Physics of Human and Superhuman Vision

Phil Nelson
Univ Pennsylvania

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

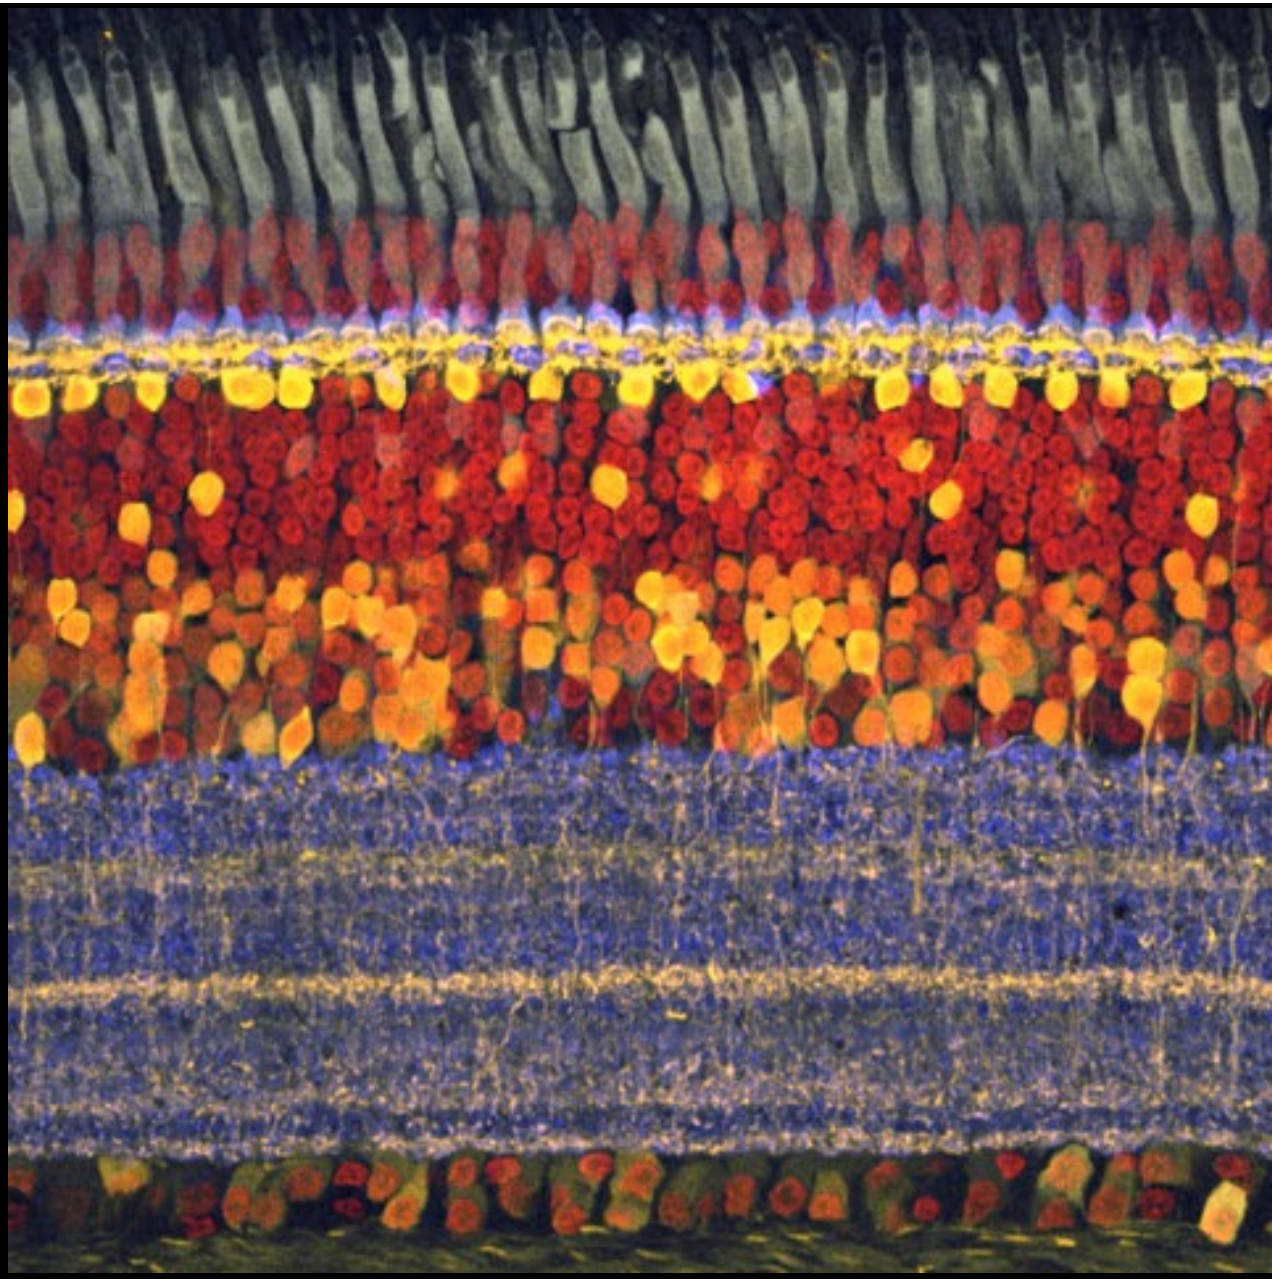

Image of chick retina by Andy Fischer, Ohio State Wexner Medical Center.

Sunshine

Compact fluorescent light bulb

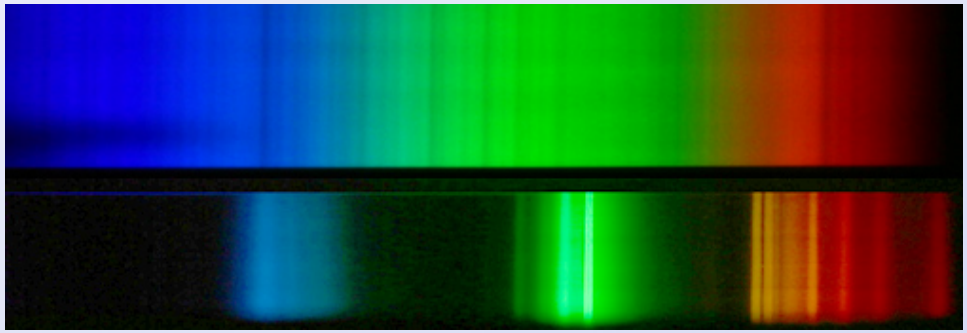

How can two such different kinds of light both look the same? ***Something doesn't fit.***

Just two topics

- What is color and how do we see it?
  - [Can we make a gadget that discriminates colors better than humans, and would that be useful?]
- What sets the ultimate limit on our visual sensitivity, and how close are we to that limit?
  - [Can we made a gadget that makes use of that insight to do something useful?]

This talk

Direct Experience

Color

Light Quanta

[3 slides for the experts]

Wrap

Just four big ideas

- Understanding your own body sometimes requires top-drawer physics ideas.

- Sometimes a simple physical measurement can give you insight decades earlier than “ought” to be possible.

- But sometimes that measurement needs to be coupled with some mathematical analysis.

- Once you understand, even partially, how Nature has implemented one of its impressive tricks (e.g. vision), often you also gain practical benefits.

[And anyway, it’s beautiful.]

Espionage

Part I: Color

Color has always fascinated humans.

Color perception is *useful* to humans, and other animals:

- Image segmentation (separating objects in a scene).
- Object recognition, including fine shades of ripe/unripe fruit.
- Sexual selection.
- Emotional signaling.

Useful, maybe, but not quite simple:

Our uncomfortable observation was that our eyes *discard a lot of information* about the spectrum of light: Perceptually, Y seems just as "pure" a sensation as R, G.

Intellectual opportunity: This doesn't fit our general notions about our (perfect?) eyes. Maybe we can learn something.

Technical opportunity: Is there some way to *not* discard that information?

Light spectrum, or color content curve

Sunlight

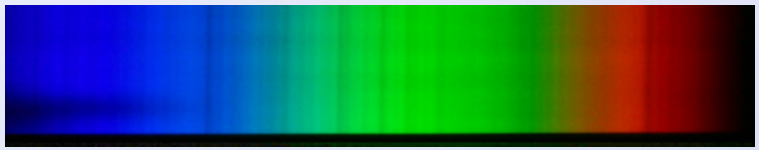

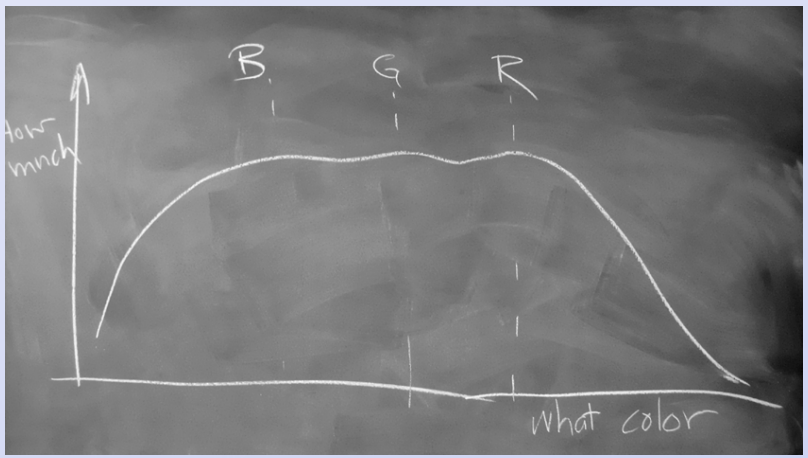

Colored light

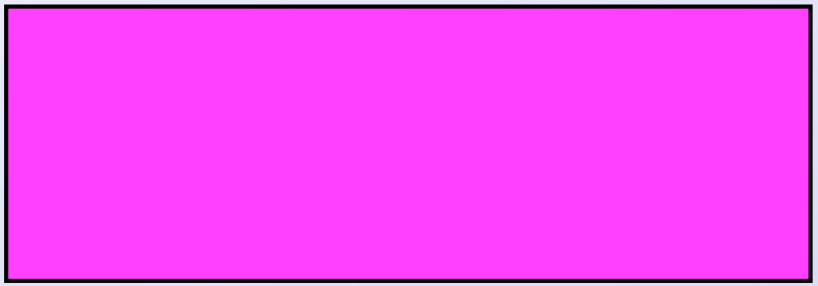

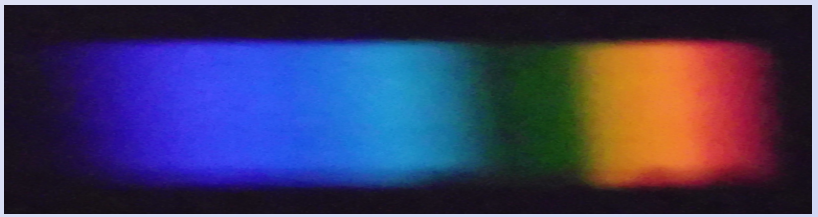

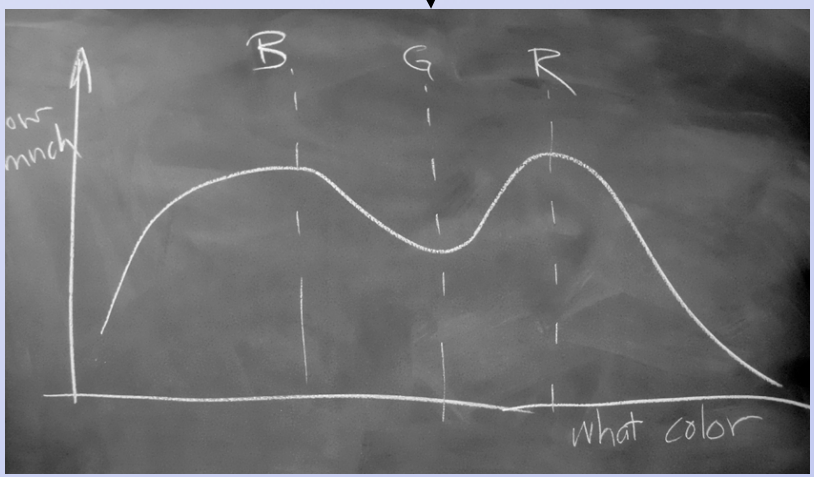

Thomas Young, 1802

An astonishingly modern chain of hypotheses:

1. Light comes in different flavors (let's call them "spectral positions").

2. Even when mixed, those flavors retain their distinct character and can be re-separated.

3. "Color" involves the relative *amounts* of these flavors.

4. Our eyes contain a mosaic of "pixels" ("photoreceptor cells").

5. *All the brain can know* about color is what it hears these cells saying.

And the key point: Each photoreceptor cell is *only sensitive to a particular range of spectral positions*: The cells are "tuned."

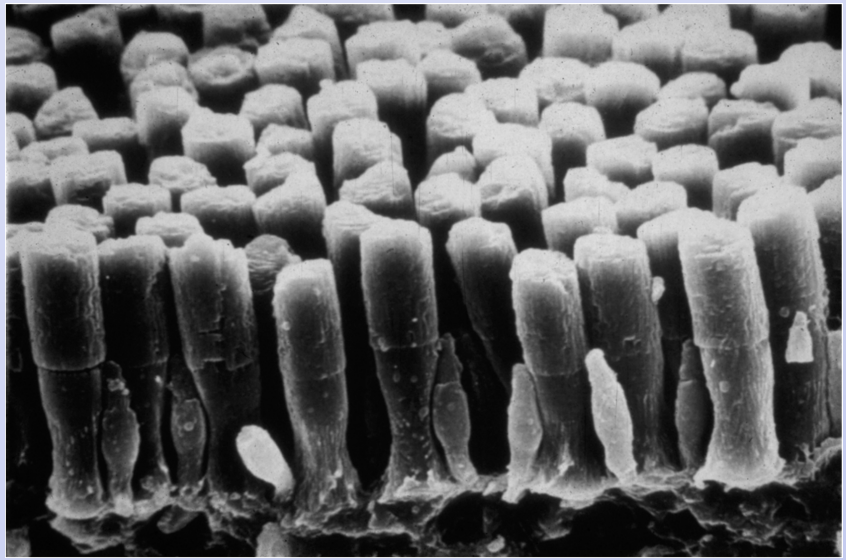

Rod cells and cone cells in the retina of the tiger salamander.
Image by Scott Mittman and Maria T. Maglio

Tuning concept

What could "tune" a receptor cell to prefer light of a particular spectral position? Young realized it could have something to do with *resonance*. Like sound, light does things that resemble what we see with waves on a pond. A wave is characterized by its frequency.

An organ pipe sings at a particular frequency, related to its *size*. Sure enough, here are solutions of "quantum dots" (nanoscale crystals), differing only in the physical *size* of the crystals, all glowing with different frequencies.

Reciprocally, a guitar string only *responds* to a particular range of frequencies.

- Maybe spectral position is a kind of frequency.
- Maybe the receptor cells in our eyes contain something that's selective because of resonance.
- Specifically, let's suppose that the receptor's response to one spectral position is the *product* of the intensity *times* the sensitivity to that color--a **linear relation.**

Photo from **Marija Drndic, U. Penn.**

Thomas Young (continued)

Continuing Young's chain of reasoning,

1. Light comes in different flavors ("spectral positions").
2. Even when mixed, those flavors retain their distinct character and can be re-separated.
3. "Color" involves the relative *amounts* of these flavors.
4. Our eyes contain a mosaic of "pixels" ("photoreceptor cells").
5. *All the brain can know* about color is what it hears these cells saying.

6. Each photoreceptor has a distinct sensitivity range.

7. They come in just 3 classes. Each cell has exactly the same sensitivity range as all the others *in its class*.

Proposed resolution of the R+G=Y paradox

This list of the sensitivities of a photoreceptor cell to light of various spectral positions can also be drawn as a graph. Unlike the light spectrum, which tells "how much is present," this **sensitivity curve** expresses "how much is *needed*" to get a response to each kind of light.

Forget about blue and consider only red- and green-sensitive cells:

If the sensitivity curves *overlap*, then sending in pure spectral yellow will excite both the green-sensitive and the red-sensitive cells equally.

But the same result can be achieved by sending in equal amounts of pure green and pure red light!

