Keeping the Physics in Biophysics -- and vice versa

Phil Nelson
University of Pennsylvania

For these slides see:
www.physics.upenn.edu/~pcn

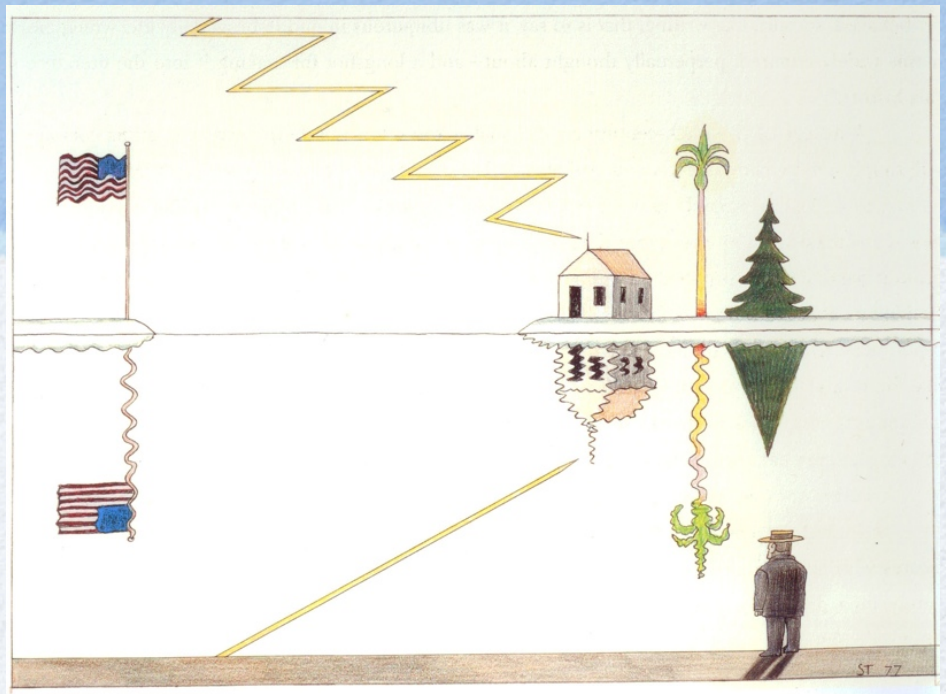

What we face

- Our Physics colleagues are *still* saying --
  - "That's not really physics."
  - "That's already offered in some other department."
  - "We don't have the resources for that."
- Our Biology colleagues are *still* not sending their students to our courses.
- "Biology students cannot/will not do math."
- "Physics students get uncomfortable with biological physics -- there are too few `Tripos questions.'"

I'll discuss those questions in the context of describing a course that works at U. Penn. If you don't face these particular questions, you may still be interested in the choice of topics.

Why do we even have classes at all?

- **To tell them facts?**
  No -- facts are now free in infinite quantity.
- **To tell them the latest, most trustworthy facts?**
  No -- facts go out of date in the blink of an eye.

Well -- *skills and habits* still matter a lot. When you walk into a room with your toolbag and encounter a problem you've never seen, which tool should you pull out of your bag? Knowing where to begin is a difficult but learnable skill.

Students need to develop the right skills and habits for that, but many (most?) courses don't really help.

A class should help them do that -- in some specific context. ***Biological physics is an interesting context for that purpose,*** regardless whether a student goes on in that field.

The interesting questions in science are those where we shake our heads and ask, ***"How could anything like that possibly happen at all?"*** And biophysics is full of such questions.

The boffins speak

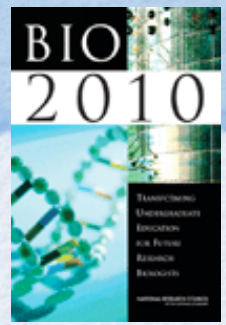

RECOMMENDATION #1.3

The principles of physics are central to the understanding of biological processes, and are increasingly important in sophisticated measurements in biology. The committee recommends that life science majors master the key physics concepts listed below. Experience with these principles provides a simple context in which to learn the relationship between observations and mathematical description and modeling.

The typical calculus-based introductory physics course taught today was designed to serve the needs of physics, mathematics, and engineering students. It allocates a major block of time to electromagnetic theory and to many details of classical mechanics. In so doing, it does not provide the time needed for in-depth descriptions of the equally basic physics on which students can build an understanding of biology. By emphasizing exactly solvable problems, the course rarely illustrates the ways that physics can be applied to more recalcitrant problems. Illustrations involving modern biology are rarely given, and computer simulations are usually absent. Collective behaviors and systems far from equilibrium are not a traditional part of introductory physics.

Students have got to learn how to get computers to do useful things *from scratch*. They also need to get used to extracting useful information from big datasets.

http://books.nap.edu/

A course that works at Penn

Physicists are pretty comfortable teaching about molecular biophysics.

- It's a good fit -- we like to talk about entropy
- It's still an exciting, opening field.
- A number of modern textbooks are now available.

But...

- Maybe you've already got one of those.
- Molecular biophysics is not so obviously connected to *medicine*, which many of our students plan to study.
- Students are intrinsically interested in *themselves*, e.g., their own visual system.
- Physicists have a lot to say about **Systems**, although we're somehow ceding a lot of the high ground to mathematicians and computer scientists.

Let's see how a Physics department could offer an useful course to a wide variety of students, including the more numerate Bio majors, **and the burgeoning group of Engineering majors interested in bio applications.**

Much inspiration from Bialek/Botstein 2004; Wingreen/Botstein 2006.

Genetic switching

Monod found something funny in the growth of bacteria in mixed medium. He asked, ***"how could that possibly happen at all?"*** And he ended up with the operon model.

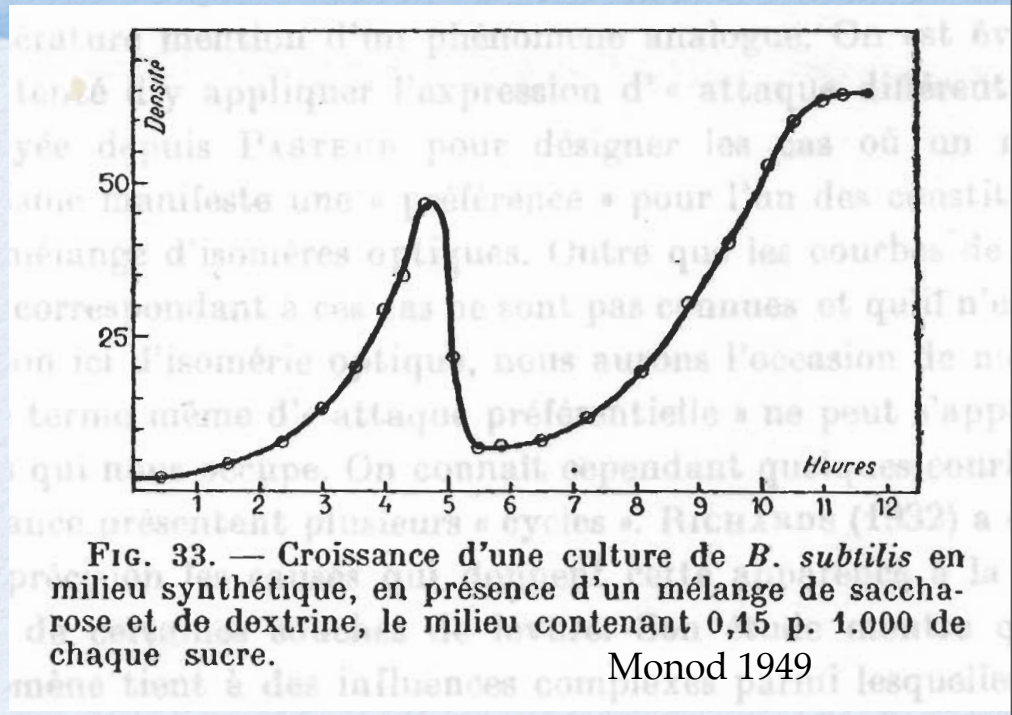

FIG. 33. — Croissance d'une culture de *B. subtilis* en milieu synthétique, en présence d'un mélange de saccharose et de dextrine, le milieu contenant 0,15 p. 1.000 de chaque sucre.

Monod 1949

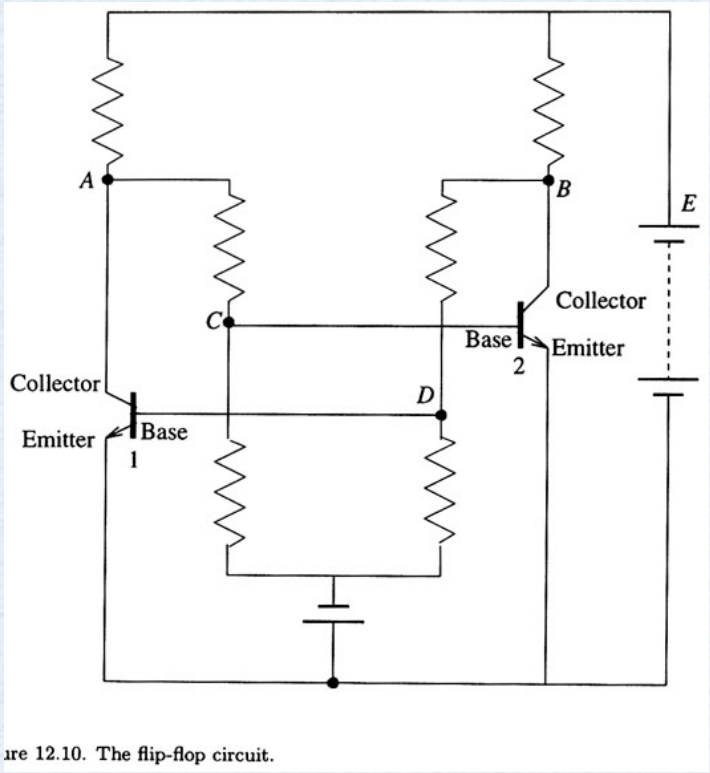

ıre 12.10. The flip-flop circuit.

Could bacteria somehow be implementing a ***two-state switch*** like the ones that changed human civilization in the mid-20th century? (Ahem -- Why doesn't this icon appear anywhere else in our Physics curriculum?)

Switching, II

Students can write a model of two mutually repressing genes, make the phase-plane analysis, and find the region of bistability in Matlab.

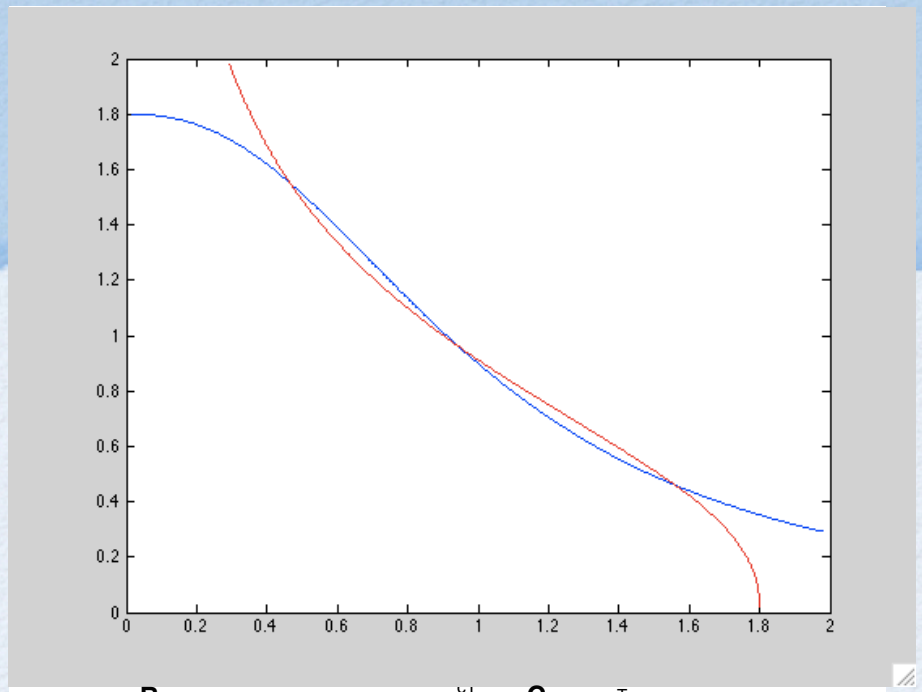

It's not speculation -- now the transfer functions of each element have been measured. The era of ***synthetic biology*** has arrived.

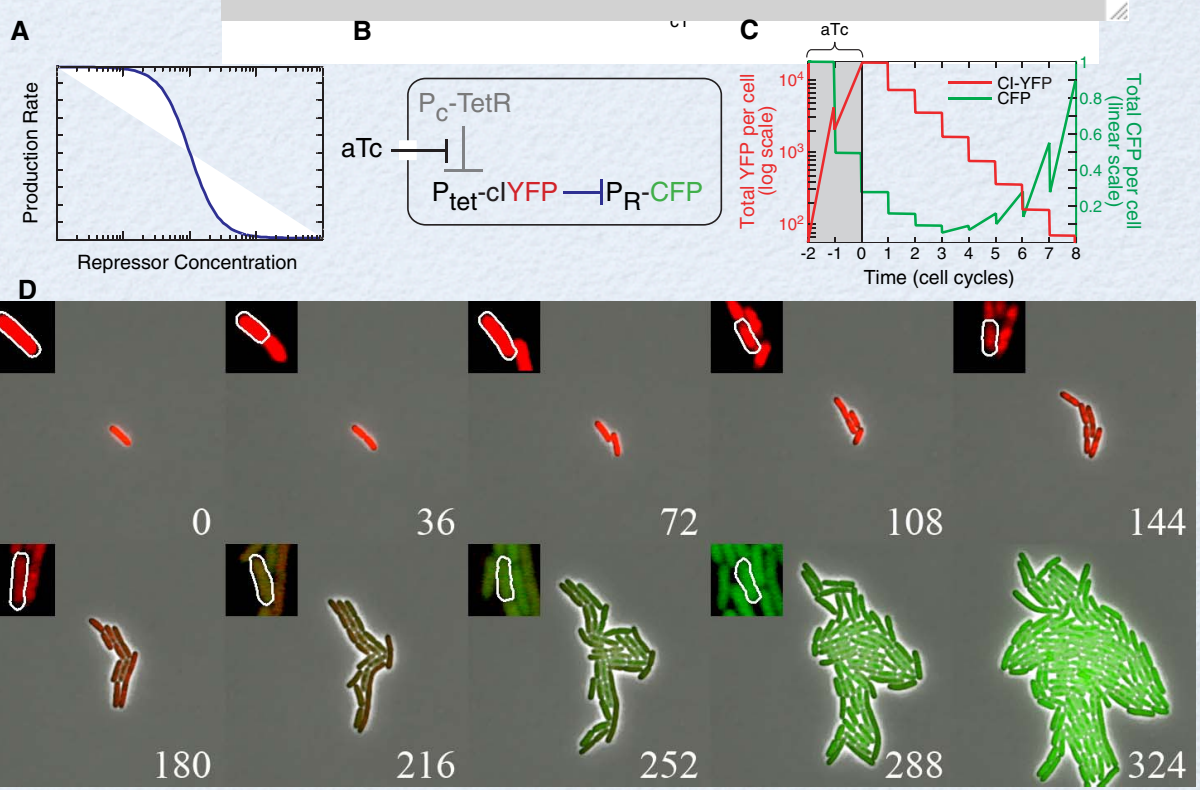

Figure S1: Snapshots of a typical regulator dilution experiment using the O_R2*-λ-cascade strain. *Panels show the same microcolony as Fig. 1D, with greater time-resolution. CI-YFP protein is shown in red and CFP is shown in green. Times, in minutes, are indicated on snapshots. Insets show a selected cell lineage (outlined in white).*

Rosenfeld et al 2005

Molecular machines

Myosin is a molecular motor that walks on actin filaments:

