

The probabilistic mind

Human brains evolved to deal with doubt

By Laura Sanders



$$P(h_i|d) = \frac{P(d|h_i) P(h_i)}{\sum_{j \in \Omega} P(d|h_j) P(h_j)}$$

$$P(h_{foe}|d) = \frac{P(d|h_{foe}) P(h_{foe})}{P(d|h_{foe}) P(h_{foe}) + P(d|h_{friend}) P(h_{friend}) + P(d|h_{family}) P(h_{family})}$$

$$P(h_{foe}|d) = \frac{(0.8)(0.5)}{(0.8)(0.5) + (0.3)(0.3) + (0.3)(0.2)}$$

$$P(h_{foe}|d) = \frac{0.4}{0.4 + 0.09 + 0.06}$$

$$P(h_{foe}|d) = 0.73$$

Humans live in a world of uncertainty. A shadowy figure on the sidewalk ahead could be a friend or a mugger. By flooring your car's accelerator, you might beat the train to the intersection, or maybe not. Last week's leftover kung pao chicken could bring another night of gustatory delight or gut agony.

People's paltry senses can't always capture what's real. Luckily, though, the human brain is pretty good at playing the odds. Thanks to the brain's intuitive grasp of probabilities, it can handle imperfect information with aplomb.

"Instead of trying to come up with an answer to a question, the brain tries to come up with a probability that a particular answer is correct," says Alexandre Pouget of the University of Rochester in New York and the University of Geneva in Switzerland. The range of possible outcomes then guides the body's actions.

A probability-based brain offers a huge advantage in an uncertain world. In mere seconds, the brain can solve (or at least offer a good guess for) a problem that would take a computer an eternity to figure out — such as whether to greet the approaching stranger with pepper spray or a hug.

A growing number of studies are illuminating how this certitude-eschewing approach works, and how powerful it can be. Principles of probability, researchers are finding, may guide basic visual abilities, such as estimating the tilt of lines or finding targets hidden amid distractions. Other behaviors, and even simple math, may depend on similar number crunching, some scientists think.

And such advanced statistical reasoning does not require paying attention in math class. New studies suggest that 1-year-olds are already tiny probabilistic machines who, in many situations, assess statistical input and perform optimally with ease.

Studying the guesstimating brain is a new enough endeavor that no one yet knows how people developed such computational abilities. Nor do scientists know the precise brain machinery behind the math.

“We’re going to continue to try to understand these processes,” says Eero Simoncelli, a computational neuroscientist at New York University. “It’s a long road. It’s going to be many decades until all of this gets worked out. But the progress is steady.”

Seeing and believing

When Pouget started studying the brain’s computations two decades ago, nobody thought that humans deal in probabilities, he says. Back then, researchers thought that if you want to catch a baseball, your brain computes the trajectory and spits out an exact answer, telling your body where to move the glove, he says. “Today, we say, ‘No, if you have a baseball flying at you, you compute the probability of where it might be and then you place your hand to maximize the probability that you’re going to catch it.’”

This shift — from studying certitude to probabilities — is largely based on the work of Thomas Bayes, an 18th century English clergyman. A claim is more reliable if initial beliefs are also included in the assessment, Bayes proposed. And these initial beliefs, known as “priors” today, can be updated as more information comes in, narrowing the range of good solutions. At its heart, the concept is simple: Learning from experience leads to better predictions.

Take a doctor faced with a medical mystery. A young boy comes into the office with a slight fever, a headache and joint pain, symptoms that could be caused by a garden-variety cold or the more nefarious Lyme disease. With no additional information, the doctor might as well flip a coin. But armed with key pieces of information — medical school

tidbits and knowledge of whether the boy played in tick-teeming woods, for instance — the physician can come up with a solid diagnosis.

Though the value of considering priors is still a matter of dispute in the statistics community (*SN*: 3/27/10, p. 26), the brain is chock-full of them. And humans constantly mediate a tug-of-war between those priors and current evidence.

By showing how assumptions can lead people astray, a new study highlights how heavily the brain leans on priors. Psychologist and computer scientist Ahna Girshick of the University of California, Berkeley, along with Simoncelli and another colleague, recently asked people to assess the relative tilts of sets of fuzzy lines on a computer screen. The task is like trying to say which way, on average, a handful of dropped toothpicks point.