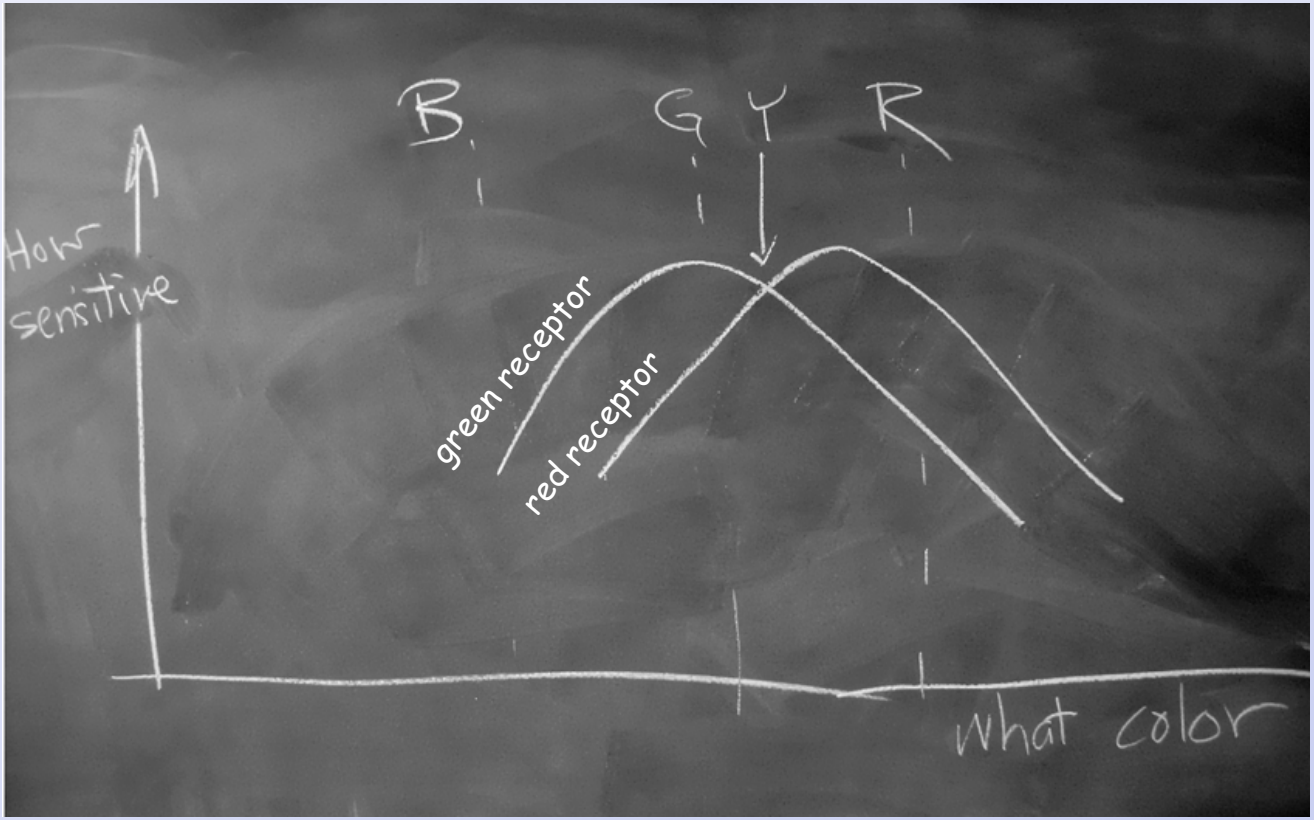

The brain can't tell the difference because *all it knows is what the receptor cells tell it*.

Summary

- Pure spectral light has a continuously-varying property (its "spectral position").
- Many kinds of mixture are possible; huge variety of spectra.
- But our eyes sample spectra with just *three* sensitivity curves, sending just *three* signals to the brain. Each signal depends linearly on the incoming spectrum.
- All the brain can know are those three signals, leading to some ambiguity of color discrimination (loss of potential information).

Young's hypothesis was way ahead of its time. Nobody had ever seen a photoreceptor cell, and when they did they all looked exactly the same in the electron microscope. Detailed confirmation came 162 years later! *That's spycraft*.

Are we done? We saw something weird; we found a hypothesis that seems to explain it. Ready to reap those golden rewards?

Well, you can take a lot of flak when you're that far ahead. Peer review wasn't built to handle it:

> "It is difficult to deal with an author whose mind is filled with a medium of so fickle and vibratory a nature...; We have searched without success for some traces of learning, acuteness, and ingenuity, that might compensate his evident deficiency in the powers of solid thinking...."
>
> -- Henry Brougham. [Criticizing Young's theory]

A missing step

And anyway...

The propositions we most desperately want to be true are the ones we must mistrust the most.

A quantitative test

OK: better *nail the case* for Young hypotheses before we call the VCs.

Quantitative, detailed, testable prediction is crucial.
That's our discipline: Don't go too far on a tangent without experimental *authority*.
Ideally we'd like a lot *more* experimental data points than unknowns (fit parameters).

"Target" light

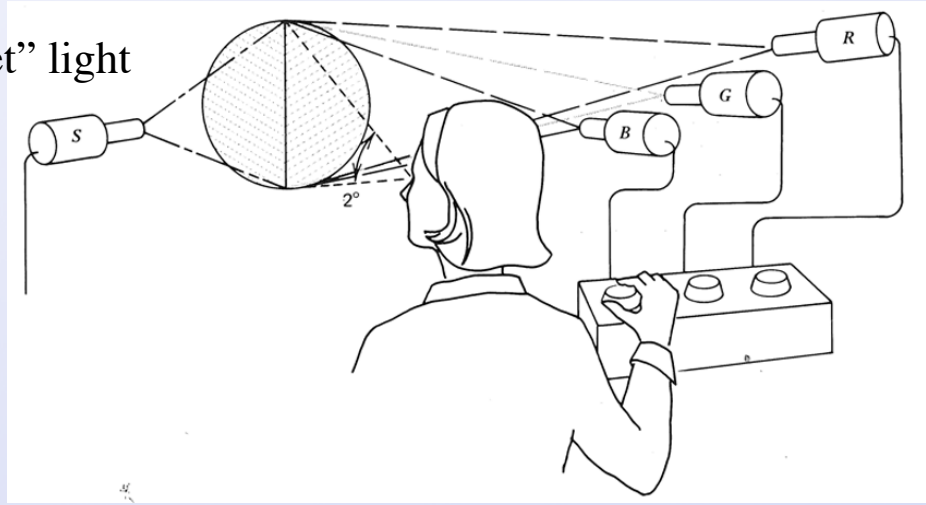

3 standard lights.

Result: three "color matching curves":

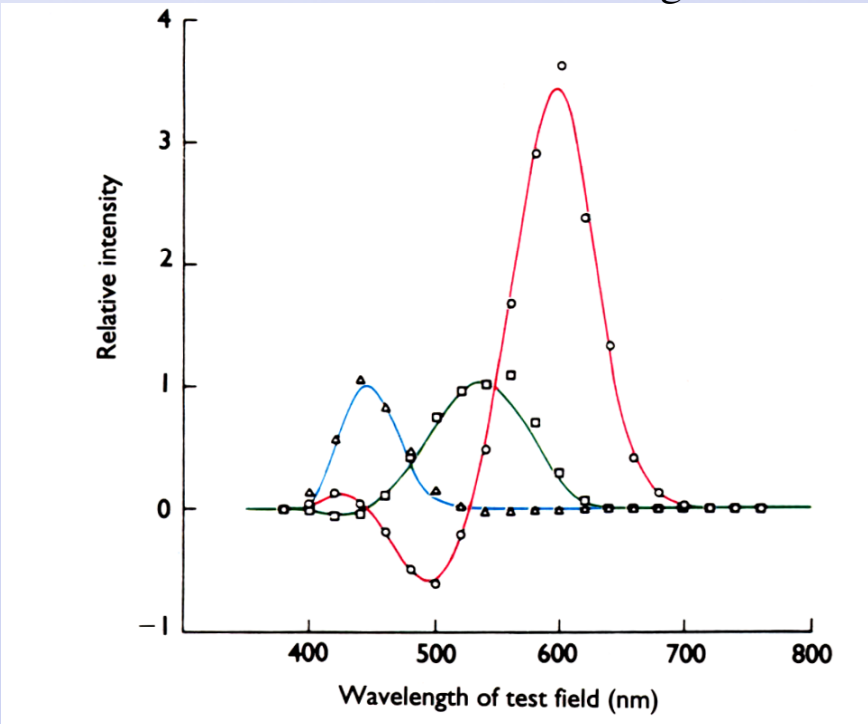

The theory makes testable predictions

Once we measure the sensitivity curves, we can predict the response of each photoreceptor to any possible light spectrum. Then we can find out how much of each of the three standard lights is needed to *mimic* the response elicited by the target light, by solving *three linear equations*.

Left: The sensitivity curves of color photoreceptors indeed fall into three well-separated classes. Notice the big overlap between the "red" and "green" curves.

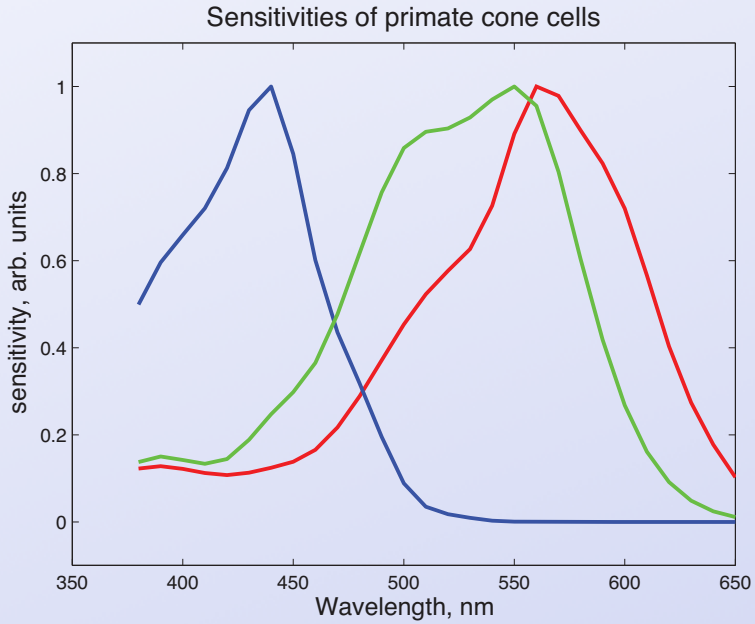

Right: Once those curves are known, the color-matching functions can be *predicted*, and *they agree with psychophysical measurements*.

Curves: data from color-matching experiments.
Dots: predictions from theory.

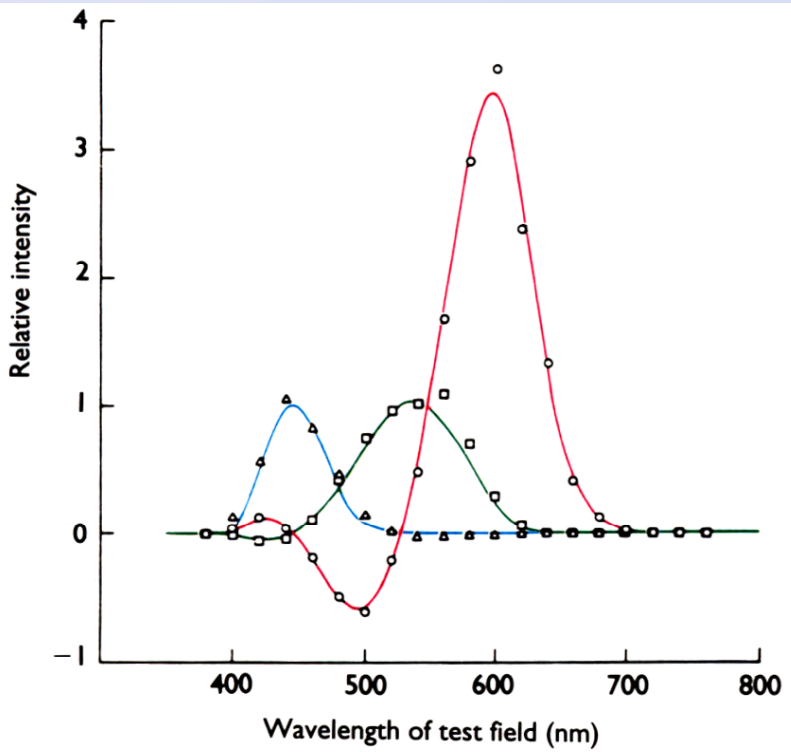

Data from Julie Schnapf and Denis Baylor 1987.

First tech payoff

You can fool the eye into thinking that a wide range of colors is present by using just three pixel types…

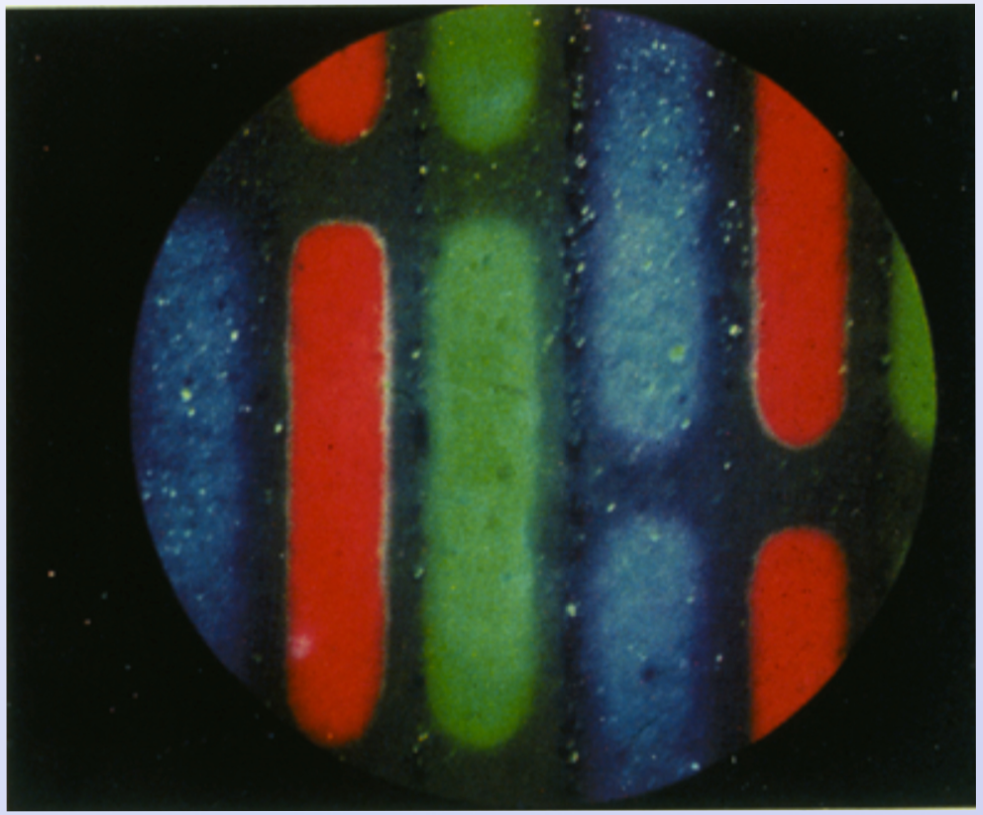

Mixing 3 colors is enough to match (almost) any color. That's good for making inexpensive computer displays.

But turning it around: Our eyes *discard* a lot of information about the spectrum of light entering any given visual field! **Can an artificial visual system discriminate *better* than that?**

Superhuman vision, 1

OK, that was a 19C phenomenon, based on 17C discoveries about light, confirmed in 1980s. Is that *all?*

Subhuman color vision

Human

Superhuman color vision

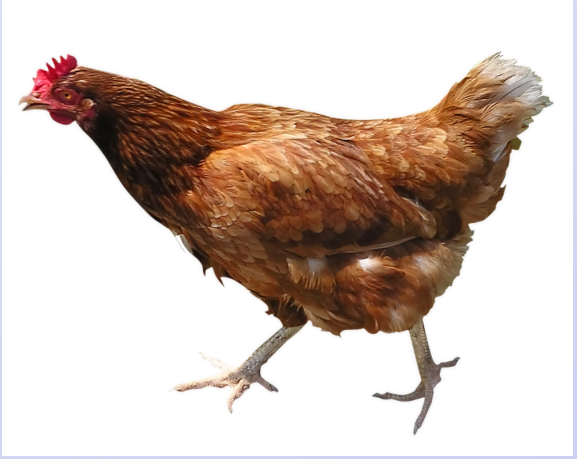

Superchicken?