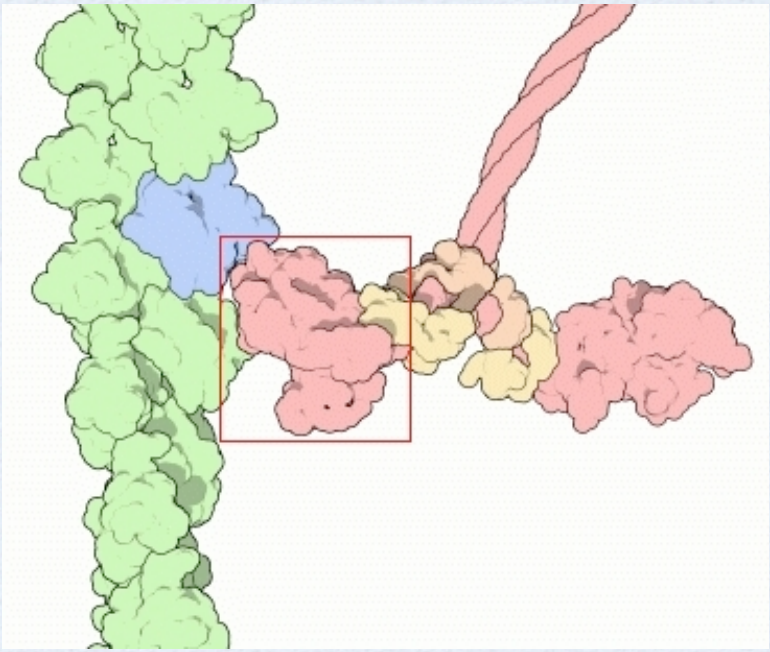

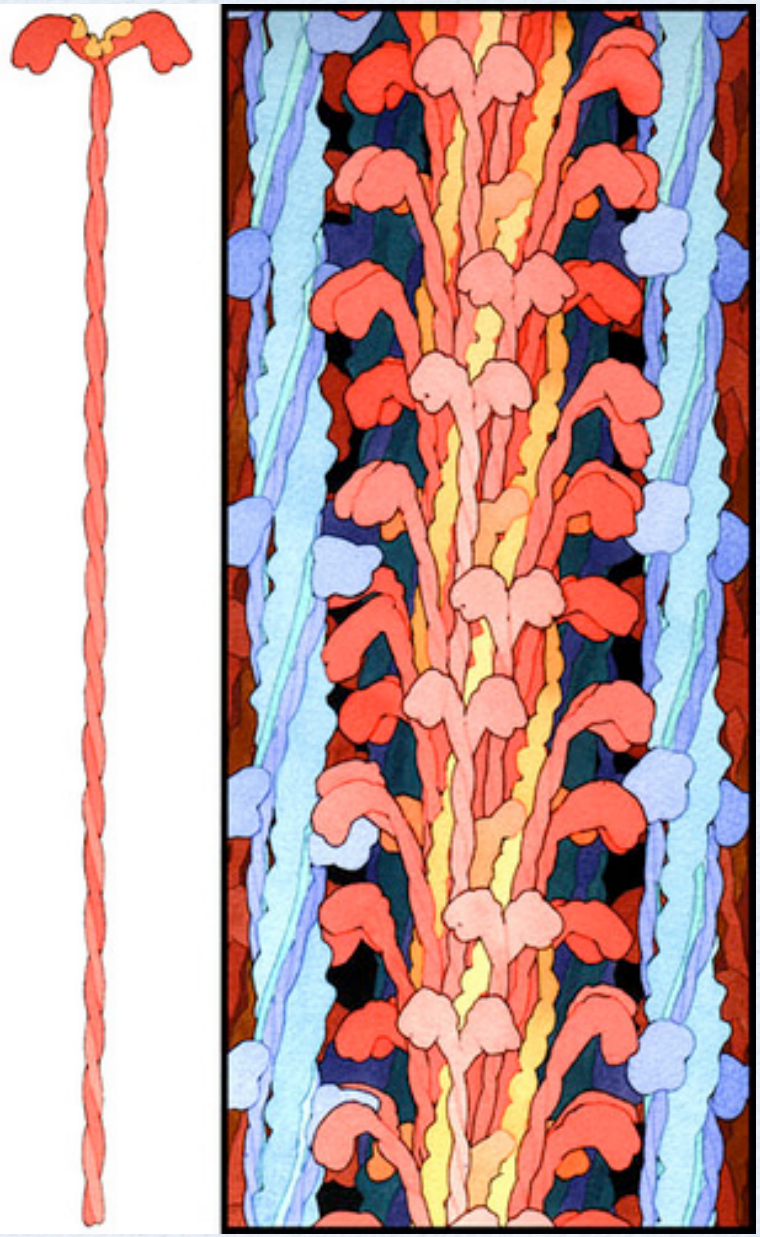

How can we get information about this invisibly small motor's mechanism?

Here is where a little probability goes a long way.

Students can test two hypotheses about the gait by analyzing the statistics of the motor's steps:

Images based on x-ray crystallography data by David Goodsell.

SCIENCE VOL 300 27 JUNE 2003

Myosin V Walks Hand-Over-Hand: Single Fluorophore Imaging with 1.5-nm Localization

Ahmet Yildiz,[1] Joseph N. Forkey,[3] Sean A. McKinney,[1,2] Taekjip Ha,[1,2] Yale E. Goldman,[3] Paul R. Selvin[1,2]*

Second, we can also detect the 0-nm steps indirectly via a kinetic analysis. If the step rate of one head is k_1 (A → B, Fig. 1) and the step rate of the other head is k_2 (B → A′, Fig. 1), then the dwell time probability distribution function for the A → B transition is $f(t) = k_1 e^{-k_1 t}$ and for the B → A′ transition is $g(t) = k_2 e^{-k_2 t}$. In the 42-33 and 52-23 cases, the total probability of dwell times is the sum of two exponentials: $P(t) = (k_1 e^{-k_1 t} + k_2 e^{-k_2 t})/2$. If each head has the same stepping rate, $P(t)$ reduces to a single exponential $P(t) = k_1 e^{-k_1 t}$. In the 74-nm case, the observable is A → A′, which is the convolution of two processes. If A → A′ takes t s, and A → B takes u s, then B → A′ will take $t - u$ s with $0 < u < t$. Integration over all possible values of u gives the convolution $P(t) =$

$$\int_0^t f(u)\cdot g(t-u)\,du = k_1 k_2 \cdot (e^{-k_1 t} - e^{-k_2 t})/(k_1 - k_2).$$

If the two rates are equal, then $P(t) = t k^2 e^{-kt}$. Note that $P(0) = 0$ and that $P(t)$ initially increases and then decreases for the 74-nm data if they contain hidden 0-nm steps, whereas the 42-33 or 52-23 data are expected to decay monotonically.

CONVOLUTION

Dwelling time distribution of the two stepping types of myosin X.
Red bars: Type 1; Blue bars: Type 2.
Assuming the same stepping rates of the two heads of myosin X.
These distributions tell a story, if you know how to convert a hypothesis to a predicted distribution (curves).

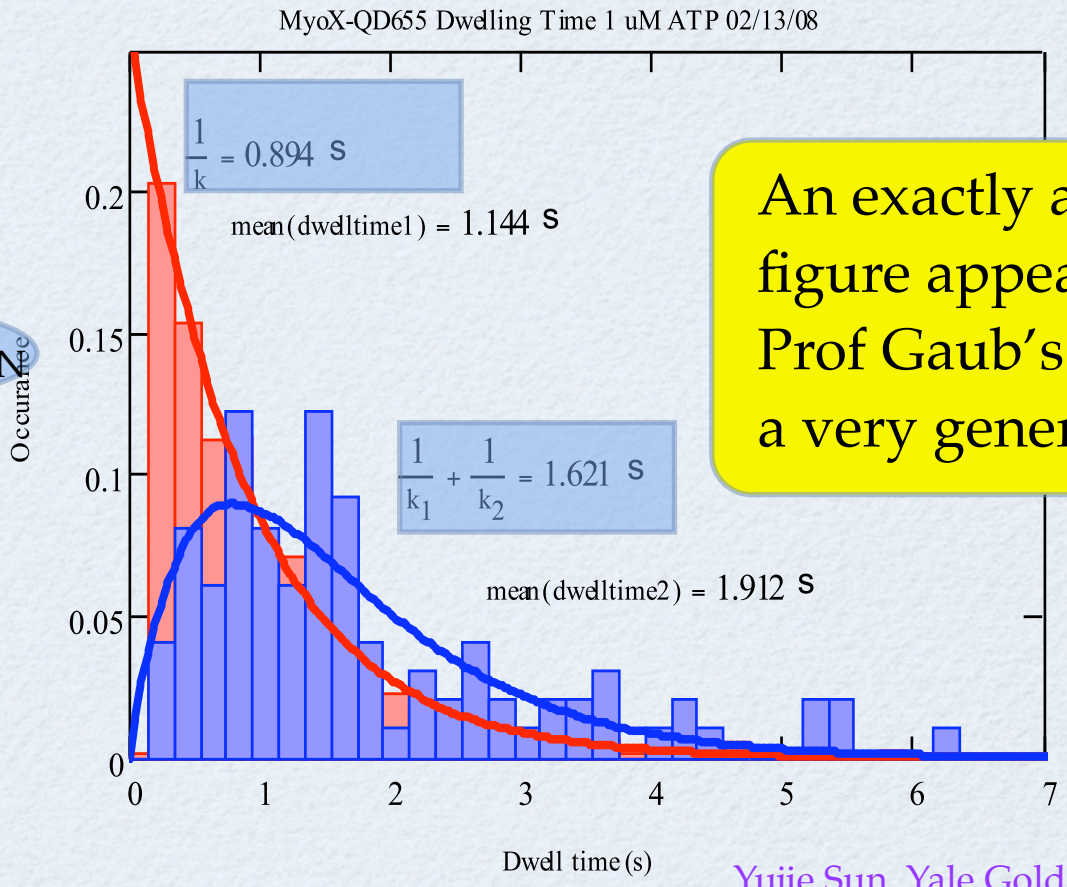

An exactly analogous figure appeared in Prof Gaub's talk--it's a very general idea.

Yujie Sun, Yale Goldman, et al.

[Subtext 1a]

Biophysical problems are an interesting road into probability theory with high-profile, current applications that can motivate students.

(Do your life-science students really understand it when they take their department's statistics course?)

Example: in the previous slide, how should we actually perform the fits to the expected distributions? Could jiggle till they look good... could hit the "fit" button on our canned sofware... or we could maximize the *likelihood*. A teachable moment.

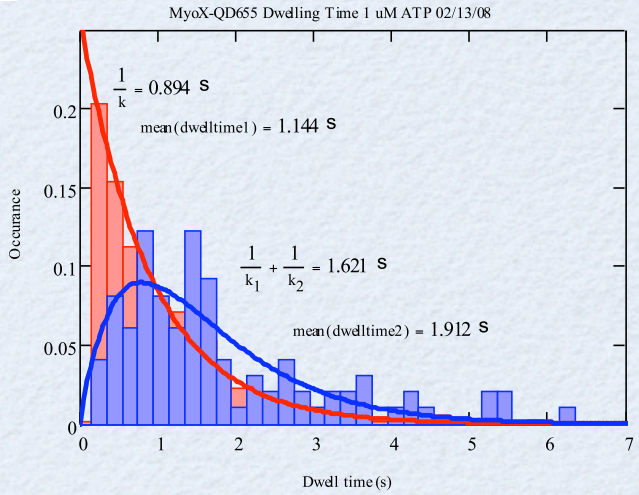

Bacterial genetics

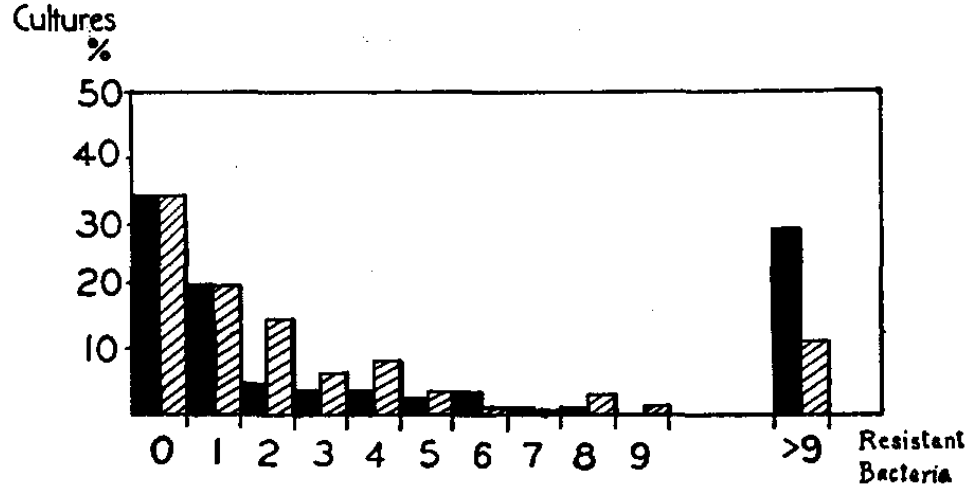

FIGURE 2.—Experimental (Experiment No. 23) and calculated distributions of the numbers of resistant bacteria in a series of similar cultures. Solid columns: experimental. Cross-hatched columns: calculated.

Luria & Delbruck, 1947

Luria and Delbruck noticed a statistical peculiarity in their data -- a huge "fat tail." They came up with a "Mendel, not Lamarck" model for drug resistance, and detailed quantitative predictions for such distributions that distinguished their model from the alternative.

They had to work very, very hard. But now it's trivial for students to simulate in Matlab.

Students can easily simulate the Luria--Delbruck model using Matlab:

Of 5000 simulated cultures, most have zero resistant mutants but a few have very many. This may be an ancient case history, but a lot of cutting-edge research is done on "fat-tail" distributions like this one. Biophysics is a good context for students to learn how to handle such things. In fact, many distributions arising in Biophysics have *infinite variance.*

Oh, also -- drug resistant bacteria are a very, very current problem.

And -- similar ideas are also illuminating when applied to retinoblastoma. ***A good physical model applies to problems beyond the one for which it was developed.***

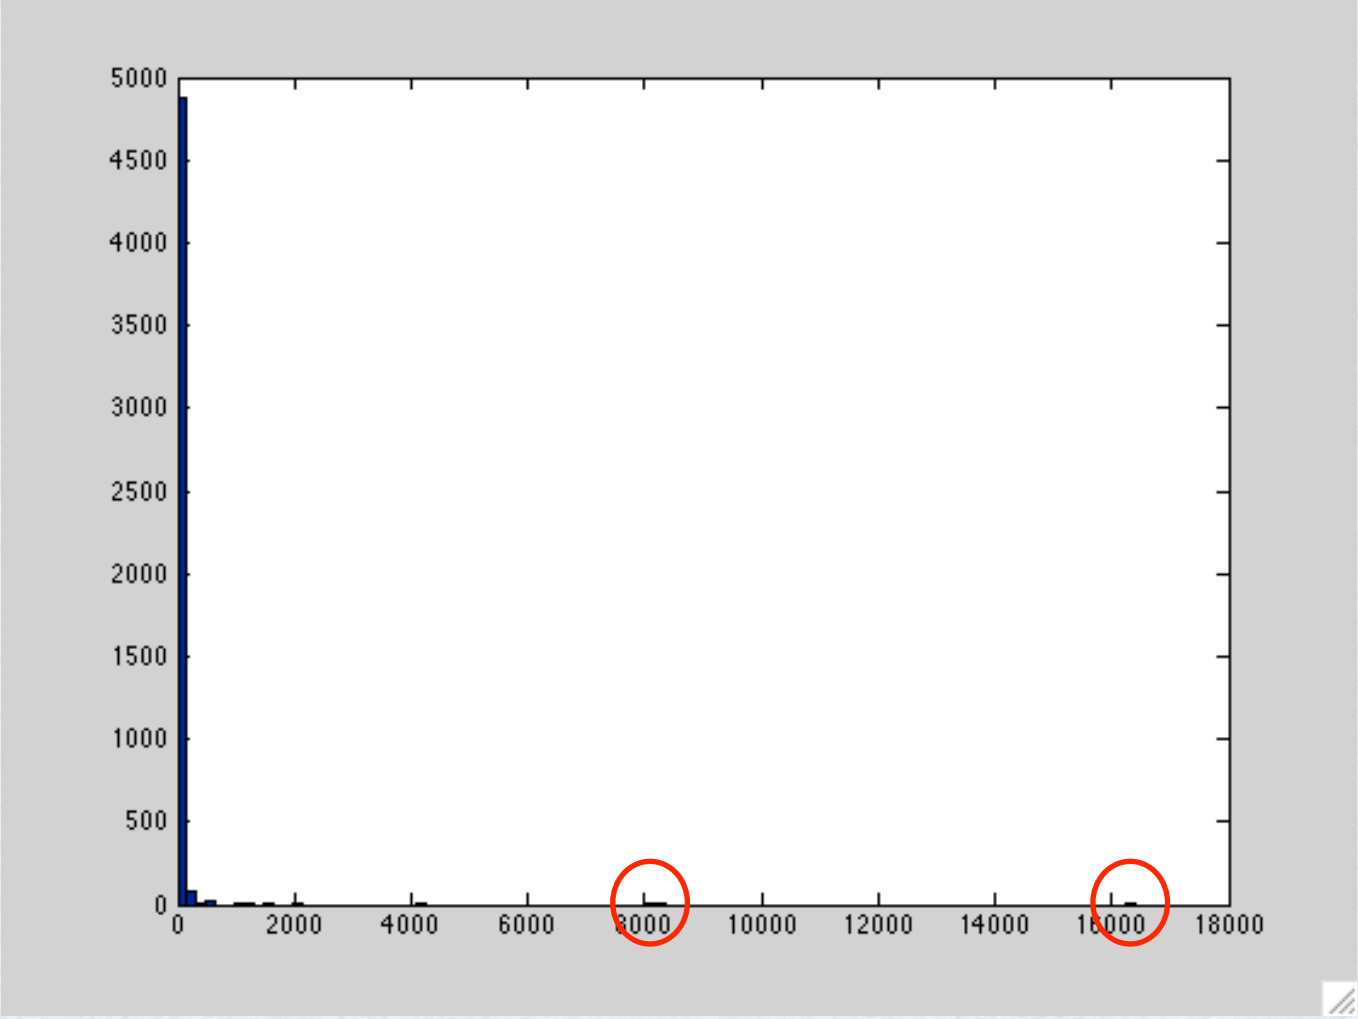

Random walks

Simple walk:

