Alec Brewer <abrewer@ourobio.com>

Curiosity Depends on What You Already Know
1 message

The Inn <theinn@mail.beehiiv.com> Mon, Jan 22, 2024 at 2:00 AM
Reply-To: The Inn <sy8005939@gmail.com>
To: "abrewer@transfoamllc.com" <abrewer@transfoamllc.com>

January 22, 2024

Curiosity Depends on What You Already Know

Humans have a drive to eat. We have a drive to drink. We have a drive to reproduce. Curiosity is no different, says George Loewenstein, a professor of economics and psychology at Carnegie Mellon University. Our insatiable drive to learn—to invent, explore, and study ceaselessly—“deserves to have the same status as those other drives.”

What’s curious about curiosity, though, is that it doesn’t seem to be tied to any specific reward. “The theoretical puzzle posed by curiosity is why people are so strongly attracted to information that, by the definition of curiosity, confers no extrinsic benefit,” Loewenstein once wrote. It makes sense for organisms to seek food, water, sex, shelter, rest, wealth, or any of the other myriad nourishing and pleasant things in life. But what is the good of deducing the nature of gravity, or of going to the moon?

A simple answer is that we never know if what we learn today might come in handy tomorrow. Take worms, for example. They’re incorrigible optimists, says Sreekanth Chalasani, a neurobiologist at the Salk Institute of Biological Sciences, in California. He studies Caenorhabditis elegans, a common, millimeter-long species of roundworm. During experiments, he will put a worm on a big patch of bacteria (its favorite food), surrounded by plenty of potential mates. “What’ll it do? It’ll leave that patch, and go looking for something more,” he says. “There’s no evidence that there’s something better outside. This is the best food you can give it. It’s craziness!”

Whether you’re simply abandoning your food patch, or flying into space, exploring does seem a little nuts—except, of course, you never really know if the food will run out. From an evolutionary perspective, Chalasani says, there’s good reason to keep looking. Information helps us make better choices and adapt to a changing environment. Maybe someday we’ll need a moon base.

> Curiosity is less about what you don’t know than about what you already do.

Curiosity is not just wanderlust, though. We’re curious about specific things, and different people are interested in different specific things. Some are hobbyists, seekers of the arcane, others jacks-of-all-trades. This divergence of interests tells us that something beyond a tendency to roam must be guiding each of our unique obsessions.

Indeed, scientists who study the mechanics of curiosity are finding that it is, at its core, a kind of probability algorithm—our brain’s continuous calculation of which path or action is likely to gain us the most knowledge in the least amount of time. Like the links on a Wikipedia page, curiosity builds upon itself, every question leading to the next. And as with a journey down the Wikipedia wormhole, where you start dictates where you might end up. That’s the funny thing about curiosity: It’s less about what you don’t know than about what you already do.

In the most basic terms, you could describe curiosity as a function of motivation plus direction. The first part is not as obvious as it might seem. Thirst, hunger, horniness—there are clear motives for our other drives. But what kindles curiosity?

The 19th-century German philosopher Arthur Schopenhauer believed that life’s chief task is “subsisting at all,” followed directly by “warding off boredom, which, like a bird of prey, hovers over us, ready to fall whenever it sees a life secure from need.” To be content is to be bored, and curiosity is our ticket out. The anthropologist Ralph Linton went even further. “It seems probable that the human capacity for being bored, rather than man’s social or natural needs, lies at the root of man’s cultural advance,” he wrote in 1936.

Humans, in other words, have managed to amass immeasurable knowledge—language, the Taj Mahal, the Snuggie—because we loathe monotony.

But boredom alone can't fully explain curiosity. "The very old view is that curiosity and boredom are opposite ends of the same continuum," Loewenstein says. The new view: bored is not to curious as hungry is to full or thirsty is slaked. Rather, boredom is "a signal from your brain that you're not making good use of a part of the brain," like the tingling of a foot you've sat on too long. Boredom reminds us that we need to exercise our minds, but there are antidotes to boredom besides curiosity—food or sex, for example. What's more, curiosity strikes even when we're not bored. In fact, we will readily give up things we want or enjoy in order to learn something new.

> Curiosity peaked when subjects had a good guess about the answer but weren't quite sure.

Like Chalasani's worms leaving their perfect patch of food, humans and other primates will consistently trade rewards for information. To measure this tendency, researchers use "bandit tasks"—a reference to slot machines ("one-armed bandits")—in which subjects must repeatedly choose between several images or other options. Different options come with different probabilities of paying out a reward (money, typically), and over time, subjects learn which options are most likely to reward them and will consistently choose those. But when an option is introduced that subjects haven't seen before, they will often choose that one, giving up a likely reward on the chance that the new option will pay out better.

Brain studies suggest that this "novelty bonus"—the additional weight we give to new options—stems at least in part from the euphoric feeling it gives us. For instance, a 2007 study found that, like Pavlov's dog salivating at the ring of a bell, the part of our brain that processes rewards like love and sweets activates when we expect to find something new, even if that expectation doesn't play out. These findings, the researchers conclude, "raise the possibility that novelty itself is processed akin to a reward."

So maybe it's true that I go to Wikipedia, as Schopenhauer said, to "ward off boredom." But then I stay there for the next three hours reading up on the Mongol invasions of Japan partly because, subconsciously, I just like the dopamine rush I get from clicking on the links—the same rush that drove my ancestors to colonize Australia and the Arctic Circle, invent pottery, and carve the Venus of Willendorf.