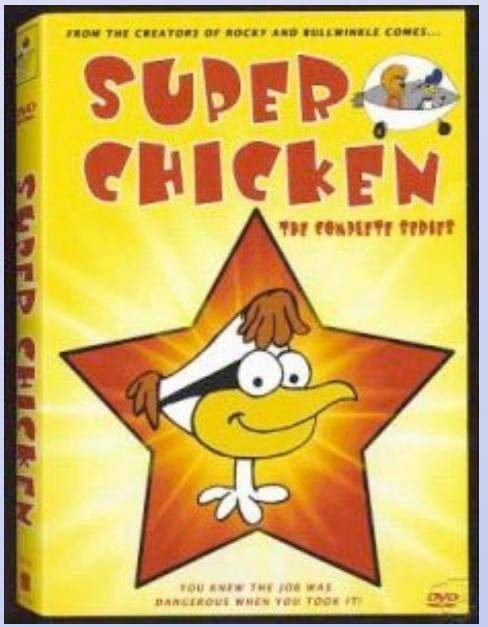

Karyotyping

Yuval Garini, Physics Department
& Bar-Ilan Institute of
Nanotechnology, Israel

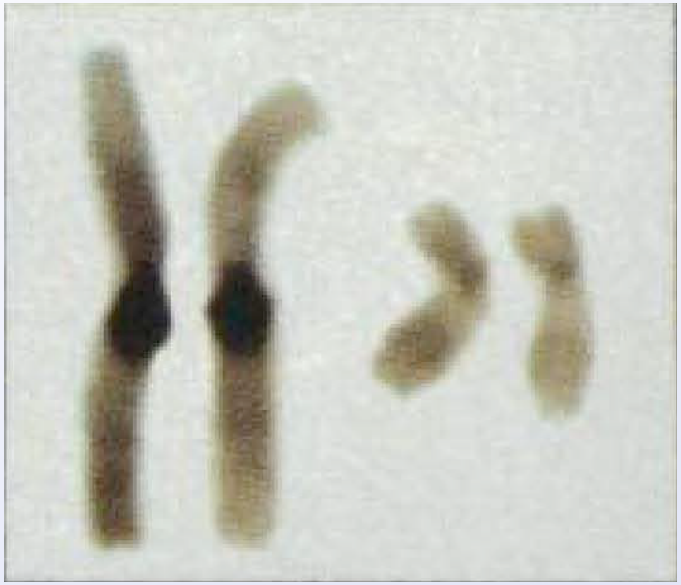

One-color (DAPI) staining can reveal, some, not all, chromosome abnormalities.

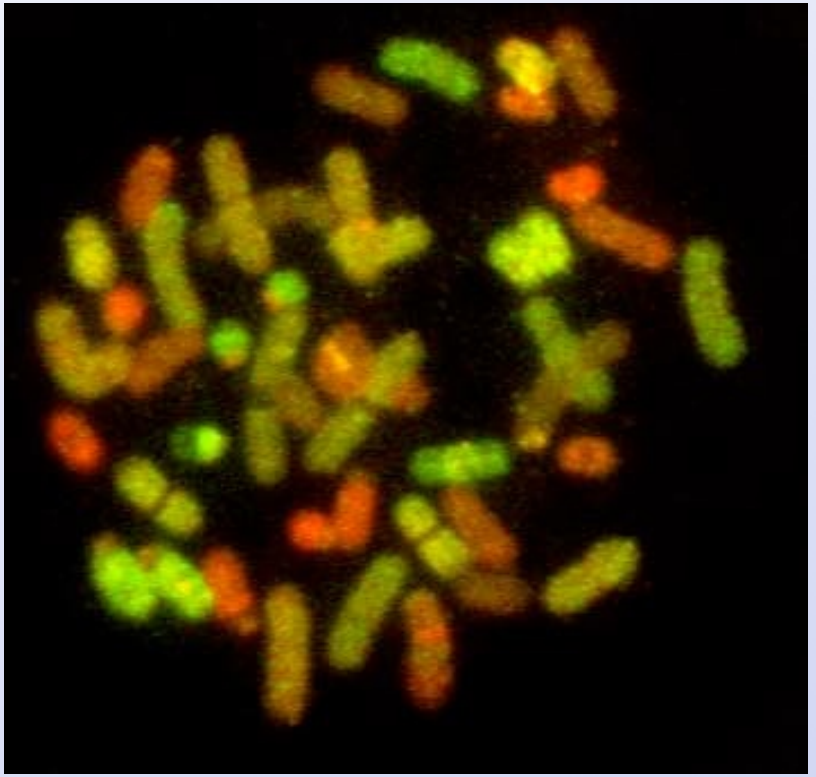

Multi-color (FISH) staining is hard to interpret when you go beyond two colors.

E Schröck, S du Manoir, T Veldman, B Schoell, J Wienberg, M A Ferguson-Smith, Y Ning, D H Ledbetter, I Bar-Am, D Soenksen, Y Garini, T Ried. Science 1996.

Spectral karyotyping, 1

Could we automatically sample each pixel with *many* sensitivity curves?

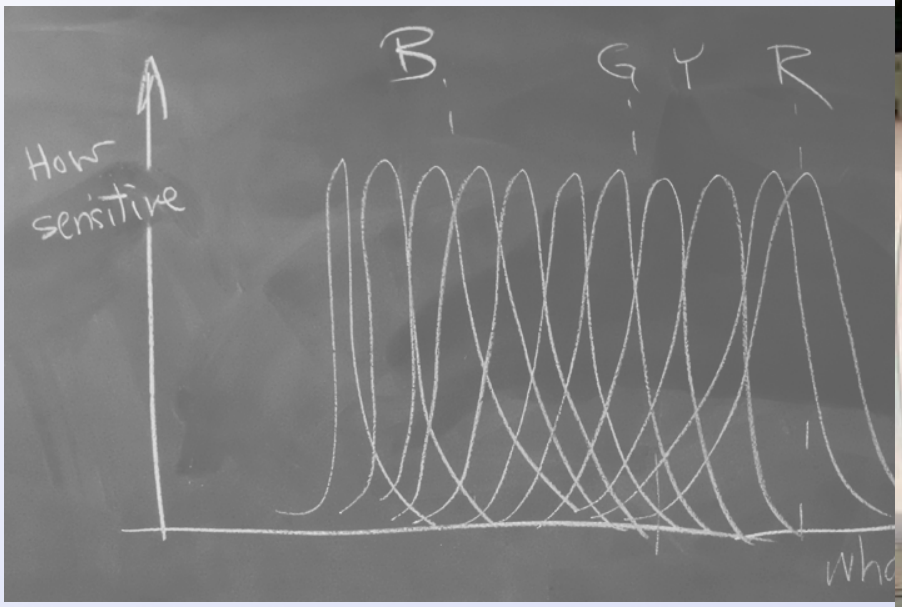

That would give us a detailed spectrum -- not just 3 numbers -- at every point in the image!

The Sagnac Interferometer is a gadget that mechanically scans through the full spectrum of every pixel in the image.

That's spycraft.

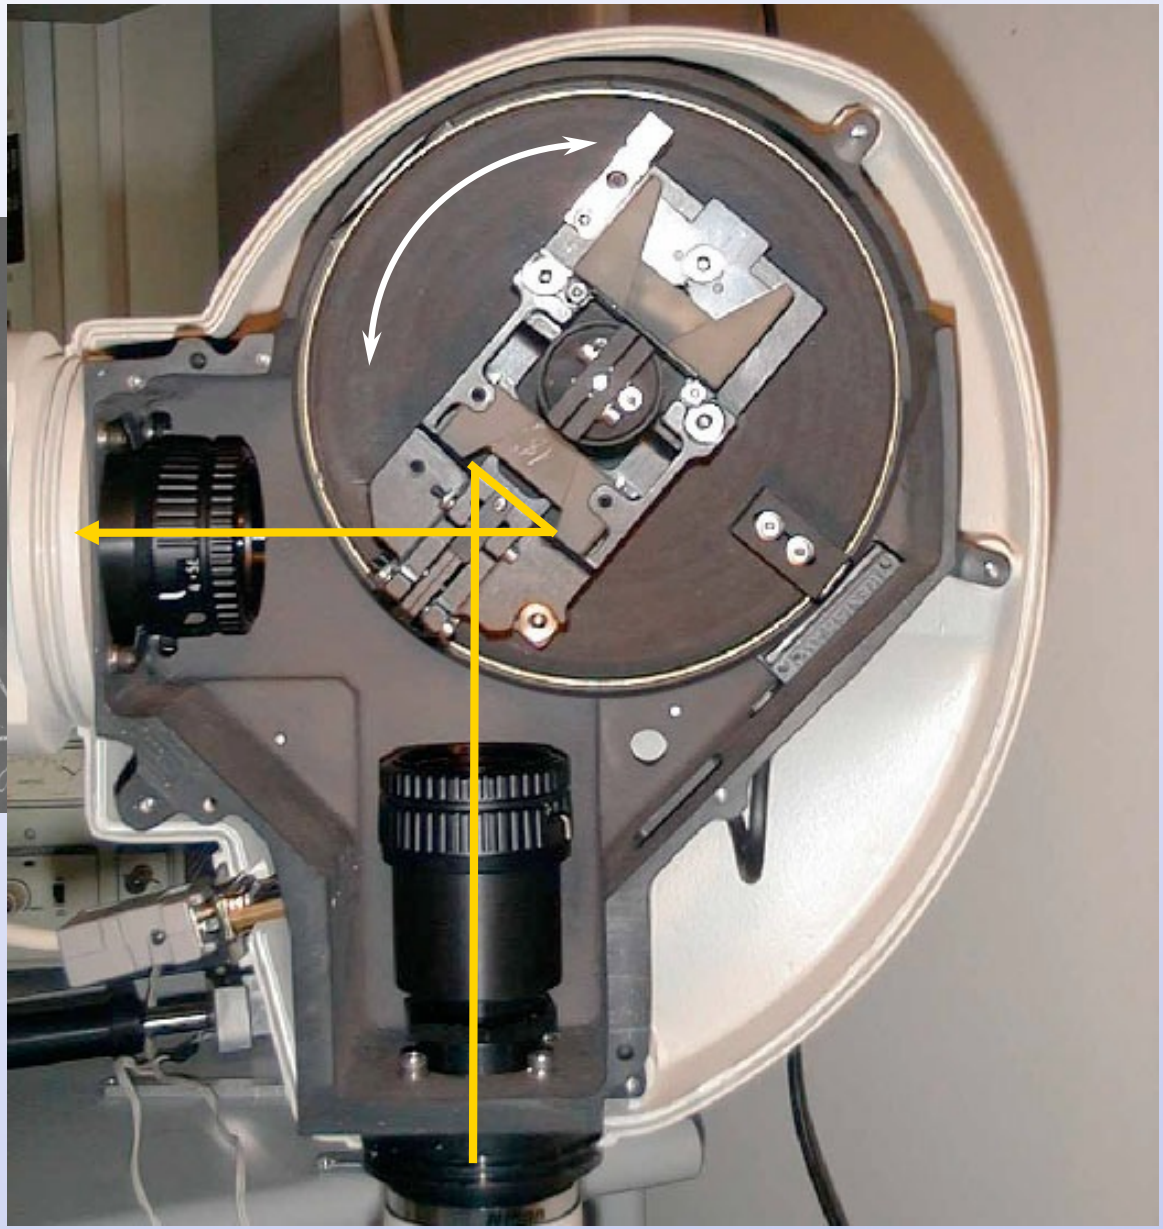

Spectral karyotyping, 2

Fluorescent dyes can be combined to give a lot of distinct spectra:

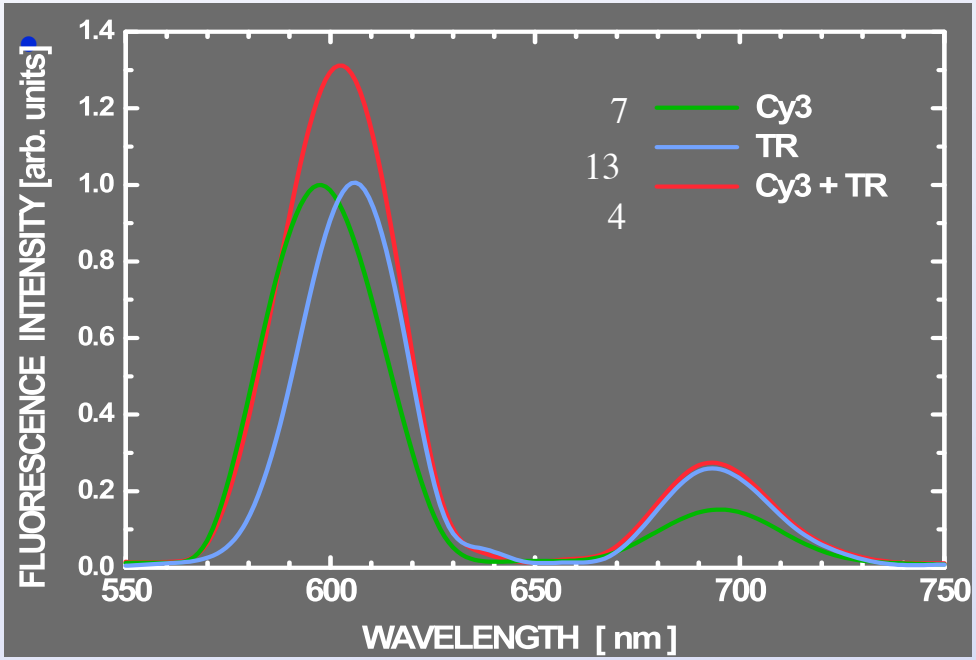

Our unaided eyes are not so good at discriminating the resulting colors:

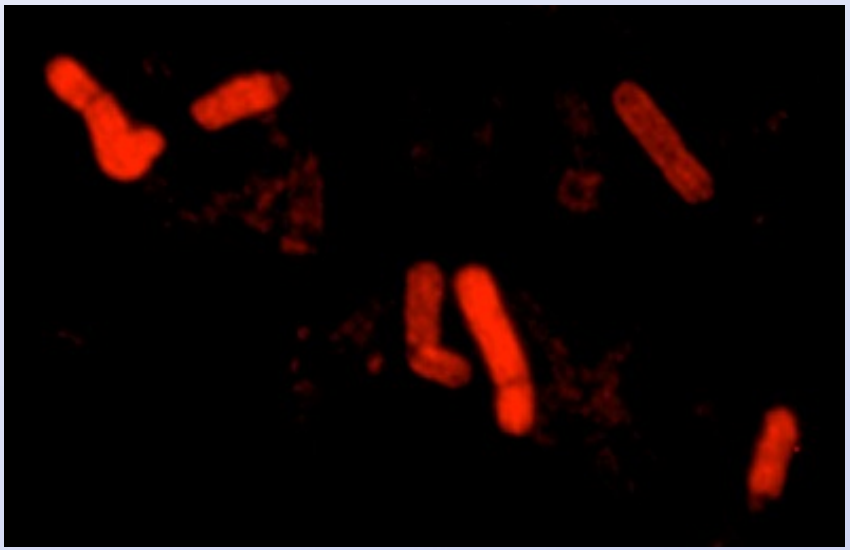

But we can compare the spectra to the known curves, make our assignments, and replace them by human-friendly false colors:

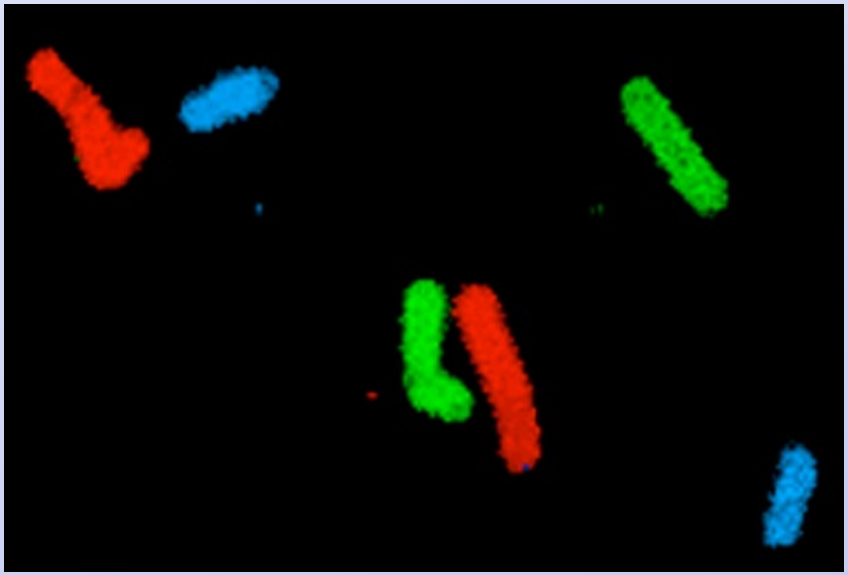

Spectral karyotyping, 3

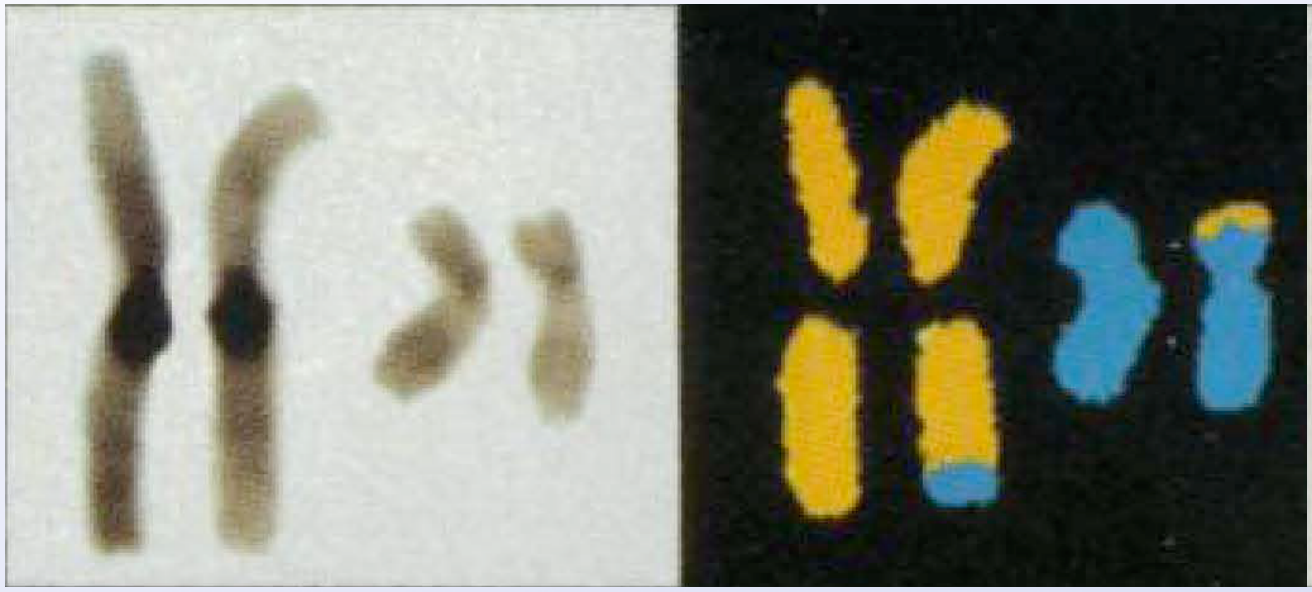

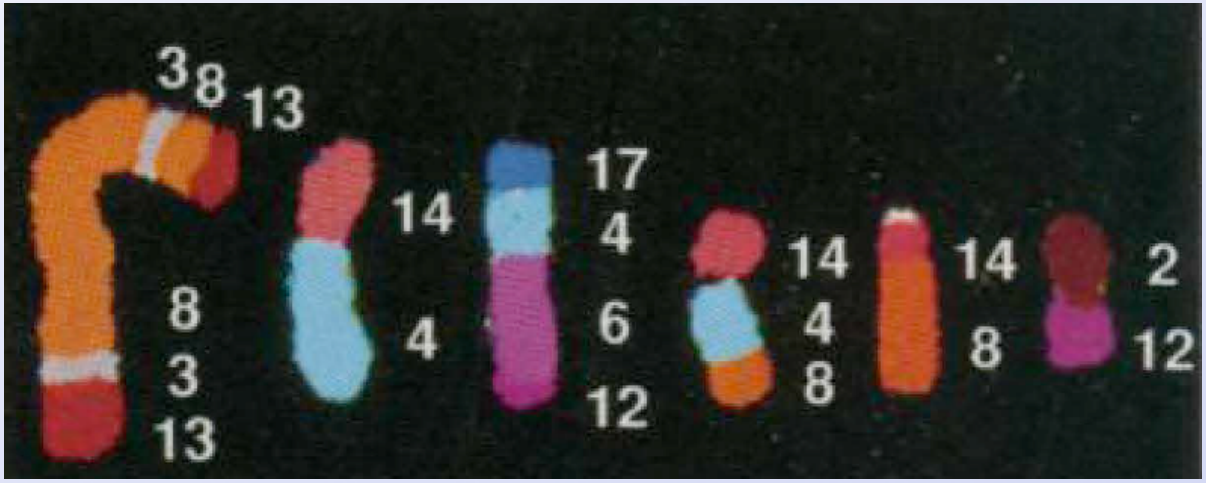

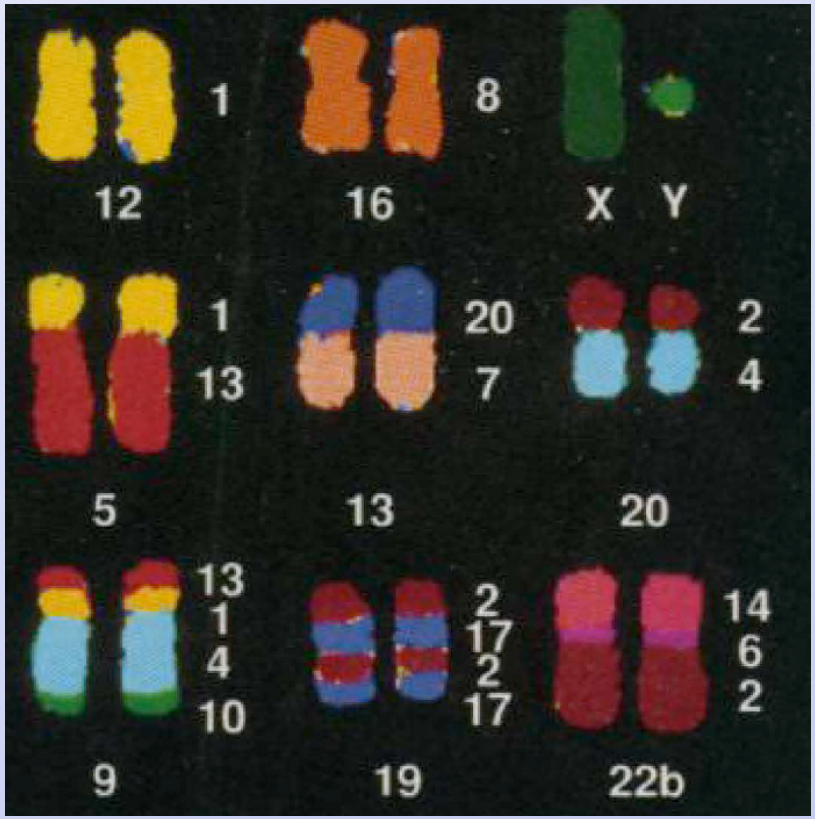

E Schröck, S du Manoir, T Veldman, B Schoell, J Wienberg, M A Ferguson-Smith, Y Ning, D H Ledbetter, I Bar-Am, D Soenksen, Y Garini, T Ried. Science 1996.

Superhuman vision, 2

If you look at a slice of brain in a microscope, all you see is a dense tangle.

Neuroanatomy dates from Golgi's invention of a way to see a complete, single neuron amid the welter of its neighbors.

Unfortunately, these magnificent images tell us nothing about the *connections* between all those neurons.

Oh, also -- the method also kills the tissue.

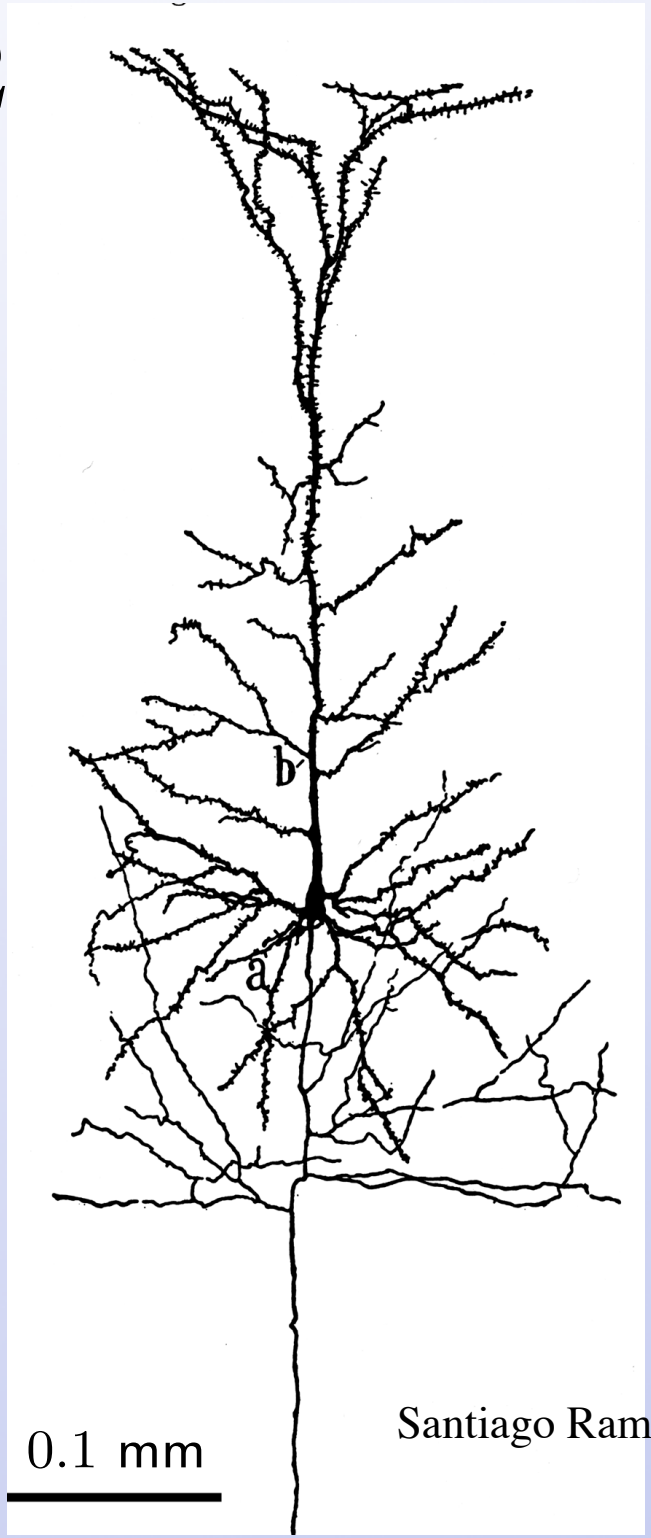

Santiago Ramon y Cajal, 1909.

Superhuman vision 2: “Brainbow” imaging

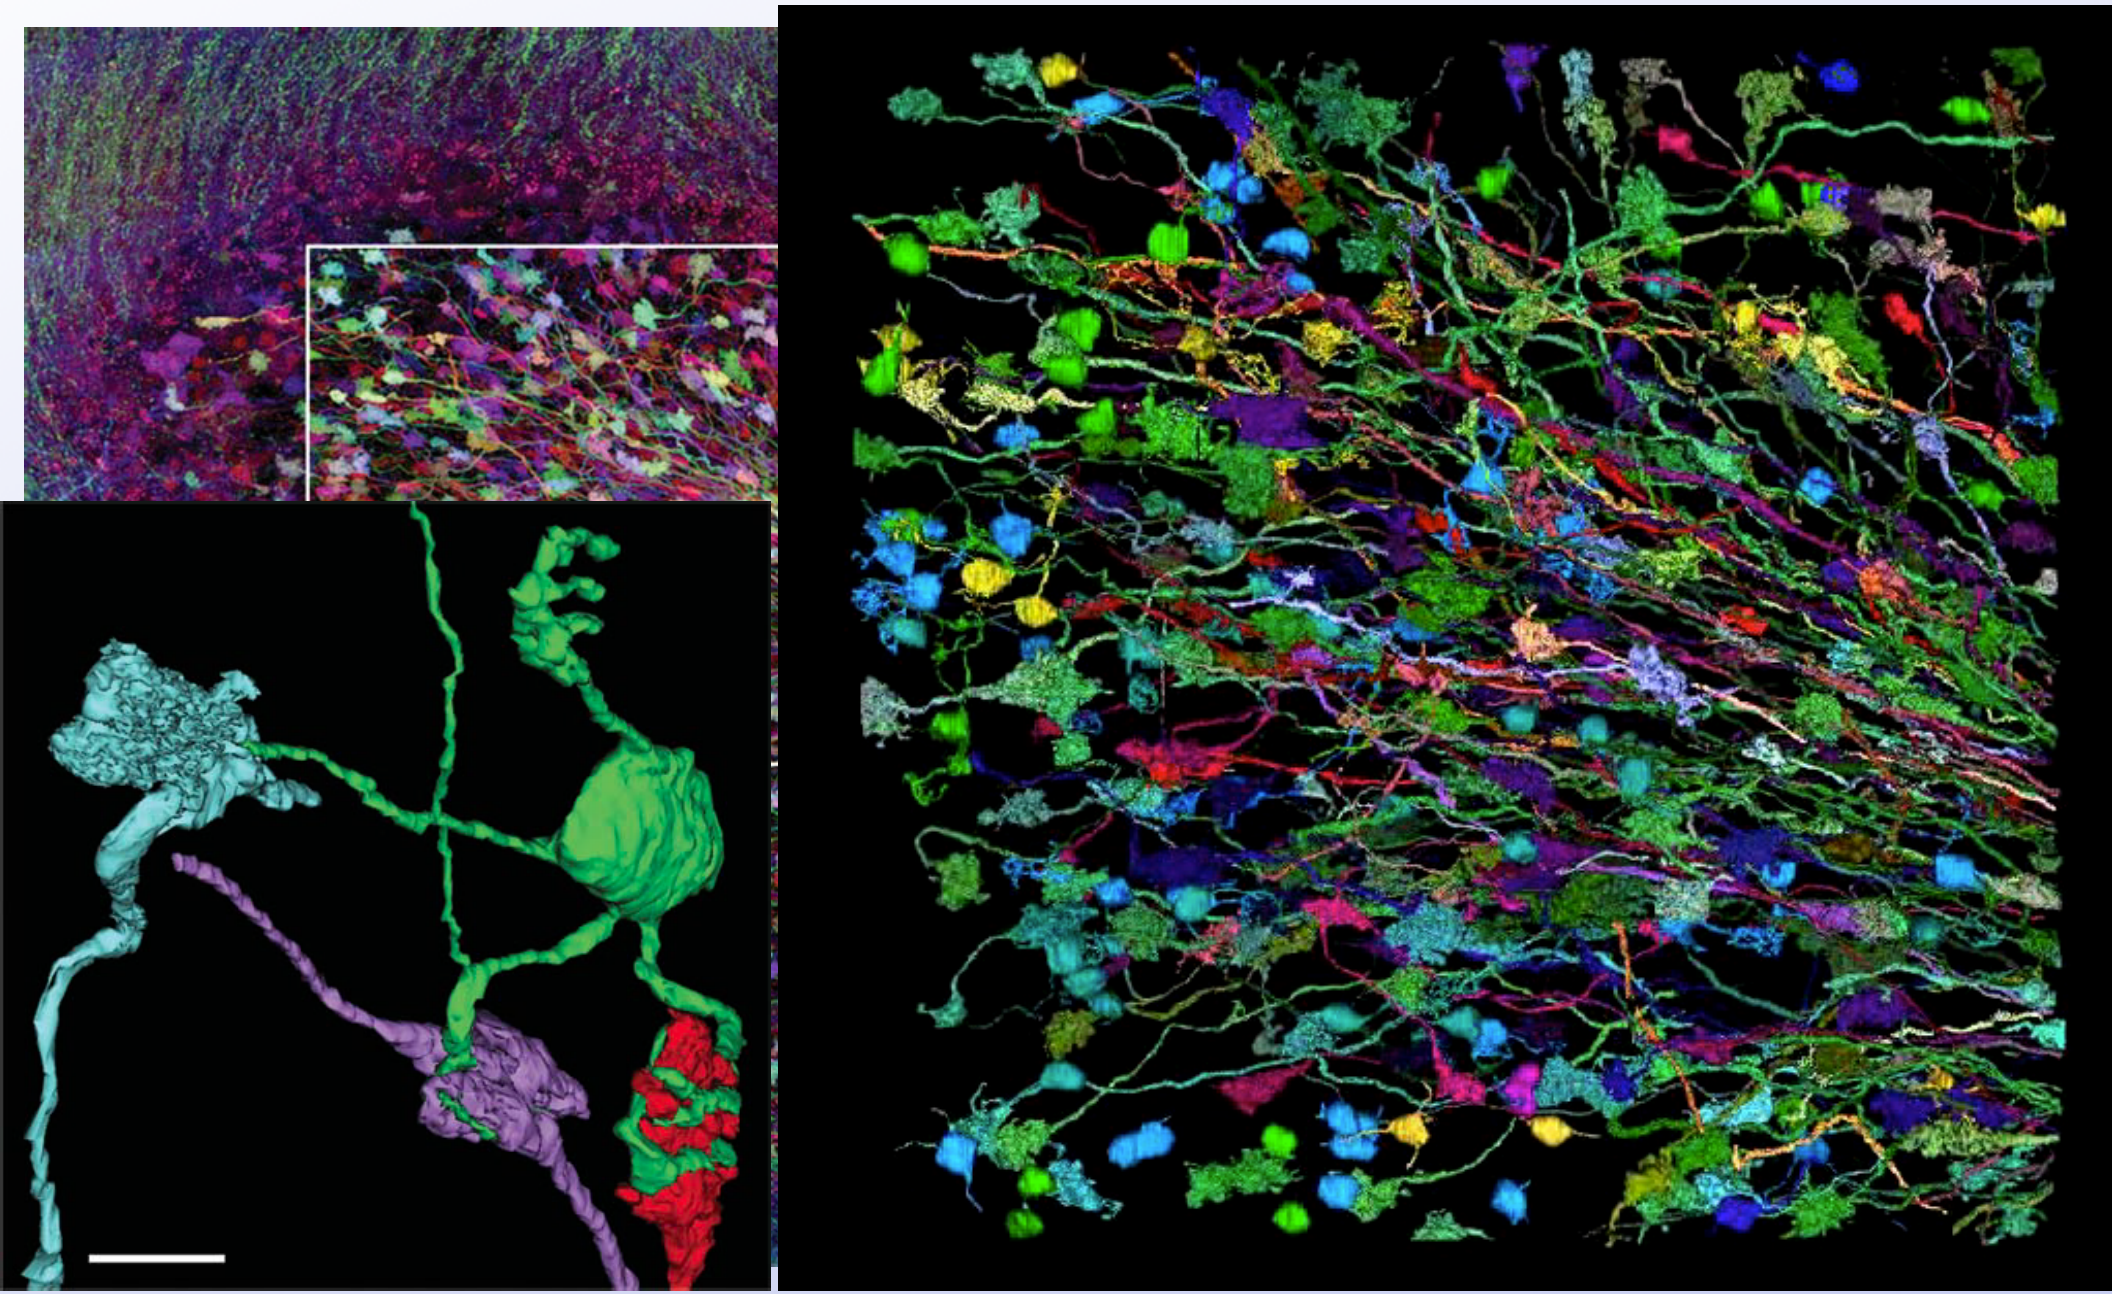

Jean Livet, Tamily A. Weissman1, Hyuno Kang, Ryan W. Draft, Ju Lu, Robyn A. Bennis, Joshua R. Sanes & Jeff W. Lichtman. Nature 2007.

