From

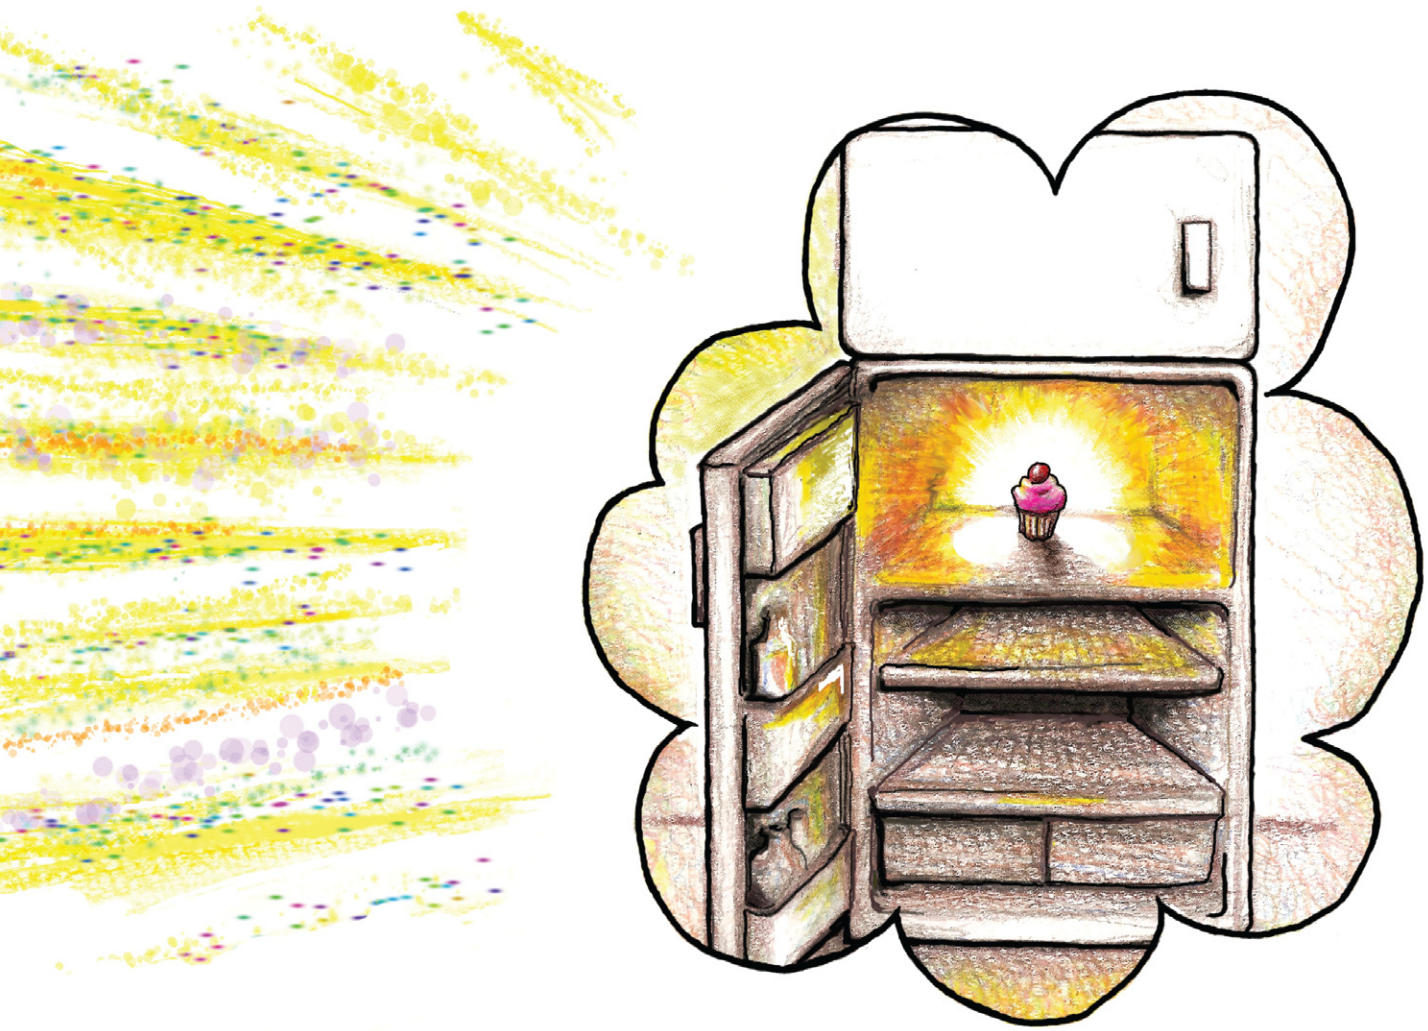

Dendrites to Decisions

By Lori Oliwenstein

The average human brain—weighing in at a scant three pounds—has, according to one estimate, upward of 100 billion neurons that connect with one another via some 100 trillion synapses.