But why follow the Mongol horde deep down the rabbit hole? Why not research Willard's sooty boubou, or any of the other theoretically interesting subjects Wikipedia's "random article" button tosses my way? Why does curiosity beckon us *this* way, and not *that*?

In a 1994 paper, Loewenstein theorized that curiosity's direction is determined by the "information gap," the sudden awareness of what you don't know and the immediate desire to fill that gap. This perceived gap can exist in the physical universe (What is this weird bug?) or the mental one (What is love?). His theory does a good job of putting into words why Upworthy headlines are so irresistible (Damnit, what *are* the 22 Reasons I'm Probably Already a Manatee Fan?), and why curiosity is viewed as both a strength and a weakness (Did you know that manatee nipples are located in their armpits?).

For the information gap to set its hook, though, it can't be too big (the headline is written in Portuguese) or too small (One Fact Is That Manatees Live in Florida). In a 2009 study, a team of researchers (including Loewenstein) put subjects in an fMRI machine, and then asked them a series of trivia questions: What is the instrument that was invented to sound like human singing? What is the name of the galaxy that Earth is part of?[1] For each question, the subjects estimated how confident they were in the answer. The researchers also asked the subjects to rate how curious they were about the question and monitored how strongly their brains' reward centers lit up—another measure of curiosity.

As expected, subjects were least curious about answers they thought they knew. But they were also uninterested in questions about which they hadn't a clue. Instead, curiosity peaked when subjects had a good guess about the answer but weren't quite sure. The sweet spot for curiosity seemed to be a Goldilocksian level of information—not too much nor too little.

Babies, too, favor new-but-not-too-new things, says Celeste Kidd, a neuroscientist at the University of Rochester. In a 2012 study, she and her colleagues sat 7- and 8-month-olds in front of a screen that showed three patterned boxes, each of which contained an object such as a cookie, a spoon, or a car. The objects emerged from the boxes in specific patterns, "like Whack-a-mole." By presenting some patterns more frequently than others, Kidd could make certain sequences seem more rare, and thus more surprising.

As the babies watched the screen, an eye-tracking device watched them. And their glances revealed a clear preference: Patterns that were somewhat surprising but not completely novel held their attention; patterns that were either very similar, or very different, from what they'd seen before did not. (Each time a baby looked away, the screen reset to an image of a laughing baby. "I don't know if you knew this, but babies love looking at pictures of other babies," Kidd says. I did not. Curious.)

The way our brains instinctively seek "just right" levels of novelty is a bit like going to a bookstore, Kidd says. "You wouldn't want to pick a children's book, or a book you've read a lot before." On the other hand, if you choose a book you can't penetrate at all, like, say, a Russian textbook on astrophysics, you hit a similar problem. "That's not going to be very interesting." To learn, you have to have something to grab onto: The next handhold can't be too far from the last—you might never reach it. So as your brain pushes you to gather information as quickly as possible, it instinctively steers you away from gaps that are too small, or too large.

Robots are good vessels for examining how this calculation may work. But because a robot lacks motivation (the primary ingredient of curiosity), first you've got to give it some. To do that, just program the robot to seek a reward, says Varun Kompella, a postdoctoral fellow studying artificial intelligence at Ruhr University Bochum, in Germany. It doesn't matter what the reward is (even a number works), so long as the robot knows that rewards exist, and that it wants to earn them. Similarly, it can't know how to earn a reward. Just as a human gets a hit of dopamine for learning something new, even if it seems completely useless, the robot's motivation system makes learning its own reward.

Kompella works with an iCub, an open-source humanoid robot with cream-colored skin, silver joints, a head, eyes, arms, fingers, even nipples, but no hair or legs. In a video he sent me, the iCub connects directly to the floor behind a table. In the center of the table is a plastic cup. The robot starts off pitching back and forth, clenching and unclenching its fists. At first, every new action teaches it something, and the rewards come quickly. Soon, though, it runs out of new things to learn about how its own body moves.

Then suddenly, lurching about in this random way, the robot knocks over the cup. The encounter earns it a reward and, more importantly, suggests a new avenue of knowledge. This is a sailor, spotting a shorebird after months at sea. This is George Mallory, first hearing of Everest. This is curiosity, no longer random, but directed.

So what next? The answer is determined by a probability algorithm that calculates what action is most likely to earn the robot another reward. In this case, the algorithm says that, because moving an arm in the region of the cup resulted in knowledge about an action (and a reward: cha-ching!), doing something similar in that same region is more likely to result in new skills than either going back to performing totally random actions and ignoring the cup, or focusing on the cup, but acting completely different. Why focus on the cup? Because it's there.

Kompella's iCub eventually learned to pick up the cup, move it over, and drop it in a certain spot on the table, which, it turns out, was the task Kompella wanted it to do. But it taught itself to become a cup dropper in large part because it was bolted to the floor in front of a table with a cup on it. It had few other options.

Similarly, Kidd's experiment was designed to track the amount of information the babies had at any given time—thus allowing Kidd to control the amount of novelty—and also to limit the babies' choices. She used 7- and 8-month-olds, she says, because they hit the sweet spot of being able to support the weight of their own head, but they're not yet learning to walk. Nothing is more interesting to a baby learning to walk than walking is, she says. "You can't compete with that."

To predict or even control curiosity would be to teach more efficiently, to better understand diseases of the mind, to entertain more consistently; life would be endlessly interesting. But the very difficulty of studying curiosity suggests its boundlessness, the near impossibility of truly directing it. For now, there are only more questions.

Unsubscribe