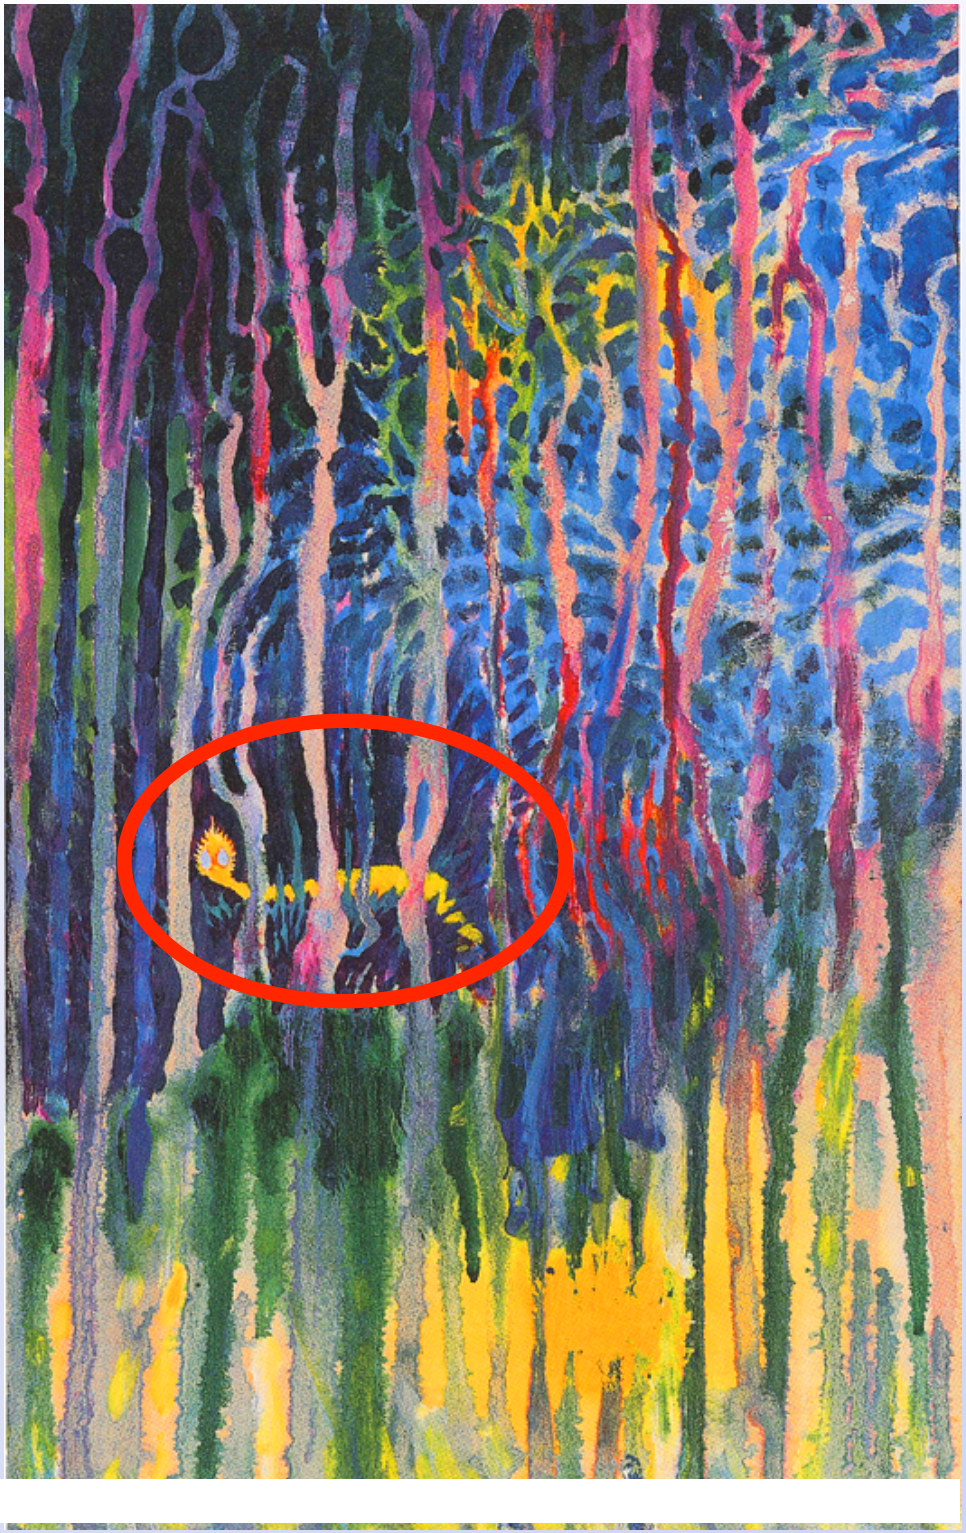

Wrap:

Without an understanding of our own (all too human) vision, we might not have imagined the possibility to do better, nor the means to do so.

That's spycraft.

"The secret art of Dr Seuss"

Part II: Quanta

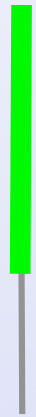

OK, great! Fun demo, fun story, good applications -- let's quit.

No, wait. A few small matters remain:

- What *is* light?
- What *is* its "color content?"
- What *is* "tuning" all about? Why is that crucial relation linear? Where did those sensitivity curves that we used come from (how do you measure them)?

Can we learn something more specific about light, and about our eyes? If so, would it have any practical value?

Uh-oh

What happens in those photoreceptor cells that translates light into nerve impulses?

We can detect very dim light with a photomultiplier tube or avalanche photodiode. Either way, light causes discrete clicks in the detector. ***Dimmer light gives equally big clicks, just less frequent***:

Dim illumination:

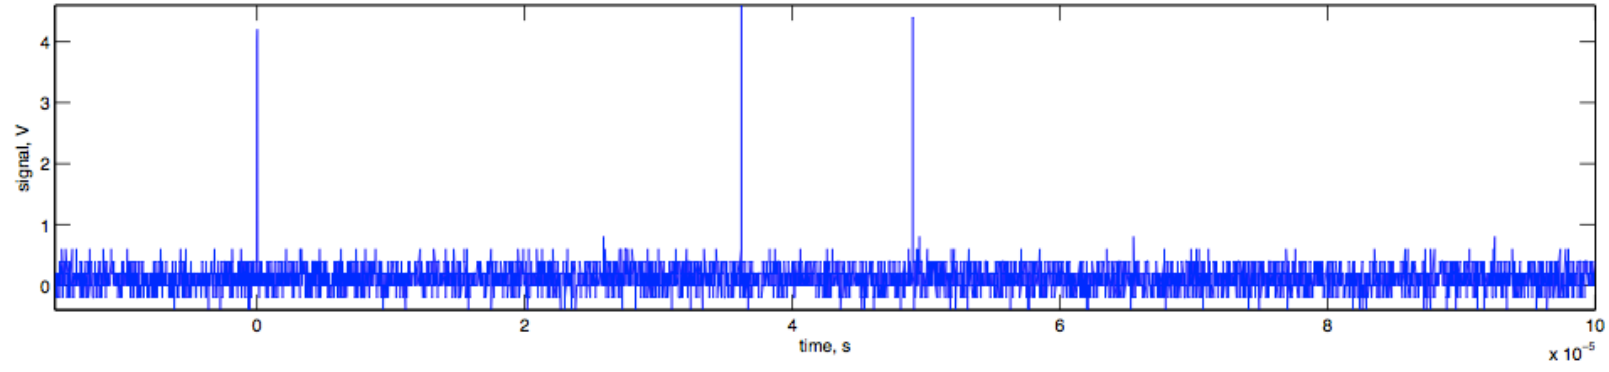

Slightly brighter (still very dim):

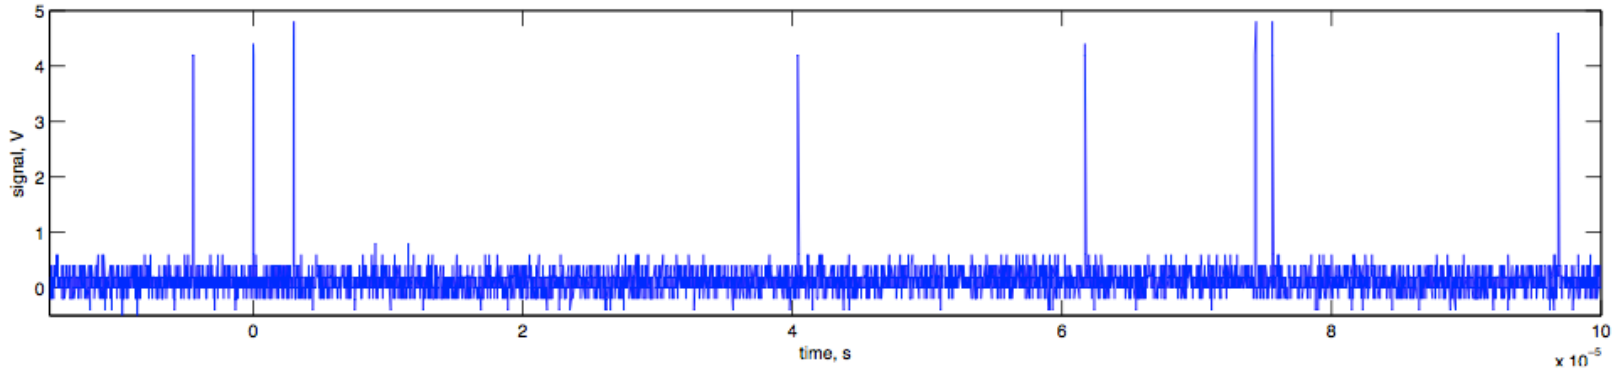

Apparently light is lumpy. Albert Einstein was forced to that conclusion in 1905, much against his (and everybody else's) will. ***Something doesn't fit*** the older picture of light.

Experimental data courtesy J. F. Beausang.

Uh-oh, 2

You might imagine a mechanism something like this device, which takes a continuous stream of water and converts it into discrete events:

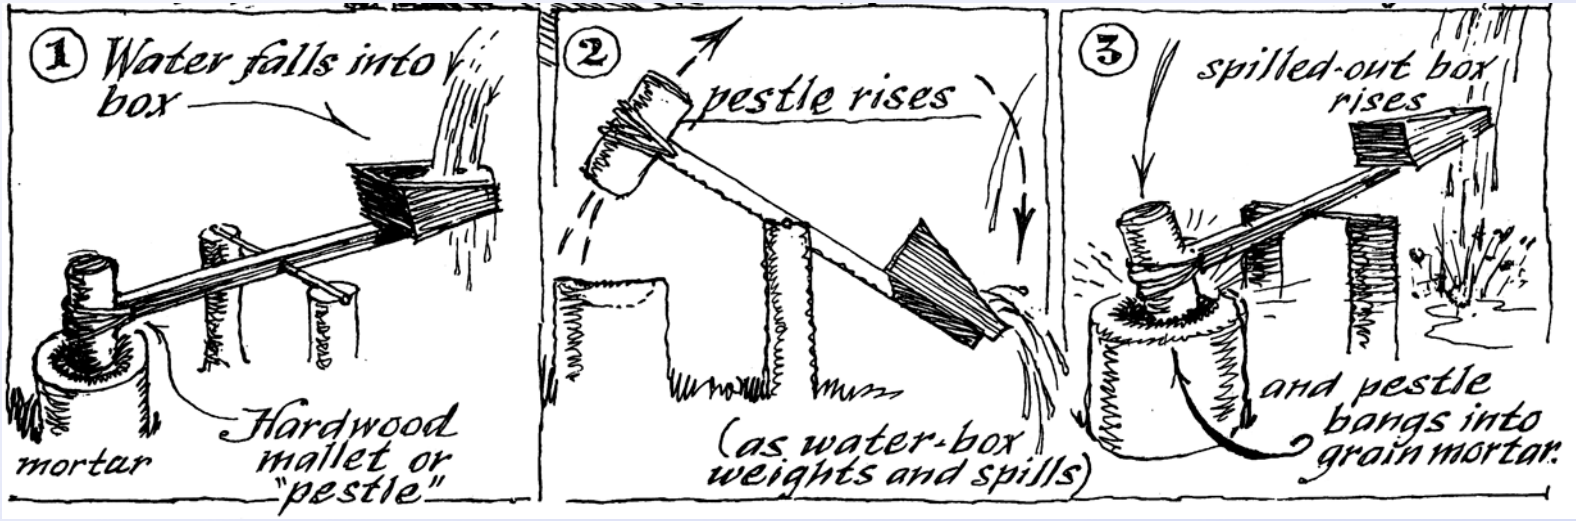

Image: Eric Sloane, *Diary of an early American boy.*

But that mechanism would give *uniformly spaced* clicks:

click for uniform clicks audio

click for actual photon recording

Instead the clicks are *as random as possible* -- they are a "Poisson process." Something about light is discrete and *intrinsically random*.

simulated Poisson process

This lumpy character of light brings out another surprise: When we get down to very few lumps, we see that individual arrivals are random in *space* as well as time:

Experimental data by Albert Rose.

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

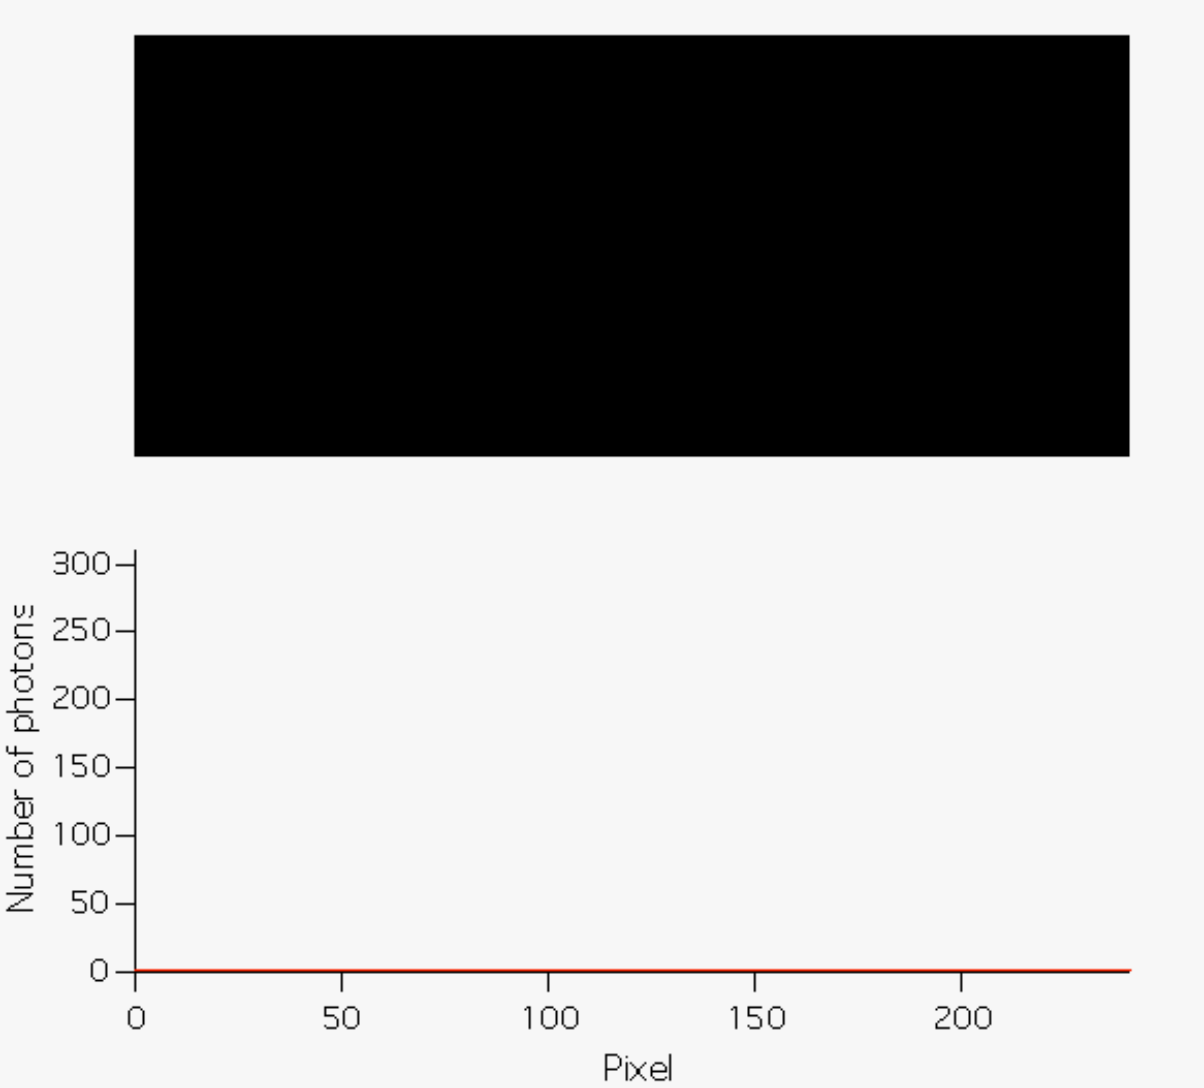

How could anything like that possibly happen at all?