A hundred trillion? Wow. You learn something new every day, right?

But wait. Was that really learning? Or was it just a bit of trivia you're likely to forget as quickly as you read it? What, exactly, *is* learning?

Ah, there's the rub—or, if you will, the learning curve—says psychologist John O'Doherty, one of the dozen or so Caltech faculty for whom learning about learning has become a scientific endeavor.

"There are lots of ways to talk about learning," he says. "The brain is always learning. It's key to our survival; we need to learn how to find things like food, water, and shelter. Equally, if not more importantly, we need to learn to avoid bad things, like getting run over by a car or being eaten by a mountain lion."

Biologist Thanos Siapas sees learning as part and parcel of memory, upon which his research focuses. "Learning and memory are two sides of the same coin," Siapas notes.

Behavioral economist Colin Camerer, on the other hand, sees learning as just another kind of decision making—a process of assigning values to objects and experiences and thus "learning" about them, and about whether you'd want to make the same choices in the future.

And if you ask Mary Kennedy, who has spent the last 30 years taking apart and exploring the inner workings of the brain's synapses, you'll get yet another take. "Learning is essentially a form of neural plasticity," she says—the ability of individual brain cells to make new connections or retune those that already exist.

LEARNING TO CHANGE

A synaptic connection has three parts: the sending neuron's axon, the receiving neuron's dendrite, and the cleft in be-