Sampled walk:

Real Brownian motion:

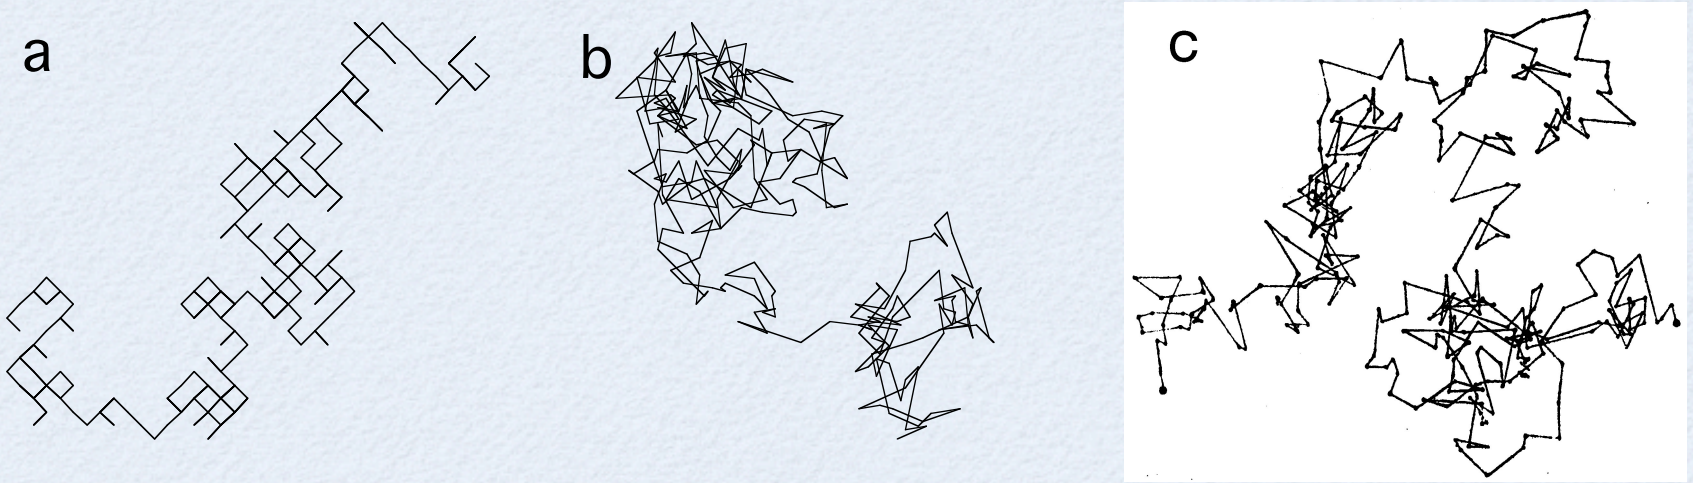

Students can simulate a random walk, then run it again and again to grasp the generic similarity of these figures despite the fact that they're always different in detail. Then they can find the mean-square displacement to confirm the diffusive law.

P Nelson, Biological Physics, updated ed 2008.

Diffusion

We can show that the diffusive flux of oxygen to a bacterium is limited by its size. But that derivation is a bit abstract. We don't "see" the oxygen molecules themselves.
And anyway, a key result (Berg and Purcell) requires solving an electrostatics problem that most of our grad students can't do!

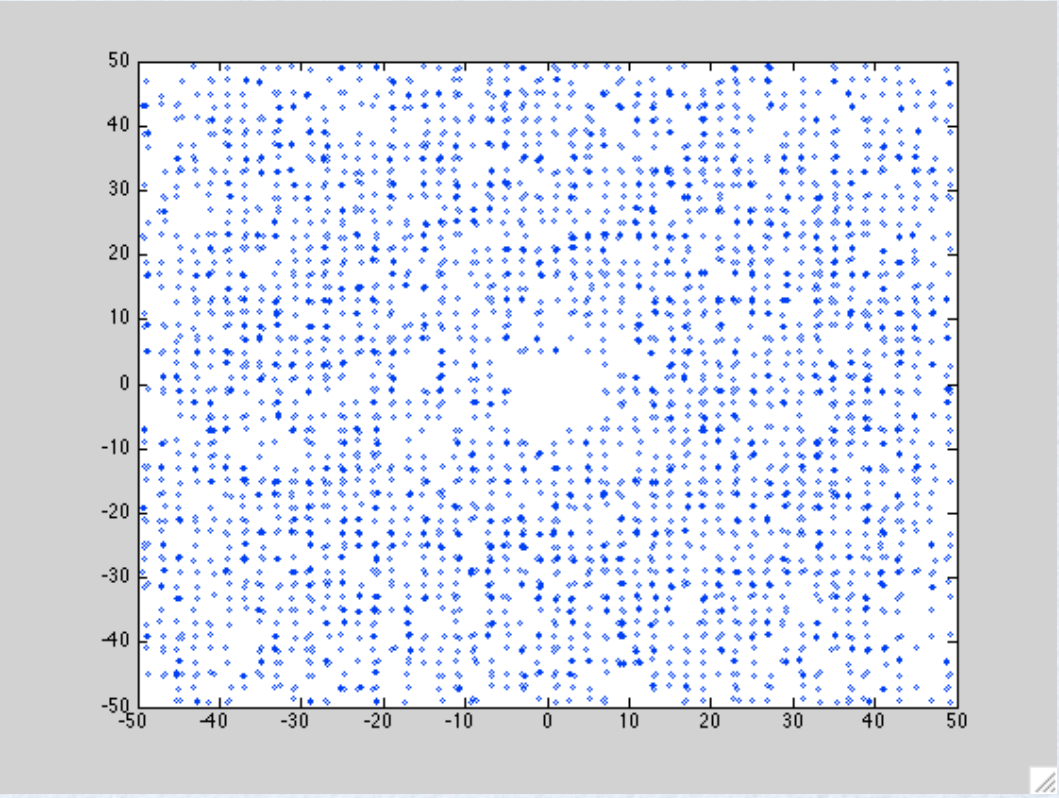

Instead, students can simulate random walks in Matlab and find the diffusive flux to an absorbing sphere. They can then find the flux to a reflecting sphere with absorbing patches -- a calculation with big implications for cell receptors.

Genetic drift

Genetic drift is also a random walk -- with non-constant "diffusion constant." Kimura had to work very, very hard to solve this model, but it led to the fundamental result that probability of fixation is proportional to 1/(population). It's incredibly important -- the basis of the molecular clock that gives us phylogenetic trees.

Students can trivially simulate this system, and obtain the key result, using Matlab.

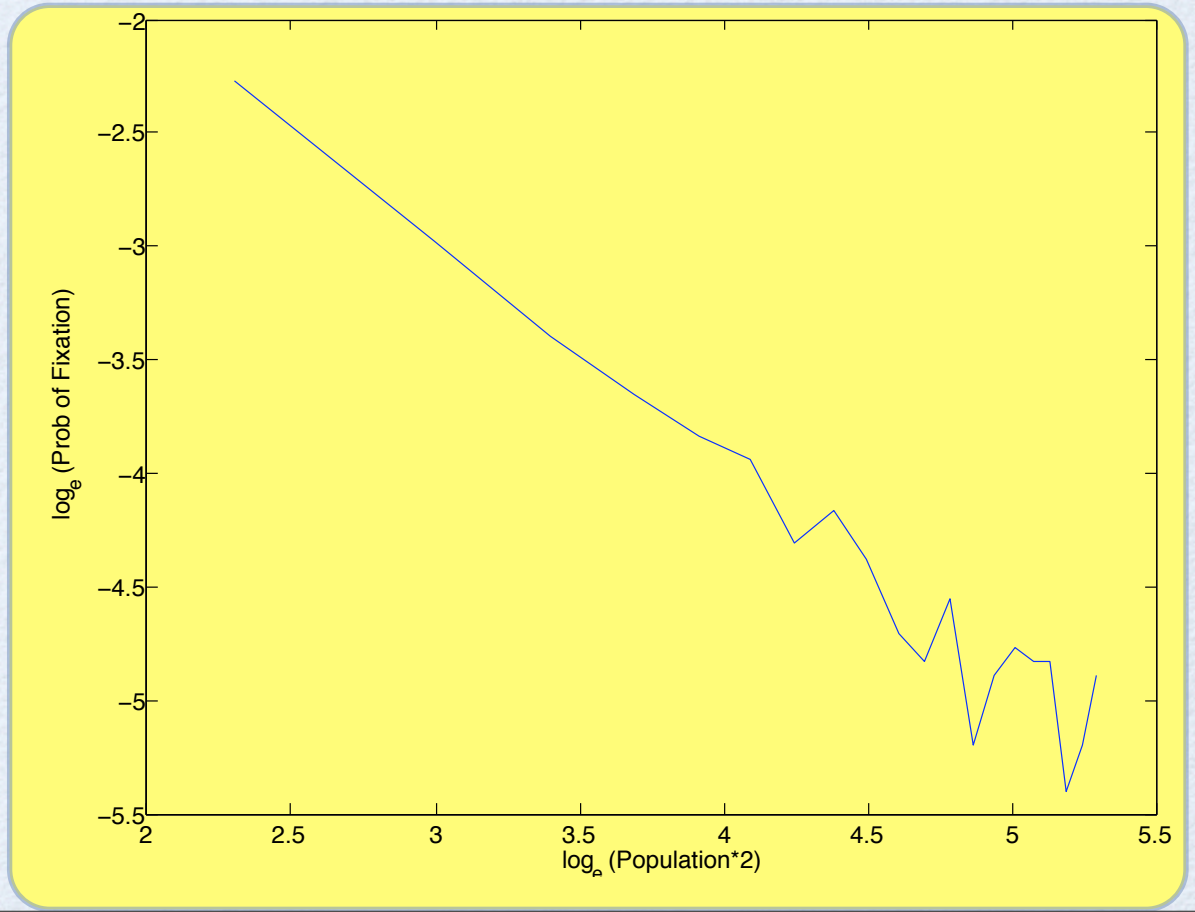

Proteomics

If you're looking at an unknown gene, and it resembles a lot of kinase genes, then maybe it's a kinase.
But resemblance can be hard to spot. Nonlocal sequence homologies can be important.
Random walks *on graphs* can tease out those nonlocal aspects.

Diffusion on graphs is also the basis of the ***Google PageRank algorithm.*** A civilization-changing discovery.

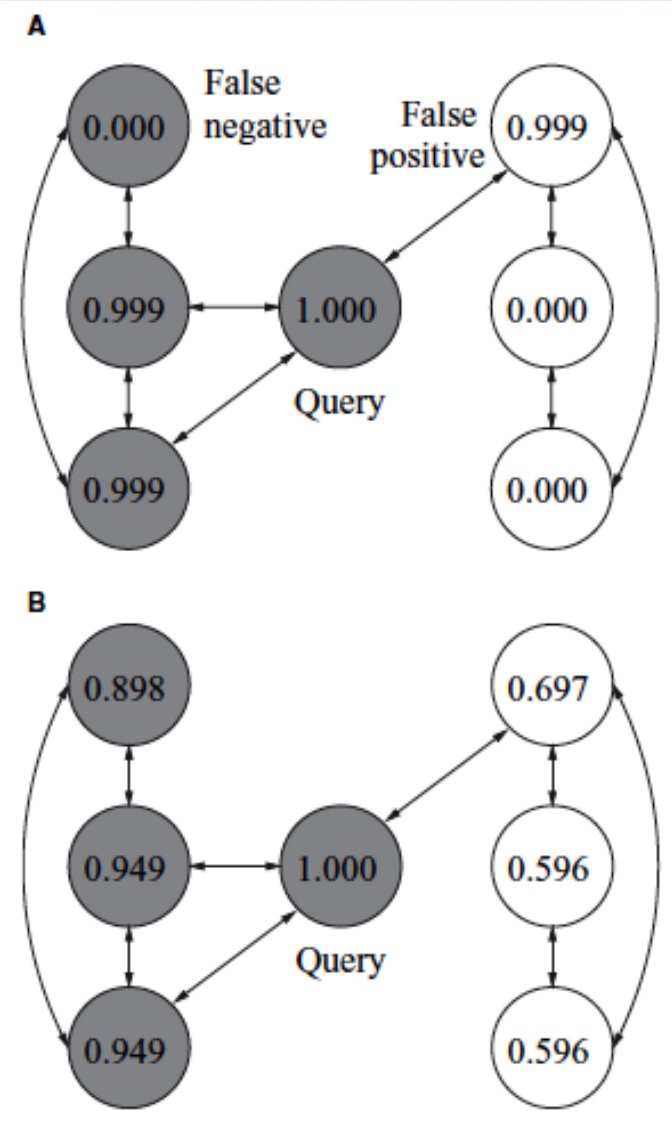

Noble et al FEBS 2005

[More subtext]

1b: Biophysical problems are an interesting road into statistical physics (including nonequilibrium) with a lot of high-profile, current applications that can motivate students.

(Are your life-science students really likely to take your Physics department's regular stat mech course?)

Some other subtexts are points that may strike us as so obvious that we need not speak them aloud -- but they're not so obvious to many life-science students, and we do need to say them:

2a: Somehow **a good physical model** (in this case random walks) **can apply to widely disparate problems.**

Eyes

Eyes are an ancient invention: Here is Trilobite, half a billion years old.

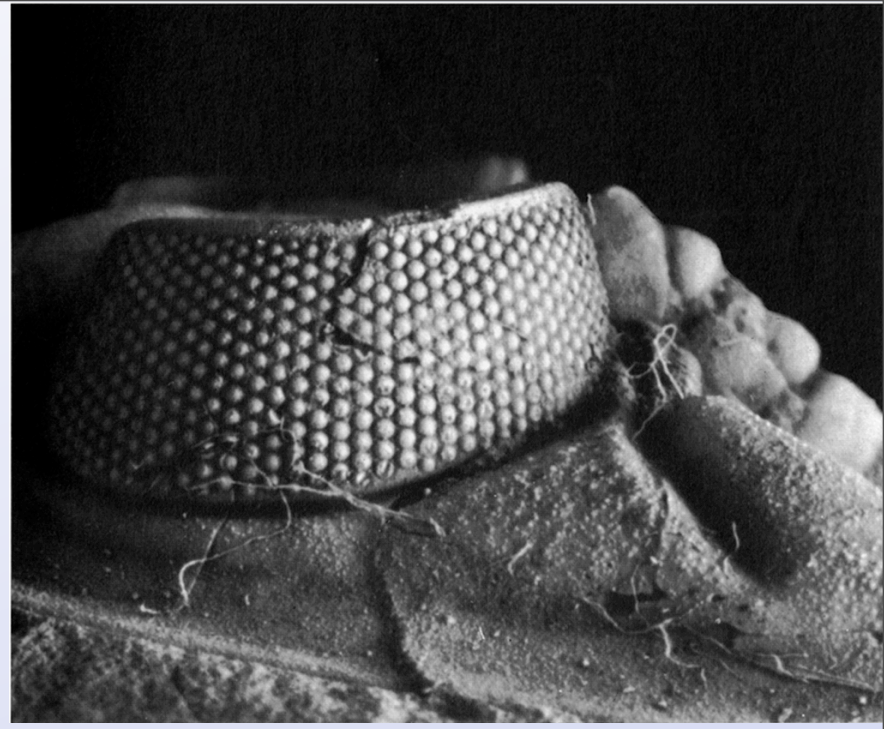

That design was successful: Here is a modern aphid.

