choose S_H if $p(x | S_H) > p(x | S_L)$
 choose S_L if $p(x | S_L) > p(x | S_H)$

Because of Bayes rule, if the prior probabilities are equal, then MAP is equivalent to maximum likelihood.

$p(x | S_H)/p(x | S_L)$ is called the likelihood ratio.

- Suppose 8 photons were counted, what is the maximum likelihood decision? (See $p(x|H)$ plots above, or recompute $p(H|x)$ using $p_H = \{1/2, 1/2\}$.)

Simple and classic OPTIMAL Rule for light discrimination: Compare the photon count to a criterion.

Let's assume the priors are equal, so we base our decisions using the maximum likelihood rule. We will use *Mathematica* to do some symbolic manipulations to find the transition point in our data measurements (i.e. photon count) where the likelihood goes from favoring a low to a high switch setting. This is where the likelihood ratio is 1, or the log of the ratio is zero.

$$\text{loglikelihoodratio} = \text{Log}[\text{poisson}[x1, b] / \text{poisson}[x1, d]]$$

$$\text{Log} \left[\begin{array}{ll} \frac{b^{x1} e^{-b}}{x1!} & x1 \geq 0 \\ 0 & \text{True} \end{array} \right]$$

$$\left[\begin{array}{ll} \frac{d^{x1} e^{-d}}{x1!} & x1 \geq 0 \\ 0 & \text{True} \end{array} \right]$$

$$\text{PowerExpand}[\text{Log}[\frac{b^{x1} e^{-b}}{x1!}] - \text{Log}[\frac{d^{x1} e^{-d}}{x1!}]]$$

$$-b + d + x1 \text{Log}[b] - x1 \text{Log}[d]$$

There is a simple point to make here--the photon or dot count, $x1$, is monotonically related to the log likelihood. This means that our intuition that we can't do any better than simply counting dots was correct. But we need to decide on an appropriate criterion.

$$\text{Solve}[\text{Log}[\frac{b^{x1} e^{-b}}{x1!}] == \text{Log}[\frac{d^{x1} e^{-d}}{x1!}], x1]$$

$$\left\{ \left\{ x1 \rightarrow \frac{b-d}{\text{Log}[b] - \text{Log}[d]} \right\} \right\}$$

If d and b are the average number of photons for the low and high switch settings, and the photon counts are distributed according to a Poisson distribution, then minimal error will result on average with the following rule:

Say "high" if

$$x > \frac{b-d}{\text{Log}[b] - \text{Log}[d]}$$

Say "low" otherwise. The photon count, x , is said to be the *decision variable*. Because it is monotonically related to the likelihood ratio, it is an optimal decision variable. The transition value is called the *decision criterion*:

$$\text{decision criterion} = \frac{b-d}{\text{Log}[b] - \text{Log}[d]}$$

So now you have all you need to know to write a program that would make optimal decisions about high vs. low--it would simply count photons, and base its decision on whether the count was greater or less than the decision criterion.

Summarizing performance using performance metrics

So far so good. But we would like to go a littler farther so that we can calculate theoretically what the average performance of an ideal observer would be.

In a yes/no task, an ideal observer can make two kinds of mistakes. It can decide the signal was there when it wasn't (a **false alarm** also called a false positive), or it can decide it wasn't there when it was (a **miss** also called a false negative). These two determine the error rate.

An ideal can also be correct in two ways. It can have a **hit** (true positive) or a **correct rejection** (true negative). We'll see later that you only need to measure two of these four statistics, because the hit + miss rate have to add up to 100% (if the signal is presented 100 times, you either say yes or no, so the sum of the hit and misses has to be 100).

In the next section you can try competing against the ideal, and have the opportunity to calculate hit and false alarm rates.

Try your luck against the Ideal Observer: Yes/no task

Set up a mini-experiment

Let's pretend that photons are like dots. In 1977, Barlow published a paper which did just that. He pointed out that humans can efficiently, but not perfectly, discriminate between patterns of differing dot density.

As above, let's define a Poisson distribution with a mean of **mean**, with a function to draw a sample from this distribution. Let **highmean** = 220 for the "high" setting, and **darkmean** = 200 for the "low" setting. Then let's initialize an array called **data**, to accumulate the human, ideal, and true answers, and the corresponding hit, false alarm, correct rejection (cr), and miss events.

```
dotsize = 0.01;

numberofphotons[mean_] := RandomInteger[PoissonDistribution[mean]];
highmean = 220; darkmean = 100;
criterion = N[-  $\frac{\text{highmean} + \text{darkmean}}{\text{Log}[\text{highmean}] - \text{Log}[\text{darkmean}]}$ ];
(*We'll learn about this in the next lecture*)
data = {"myanswer", "idealanswer", "trueanswer",
        "ideal hit?", "my hit?", "ideal false alarm?", "my false alarm?",
        "ideal cr?", "my cr?", "ideal miss?", "my miss?"};
```