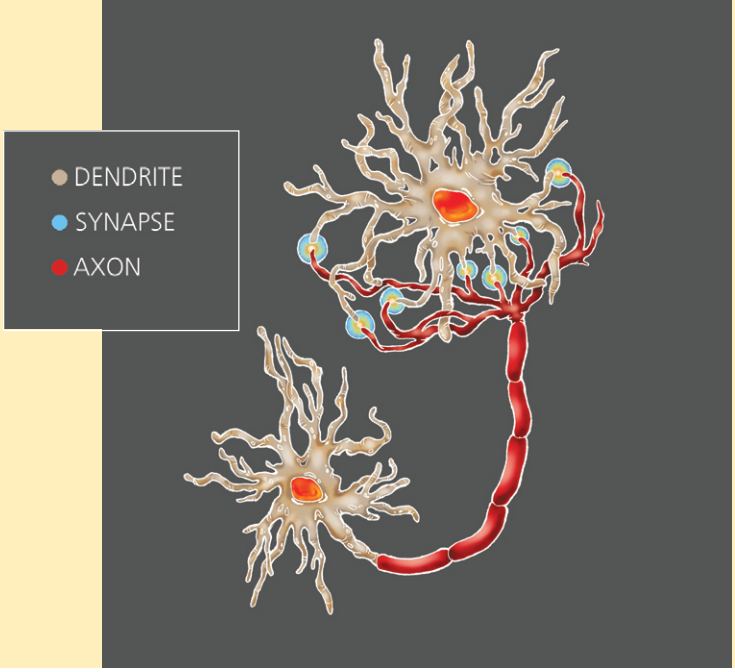

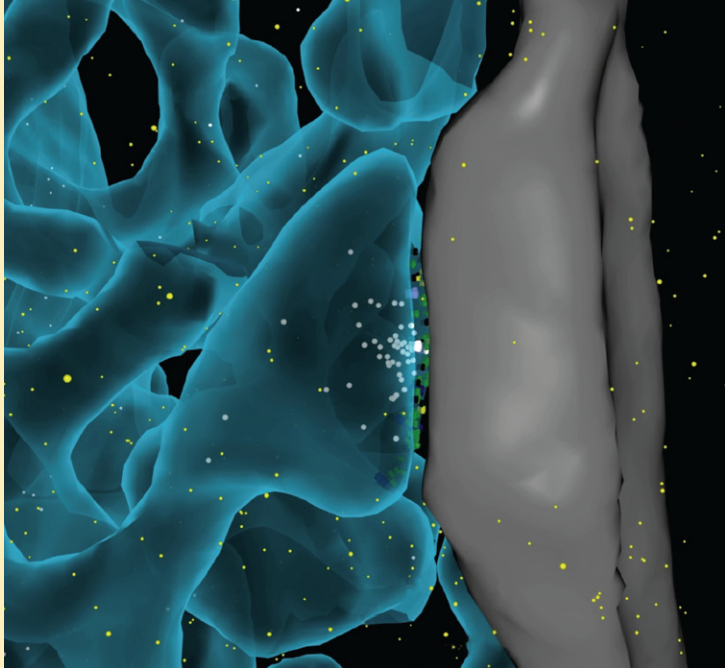

tween where axon and dendrite almost, but don't quite, touch. In an amazing feat of sleight of hand (sleight of synapse?), an electrical impulse reaching the axon's tip is transformed into a burst of chemicals that cross the synaptic cleft, only to be changed back to an electrical impulse again in the dendrite. This impulse then travels down the dendrite, through the neuron's cell body, and out along the axon to another set of synapses.

Not all synapses are created equal, however. Although each neuron has thousands of synapses, most of them are small and weak, and have little if any influence on the next nerve cell in line. Learning something new—opening the lines of communication between neurons that previously wanted nothing to do with one another—requires pumping up the volume in the synapses connecting them, so that the intended message comes through loud and clear.

It is this volume adjustment that Kennedy studies in the most minuscule detail. Specifically, hers is one of a handful of laboratories in the world that focus on what is known as the postsynaptic density, or PSD. The PSD, as the name implies, is at the receiving end of the synapse—it's that part of the dendrite that includes the cell membrane and the area just beneath, where the chemical signal is plucked from the cleft and converted back to electrical form.

Recreating the electrical impulse in the dendrite requires an influx of ions—and calcium, Kennedy says, does a memory good. The intercellular soup in the synaptic cleft is heavily seasoned with calcium ions, and the surface of the PSD is studded with proteins called NMDA receptors that, when activated, open to let calcium ions into the dendrite.

It's an almost impossibly complex process, but it all comes down to this: The more calcium that comes into a dendrite through the NMDA receptors, the more that dendrite's internal skeleton branches and expands. The more the skeleton branches, the more the PSD's membrane sprouts another kind of receptor—called an AMPA receptor—that ultimately causes the neuron to fire.

And all of that—the influxing of calcium, the expanding of the skeleton, the adding of AMPA receptors—is what defines neural plasticity.

"Neurons that fire together wire together," Kennedy quips. But it's no joke: That calcium cascade tightens and strengthens the connections between neurons. Without that cascade—without a robust, well-connected PSD—learning comes to a screeching halt.

"When I started at Caltech 30 years ago," says Kennedy, "we didn't know any of the molecules in the postsynaptic density; none of them. We knew it as a dark thing that we saw in the electron microscope. Now we know most of the proteins that are there, and we know quite a bit about how they work and how they respond to calcium."

Having dissected the PSD and identified its components, Kennedy's group is now trying to figure out the system's dynamics. "We want to try to understand how subtle differences in calcium flux lead to strengthened synapses—or, sometimes, to weakened ones," she says. "The precise pattern of calcium flux into the synapse is what controls whether it strengthens or weakens; it's at the core of what happens during neuroplasticity. And yet, nobody really knows how it works."

REMEMBERING TO LEARN

What we do know is that learning is an oh-so-deliberate process, a sweat-and-tears and all-night-problem-set endeavor. In other words, learning takes time, says Thanos Siapas.

Siapas studies just how the brain takes incoming information, shunts it around, and finally lays it down in ways that will allow it to be retrieved quickly and easily at a later date—that will allow the information to be remembered, to be learned.

"It's a process," says Siapas. "When you learn something, your brain continues to make changes for a long time after. You can't just store a new memory; you need to integrate it with the other things you already know. And that's much trickier than you might think."

While Kennedy and her group are peeking inside individual neurons, Siapas and his colleagues are taking a

Far left: The brain's synapses lie in the blue-encircled areas where the branching tips of the sending cell's axon (in red) almost-but-not-quite meet the receiving cell's dendrites.

Left: Calcium ions are dissolved in the soup that surrounds the neurons and fills the synaptic cleft between axon (gray) and dendrite (blue). Here, calcium ions (the clustered white specks) rush through activated NMDA receptors into the dendrite, where they will interact with proteins in the postsynaptic fluid to help strengthen the synapse.

Right: The prefrontal cortex, outlined in red, is part of the cortex, the brain's outer layer, and is involved in planning, decision-making, and other higher-order functions; the two major subparts shown here play roles in goal-directed learning. The intraparietal sulcus, another part of the cortex, also contributes to action-planning and decision-making. The hippocampus, which helps process memories for storage, is buried deep in the brain's interior.

