Physicists Take a Hard Look at Vision

To these cross-disciplinary migrants, how the eye extracts a clear signal from the sea of noise generated by light and nerve impulses is an irresistible puzzle

At first glance, physicist Rob de Ruyter's lab at the NEC Research Institute in Princeton is just the kind of working environment you would expect for a physicist: crisscrossing wires and cables, racks of electronics, a vibration-damped table. But in the background of this scientific still life are a few details that don't quite fit. Look closely at a tangle of equipment and you'll spot a fly, perched on a pedestal, immobilized by wax and penetrated by tiny electrodes, watching a television screen that features a display of moving vertical lines. This is de Ruyter's approach to understanding a highly atypical subject for a physicist: the visual system.

For the 3 or 4 days that the fly survives, de Ruyter taps into the firing of individual neurons in the creature's brain. He's particularly interested in the H1s, two of a handful of neurons that register horizontal motion and enable the fly to keep track of its direction of flight. De Ruyter has been measuring how these neurons fire their quick pulses in harmony with the racing lines on the screen. They do so, he has found, with a reliability that an engineer could only marvel at. In the jargon, the fly's visual system is "optimized." Says de Ruyter's NEC colleague William Bialek, "What is amazing about the fly's performance is...that the precision of one neuron's response is just equal to the physical limits" set by the nature of light, the noise in neuronal connections, and the limited capacity of the pathways connecting eye to brain.

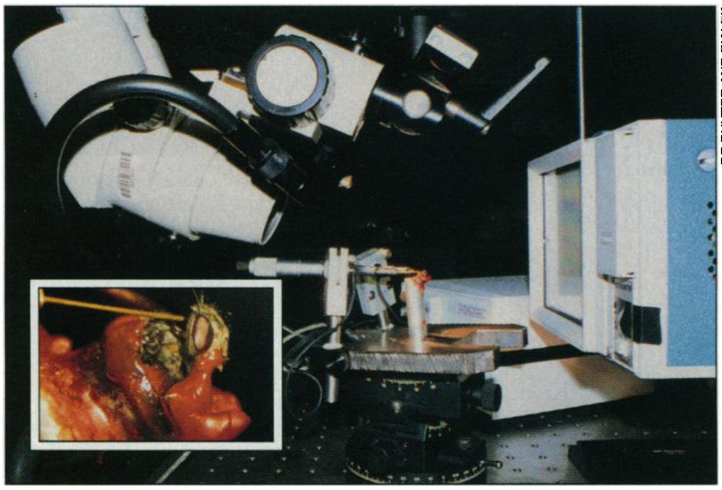

DE RUYTER AND BIALEK

Glued to the tube. Embedded in orange wax, a fly watches TV while electrodes monitor how faithfully neurons respond to the video display.

To de Ruyter, Bialek, and a growing band of other physicists, how an animal's eye manages to deliver information of such high quality to the brain is a puzzle—and the best way to think about it is to begin with the physics of the problem. One of the biggest challenges for the fly, as for other creatures, is dealing with a flood of data corrupted by the random "noise" resulting from the nature of light and the properties of the eye. How does the retina of a fly or any other animal avoid losing the signal in the noise as it distills the incoming data torrent—billions of photons showering the retina every second—to a tiny electrical trickle that can fit down the threads of the optic nerve? Or, as Harvard University physicist Marcus Meister phrases the question, "What is the language the retina speaks when it conveys a visual image to the brain?"

That's a problem of information and coding, topics familiar in computer science. It also entails pulling signals out of noisy backgrounds, a problem physicists battle in everything from the search for quarks to the study of cosmic microwaves. And so, say Meister and his physicist colleagues, understanding vision is a physics problem as well as a biology problem. At last June's Princeton Lectures on Biophysics, sponsored by NEC, physicists and like-minded biologists came together with a group of 50 students to explore the state of their interdisciplinary efforts. The message to the students was clear: "We don't need physicists to go into biology and do biology," as they have in the past, says Rockefeller University physicist Joseph Atick. "We need physicists to use a physical approach." The hope of these cross-disciplinary migrants and their like-minded biologist colleagues is that the principles they uncover by studying flies, salamanders, and other creatures will constitute a universal code for visual processing in all animals, just as there's a single array of particles and forces underlying the behavior of all matter.