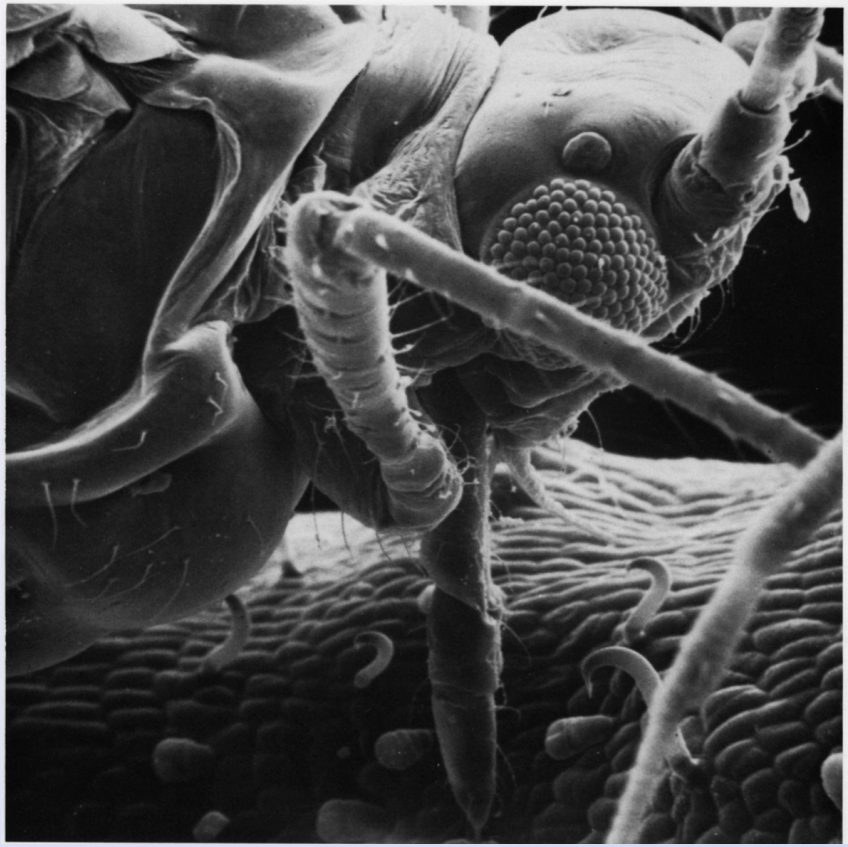

But if you can afford to carry more weight around, here is a better design:

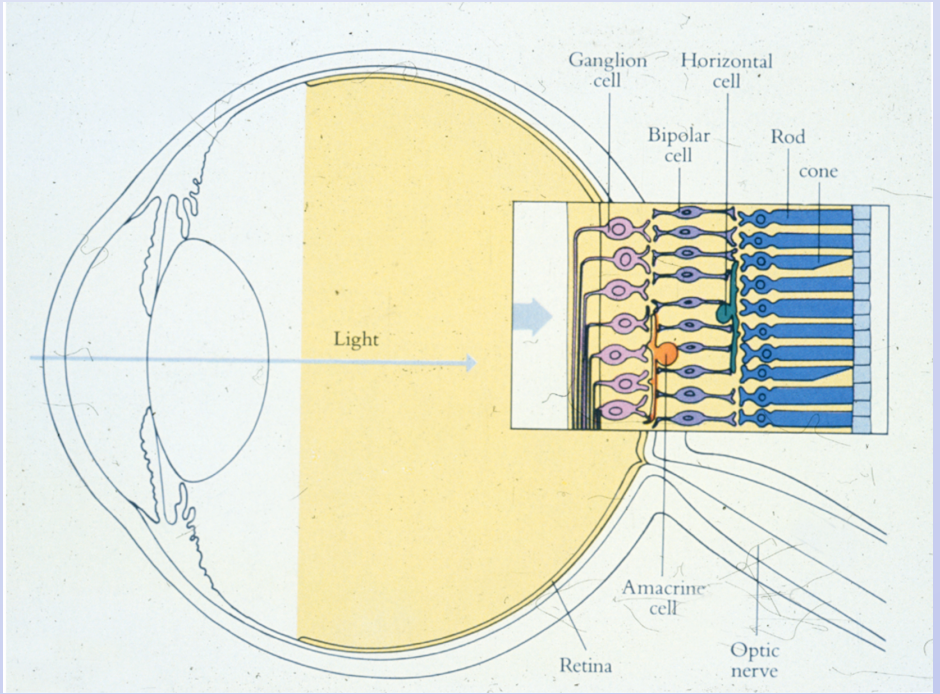

False sense of security

The human eye also has a lens-based focusing system, again like a camera. It seems to make sense in terms of Snell's Law, a consequence of the wave theory of light.

A key realization is that the finite size of our pupils limits the resolution we can get at the retina, due to diffraction. There's no point having a pixel size smaller than this resolution limit, and remarkably, *our photoreceptor cells really are about this size*.

Looks like the wave theory of light explains everything.

Uh-oh

But what happens *next?* What happens in those photoreceptor cells that translates light into nerve impulses?

We can detect very dim light with a photomultiplier tube, or a photodiode. Either way, light causes discrete clicks in the detector. *Dimmer light gives equally big clicks, just less frequent.* You might imagine a mechanism something like this:

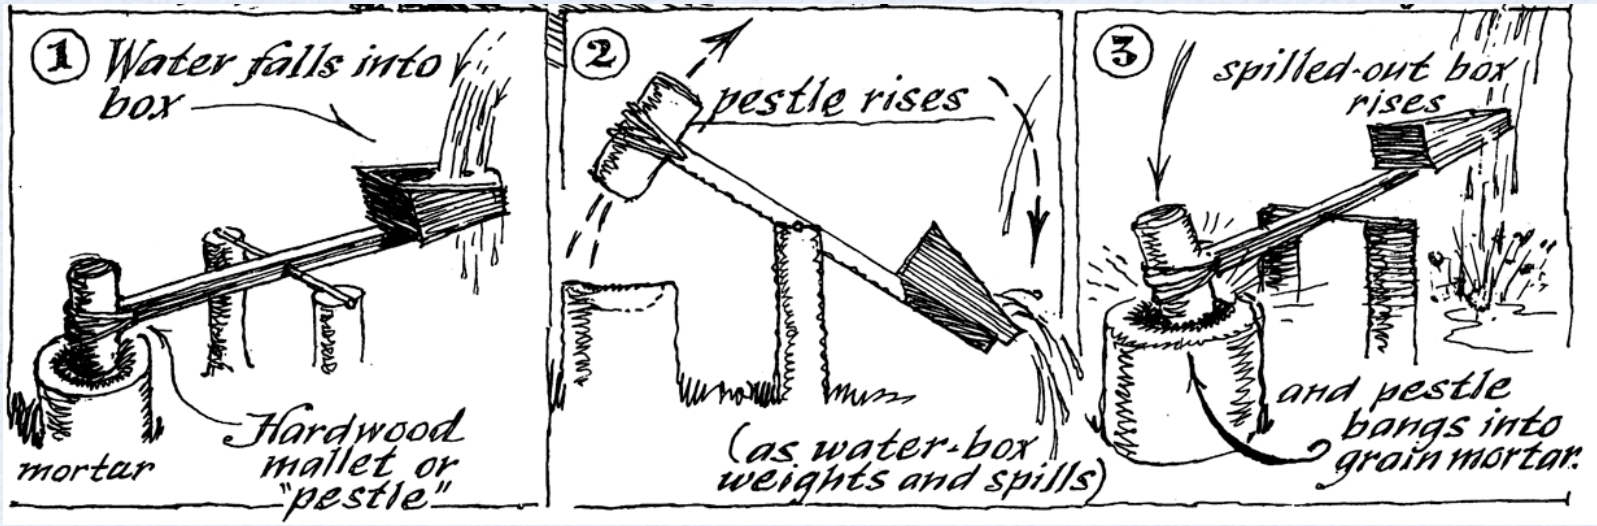

But that mechanism would give *uniformly spaced* clicks. Instead the clicks are *as random as possible* -- they are a "Poisson process." Something about light is intrinsically random.

click for uniform clicks audio

click for shot noise

Moreover, when we shine dim light on *several* photodetectors, they *never respond in unison*: Each click comes from just *one* detector, even if the beam of light is spread out to cover them all.

[Still more subtext]

2b: By the way, somehow **a good physical model** (in this case the Poisson process) **can apply to widely disparate problems.**

Both digital and film cameras also expose one pixel at a time, at random:

Even classic diffraction effects turned out to be particulate in character.

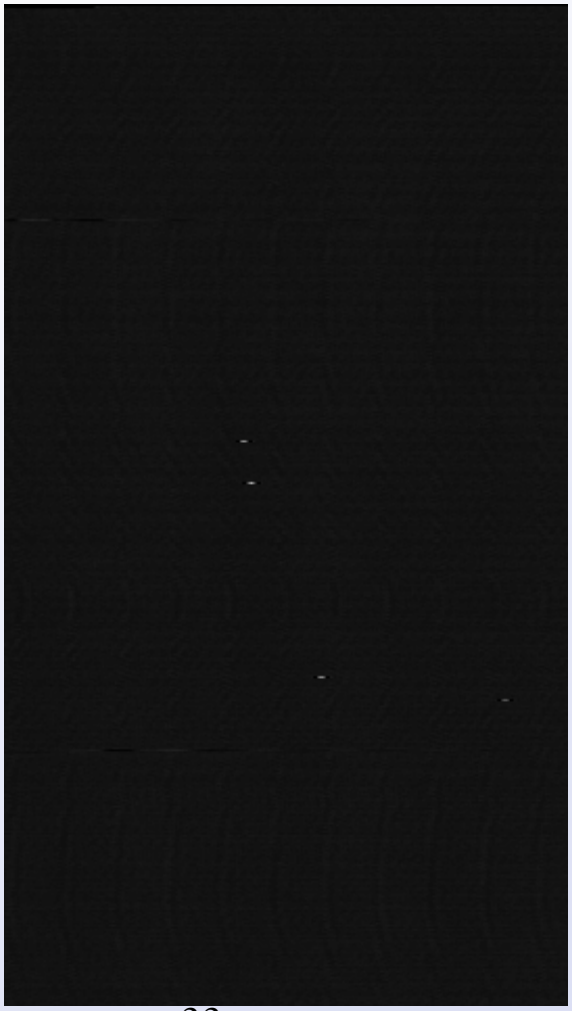

33 msec

1 second

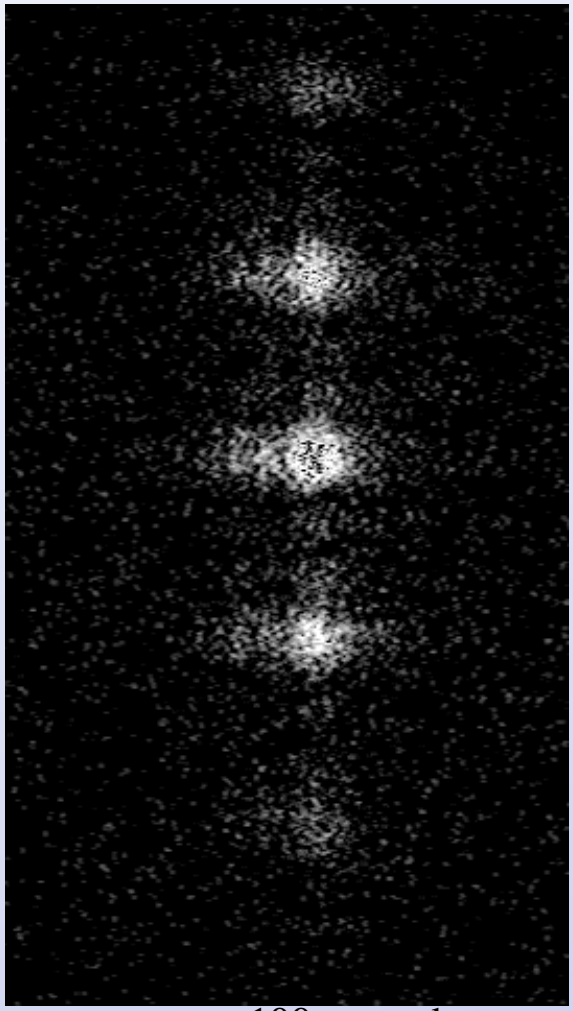

100 seconds

Einstein found he could only understand phenomena such as the photoelectric effect and thermal radiation by postulating that light consists of tiny *lump*s -- the "photon hypothesis."

[Aside -- how they see us]

All this way-out stuff about quanta sounds a good tale, but what's it to do with Life? Students could be forgiven the suspicion that, like Calvin, we are taking a fundamental idea and found a naive/irrelevant application for it.

But on the contrary: If your competitor/threatening predator/prey has better night vision than you, that could be a problem. That's the mother of invention, big time. If there's a fundamental limit to vision, you want to *be there*.

Human vision

"What's all that theoretical stuff got to do with *vision*, a real biological process? Surely vision is a terribly complex system, impossibly difficult to understand?"

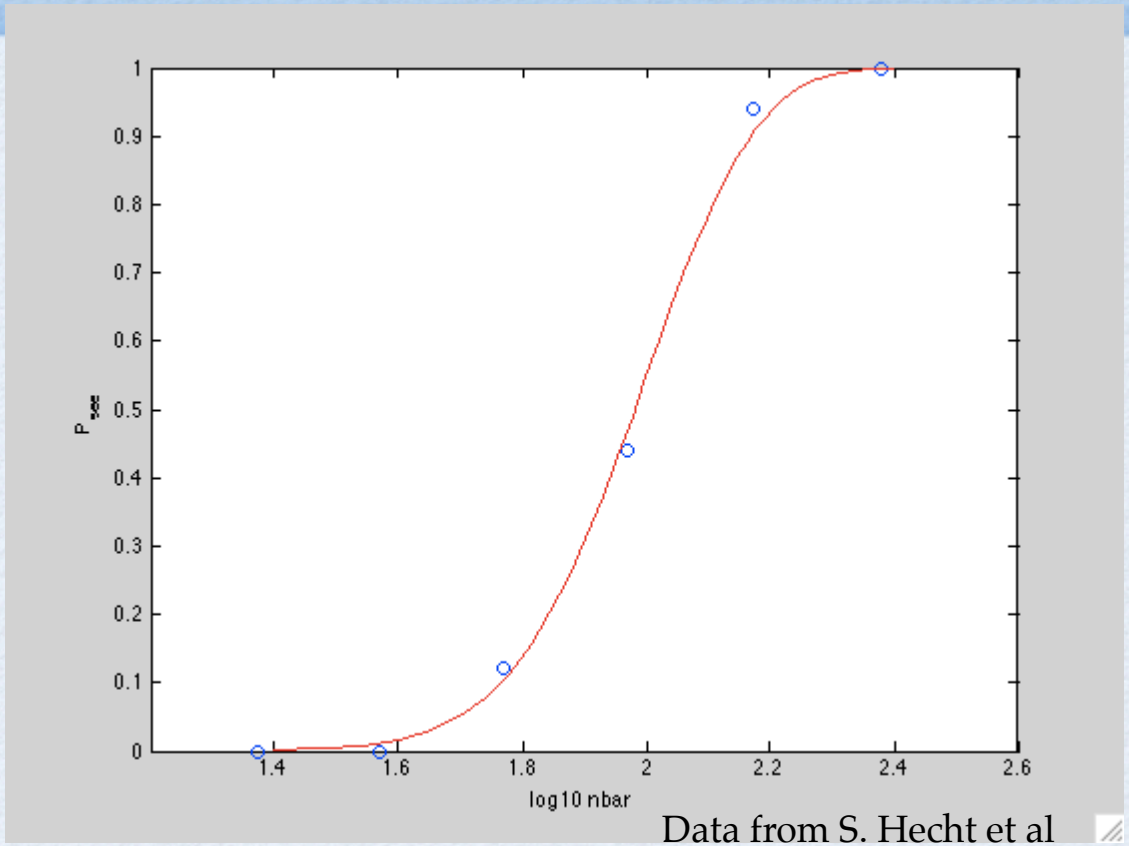

Data from S. Hecht et al

Hecht et al measured the probability of seeing a flash vs intensity. They found a simple **physical model** predicting the form of this "probability of seeing curve."
Then they were led from this information to the conclusion that ***single photons can excite rod cells***, and that a quorum of simultaneous rod-cell firings is registered consciously.

```
%% pcn 10/07  hecht.m
% hecht experiment
% data from hecht et al page 835 table V 4th column
p=[0,0,0.12,0.44,0.94,1]; nphoton=[23.5,37,59,93,149,239];lnphoton=log10(nphoton);
plot(lnphoton,p,'o');
xlabel('log10 nbar','FontSize',16);ylabel('P_{see}','FontSize',16);title('probability of
seeing','FontSize',16)
%%
figure
plot(lnphoton,p,'--');hold on
for mstar=1:2:12,
   q=0.12;
for j=1:46,
   photons=20+5*j;
   mbar=q*photons;
   total=0;
   for i=mstar:50
      total=total+exp(-mbar)*(mbar^i)/factorial(i);end
   ptheory(j)=total;lphotontheory(j)=log10(photons);end
plot(lphotontheory,ptheory,'r');
end
xlabel('log10 nbar','FontSize',16);ylabel('P_{see}','FontSize',16);title('probability of seeing,
q=.12, various mstar','FontSize',16)
%%
figure; plot(lnphoton,p,'--');hold on
for q=.04:.01:.10,
   mstar=7;
for j=1:46,
   photons=20+5*j;
   mbar=q*photons;
   total=0;
   for i=mstar:50
      total=total+exp(-mbar)*(mbar^i)/factorial(i);end
   ptheory(j)=total;lphotontheory(j)=log10(photons);end
plot(lphotontheory,ptheory,'r')
end
xlabel('log10 nbar','FontSize',16);ylabel('P_{see}','FontSize',16);title('probability of seeing,
mstar=7, various q','FontSize',16)
```

Students can fit Hecht et al.'s data in Matlab and find the critical quorum size.