We'll define a blank screen to turn the display on and off. Set the variable **flash** to blank.

```
blank = Graphics[{PointSize[dotsize], Black, Point /@ {{}, {}},
                 AspectRatio -> 1, Frame -> False, FrameTicks -> None,
                 Background -> GrayLevel[0.0], PlotRange -> {{-0.2, 1.2}, {-0.2, 1.2}}];
flash =
  blank;
```

Now create a window whose contents contain the graphics structure "flash". We make **flash** dynamic so that the display will get automatically updated whenever we change the value of **flash**.

```
nb3 = CreateWindow[DocumentNotebook[Dynamic[flash]],
                  ShowCellBracket -> False, WindowSize -> {300, 300}, WindowElements -> {},
                  Background -> Black, NotebookFileName -> "Yes/No Flash Display",
                  WindowMargins -> {{600, 200}, {120, 120}}];
```

Define a trial that shows one flash, either high or low

Now, randomly turn the switch to "high" or "low", draw a sample. We put in a dialog box so that you can input your response (click "yes" for "high" and "no" for "low").

```

oneflash := Module[{sample, idealanswer}, whichflash = RandomInteger[{0, 1}];
If[whichflash == 1, numsample = numberofphotons[highmean],
  numsample = numberofphotons[darkmean];
sample = Table[RandomReal[{0, 1}, 2], {numsample}];
flash = Graphics[{PointSize[dotsize], Red, Point /@ sample},
  AspectRatio → 1, Frame → False, FrameTicks → None,
  Background → GrayLevel[0.0], PlotRange → {{-0.2, 1.2}, {-0.2, 1.2}}];
Pause[.5];
flash = blank;
myanswer = ChoiceDialog["High dot mean?",
  {"yes" → 1, "no" → 0}, WindowMargins → {{570, 100}, {120, 120 + 300}}];
idealanswer = If[numsample > criterion, 1, 0];
trueanswer = whichflash;
data = Append[data,
  {myanswer, idealanswer, trueanswer, If[idealanswer == 1 && trueanswer == 1, 1, ""],
  If[myanswer == 1 && trueanswer == 1, 1, ""],
  If[idealanswer == 1 && trueanswer == 0, 1, ""], If[myanswer == 1 && trueanswer == 0,
  1, ""], If[idealanswer == 0 && trueanswer == 0, 1, ""],
  If[myanswer == 0 && trueanswer == 0, 1, ""], If[idealanswer == 0 && trueanswer == 1,
  1, ""], If[myanswer == 0 && trueanswer == 1, 1, ""]}];];

```

Practice run

Let's run 20 practice runs with highmean = 220; darkmean = 100;

```
For[iter = 0, iter < 20, iter++, oneflash]
```

Display the data

```
Style[TableForm[data], FontSize → 9]
```

myanswer	idealanswer	trueanswer	ideal hit?	my hit?	ideal false alarm?	my false alarm?	ideal cr?	my c
0	1	1	1					
0	1	1	1					
0	1	1	1					
1	0	0				1	1	
1	1	1	1	1				
1	1	1	1	1				
1	1	1	1	1				
0	0	0					1	1
1	1	1	1	1				
0	0	0					1	1
1	1	1	1	1				
0	0	0					1	1
1	1	1	1	1				
1	1	1	1	1				
1	1	1	1	1				
1	1	1	1	1				
0	1	1	1					
0	0	0					1	1
1	1	1	1	1				

Analyze the data

Let's drop the table heading stored in row 1, and then transpose the matrix so that the columns become the rows:

References

- Applebaum, D. (1996). Probability and Information . Cambridge, UK: Cambridge University Press.
- Barlow, H. B. (1962). A method of determining the overall quantum efficiency of visual discriminations. Journal of Physiology (London), 160, 155-168.
- Barlow, H. B. (1977). Retinal and central factors in human vision limited by noise. In B. H. B., & F. P. (Ed.), Photoreception in Vertebrates Academic Press.
- Duda, R. O., & Hart, P. E. (1973). Pattern classification and scene analysis . New York.: John Wiley & Sons.
- Geisler, W. (1989). Sequential Ideal-Observer analysis of visual discriminations. Psychological Review, 96(2), 267-314.
- Green, C. S., Benson, C., Kersten, D., & Schrater, P. (2010). Alterations in choice behavior by manipulations of world model. Proc Natl Acad Sci U S A.
- Green, D. M., & Swets, J. A. (1974). Signal Detection Theory and Psychophysics . Huntington, New York: Robert E. Krieger Publishing Company.
- Poor, H. V. (1994). An introduction to signal detection and estimation (2nd ed.). New York: Springer-Verlag.
- Van Trees, H. L. (1968). Detection, Estimation and Modulation Theory . New York: John Wiley and Sons.
- Watson, A. B., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method. 33, 113-120.