Jean-François Roch
François Treussart
Philippe Grangier
http://www.physique.ens-cachan.fr/old/franges_photon/interference.htm

Light hypothesis, 1

- Light comes in lumps ("photons").
- Each lump has one distinguishing quality: its spectral position.
- These lumps arrive at random, no matter how hard we try to make a steady light. Their *average rate* (probability per second to arrive) corresponds to what we think of as "brightness."
- The light spectrum/color content is the list of these average arrival rates, for each type lump. **That's the meaning of the light spectrum** (color content curve).
- An "image" is a spatial modulation in those average rates.

Light hypothesis, 2

- Light comes in lumps ("photons").
- Each lump has one distinguishing quality: its spectral position.
- These lumps arrive at random, no matter how hard we try to make a steady light. Their *average rate* (probability per second to arrive) corresponds to what we think of as "brightness."
- The light spectrum/color content is the list of these average arrival rates, for each type lump.

- Some single molecule in the photoreceptor cell can flip like a toggle when a photon comes by, absorbing it. Or, the photon can pass right by with no effect. The choice is random.
- The *probability* to be absorbed depends on the type of molecule and the spectral position of the photon. **That's the meaning of the sensitivity curve** (tuning spectrum).
- The three kinds of photoreceptors are each packed with just one of three kinds of sensitive molecule.
- Some cellular apparatus *counts* how many molecules flipped and reports that to the brain. Each receptor type gets reported separately.

So there, you go - answers to all those vexing questions. But... that's a lot of new and crazy ideas! What other kind of experiments could confirm (or demolish) such a story?

Objection

"Even if it's "true," surely this photon stuff has got nothing to do with our eyes: Each photon carries an inconceivably small energy."

Maybe, like Calvin, we are taking a fundamental idea and naively shoehorning it onto an application for which it's irrelevant!

Well... but the evolutionary payoff for good night vision is huge. Nature will certainly have tried very hard to get all the way to the absolute limit of sensitivity... which is *one photon*. (There's nothing between 0 and 1!) Maybe we should look into this unlikely-sounding idea.

By the way, Calvin does have a key idea! Let's follow his lead.

Uh-oh, 3

As before: First nail the case. Then move to the juicy and societally relevant applications.

As before: Quantitative, falsifiable prediction is good.

But... “How can we make a quantitative test of a theory that freely admits that every trial will be different? **Doesn’t Science demand repeatability?**”

It’s a key point. *Sometimes your prediction is statistical*. You *can* test such predictions, but not with a single trial. As you make more trials, you gain greater conviction that the hypothesis is right (or wrong).

The Light Hypothesis, and a simple simulation, lets us predict mathematically *how much randomness* to expect, and compare to the performance of real eyes.

A weird kind of prediction

Suppose that our eyes really could respond to individual photon absorption events, crazy though that may seem. Then if we place a photoreceptor cell in total darkness, and present very dim flashes of light, the Light Hypothesis predicts there must be some *unavoidable randomness* in the response, stemming from the very character of light.

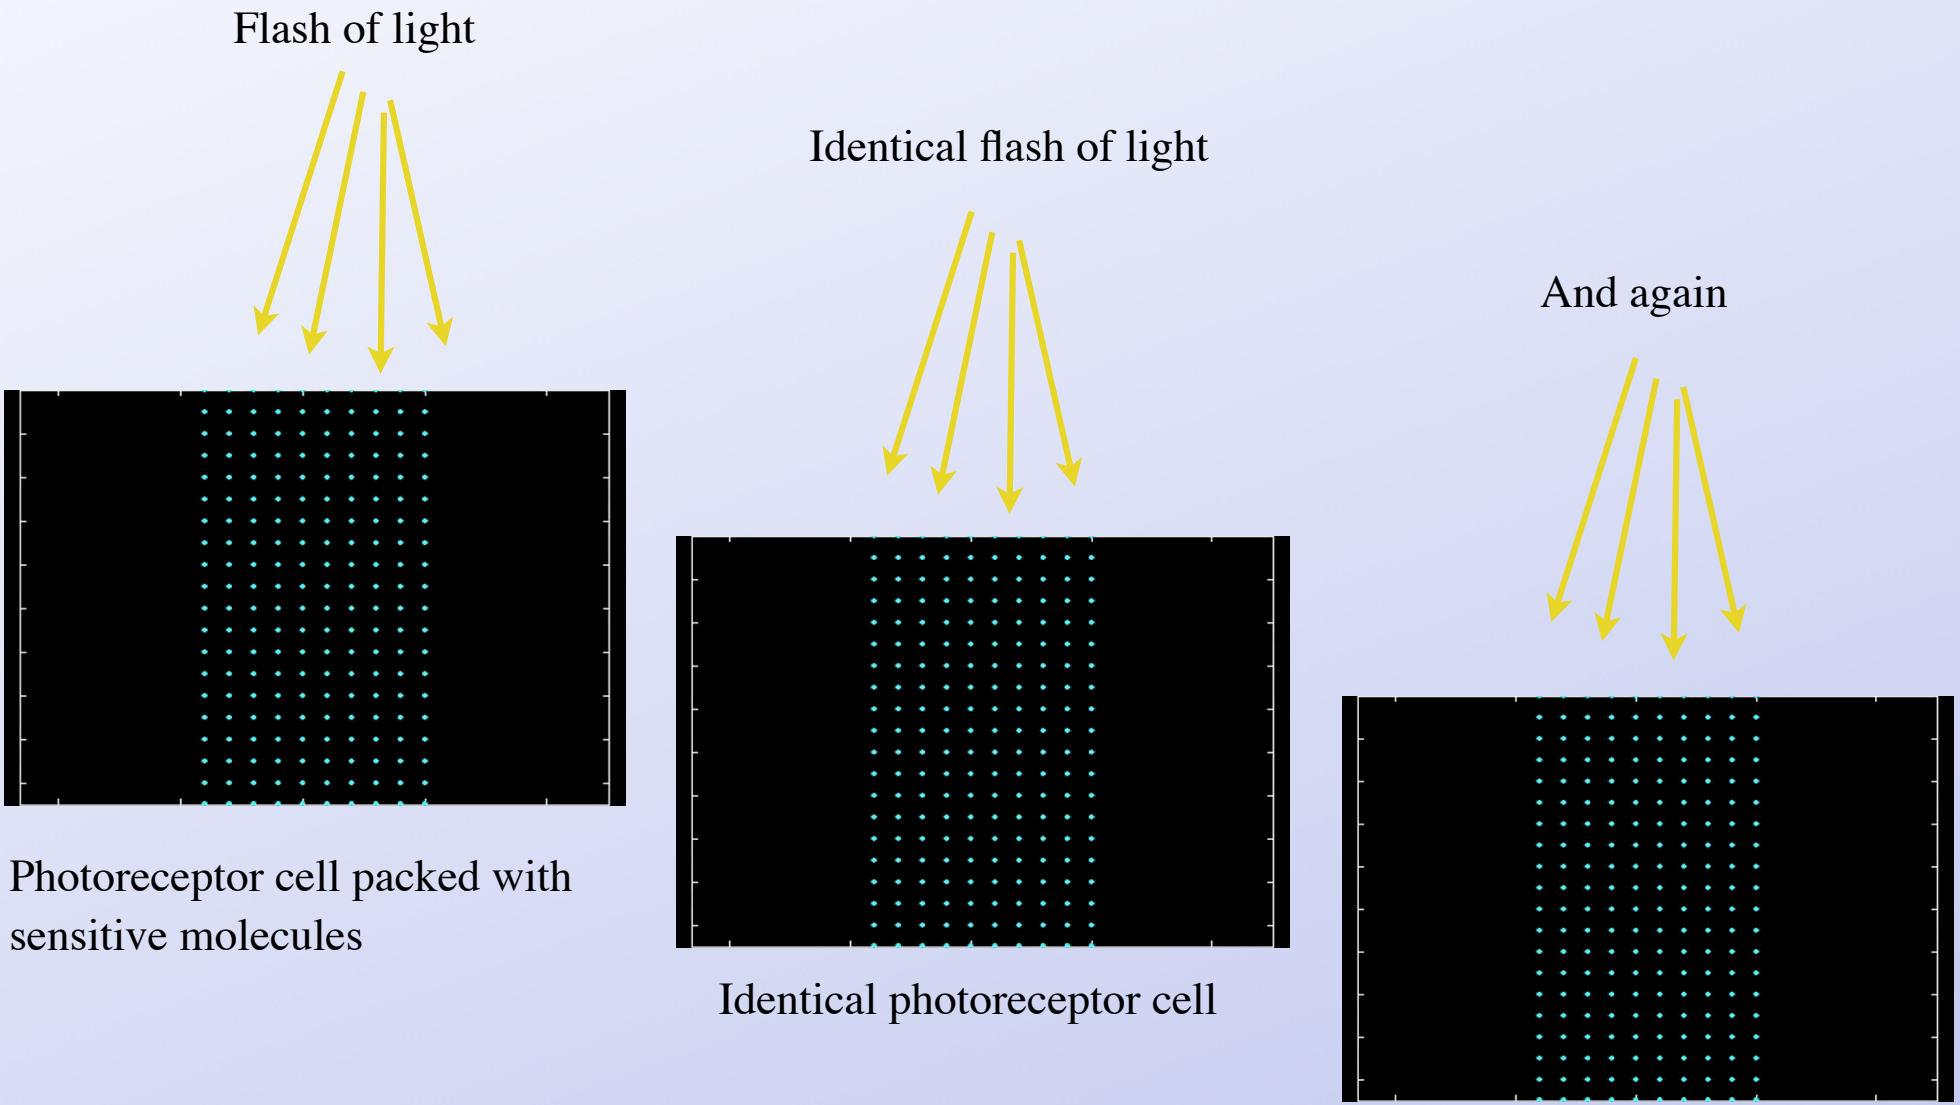

Image from King-Wai Yau. Data and sketch from Fred Rieke.

Receptors respond to single photons

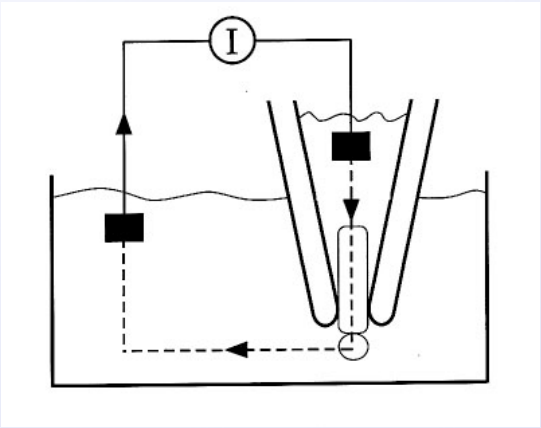

An individual rod or cone cell's response can be measured by gently sucking its outer segment into a pipette electrode and stimulating it with a flash of light (green).

Inset: Flashes of light give rise to discrete current blips. Most observed blips fall into clearly separated categories: Here we see four 1-pA blips, two 2-pA blips (and one that's hard to classify).

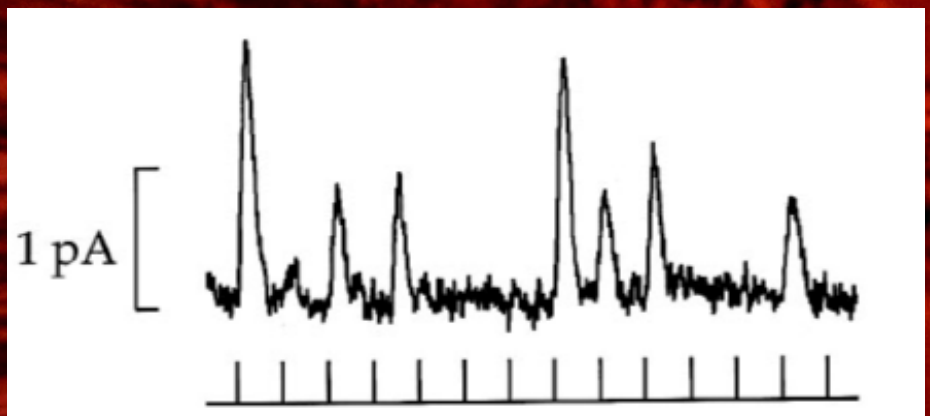

Test a quantitative prediction

It's tempting to imagine that each blip is the response to a single photon absorption. But we cannot relax till we've checked the detailed prediction of the theory: Identical flashes of light should make a *distribution* of blip types that agrees with the Light Hypothesis.
We also get a hard prediction for how that distribution should *change* when we supply a stronger flash.

If those predictions pan out, then we'll have strong evidence not only that the Light Hypothesis is right, but also that all this matters for our vision.

OK: looks promising. Baylor et al even confirmed that the average number of absorptions varies linearly with the flash strength.

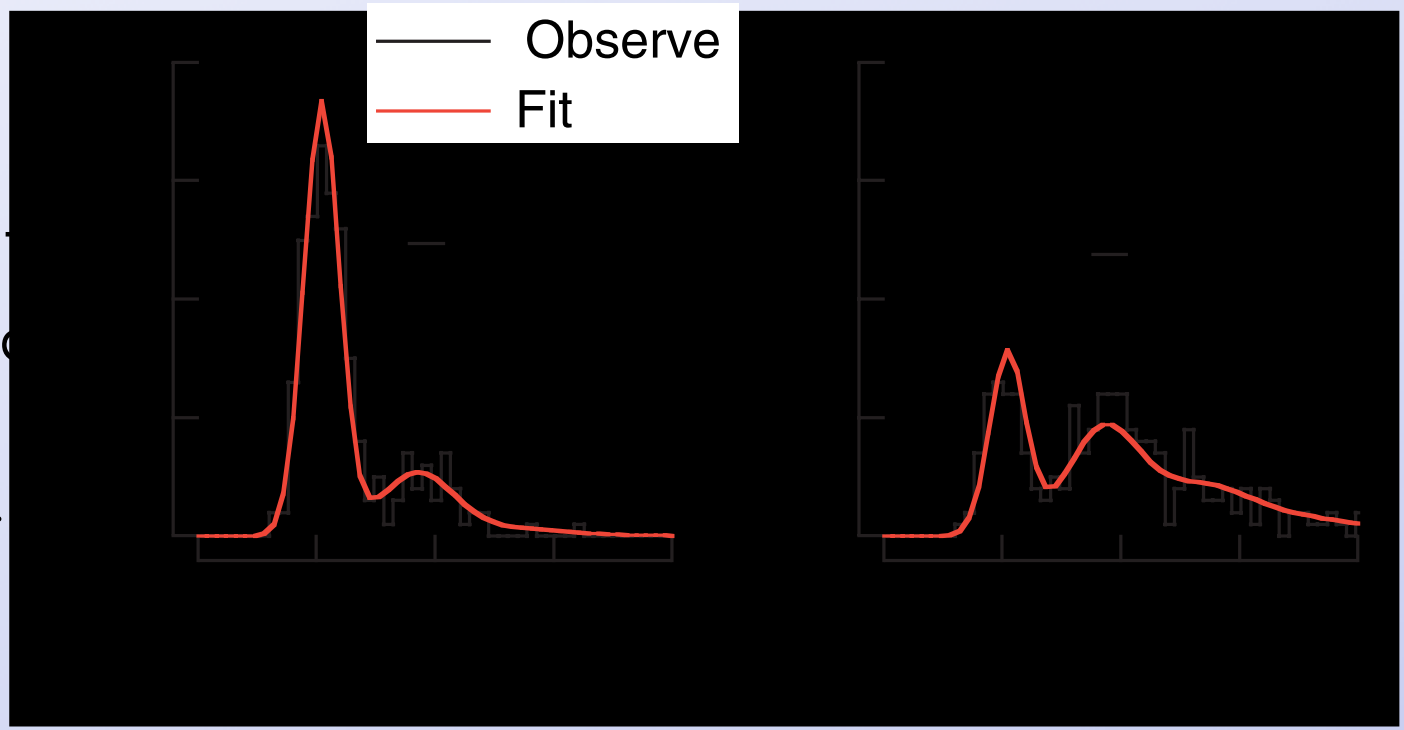

From F. Rieke, in *The Senses: A Comprehensive Reference*; experimental data from G. Field.