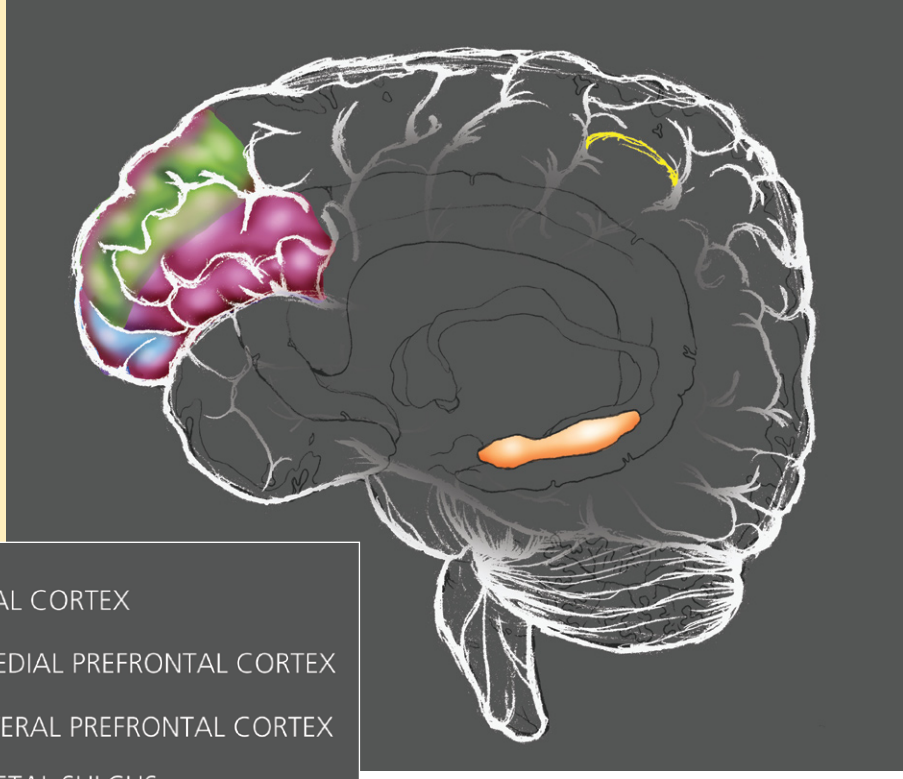

step back to look at large conglomerations of brain cells and their relationships with one another. At the center of consideration is the hippocampus, a curved ridge of gray matter known to be essential for learning and memory formation—though it is not, Siapas points out, where those memories are ultimately stored.

"The hippocampus helps establish memories, helps consolidate them," he says. "But it consolidates them somewhere else. We want to understand how the hippocampus is activated during learning, and how it communicates with other brain areas during this process."

To do that, Siapas says, requires monitoring many brain areas over long periods of time. "We're talking about months or even years in humans, weeks in mice," he notes.

But that long-term effort has paid off. For instance, using high-tech recording and computational techniques, Siapas and Casimir Wierzynski (PhD '09), now a postdoctoral scholar, were able to pinpoint a number of synchronized neuron pairs in which a hippocampal neuron's firing was followed within milliseconds by the firing of a neuron in the prefrontal cortex.

"This is exactly the kind of relationship that would be needed for the hippocampus to effect changes in the neocortex—such as the consolidation, or laying down, of memories," says Wierzynski, who was the lead author on the 2009 *Neuron* paper reporting the work.

The scientists also found that these bursts of neuronal chatter happen only during slow-wave sleep—the deep, dreamless portion of your night's rest. During the dream-laden periods of REM sleep, it seems, you may well be too busy fighting zombies or wandering naked through your high-school hallways to get much real brain work done.

But what *do* kick in during REM sleep, Siapas says—as well as when we're up and about during the day—are the theta oscillations. These are prominent brain rhythms that orchestrate the activity of neurons in the hippocampus, and for decades they were thought to pulse in sync across the entire structure—acting as a master clock, a centralized pacemaker. It wasn't until postdoc Eugene Lubenov, now a senior research fellow, and Siapas took a more detailed look at these biological cadences that the pacemaker paradigm was dealt its death blow. It turns out instead that theta oscillations sweep across the hippocampus in a traveling wave, moving steadily from one end of the structure to the other. "In other words," says Siapas, "the hippocampus has a series of local time zones, just like the earth has."