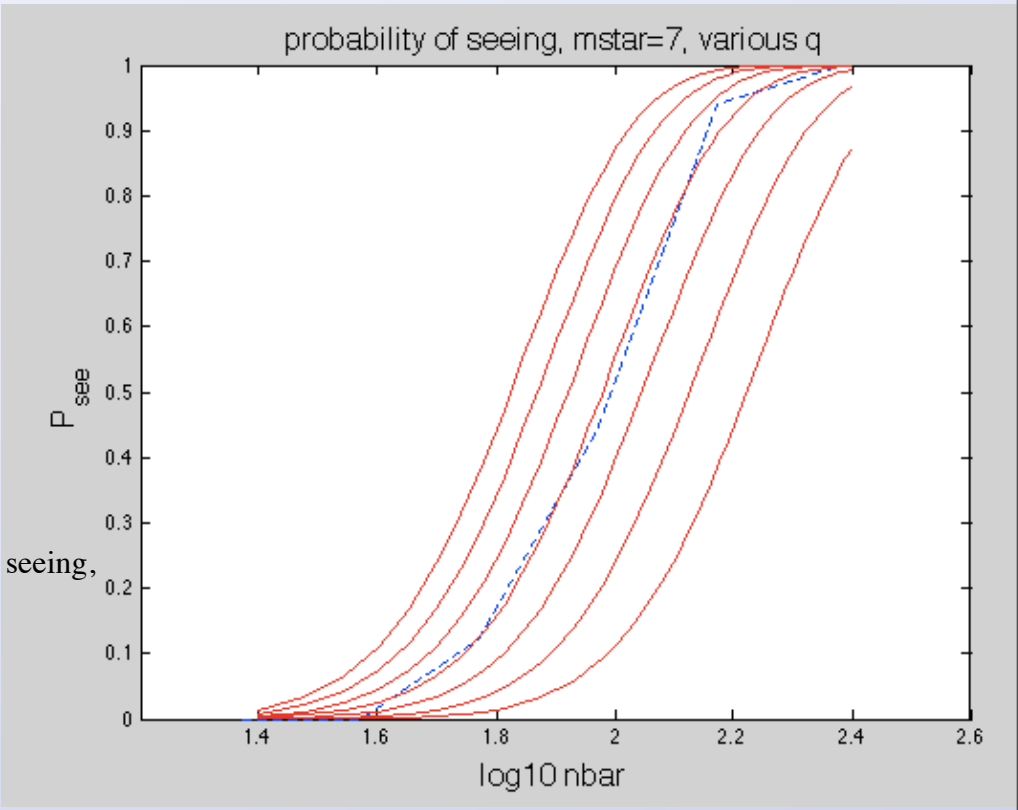

Above: threshold=7 fits the data well.

[Subtext 3]

Sometimes you can roll up your sleeves, scribble down some symbols or key in some code, and **extract new knowledge** from an experiment that wasn't going to divulge that knowledge to your unaided eye. Those scribbles, or that code, often embody some **physical model**.

Modern methods

Images D. Baylor; F. Rieke and D. Baylor

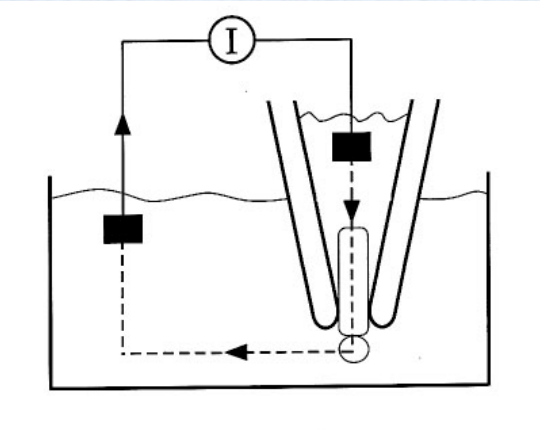

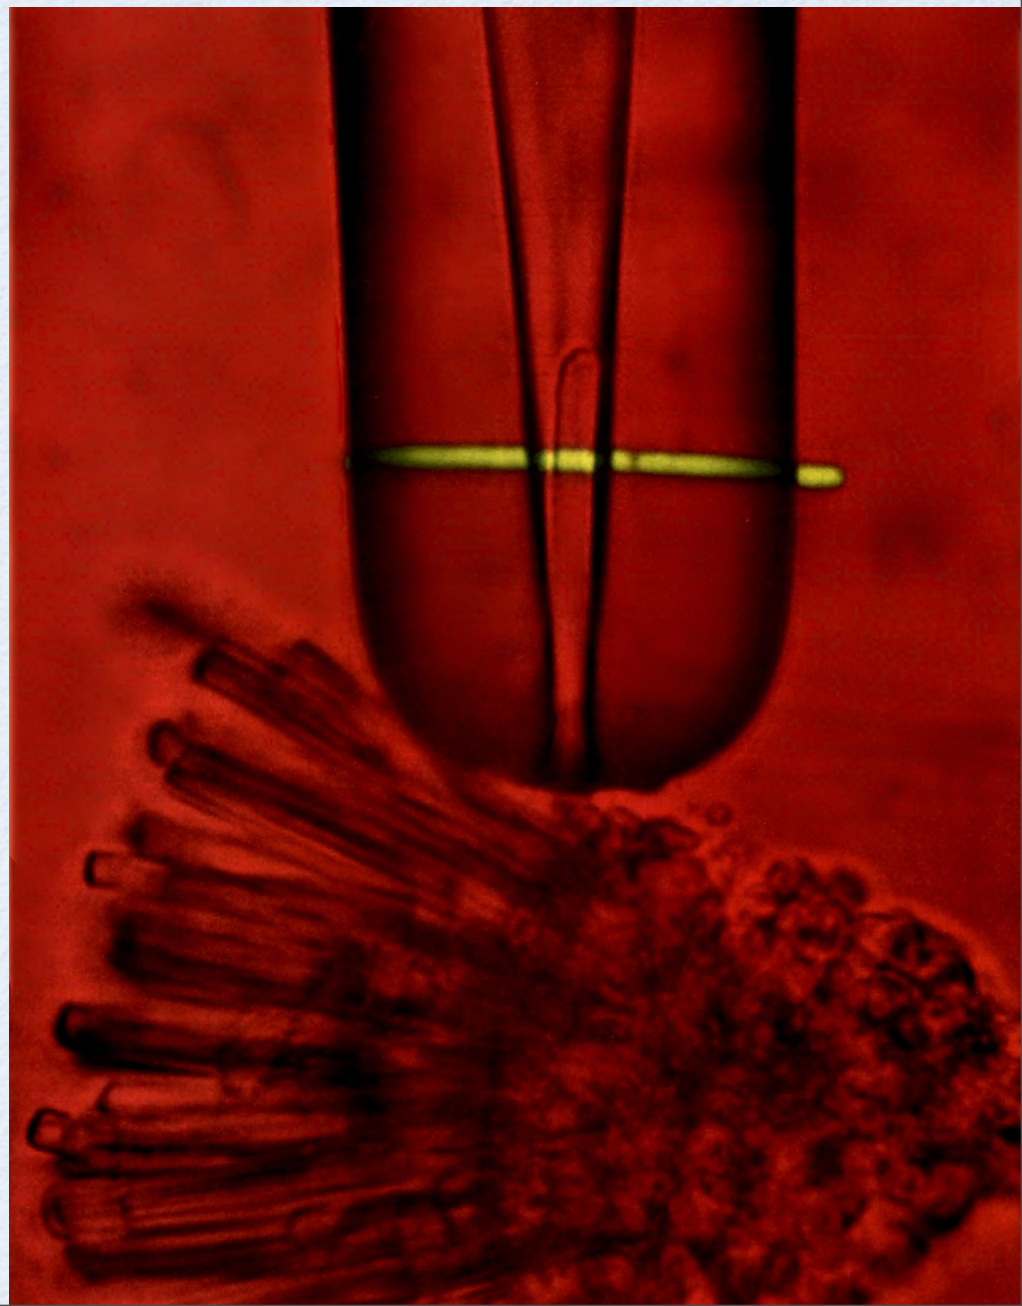

Direct measurements on single rod cells confirm that they can respond to single photons, and confirm the inherent randomness of the response.

An individual rod or cone cell's response can be measured by gently sucking its outer segment into a pipette electrode and stimulating it with 500 nm light (green).

Scale: outer diameter of pipette about 6 micrometers.

Intolerable

"OK, so light comes in tiny lumps. Is that all?"

And yet, I mentioned earlier that light *also* shows many other properties long thought to be slam-dunk evidence of *wavelike* behavior, much of it critically important for the design of visual organs. How could any of that ***possibly happen at all*** in the particle picture? Einstein didn't know.

Now, in physics we often put a box around a set of issues and say, "We can't understand that today," and move on. But this is an ***intolerable contradiction*****.** It's too big to put a box around it. We have to understand it before we have any business moving on.

Generally professors say, "You're not ready for that. You'll understand that some day."

(They really mean, "Shut up.")

Is that really an adequate response? Students would have to wait till they were halfway through a PhD in high-energy particle physics (which they're not going to do anyway) before we'd get around to telling them.

Can't we tell them something we actually believe is *true*? Can't we have them do a calculation *for themselves* that illuminates this apparent paradox?

Get serious

Again: How *can* little bullets display the diffraction and refraction needed to explain physiology (and much more)?

Photon hypothesis, continued:

- The probability to observe a photon is the length-squared of a certain complex number ψ.
- If there are multiple routes (or processes) by which a photon could make the trip, and we don't directly observe which one was taken, then they *all* make additive contributions to the total amplitude. These contributions are all complex numbers with equal *lengths*, but different *angles*, so they may reinforce or cancel.
- The angle (phase) of any one contribution equals the angular frequency of the light (related to its color) times the transit time for the path. In vacuum, light travels at the fixed speed c, so transit time may be written as (path length)/c.
- In a transparent material, it gets complicated by all those electrons, but for some purposes it's a good approximation to say that in water, etc. the speed is reduced to c/n, where n is the "index of refraction."

That's it.

See Feynman, *QED*

So how does that help?

It's incredible, but with that additional info we can reproduce all of the familiar classical optics results, including focusing by the lens of the eye. And you don't need to trust some authority figure on that -- ***students can do the calculations for themselves.***

The key question is, ***Why does light (usually) (seem to) go on (pretty) straight lines?***
After all, a penny held in the sunlight casts a sharp shadow. And yet our photon hypothesis says that photons take *all possible paths* between source and detector!

To answer that question, students can approximate the integrals in Matlab as sums, drawing little arrows to represent each term in the sum. The full integral (red arrow) is the vector from one end of the chain to the other end, times dx.

$$\int_{-5}^{5} e^{ix^2}\,dx$$

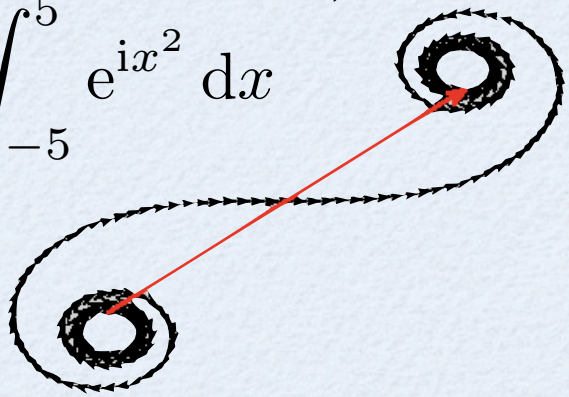

$$\int_{-1}^{9} e^{ix^2}\,dx$$

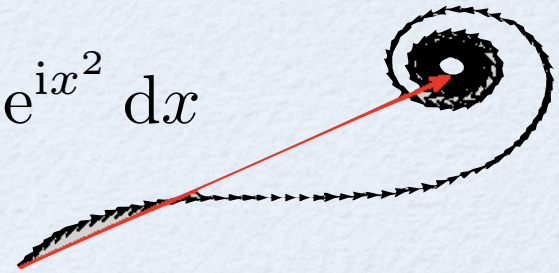

A similar integral whose range of integration contains *no* stationary-phase point will have a small total value:

Again: Near $x = 0$ the arrows point mainly to the right and they add up to something significant. If the range of integration contains that stationary-phase point, we'll get a large total (long red arrow). Otherwise, we won't.

$$\int_{1}^{11} e^{ix^2}\,dx$$

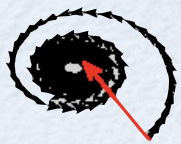

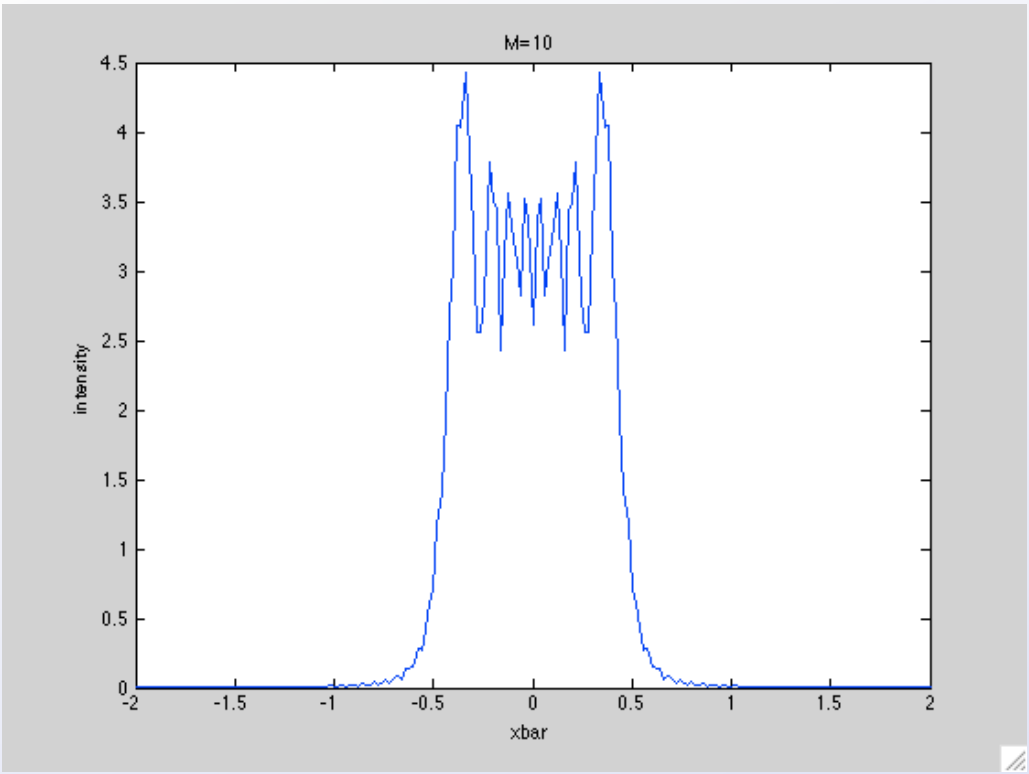

The concept we found -- "stationary phase" -- explains why we sometimes get sharp edges, other times not. Students can approximate the integrals in Matlab, then see the phenomena themselves with a laser.

Left: wide slit, sharp edges.
Lower left: medium slit, medium edges.
Lower right: narrow slit, fuzzy edges.

That gives us the "diffraction limit" on resolution (how well we can see) of an optical system. You can now understand other optical phenomena (like focusing) with similar principles.