[A more detailed measurement]

H Barlow** proposed, and Barbara Sakitt then performed, a classic experiment: She measured a function of ***two*** variables, the probability of a subject assigning each subjective rating to the strength of a flash as a function of that flash's nominal strength. That gives a rich dataset:

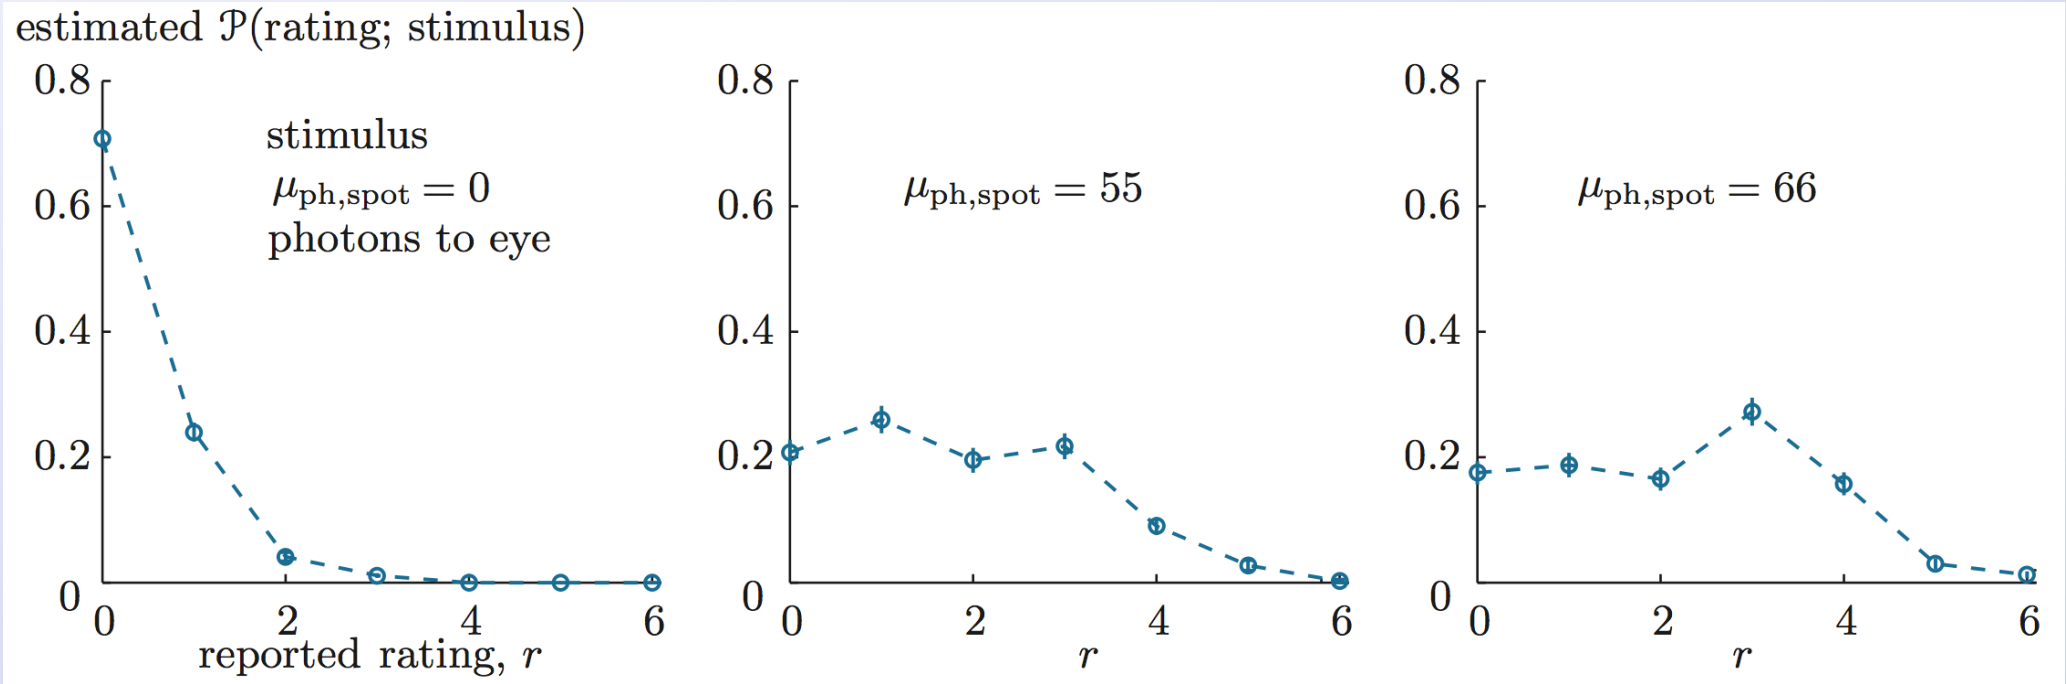

B. Sakitt, ***J Physiol*** (1972).

[Absurdly simple model]

It may be crazy-ambitious to attempt to model human psychophysics, corrupted as it is by all the complexity of the conscious and unconscious mind, with any ultra-simplified physical model – but let's try. Barlow proposed a model involving just two modules:

Poisson processes thinned by various losses

Assume no random losses, just pool signals and apply network thresholds.

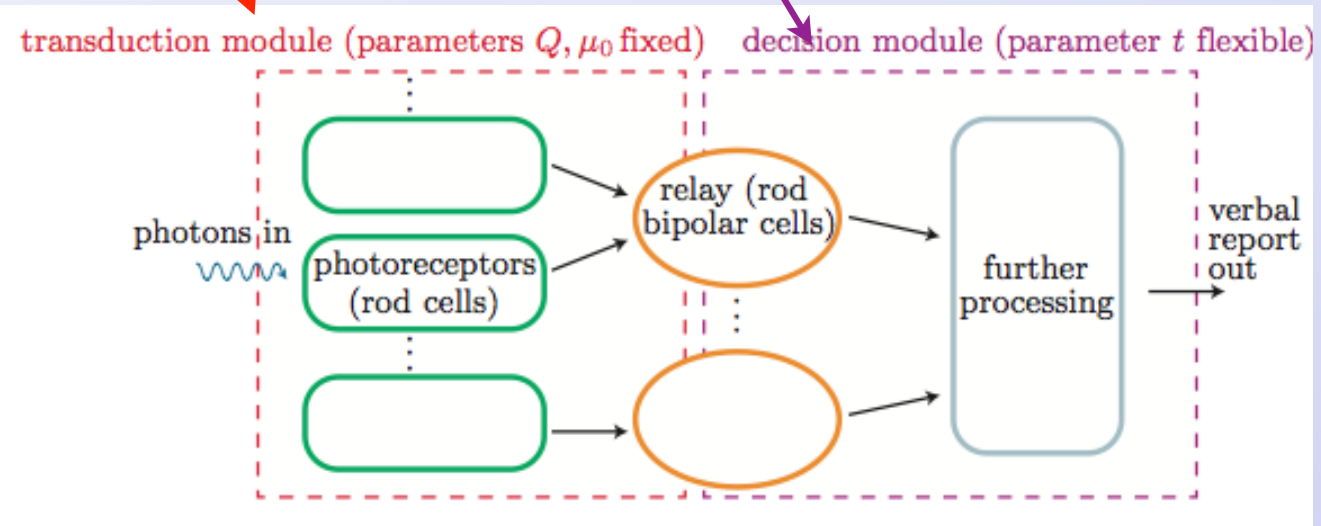

The model was ahead of its time – too many parameters were not yet measured. But today, we know everything about the transduction module from single-cell measurements. And the decision module has very few parameters. And Sakitt's experiment has been redone with modern apparatus, yielding a rich dataset.

So the model is highly falsifiable. *Can such a simple picture really succeed?*

From P. Nelson, ***From Photon to Neuron: Light, Imaging, Vision*** (Princeton, 2017).

[Detailed measurement meets theory]

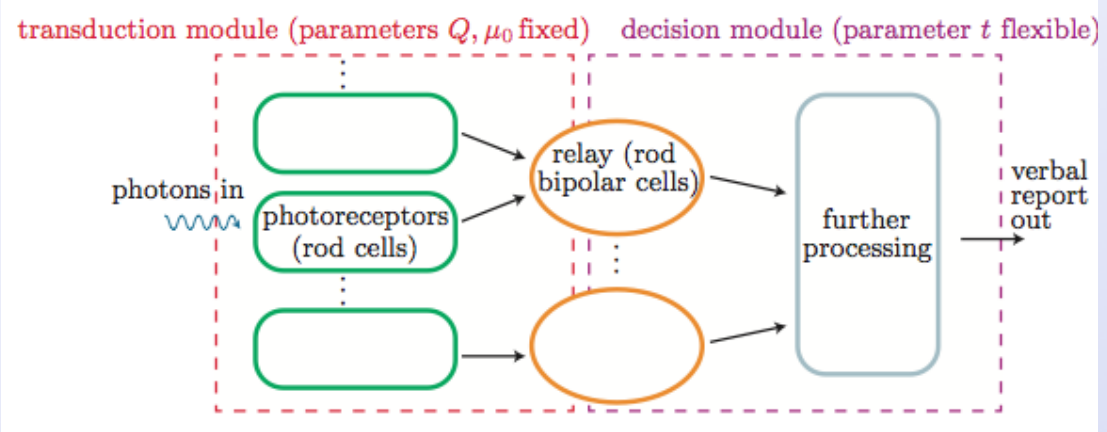

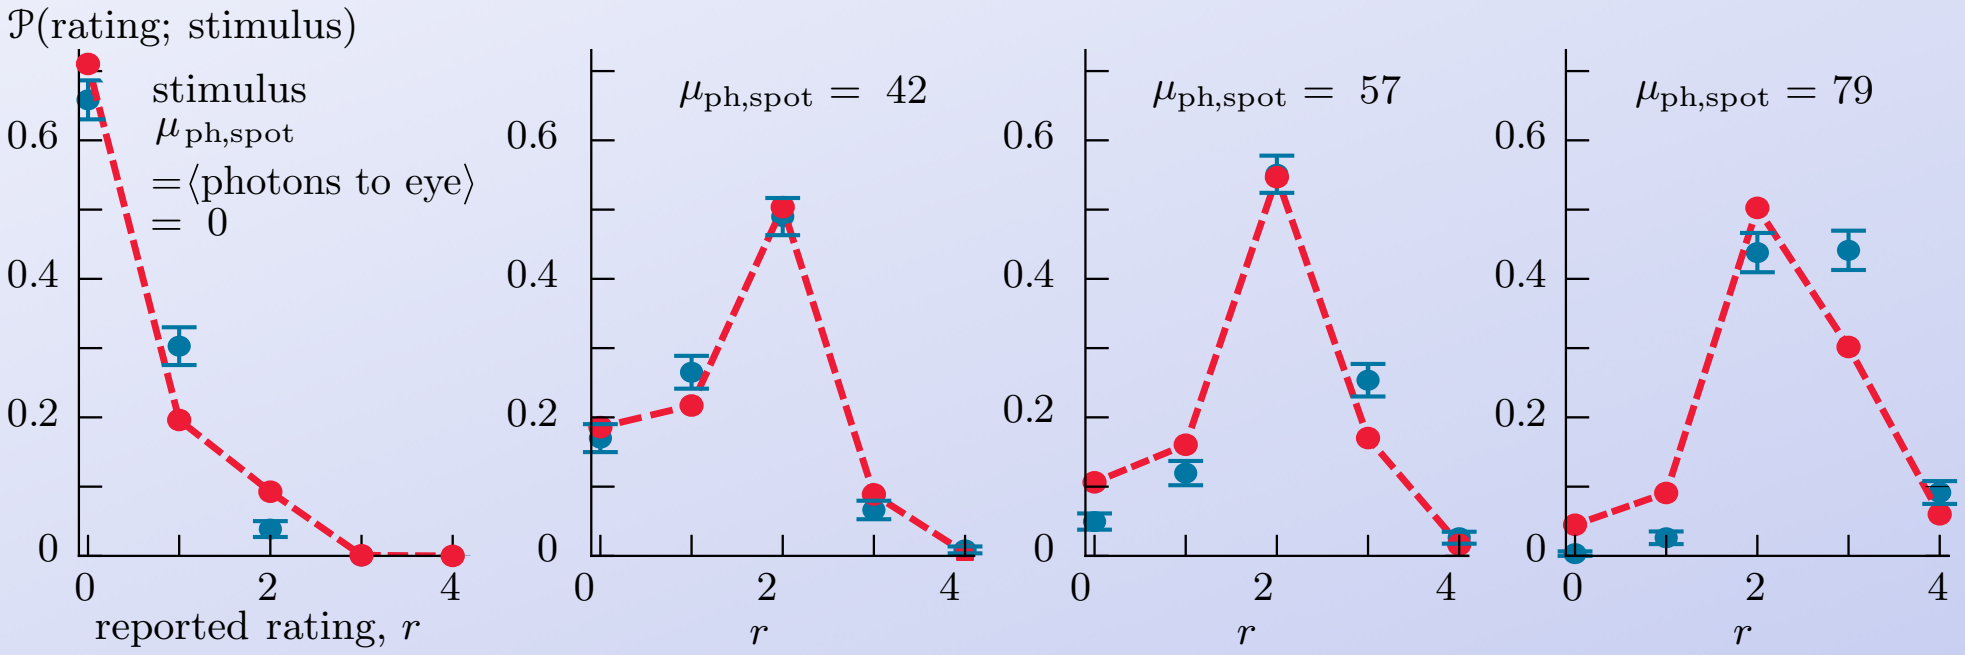

From P. Nelson, ***Physical Biology*** (2016). Data courtesy Heidi Hofer; see also Koenig and Hofer , ***J. Vis.*** (2011).

Superhuman vision revisited

Baylor et al. also varied the color and measured directly the sensitivity curves of over a hundred primate cone cells. They confirmed that there were three classes of cones, distinguished by their sensitivity curves. (Remember that Young had guessed this in 1802!!!) Then they used the curves to make the predictions of color matching discussed earlier.

Great -- we tied up some loose ends in Part I. Is that *all?*

No; we get more. Once we really believe the lumpy nature of light, we can see how to make another big step forward.

Any microscope can give you "superhuman vision." But even the most expensive light microscope can't resolve objects closer than a certain minimal distance (the wavelength of light), again because of that randomness business.

Sadly, nearly all the key machinery inside cells is smaller than this "diffraction barrier."

So superhuman isn't enough. We'd like
superresolution microscopy.

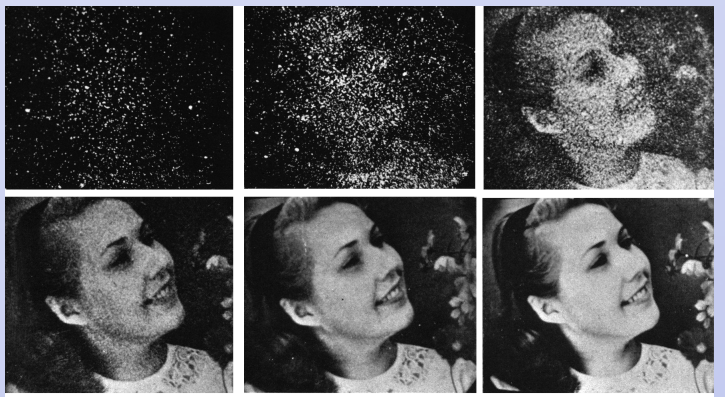

Superhuman 3: Beyond the diffraction limit

For example, how does one measure the steps taken by a molecular motor using visible light? The diffraction-limited spot is at least 200 nm wide!

The key points are to realize that

- Although we cannot resolve *two* spots closer than this, sometimes all we want is to detect *motion* of *one* spot.
- Although the spot may be smeared out by randomness, nevertheless we can find its *center*, using easy statistics.
- If we collect enough photons, then we can find that center with very high accuracy.