The volunteers’ performance suggested that they thought the lines were more aligned with the horizontal or vertical axes than they actually were, the team reported in the July *Nature Neuroscience*. That assumption may exist for a very simple reason, says Girshick. “In nature you see these very strong verticals because of trees, and you also see horizon lines and flat surfaces to walk on,” she says. “We’ve all been raised on planet Earth, and there are mathematical structures to the world around us that you can measure.”

What’s more, the researchers could strengthen the misperception by changing the conditions: When the arrays of lines varied more, people showed an even greater bias toward the horizontal or vertical directions. Greater doubt led to a stronger reliance on preconceived ideas.

Scientists don’t yet know what physical hardware in the brain might be performing such Bayesian reasoning, but simulations suggest variations in nerve cell behavior might be responsible for these seemingly complex calculations. “It seems like sophisticated math,” Girshick says. “But it could be quite simple.”

Some nerve cells respond strongly to horizontal or vertical lines, while others don’t give those orientations special attention. “You get this Bayesian-like behavior simply by the fact that you have this nonuniformity in the brain,” Girshick says.

Bomb amid batteries

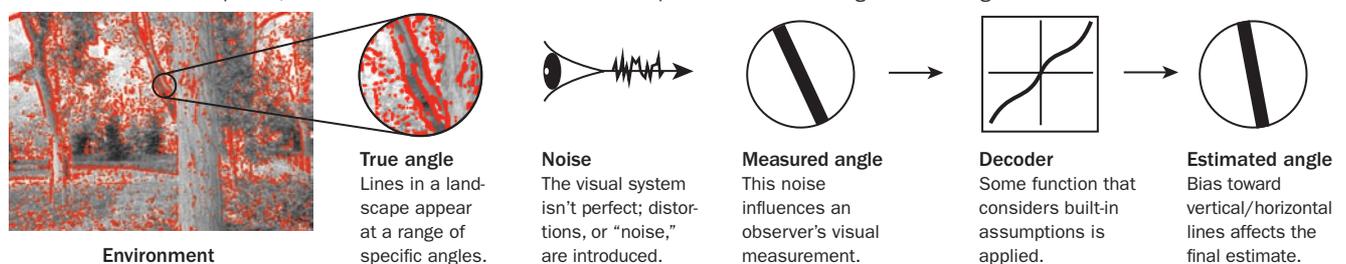
As any airport security screener knows, spotting a bomb among a steady stream of computer batteries, alarm clocks and blow-dryers is notoriously difficult. But in the case of this visual challenge, called a visual search, the Bayesian brain appears to perform surprisingly well.

Given the incomplete information that humans get from their retinas, people’s visual search skills are remarkable, Pouget says.

“Visual search happens absolutely all the time,” he says. “We thought this is exactly the kind of task where a probabilistic approach would be great.” In a recent study, he and his team had participants watch a computer screen for a quick flash of a target — a previously seen line tilted at a particular angle. On the screen, this line was surrounded by distracting objects. Participants reported whether the target was there or not, and how confident they were in the answer.

When the target blended in with the background and the distracters were nice and bright, people grew worse at

Bayesian-based brains Though they don’t yet have a clear idea of how the brain does the calculations needed to compute probabilities based on built-in assumptions, scientists do have some sense of the steps involved in encoding and decoding an environmental stimulus.



recognizing the target, assuming that it was simply not there. But they grew worse in a very particular way. People's behavior closely mirrored what Bayesian math predicted, the team reported in the June issue of *Nature Neuroscience*.

"A visual search starts involving pretty complicated mathematics," Pouget says. Yet in the study, the human subjects were "as good as they could possibly be."

Now the team is wondering just how good humans' Bayesian thinking can get. "The lab is on a quest to find out, 'OK, where do we break down? How much complexity do we have to put in the task before we can no longer come up with the optimal solution?'" Pouget says. "And so far we haven't found where that boundary is."

Psychologist Wilson Geisler of the University of Texas at Austin prefers an approach that starts with the outside world. His team uses carefully calibrated cameras to capture a scene and range-finders to measure the distance from the

cameras to each point in the scene and the brightness of the light coming from each of those pixels. These tools allow the researchers to construct an exact mathematical description of the natural world.

"We try to measure the actual 3-D world, and then we try to learn how you would estimate the shape or distance of an object," Geisler says. With this precise mathematical description of the world, Geisler then builds a theoretical tool that mimics the behavior of a perfect Bayesian-thinking human inhabiting that world — an "ideal Bayesian observer."