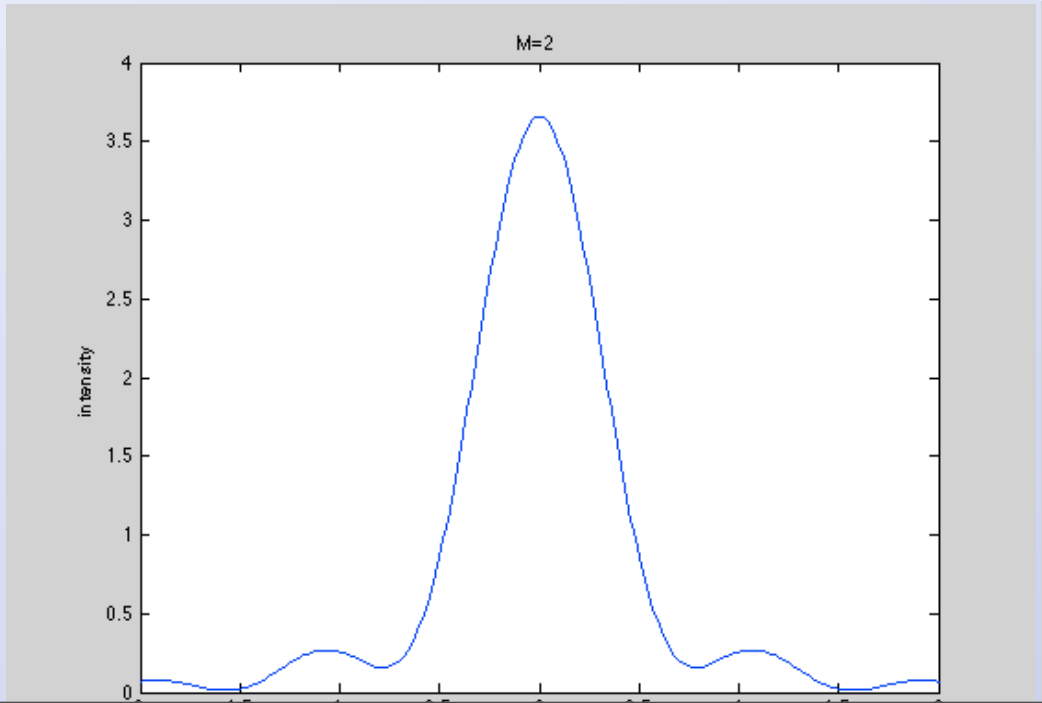

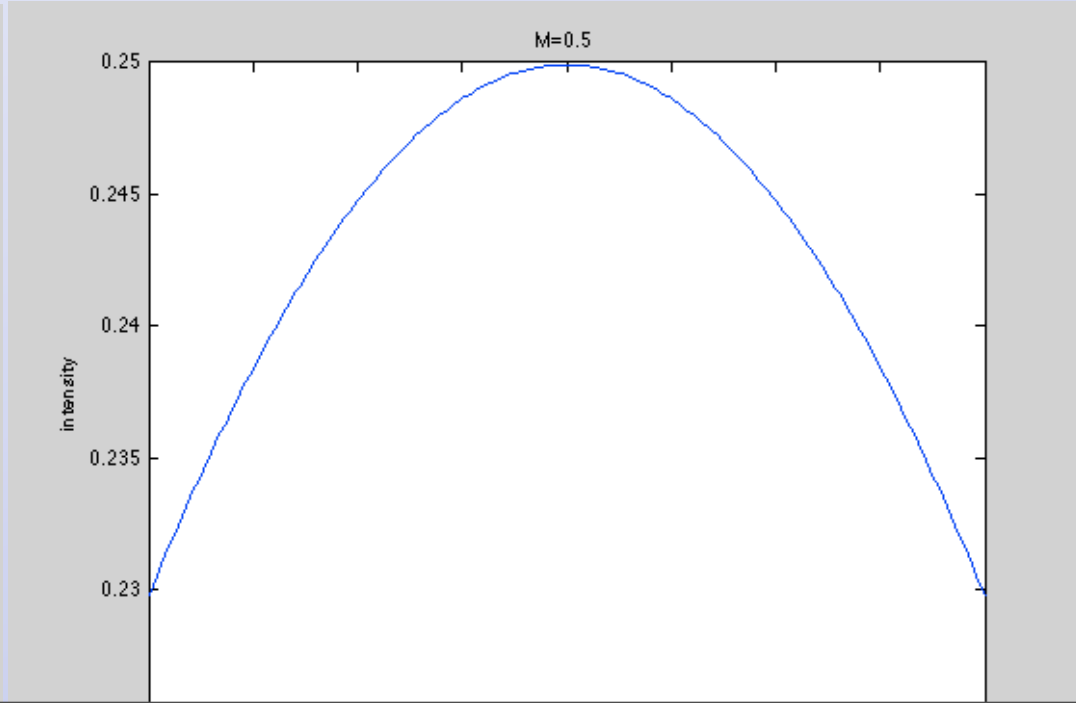

[Sorry to be so repetitive]

Subtext 1c: **Biophysical problems are an interesting road into quantum physics with a lot of high-profile, current applications that can motivate students.**

(Are your life-science students really likely to sit your department's regular quantum mechanics course?)

2c: By the way, somehow **a good physical model** (in this case the sum over trajectories) **can apply to widely disparate problems.**

"OK, OK... but"

I guess I believe that this photon hypothesis is possible. So OK, we learned about the true character of light, we reconciled its wave and particle aspects... **Is that *all*?** Is this just wiggling out of some paradox, or can we learn something *new* about biophysics?"

Yes. Somehow or other, the reception of a single photon gets converted to a neural impulse. By the time that impulse reaches the optic nerve, it has become a *spike*.

Spiking neurons, e.g. in the optic nerve, have limited dynamic range. They represent their signal by the times of individual spikes, each of which is exactly like any other. There are upper and lower bounds on this rate. The eye must make the best possible use of this limited range of representing the intensity of light at each pixel.

And the problem is even more acute than that: you have 10^8 photoreceptors, but only 10^6 optic nerve fibers!

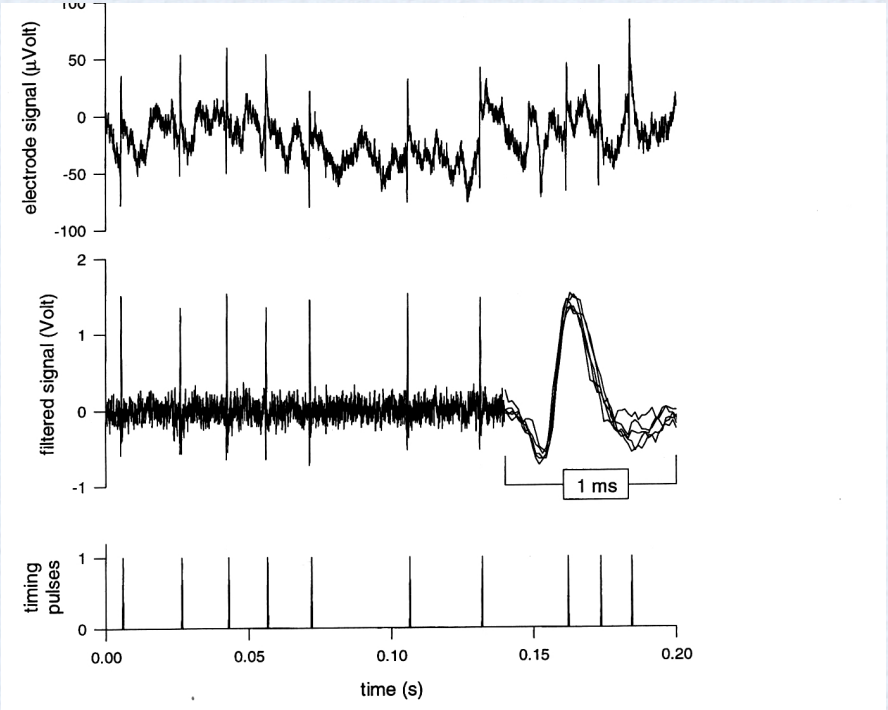

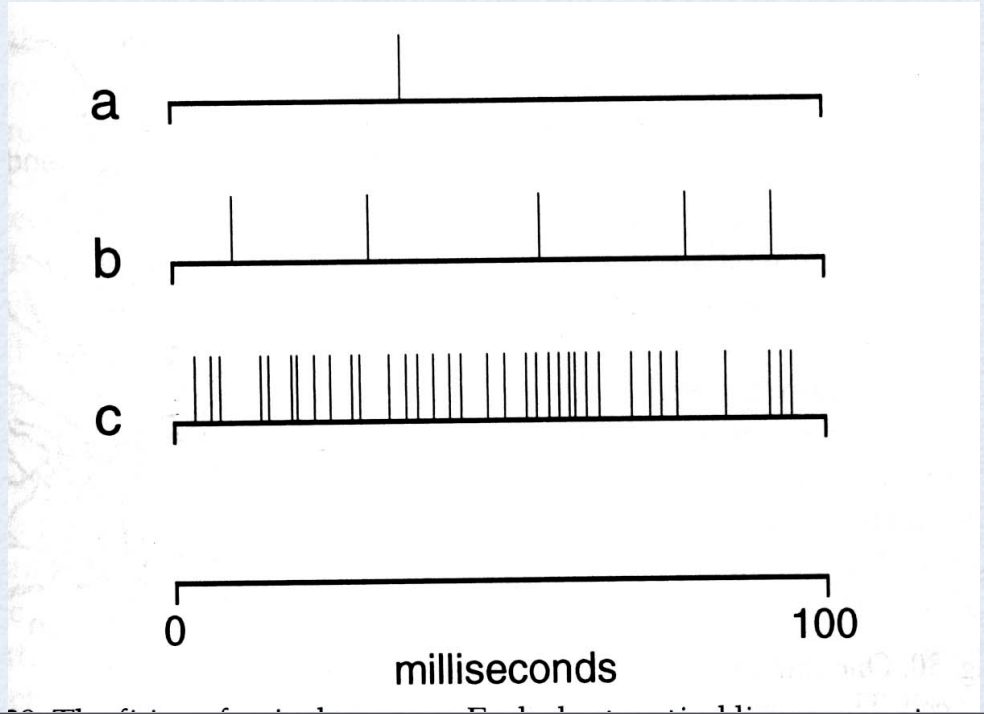

At high illumination, truncation to 1-bit depth destroys a lot of detail (top right). But we can do much better than this, at the same high level of compression.
To do so, students can apply a filter (bottom). The filter is called "**center/surround**"; it is the difference of two concentric Gaussians. And by the way, it's a convolution -- an idea we met in a very different context in the start of the course.

Truncating in this way vastly reduces the image size, or the bandwidth needed to transmit it in a given amount of time. (In this case we went from 8 bits per pixel to just one.)

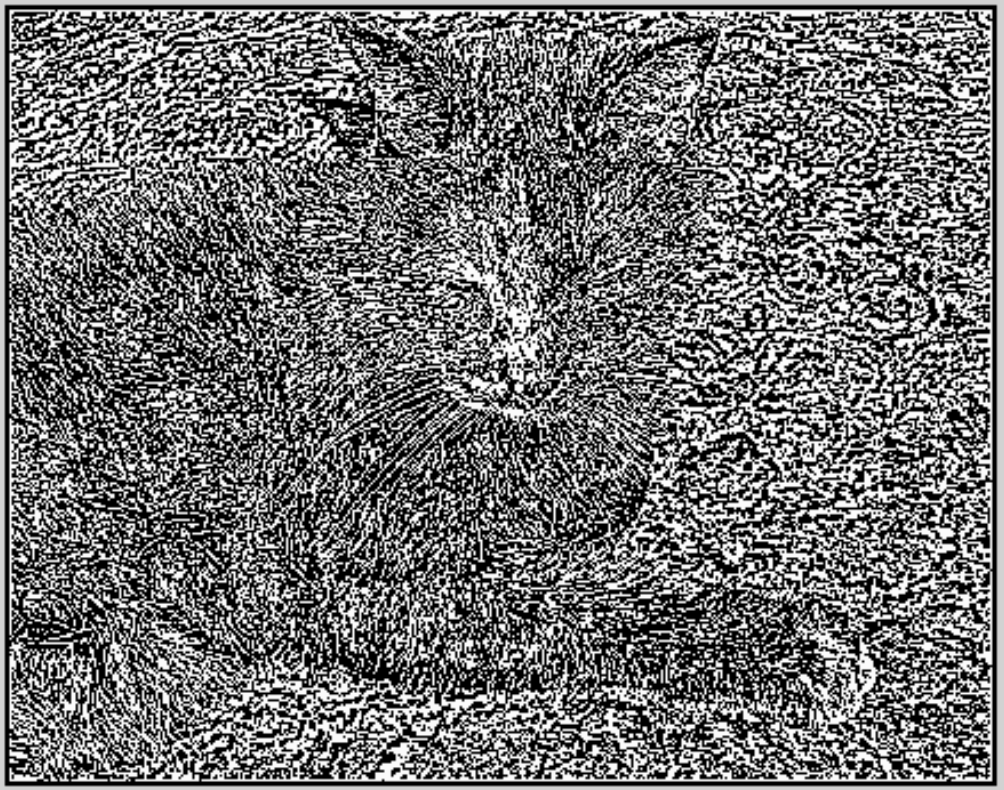

"Well, cool, but... ***Do our eyes really do that?"*** First look at psychophysical clues.

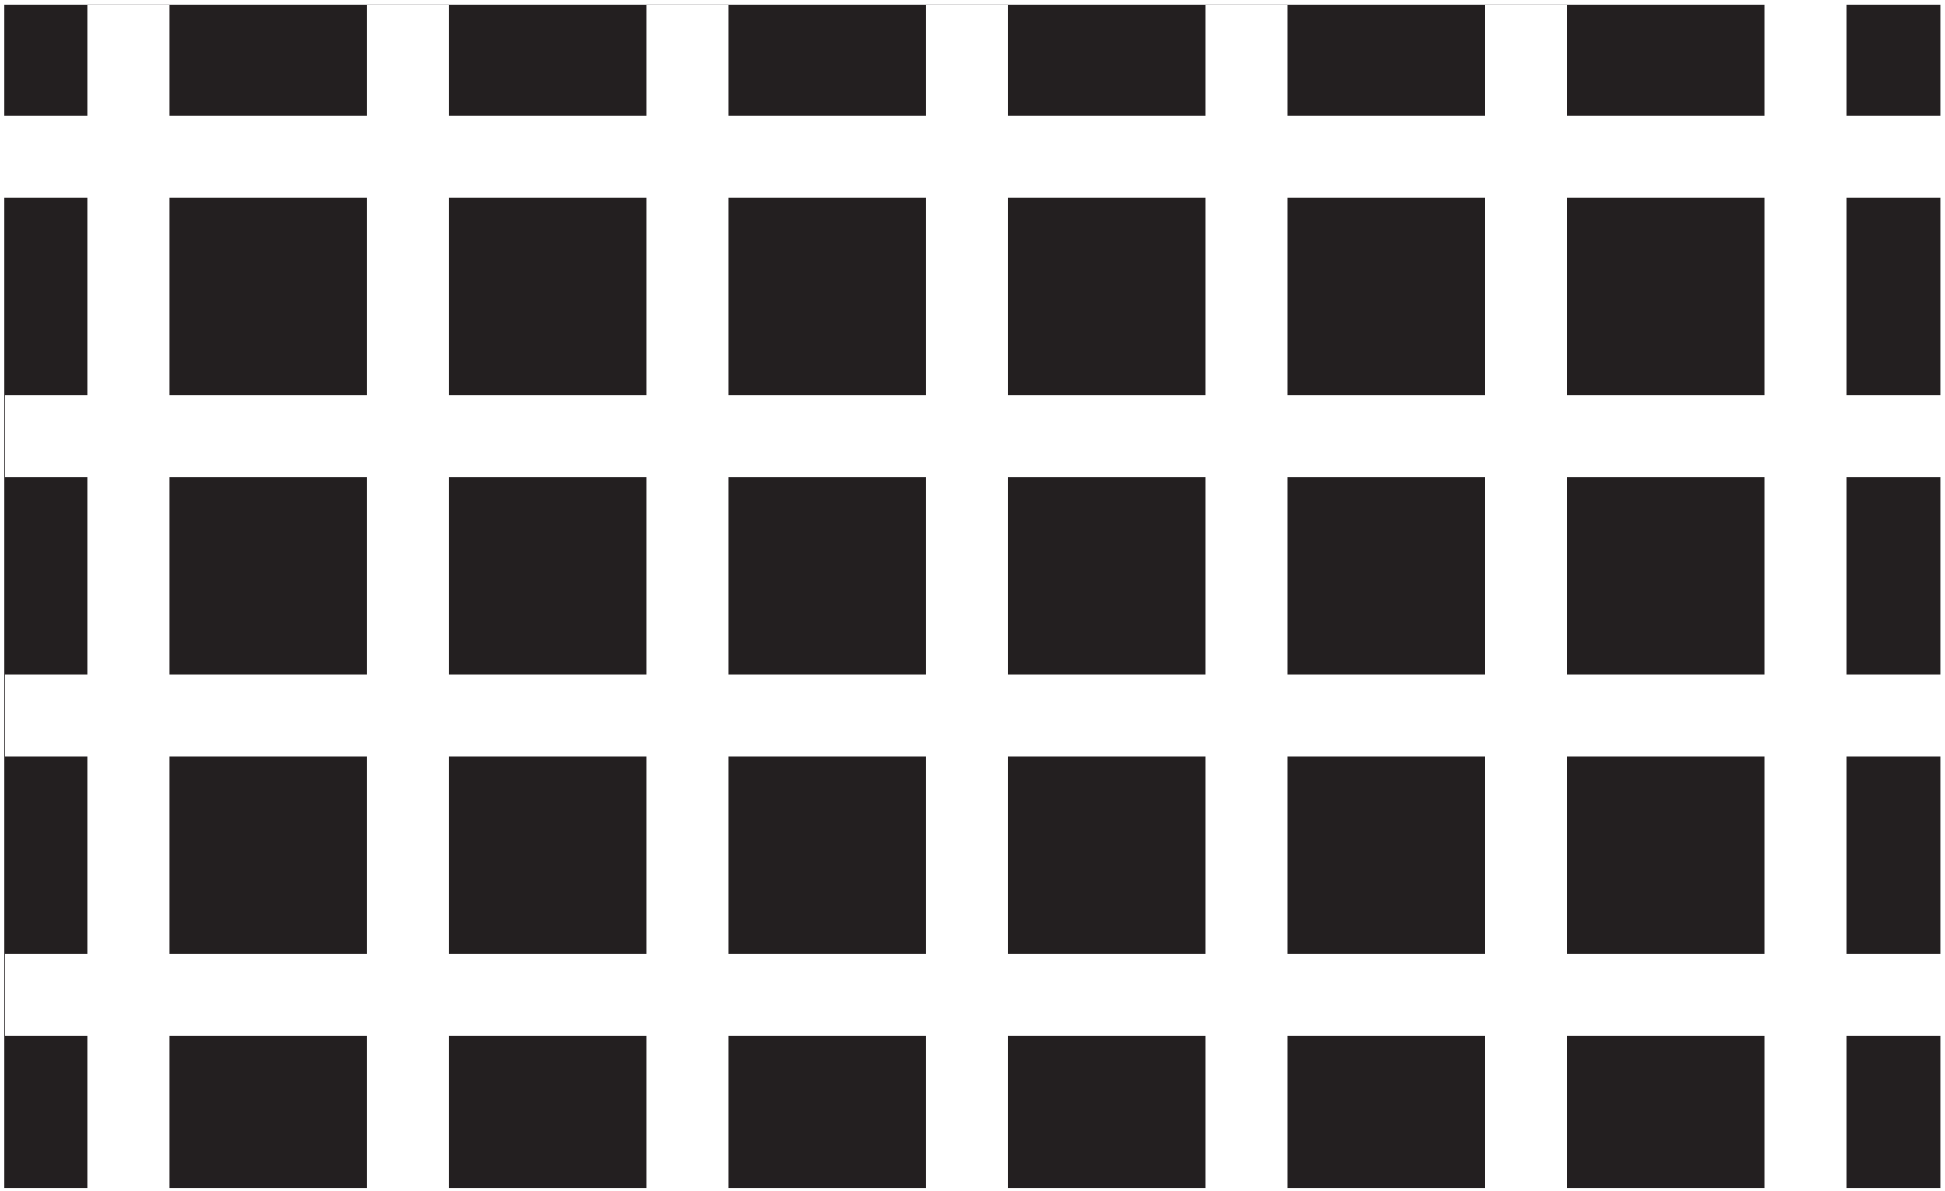

If you fixate on one intersection, other intersections appear gray, darker than the streets.