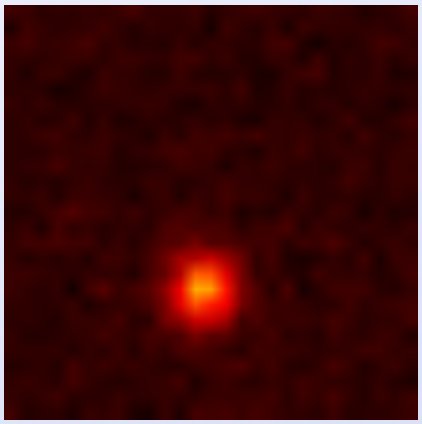

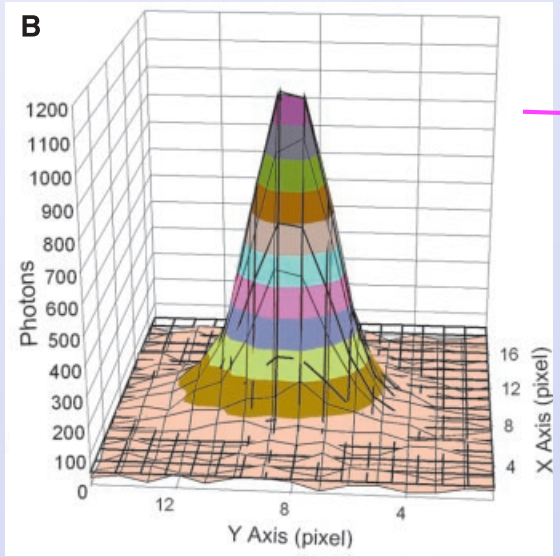

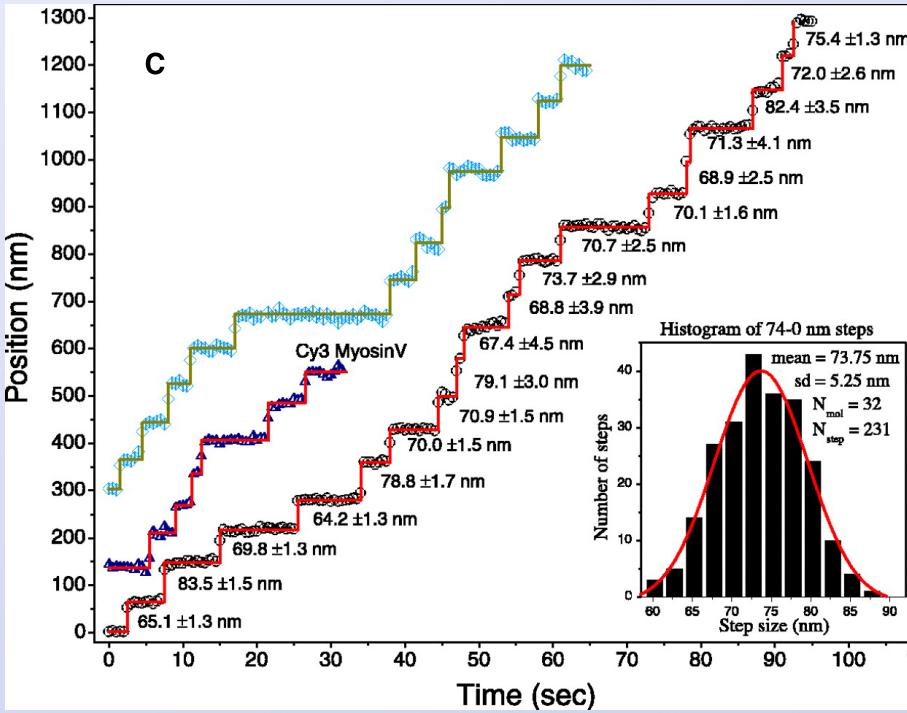

Fluorescence Imaging at One Nanometer Accuracy...

A Yildiz, J N Forkey, S A McKinney, T Ha, Y E Goldman, P R Selvin.
Science (2003) **300**: 2061

F.I.O.N.A.

Superhuman 4

Is that *all?* Is anything *newer* going on?”
Well, usually we want an image, something a lot more structured than one point of light. But if the object consists of points that blink on and off, we can just apply localization to each one and accumulate the results:

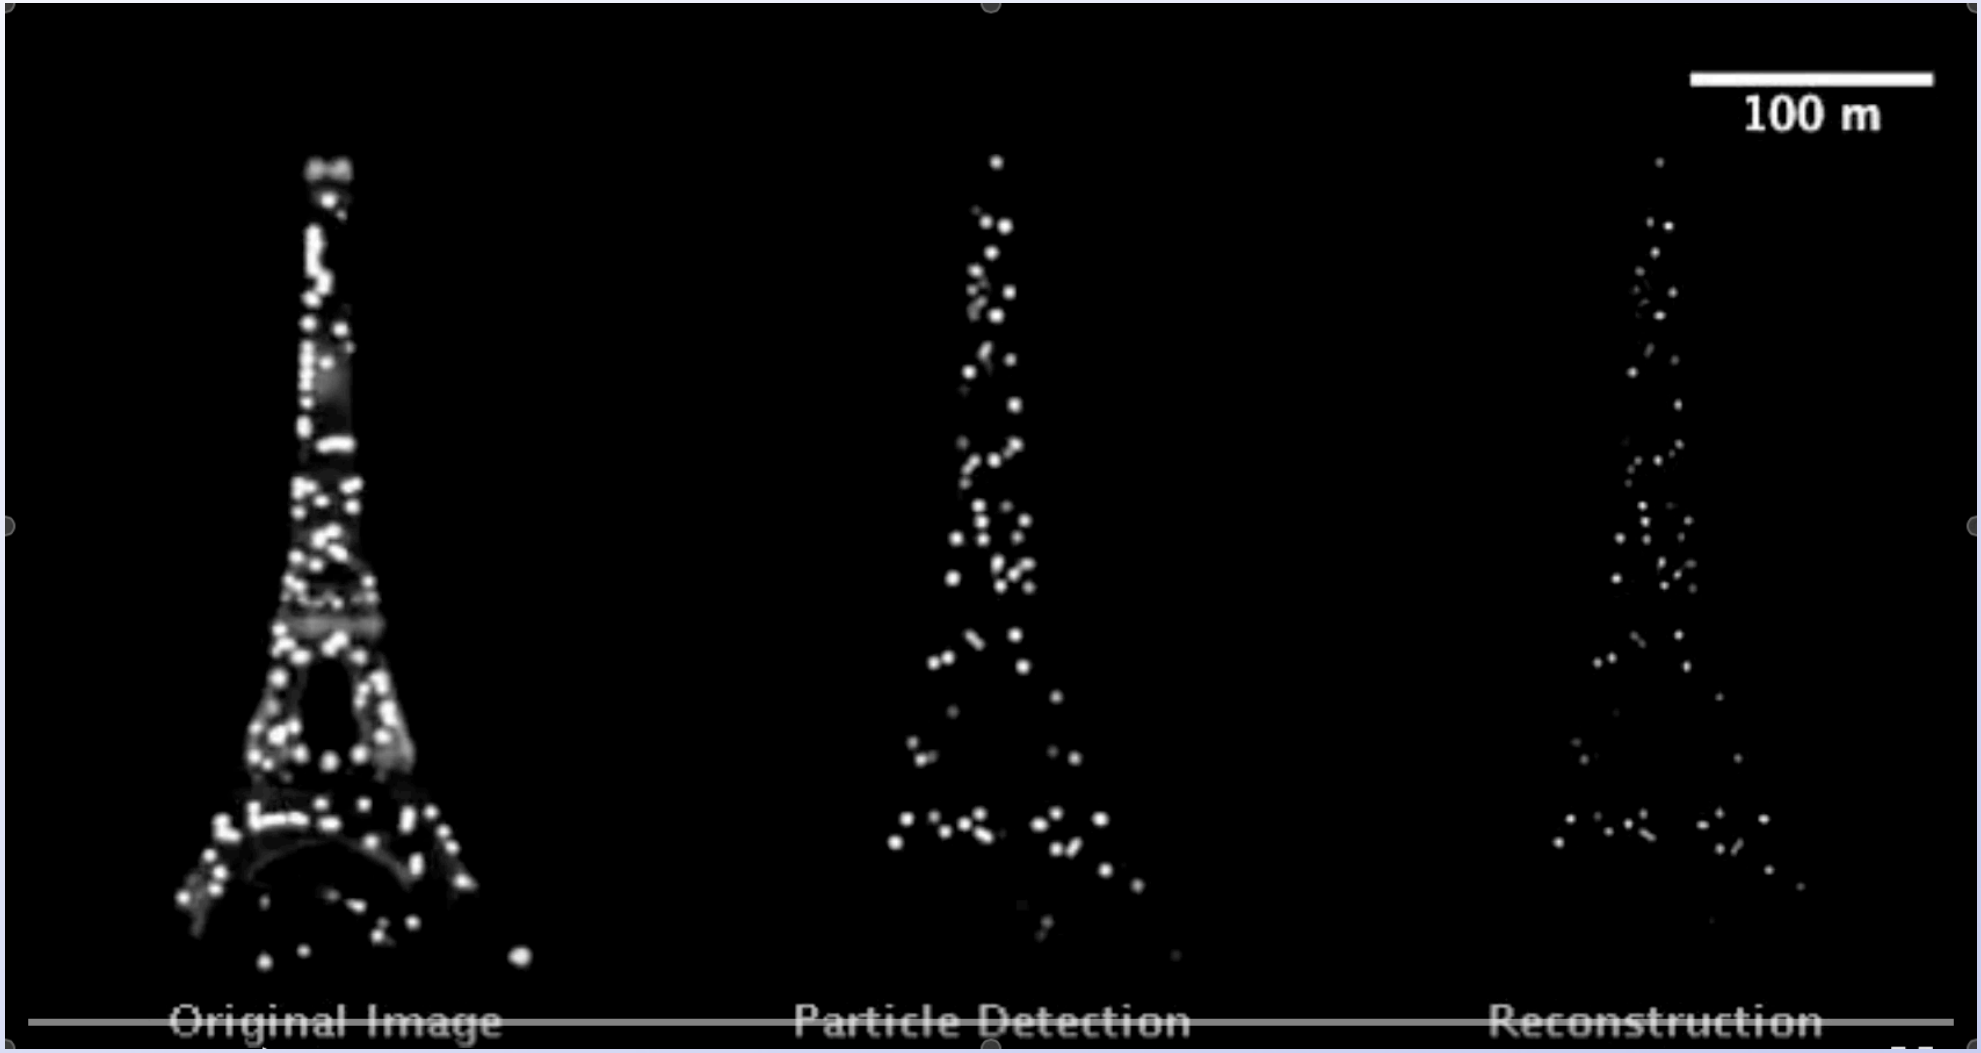

https://www.youtube.com/watch?v=RE70GuMCzww

Superhuman 4

The resulting family of techniques got named **"Method of the year 2008"** by *Nature Methods*:

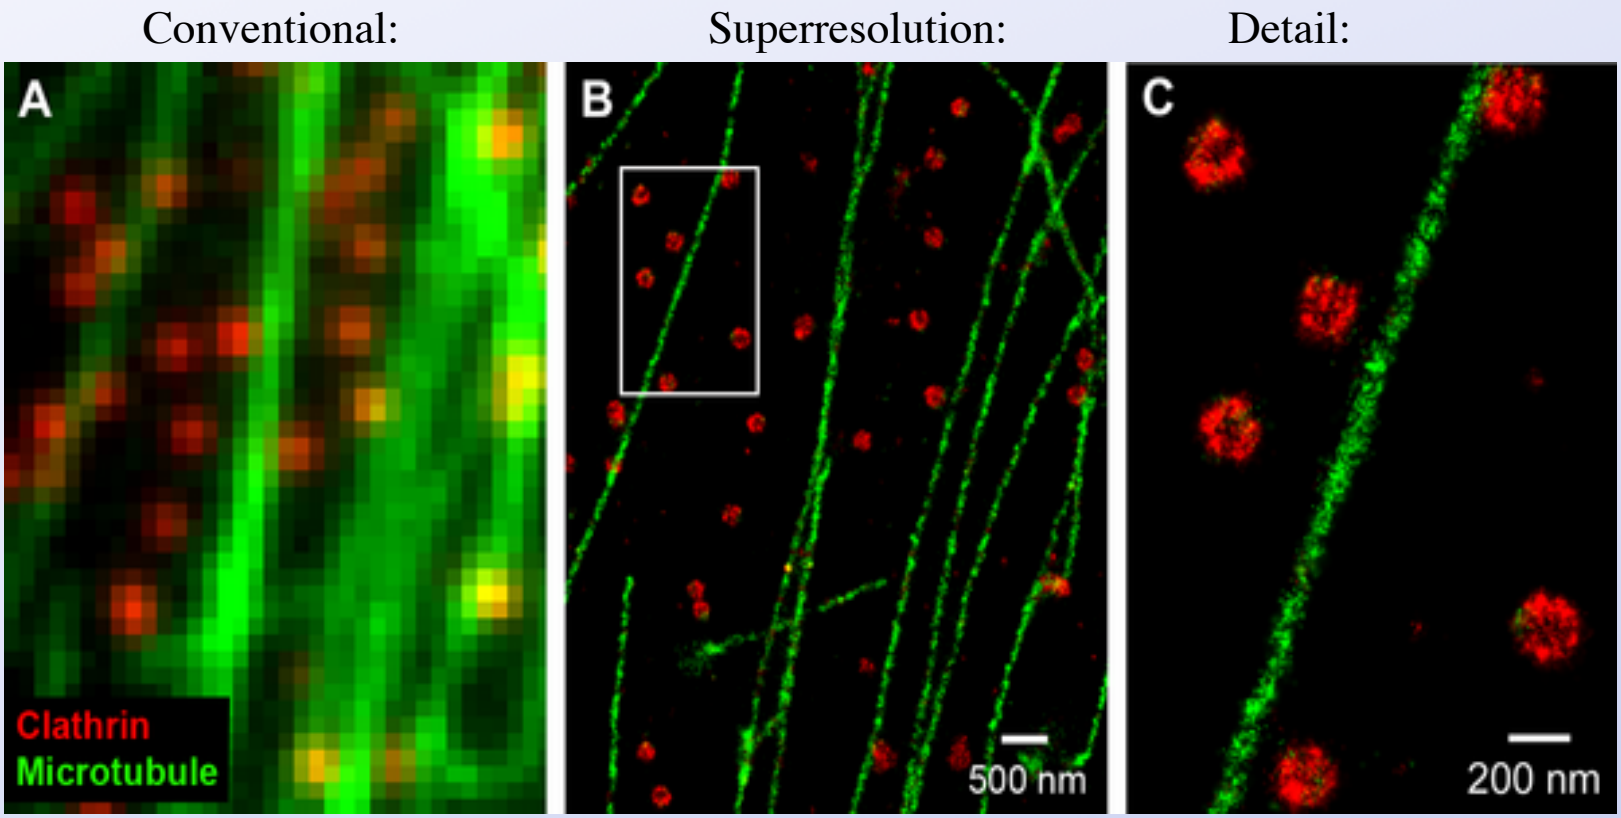

Understanding the lumpy, statistical character of light has led to microscopy methods like PALM, STORM, STED... *That's* spycraft.

Images: Bo Huang, Mark Bates, Xiaowei Zhuang. Annu Rev Biochem (2009) vol. 78 pp. 993-1016.

Wrap part II:

Many clues, including our own vision, led us to a surprising conclusion about light itself.

Without that understanding, we would not have been able to imagine how to break the resolution barrier.

That's spycraft.

OK, I admit I haven't told you everything about light. But I *have* told you most of what you need to understand a lot of biological physics!

Wrap

How did these stories differ from occult voodoo?

- We started out with some real reality.
- It led to a fruitful paradox.
- We discussed how to test a theory--and why. It takes effort. (It also takes money.) Much of my course is dedicated to giving students the skills and frameworks needed for this step.
- We learned some lessons that translated into methods that have paid off in unexpected ways.

Why I like biophysics

What did you learn in this talk?

Well, strictly speaking... *nothing*. I believe you don't learn till you do things yourself. But many of the most important calculations in biophysics really are things you can do for yourself, using modern tools unavailable to the Ancients. *I like that*.

It turned out we could not understand vision at all without some top-drawer ideas from fundamental physics (like quantum theory).

When properly fleshed out, the discussion also makes use of probability theory, biochemistry, evolution... (plus a little information theory, physiology, kinetic theory, physical chemistry, cell biology, neuroscience...). *I like that too*.

Well, I've done my best to share with you my conviction that Biophysics is a unified whole, best approached without artificial, outdated discipline boundaries. That's the Deep Program.

“Learn from science that you must doubt the experts.”
-- Richard Feynman

“I've learned how to live without knowing… I think my life is fuller because I don't know what I'm doing.” -- Richard Feynman

Read More

I realize this was a whirlwind tour. You will enjoy reading…

Light:

R. P. Feynman, ***QED: The strange theory of light and matter.***

Vision:

David Hubel, *Eye, brain, and vision*, also available free online:

http://hubel.med.harvard.edu/index.html .

Sean Carroll, *Making of the fittest.*

Jeremy Nathans lectures: http://ibioseminars.org/nathans/nathans1a.shtml .

For a deeper dive into light, imaging, and vision, I wrote:
From photon to neuron (Princeton, 2017).

Thanks

University of Pennsylvania NSF PHYS NSF NSEC

These slides are available at www.physics.upenn.edu/~pcn