By comparing flesh-and-blood humans against this "ideal observer," Geisler and his colleagues are getting a sense of how people stack up. They are "almost perfect," Geisler says. But before cockiness sets in, Geisler points out that in these studies perfection doesn't mean always being correct. For instance, people judging whether two patches of green behind a mushroom belonged to

the same leaf or different ones would get worse at the task as the mushroom grew bigger and hid more of the scene. Unreliable information leads people astray in a way that Bayesian math predicts.

In a study published last year in the *Journal of Vision*, Geisler and colleagues showed participants close-up pictures of leaves photographed at a nearby botanical garden. People's performance at distinguishing two overlapping leaves in patches from a two-dimensional image mirrored the performance of an ideal observer. Participants seemed to operate with existing knowledge of how to visually unjumble a pile of leaves.

In a way, it's self-evident that humans rely on existing knowledge. A brain that didn't rely on its experiences would be a pretty pathetic brain. "You could argue that it would be a little strange if we were bad at it," Geisler says. "It's something that we have enormous experience with, evolutionarily. The same problem has been there for a billion years. But nonetheless, the statistics are complicated."

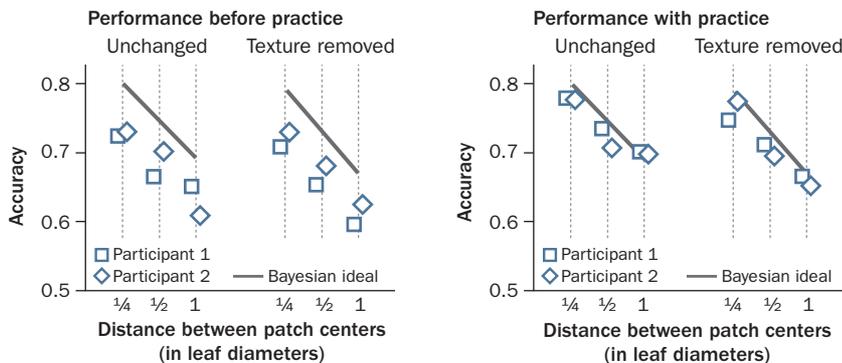
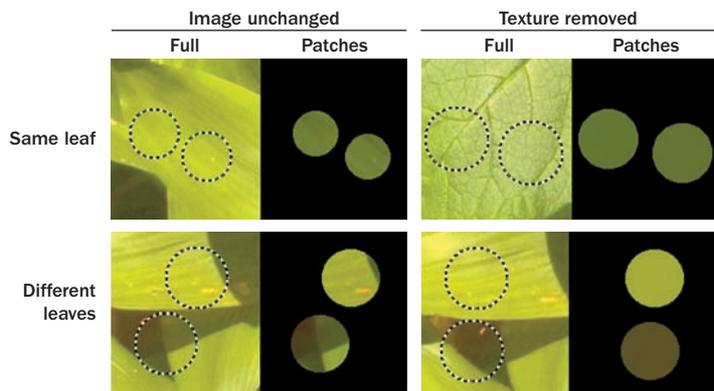
Parsing these statistics isn't just a task for the visual system. So far, some scientists have turned up hints that movements, smells, hearing, cognition and the ability to perform easy addition problems may be based on Bayesian techniques. And these abilities might be present well before a child learns $2 + 2$.

A's, Bayes, C's

By studying babies and young children, scientists can test whether probabilistic reasoning is present before life experiences begin sculpting the mind. Babies haven't been alive long enough to develop strong beliefs about how the world works. If babies act Bayesian, then they may have been born that way.

Sixteen-month-olds can make correct assumptions when faced with complicated data, cognitive scientists Laura Schulz and Hyowon Gweon of MIT reported June 24 in *Science (SN Online: 6/28/11)*. In the study, babies watched as experimenters pressed a button on a toy, causing music to play. In some cases, the toy worked beautifully the first time each experimenter pushed

Optimal performers Study participants who were shown intact or distorted patches of leaf images (top panel), and were asked to determine whether the patches came from the same leaf or different leaves, performed nearly as well as an ideal Bayesian observer (gray line in graphs).