"Oh, please...

... Spare us the psychophysics. Is there some *objective* data on this?"

Meet *Limulus:*

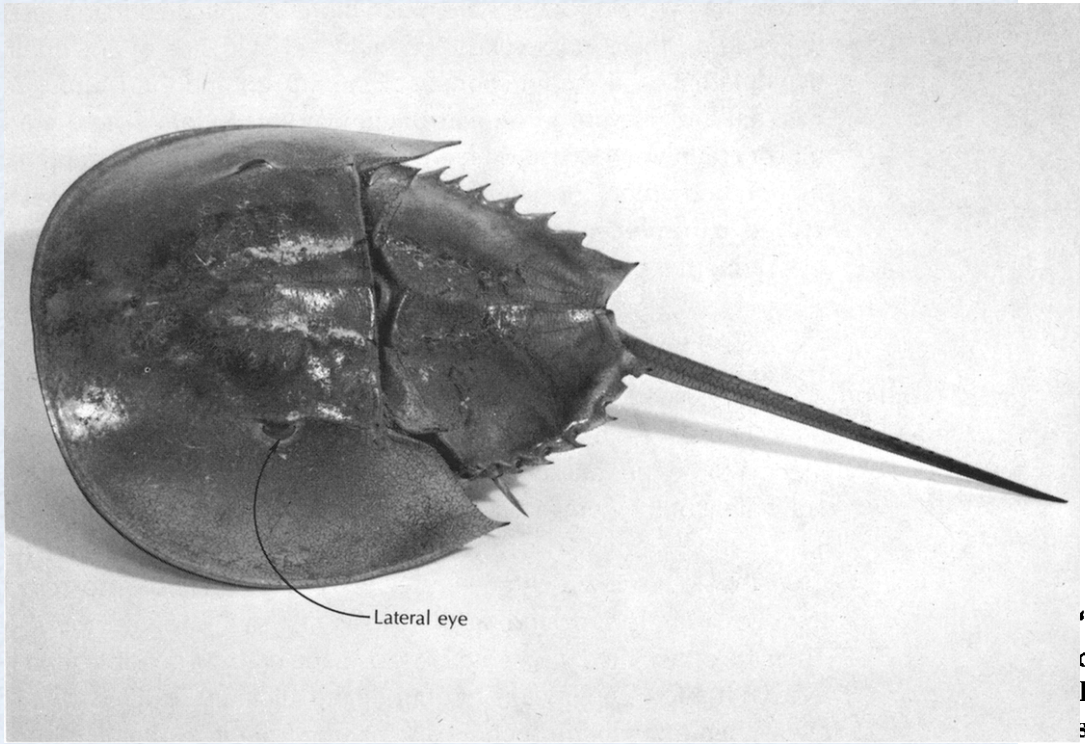

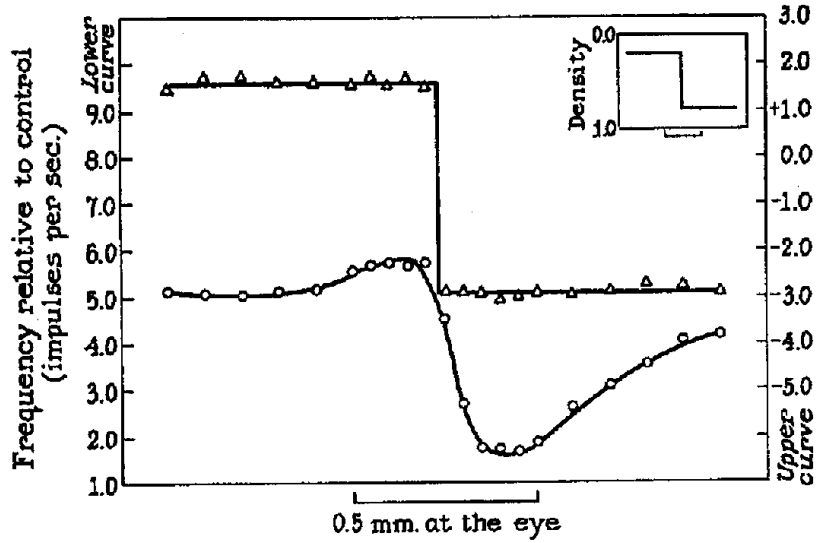

FIG. 4. The discharge of impulses from a single receptor unit in response to a simple "step" pattern of illumination in various positions on the retinal mosaic. The pattern of illumination was rectangular, covering an area 1.65 mm. × 1.65 mm. on the eye. It was obtained by projecting the demagnified image of a photographic plate on the surface of the eye. The insert shows the relative density of the plate along its length as measured, prior to the experiment, by means of a photomultiplier tube in the image plane where the eye was to be placed. The density of the plate was uniform across its entire width at every point. The measurements illustrated were made over the central 1.5 mm. of the image on the eye.

The upper (rectilinear) graph shows the frequency of discharge of the test receptor, when the illumination was occluded from the rest of the eye by a mask with a small aperture, minus the frequency of discharge elicited by a small "control" spot of light of constant intensity also confined to the facet of the test receptor. Scale of ordinate on the right.

The lower (curvilinear) graph is the frequency of discharge from the same test receptor, when the mask was removed and the entire pattern of illumination was projected on the eye in various positions, minus the frequency of discharge elicited

Students can find a simple solution to this model, on a grid of just 8 points:

```
a=0.7; b=a^4; c=a^9;
M1=[a b c 0 0 0 0 0 0 0;...
    b a b c 0 0 0 0 0 0;...
    c b a b c 0 0 0 0 0;...
    0 c b a b c 0 0 0 0;...
    0 0 c b a b c 0 0 0;...
    0 0 0 c b a b c 0 0;...
    0 0 0 0 c b a b c 0;...
    0 0 0 0 0 c b a b c;...
    0 0 0 0 0 0 c b a b;...
    0 0 0 0 0 0 0 c b a     ];
E=[0 0 0 0 0 1 1 1 1 1 ]';
F=(eye(10) + .5*M1)\E;
```

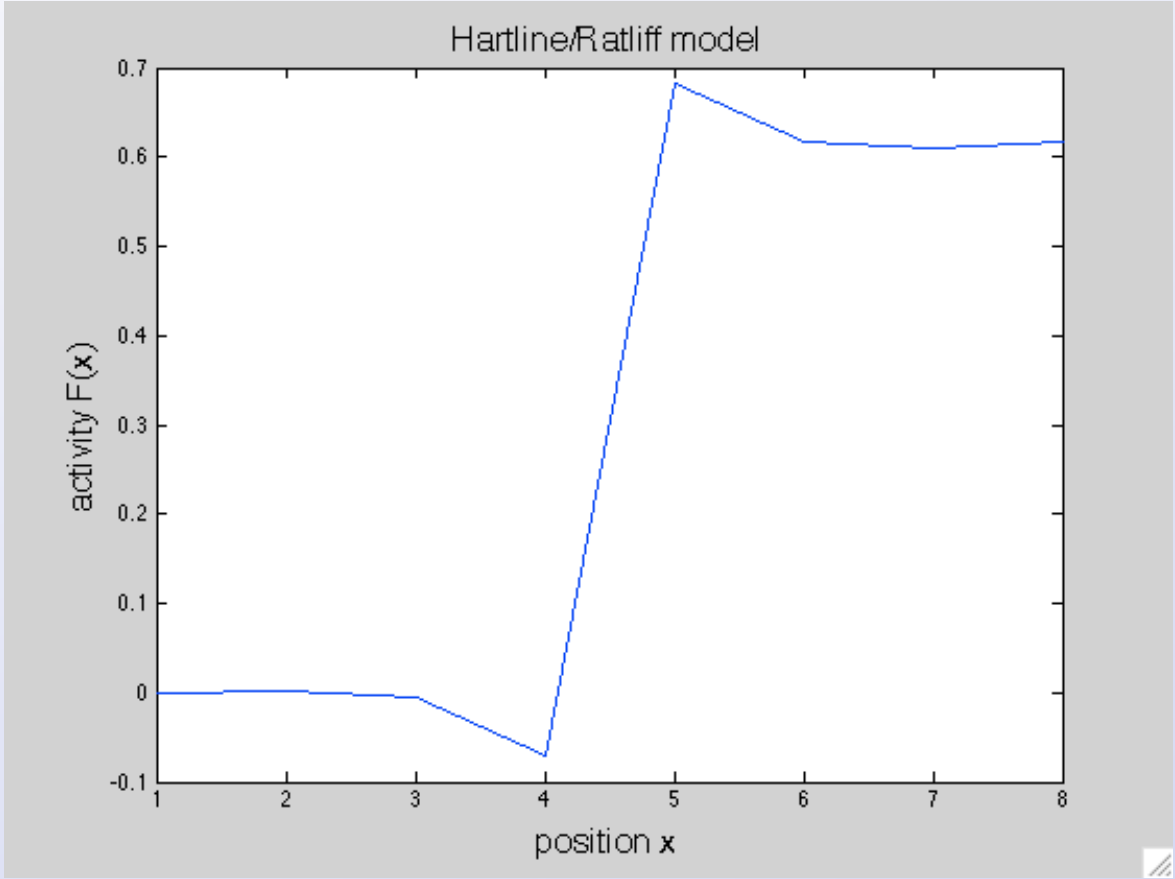

Hartline and Ratliff's model was a bit more elaborate than that -- but not a lot more. We have exploited an analogy to a physics subject: antiferromagnetism.

Effect of randomness

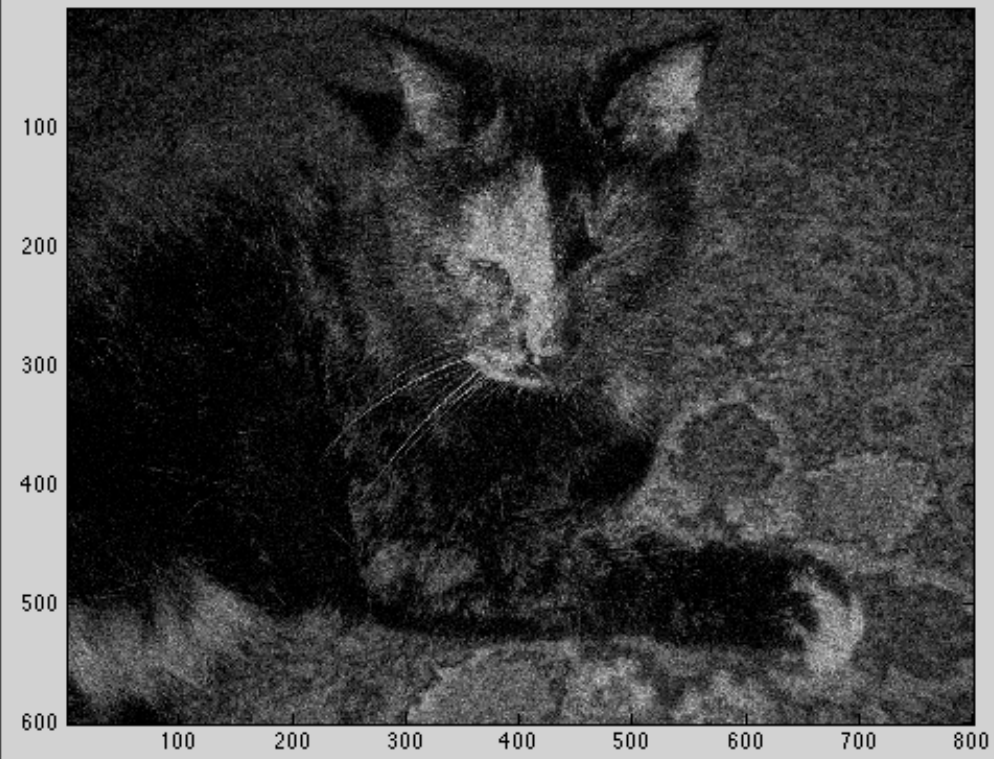

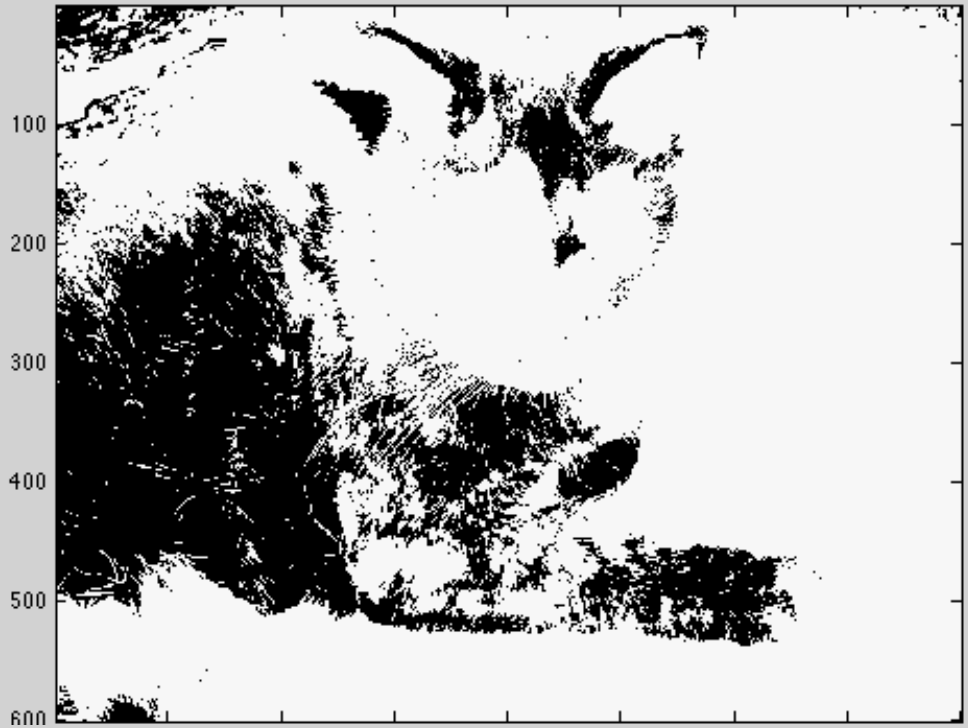

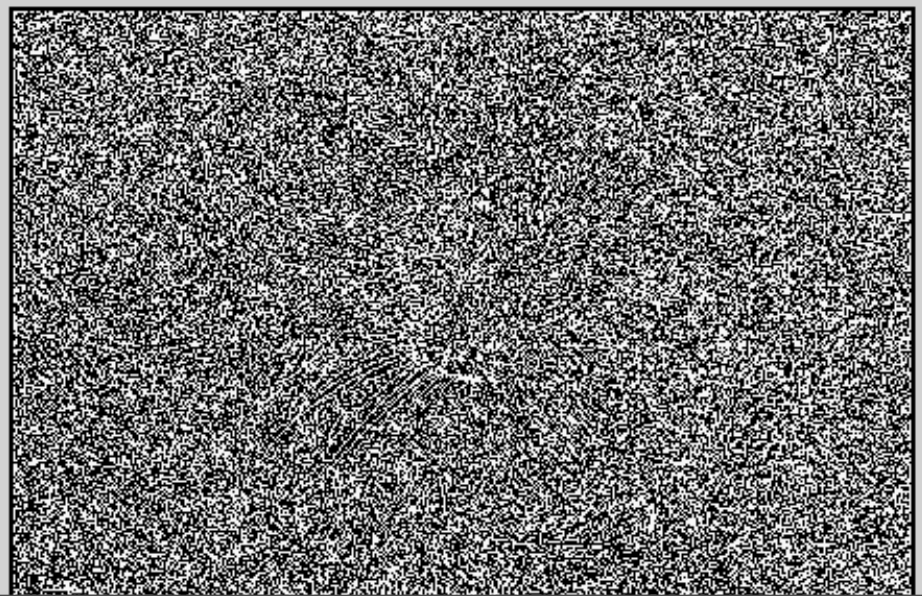

At low illumination we have a lot of Poisson noise (top left), but we can still see detail. Truncation to 1-bit depth again destroys detail (top right), but this time it's *even worse* if preceded by the same filter that was so helpful before (bottom).