To find out how the retina has optimized its response to the noise and data overload, some physicists, including de Ruyter, are collaborating with biologists to monitor how predictably individual neurons in the retina and the brain respond to different stimuli, with the aim of pinning down the sources of noise and figuring out how the system contends with it. Others are looking at many neurons at once in the retina, shining in patterns of light and trying to correlate them with the patterns of electrical pulses that result. In the meantime, theorists are starting to develop mathematical models of information coding in the retina, trying to describe how it might walk the fine line between overloading the brain with information and losing data in the noise.

The physicists' working assumption—that the eye does the best job possible within physical limits—actually originated with a biologist, Horace Barlow of Cambridge University. In the 1950s and 1960s, working on the multi-lens compound eyes of insects, Barlow showed that they have evolved just the way an astute instrument designer would have made them. For a given eye size, a larger number of smaller lenses yields a visual image consisting of finer "pixels," just like those in a computer screen or any digital image, and hence produces finer detail. But as lenses get smaller, the light passing through each lens feels more distortion from diffraction—the process that causes waves rounding obstacles or passing through small openings, such as a lens, to fan out. At a certain point, which varies depending on the size of the eye, diffraction overwhelms the advantage to be gained by having more lenses. Based on these simple physical principles, Barlow calculated the perfect compromise size for a given eye area—and then measured just that optimal size in insect after insect.

Data distillation. Encouraged by the success of this prediction based on simple physical analysis, physicists are looking for deeper manifestations of optimization. They are hoping this approach will guide them to an understanding of how the retina sorts the information that pours into the eyes at a torrential rate. "You receive about 100 megabytes of data each second," says Rockefeller physicist Atick. "Even computers of today's power couldn't handle it." Before the data continue down the optic nerve, the stream has to be reduced by a factor of at least 20.

To compress the data so dramatically while retaining enough information for the brain to reconstruct the world, the retina (and later the brain) must do some exceptionally smart processing.

One way to shrink the data flow would be to divide a scene into pixels and average all the data within a pixel to a single value. "That would be inefficient," says Bialek. There's no way the eye could do so and preserve any reasonable resolution. A smarter way, say the vision researchers, would be to separate real information from what they call "redundancy," which is the part that can be predicted once you understand an overall pattern. As a simple example, in most natural scenes bright or dark spots are strongly correlated: Next to one bright or dark spot the probability of finding a similar spot is high. Such predictable features can be eliminated without affecting perception; most of the information needed to define a scene is concentrated in the borders of bright and dark regions. At least in theory, says Atick, "it is only the deviations that [need to be] signaled to the brain."

To see if the eye really has optimized its coding strategy by weeding out such redundancy, Harvard's Meister has developed a technique for measuring redundancy at different levels in the visual system of a salamander. He dissects the salamander eye, keeping together the photoreceptors and neurons of the retina. Then he lays the retina into an array of 61 electrodes, each of which can register signals from up to 100 neurons at once, and exposes it to a flashing light.

By looking at so many cells at the same time, Meister can watch for correlations in these firing patterns and calculate the amount of redundancy to see if it is indeed being reduced. Is each neuron independent or can the firing of one be predicted from its neighbor? Is the representation of information efficient? To answer those questions, Meister is planning to try to figure out the coding scheme that connects light signals with electrical pulses. Ultimately, he says, he'd like to be able to work backwards to "decode" the patterns of spikes and reconstruct the visual image. "This is what the animal's brain has to ask," he says. "Given I'm seeing this pattern of spikes, what is going on in the outside world?"