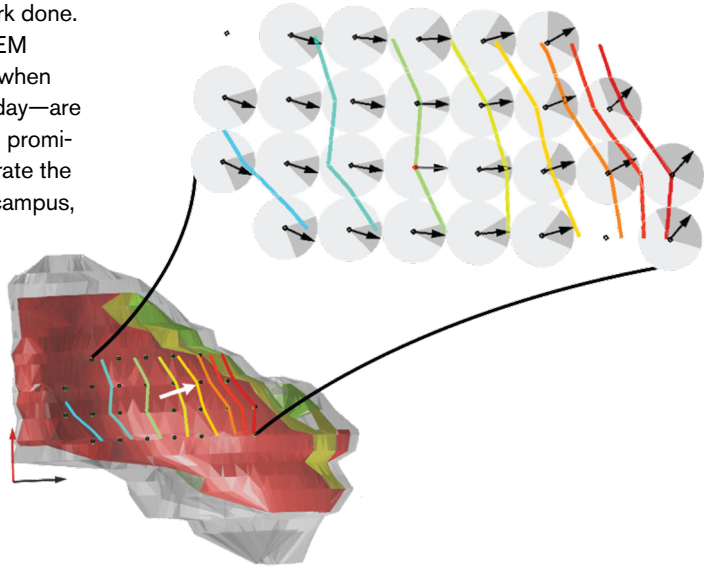

Lubenov and Siapas showed that brain rhythms called theta oscillations move across the hippocampus in waves. In this diagram, each colored line represents a "time zone" in which the oscillations are in sync, as shown by the clock hands in the inset.

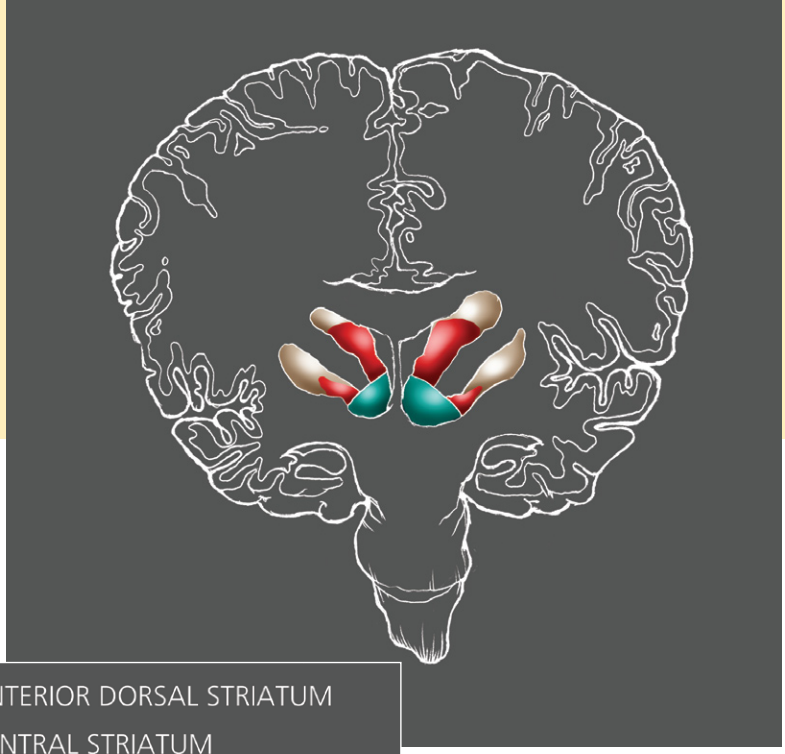

The striatum, buried deep in the brain, takes inputs from various parts of the cortex. The anterior dorsal striatum (red) receives inputs from the ventromedial prefrontal cortex and is part of a circuit for goal-directed learning. Both model-free learning, which is learning without a map, and fictive learning, which is learning from what others do, involve the ventral striatum (green). Meanwhile, the posterior dorsolateral striatum (tan) controls habitual behaviors.

DECIDING TO LEARN

Learning may be a slow-and-steady turtle of a process, but at its inception, it can be a hare-like burst of action that begets an unexpected response—a fleeting meeting of instinct and serendipity.

The most basic of instincts—flinching at a really loud noise, wincing in pain when you stub a toe—are reflexive rather than learned, notes psychologist John O'Doherty. But everything above that level involves some form of knowledge acquisition.