(In these pictures the whitest pixels are taken to contain just 30 absorbed photons.)

Rescue

The problem is a familiar one in statistics: “The difference of two noisy variables is itself a *very* noisy variable.”

We can gain back some detail by using a filter with slightly wider center and much wider surround, which averages out some of the noise.

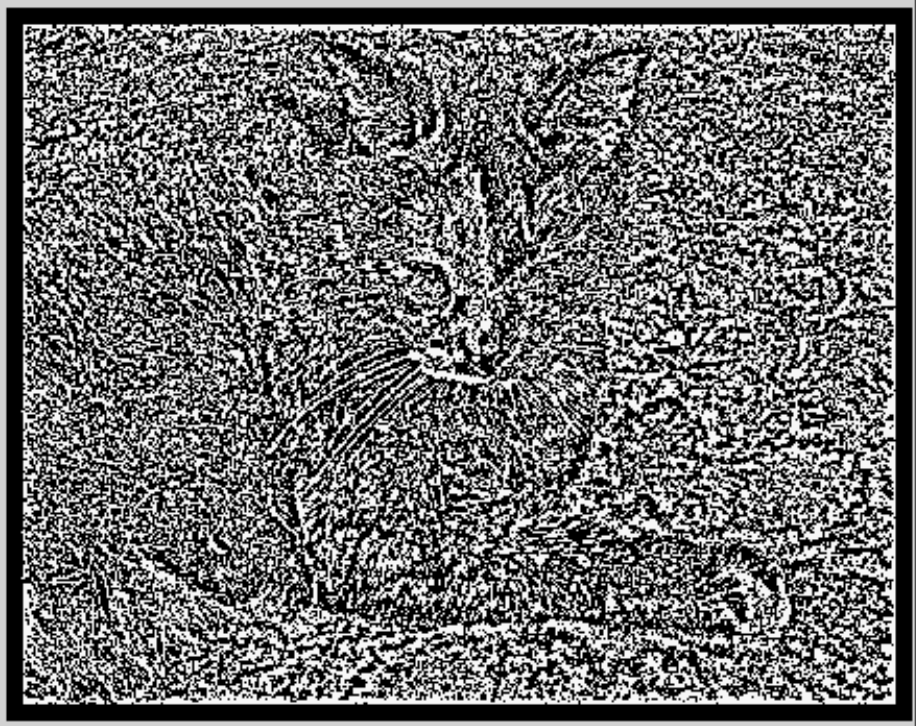

Looks like lateral inhibition is not such a good idea after all, if we want to see in dim light. But an engineer might say, “Simply turn it on only at high illumination.”

Can our eyes do *that?*

Why, yes. Even *fruitflies* can do that.

Students can do this demonstration too, and play with different center/ surround options.

Below left, the center-surround structure of fly visual field (solid curves) changes with changing S/N ratio: At high signal the most inhibition is at 2.5 degrees, whereas at lower signal it moves out to 3 degrees.
Below right, at high illumination (top curves) there's a peak in sensitivity (circles) at 5 cycles/degree. At lower illumination this peak moves to lower spatial frequency (wider center), until at very low illumination there's no peak at all.

Fly

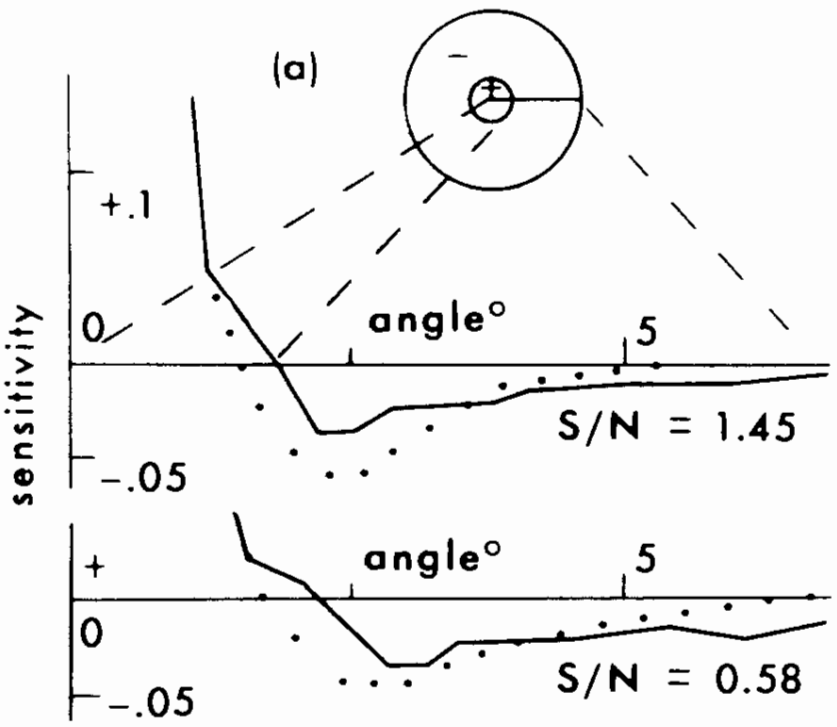

Mammal

Primate

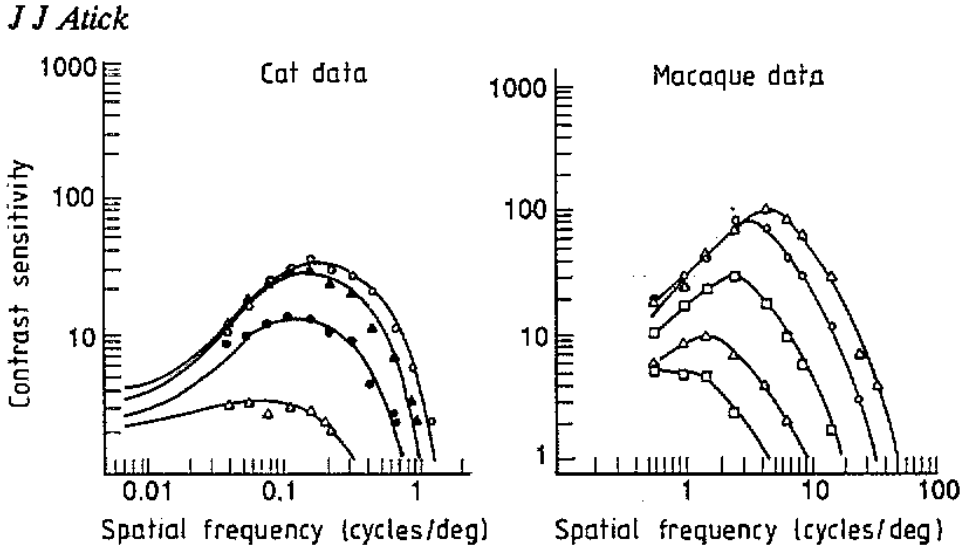

Figure 4. Measured contrast sensitivity. The data in the left figure are reproduced from Enroth-Cugell and Robson (1966), while that on the right are from De Valois *et al* (1974). In both cases, the luminance level I_0 decreases by one log unit each time we go to a lower curve.

Is *that* all?

Is anything *new* going on?"

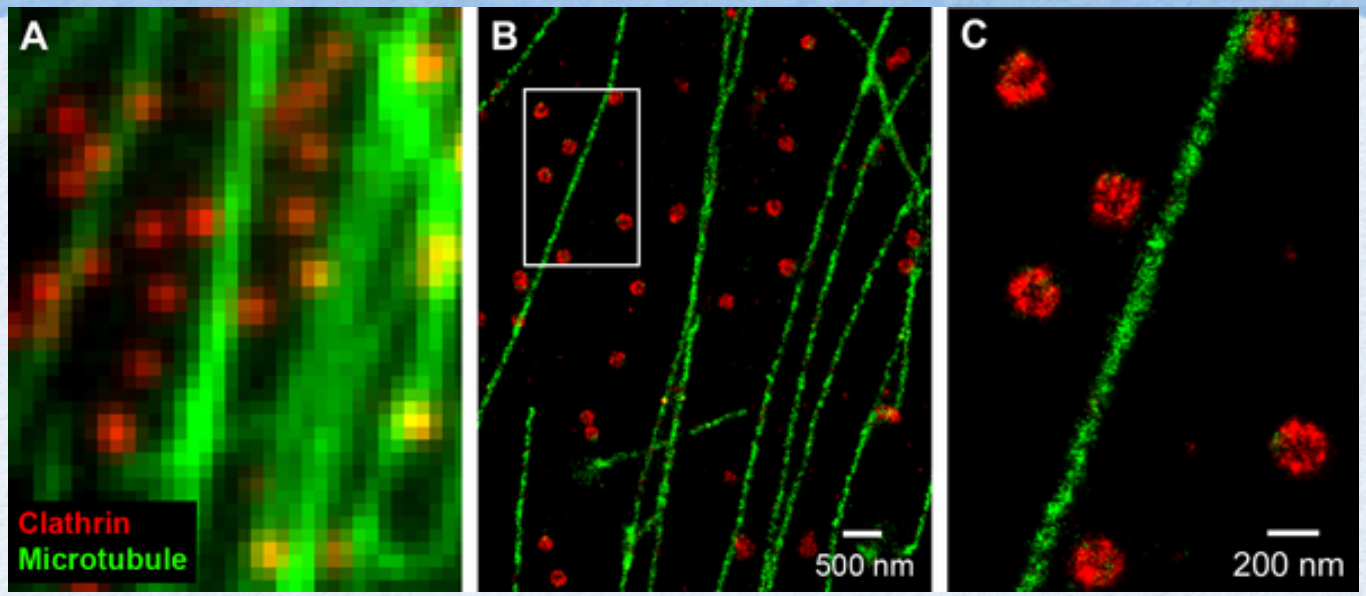

Oh, yes. Perhaps you are tired of the diffraction limit? Perhaps you want to image things smaller than the wavelength of light? Perhaps you want to invent a method that gets named "**Method of the year**" by Nature Methods?

Understanding the statistical character of light has led to subdiffraction microscopy methods like PALM, STORM, FIONA... Now we can see.

In fact, every few months somebody comes up with a new mode of microscopy, which then reveals something important. One might have expected this to be rather dull engineering, but on the contrary many of these are based on truly deep physics ideas.

[Subtext 4]

Besides superresolution, there is 2-photon microscopy... **You cannot play that game if you don't have the right physical model of light.**

(By the way, some students are already using such instruments in their research. They get excited when we help them to understand those black boxes.)

And once you've understood photons, you can move on to fluorescence... then Brainbows... then to FRET... then to the photosynthetic antenna complex...

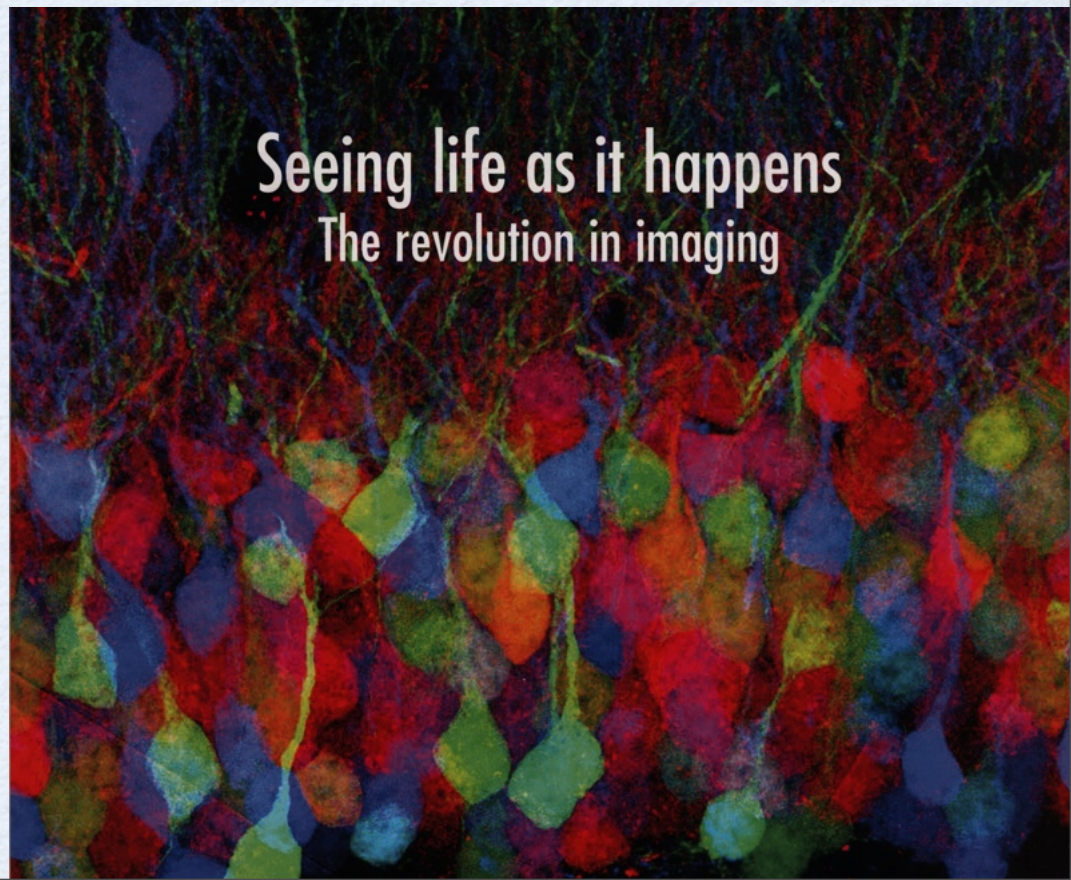

Last word

It turned out we could not understand vision at all without some top-drawer ideas from fundamental physics (like quantum theory). Other cool ideas entered too (stationary-phase, antiferromagnetism...).

Now let's step back from those specifics.

There is a cohort of students who are very interested in Physics, but who don't pursue it. They (or their parents) think of "physicist" as something akin to "professional poet." ***What would you do with that???*** They feel they must study something like Biochemistry or Medicine if they want to be employed some day.

- If we can meet these students halfway, show them a bit of what's happening in Biological Physics, everyone can win. For example, we can use biophysical ideas to introduce statistical physics, or even quantum physics. ***Keep the biophysics in Physics*****.**
- But we shouldn't be ashamed of, or downplay, fundamental physical ideas. Those ideas really matter in the latest biomedical research. ***Keep the physics in Biophysics*****.**
- What's more, **physical models** are often weirdly, *unreasonably effective* in stripping away the inessential from a biological system -- and in displaying connections between things that seemed not obviously connected to our untrained imagination.

Thanks

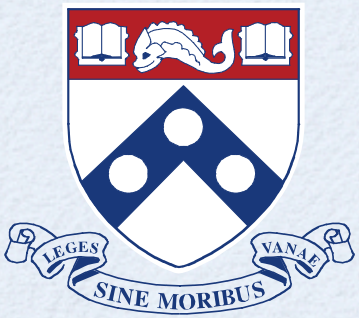

University of Pennsylvania

NSF DMR, BIO

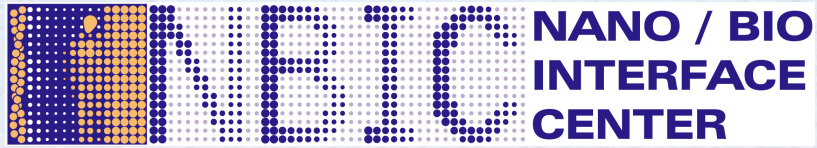

NSF NSEC

For a syllabus and course materials, please contact me:

`nelson@physics.upenn.edu`

For the slides of this talk:

`www.physics.upenn.edu/~pcn`