But reducing redundancy can only be taken so far. The patterns and regularities that give rise to redundancy are also essential for distinguishing a real signal from the outside world from optical and electrical static, says Bialek. "These sorts of [redundant] structures are, presumably, what distinguish the real world from an array of random pixels, and hence what distinguish signal from noise," he says. Rockefeller's Atick agrees. "Statistical structure is what enables you to see." Before redundancy is eliminated, he and his colleagues agree, the eye has to exploit it somehow to combat noise.

Dark noise. There's plenty of noise to be sifted out. As the Princeton conferees repeatedly pointed out, every step in the eye, from a photoreceptor's response to a photon of light to the transmission of a signal down the optic nerve, involves the opening of only a few channels in a cell membrane or the release of a few neurotransmitter molecules, and processes involving only a few molecules are always somewhat random. The jostling of molecules in the photoreceptors and neurons might not seem "loud," but from the point of view of an individual photon, it becomes a raging storm of "dark noise"—noise originating not from the outside world of light but from the dark inner connections in the retina and brain. Neurobiologist Charles Stevens of the Salk Institute, for example, has measured the reliability of connections between pairs of nerve cells in culture, and the poor quality of any individual connection surprised the conferees. "If you stimulate Cell 1," he says, "Cell 2 gets a signal [only] about half the time."

The nature of light itself adds another source of noise, especially at very low light levels. Because light is grainy, consisting of discrete photons, it makes an uneven impression on the photoreceptors of the eye. In a short enough period, one receptor might capture two photons, another might pick up three, a third none at all. The uncertainty introduced by the graininess is known as "photon shot noise." When many photons are pouring into the retina, such noise has a smaller effect, but experiments dating from the 1940s show that people in dark rooms can detect individual photons, and Dennis Baylor of Stanford and colleagues later measured single-photon responses in individual photoreceptors.

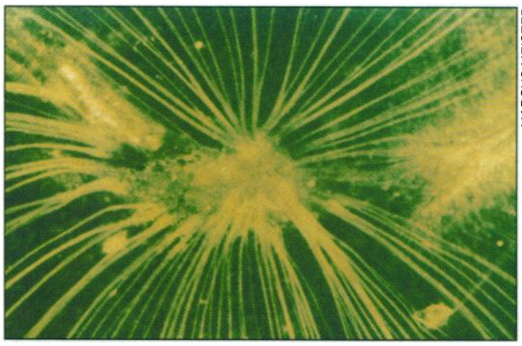

MARCUS MEISTER

Information highways. In a salamander retina, axons converge on the optic nerve.

With so much noise threatening to swamp these subtle signals, how can the retina cut out redundancy without losing the signals altogether? Atick says that a single coding scheme may explain how the visual system, beginning with the retina, accomplishes the feat. In his theoretical computer model of how the process might work, the retina first relies on redundancy to cut out the noise. It does so by pooling data from different photoreceptors. The redundancy marks real signal, which gets reinforced, while the noise gets weeded out. "You build up some kind of consensus," he says. "It's like making one person represent groups of five—you agree on something and send a spokesman," he says. But this same process of pooling also reduces the data rate, in effect pruning the redundancy. This process continues, he points out, beyond the retina and into the deeper layers of the visual pathway in the brain.

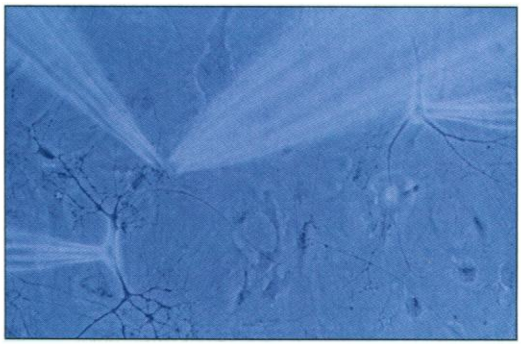

CHARLES STEVENS

Listening in. Tiny electrodes record the cell-to-cell chatter of neurons in a culture medium.