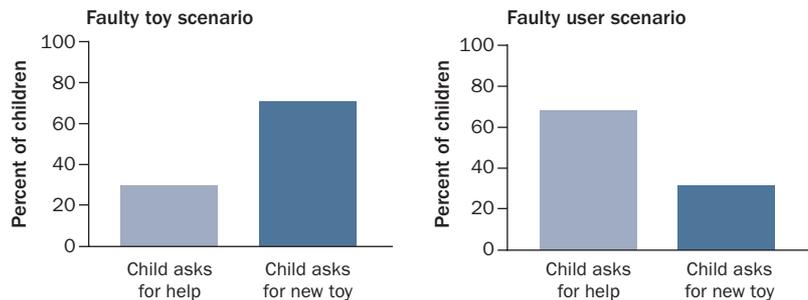
SOURCE: A.D. ING ET AL/JOURNAL OF VISION 2010

LEAF IMAGES: A.D. ING ET AL/JOURNAL OF VISION 2010



Baby rationale Babies can think in probabilistic ways. After seeing two researchers both succeed and fail in using a sound-making toy, a baby who then fails with the same toy is likely to think the toy is faulty and tends to go for another (below, left). But when one researcher fails and a second succeeds, the baby more often takes the blame and asks for help (right).

SOURCE: H. GWEON AND L. SCHULZ/SCIENCE 2011



the button, but fritzed out the next time. This created the semblance of a faulty toy. In other cases, the toy worked well for one experimenter but never worked for another, suggesting that the toy was fine but the second experimenter was a poor operator.

When the babies were handed the toy that seemed like it was faulty, they quickly reached for a different toy. But when the babies thought they themselves might be to blame (when they witnessed the second experimenter fail with the toy and then they failed themselves), they handed the toy to a nearby parent in a plea for help.

By assessing others' toy travails and applying that knowledge to their current problem, babies displayed very sophisticated reasoning, Schulz says. "As early as we can test, babies are using things that are consistent with probabilistic models," she says. "Babies are sensitive to the statistics of the environment."

Instead of looking for signs of probabilistic reasoning in young humans, some scientists are looking for signs in other species. A recent study in owls suggests that aspects of their brains also follow Bayesian rules.

Though owls are admirable hunters, they typically don't hear sounds that come from areas in the periphery as well as sounds coming from the front. To explain this deficit, Brian Fischer of École Normale Supérieure in Paris and José Luis Peña of the Albert Einstein College of Medicine in the Bronx, N.Y., turned to Bayesian math.

The team devised a statistical model

of auditory processes with the assumption that owls may have evolved to assign less importance to signals coming from the periphery because hunting something at their backs might be too costly. A turning motion might scare prey away, for instance. In tests, Bayesian models closely predicted this actual owl behavior, the researchers reported in the August *Nature Neuroscience*.

In the owl's auditory system, this bias toward hearing objects right in front may come preinstalled. Likewise, babies may be hardwired to quickly infer whether they are to blame for a nonworking toy.

Where, when and how these pieces of prior knowledge get filed away in the brain is still a mystery. Some scientists think priors — and the ability to use them — were built into brains over the course of evolution.

"Biological systems are not accidental," Simoncelli says. "We believe that evolution shaped them, and shaped them to be good at what they do. And we have a lot of evidence that that's true."

'Prior' engineering

Whether or not evolution designed Bayesian brains, some of those very brains are now intent on passing their Bayesian abilities on. Trained as an engineer, Simoncelli says that the same principles at work in the brain could be incredibly useful elsewhere. "My belief is that when we finally figure out how some of these circuits operate in brains in order to accomplish these feats, we're going to change engineering," he says.

"We're going to revolutionize the way we think about designing systems."

Many of today's robots, for example, excel at precise tasks but are totally inflexible. Robots that install windshields on new cars perform the job flawlessly each and every time. "We can make that robot be fantastically good at putting that windshield on," Simoncelli says. "They're beautifully engineered systems." But those paragons of windshield installation would have a complete meltdown if they were handed a glass sheet of the wrong size. A similar robot based on the human brain, though, might easily adapt to changing circumstances and even file away some priors of its own.

Nerve cells exhibit enormous flexibility. Constantly readjusting to input, interacting with neighbors and changing firing rates can lead to incredible adaptability, a prerequisite for Bayesian learning, Simoncelli says. The more scientists understand about nerve cell function, "the more we find they're not fixed, dedicated devices that operate the same way throughout your lifetime," he says.

Cracking the brain's Bayesian operating system might lead to a new set of engineering principles. "We don't know how to engineer systems that are more flexible, and we don't know how the brain works. And we're going to figure both those things out," Simoncelli says. "And I believe that we're going to do it at the same time." ■

Explore more

■ For Alexandre Pouget's papers: www.bcs.rochester.edu/people/alex/