Some preliminary evidence supporting Atick's scheme comes from experimenters who have shown that in a variety of creatures, signals from large numbers of photoreceptors do get pooled in deeper layers of the retina. And, in keeping with the idea that redundancy helps distinguish signal from noise, the pooling often increases in dimmer light, where noise becomes a larger factor.

Even if Atick's scheme or something like it can successfully explain common features of information coding in the retina of every animal, the concept of optimization predicts that animals should have evolved variants of this processing strategy, depending on the physical constraints of their environment or lifestyle. Many insects, for example, probably trade off some noise-reducing ability for speed, says biologist Simon Laughlin of Cambridge, because they have to gather information literally on the wing. Fast-moving insects such as the common house fly, he's found, use a special mechanism to speed up the transmission of signals across each neuron: an unusual set of electrically conductive channels in their nerve membranes that lower the membranes' electrical resistance.

A lower resistance allows voltage across the membrane to change more rapidly. This increases the neuron's speed of response but also makes it more vulnerable to photon shot noise because it responds to a smaller number of photons. Slower insects, such as the crane fly, don't have this mechanism, he says. They ignore rapid variations and use the extra time to gather the redundant information that combats noise. "This is [a] way in which the neural circuits are microengineered for specific purposes," says Laughlin.

Results of this sort reassure physicists that they are on the right track in viewing the visual system as an optimal response to physical limits. Says Bialek: "In the 70s, I thought that the central nervous system was a mess, and physicists shouldn't study it...but now people can do these unbelievably quantitative and reproducible experiments." Adds Barlow, "It's easy when one starts looking at biological tissues to see them as hopelessly inefficient." But as people have probed deeper over the past few years, he says, "it gives us a source of perpetual amazement."

–**Faye Flam**

MATERIALS SCIENCE

Magnetic Films Get Sensitivity Training

It's easy to get excited about discoveries in high-profile fields with catchy names like "chaos" or "superconductivity," but many other areas of science promise at least as great an impact on our lives yet don't have nearly the same media glamor. Take giant magnetoresistance. That's an inelegant name for a rather elegant physical effect—a change in the electrical resistance of a material when it is placed in a magnetic field. A community of researchers is hoping to make this jawbreaker a small but important part of everyday life, and a paper in this issue of *Science* moves that goal another step closer to reality.

A group at IBM ADSTAR in San Jose, California, has found a way to create materials that show giant magnetoresistance (GMR) at low magnetic fields (see page 1021). The discovery, by Todd Hylton, Kevin Coffey, Michael Parker, and Kent Howard, could greatly increase the capacity of magnetic data storage devices, such as disk drives, by making it feasible to read and write smaller bits of magnetically coded information. "A lot of people have been trying to do this," says John Barnard, a materials scientist at the University of Alabama in Tuscaloosa. The work is a "proof of principle," he says, because it shows it is possible to create GMR at fields low enough to be practical.

Though it isn't a staple of popular science books or newspaper accounts, GMR has been the hottest area of magnetic materials research almost since its discovery 5 years ago. Part of the reason has been excitement over a new physical phenomenon, but part of it is GMR's commercial promise. In fact, a first cousin of GMR—ordinary magnetoresistance (MR)—is already being used in some high-performance computer data storage devices. Most of the "read" heads in today's disk drives depend on induction—the creation of an electrical current by a changing magnetic field—but induction heads are fast approaching their physical limits, where the bits of data can get no smaller and still be read. MR, an effect created in certain magnetic metals when a magnetic field aligns the spins of the conduction electrons and affects the material's resistance to a current, yields heads that can read smaller bits than conventional heads can.

But MR has limits of its own. The resistance of typical MR materials generally changes by no more than about 4%, and by only 1% to 2% in the practical, mass-produced form used in read heads. If the resistance varied more, the magnetic bits could be made still smaller. So companies like IBM were excited 5 years ago when researchers at the Université de Paris-Sud reported a more dramatic effect, which they called giant magnetoresistance: resistance changes of up to 50% at a temperature of 4.2 K.

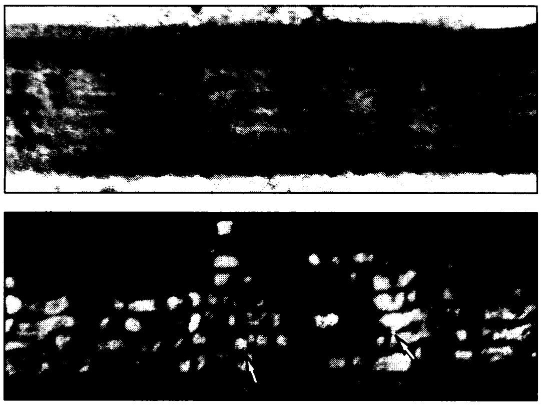

MICHAEL ANDREW PARKER/IBM ADSTAR

Slices of the sandwich. Electron micrographs of the new magnetoresistive material show continuous magnetic layers *(light bands in top image)* that break up *(arrows)* after annealing.

While MR is a "bulk effect" that appears in materials without specific atomic-scale structures, GMR arises in very structured materials consisting of atoms-thick layers of magnetic materials, such as iron or nickel, interspersed with layers of nonmagnetic materials, such as silver. Each of the magnetic layers can be thought of as a little sliver of a bar magnet, with the magnetizations of alternating magnetic layers pointing in opposite directions. When electrons of an electric current pass through one of these layers, half of them find it easy going and half are slowed down, depending on the direction of their spins. When they pass into the next layer, the roles are reversed. The upshot is that none of the electrons has an easy time passing through the entire layered structure, and the resistance is relatively high.

Applying a magnetic field to this layered structure, however, causes the magnetizations to line up, so that half of the electrons get a smooth ride through the whole material. Although the other half finds it even tougher going, the end result is that the total resistance drops. Since the phenomenon was first observed, researchers have made GMR materials whose resistance changes by as much as 65% at room temperature.

The problem is that achieving such dramatic resistance changes takes very high magnetic fields—typically 250 oersteds or more, or at least 500 times Earth's magnetic field—perhaps because the opposite magnetizations of the layers tend to "couple," or lock together, says Hylton. These fields are too high to be practical for magnetic data storage. The IBM ADSTAR group has now managed to create giant magnetoresistance in much smaller magnetic fields—only 5 to 10 oersteds. So far they see a resistance change of only 4% to 6%, but the group hopes to keep improving on it.

To achieve the effect, the researchers created alternating layers of nickel-iron and silver only a few nanometers thick, then "annealed" the material by quickly heating and cooling it. The GMR appears only after annealing, so the group members believe this step holds the key to GMR at low magnetic fields. The reason is still unclear, but the group offers a hypothesis: When the material is heated, the silver starts to diffuse into the nickel-iron layer, working its way into the cracks between the grains of nickel-iron. This creates a collection of disk-like islands of nickel-iron in a three-dimensional silver sea, with the islands in alternate layers having opposing magnetizations. The shape and size of these islands, Hylton says, make it easier for relatively weak magnetic fields to align their magnetizations, perhaps because the magnetic coupling between islands is weaker than it is between intact layers.

It's too soon to know for sure whether the GMR films made by this technique can be adapted to the high-volume production necessary for electronic equipment, but Hylton and his colleagues are optimistic that they can be. At the University of Alabama, Barnard agrees that some variant of the IBM ADSTAR technique—if not this particular material—should find its way to commercial application. One way or another, GMR may overcome the handicap of its name and find a place in the technological spotlight.

–**Robert Pool**