"Learning comes in even when you're talking about Pavlovian conditioning," O'Doherty says. "It's one of the most basic forms of learning—associating a cue like a buzzer in the lab with food, or a rustling in the bushes with a mountain lion. We *learn* to associate those cues with something significant, and then to respond physiologically to the cue—to salivate even before we see the food, or to feel a fear-based rush of adrenaline before the bushes part and the lion is upon us. By learning to anticipate significant events based on past experience, we buy ourselves time so that we're better prepared to eat, or fight, or flee, or whatever the appropriate action is."

We kick it up a notch when we learn how to interact with our environment and make things happen to suit ourselves—the so-called instrumental conditioning that drives mice to push a lever to get food, or tells us to shake a tree to make the fruit fall down. We've learned that an action will give us what we want, and we go for it.

In general, O'Doherty says, instrumental conditioning can be divvied up into two general categories: habit learning and goal-directed learning.

Imagine you're a toddler walking past the refrigerator. For absolutely no reason, without giving it a second (or first) thought, you open the door. And there, in a spot of fridge-lit glory, sits a single, perfectly frosted cupcake. You grab it—of course you do—and eat it. The moment you stuff that cakey goodness into your mouth, you will have learned something: opening a fridge door can have positive consequences. "And the next time you walk past a refrigerator," O'Doherty notes, "you're much more likely to open it."

That, in a nutshell, is habit learning.

Goal-directed learning, on the other hand, is a tad more sophisticated, and less dependent on dumb luck. "You're thinking about the consequences of taking a particular action," says O'Doherty. "You think, 'I want a cupcake. Where would it be? Ah, let me try the fridge.'" Rather than just opening random dresser drawers or your toy box, you stop before you act, before you waste your energy, and evaluate the possible outcomes.

And this, he notes, can motivate less physical actions as well. It's goal-directed learning that tells you to sit down and study for a test, because you're more likely to get a positive result—a good grade—than if you blow it off. No gold star or Saturday-night use of the car on the line? No real reason to pick up that textbook.

"Habits are things we do without thinking of the consequences," O'Doherty says. Being goal-directed, on the other hand, is all about reaping what you sow.

Which is not to say that goals and habits aren't linked. In fact, they are—intimately so, in many cases. Take bike riding. When you first get on a bike, you are completely goal-directed; you have to think about every movement your body makes in an attempt to keep yourself upright. But, after a while, you don't have to think any more. Your responses become habitual, reflexive.

Sounds obvious, right? And yet it was only in the last couple of years that O'Doherty, postdoc Elizabeth Tricomi (now an assistant professor at Rutgers), and Bernard Balleine at UCLA actually showed experimentally that—over time and with training—goal-directed behavior in humans can indeed become habit.

"People had skirted around the issue before then," O'Doherty admits. "They'd assumed certain behaviors

were habitual. But no one had actually done what we did."

In addition, the scientists were able to pinpoint, for the first time, the control of habitual behavior to a specific area of the brain—the posterior dorsal striatum.

O'Doherty says such insights are critical. "We want to know which parts of the brain are involved in learning, and what are the algorithms—what programs does the brain run—to allow these different kinds of learning to take place," he says. And the only way to get those sorts of insights is to actually watch the brain at work.

In the functional magnetic resonance imaging (fMRI) machine at Caltech's Brain Imaging Center, almost any type of mental gymnastics is fair game. This fMRI is the same sort of whole-body scanner that an orthopedist might put you in to look for a torn ligament; here, the volunteer lying in the machine performs a predefined task—placing a bet, for example—and as the thought process unfolds, the scanner tracks the brain's active areas in 3-D. Says O'Doherty, "Not only can we identify what parts of the brain are active when, we can also figure out what algorithms are being implemented when you do one task or another."

Caltech-led fMRI studies have confirmed that goal-directed learning tends to begin in a part of the brain just above your eyeballs called the ventromedial prefrontal cortex (VMPFC); this is the area made famous by railway worker Phineas Gage in the mid-1800s, when a large iron rod pierced his brain and robbed him of his decision-making and social skills. The VMPFC, says O'Doherty, talks to a region in the center of the brain called the anterior dorsal striatum. Eventually, control of the behavior in question passes to the posterior dorsal striatum, which is closely connected to the motor cortex and thus plays a larger role in habit learning than it does in the thinking-things-through process of goal-directed learning. At this point, what used to require mental effort is starting to go on autopilot.

Learning typically involves updating your expectations continuously as things change around you, applying what you've learned in the past to figure out what to do in the future. You might look at what the stock market did around this time last year, for instance, before deciding whether to throw a little extra money into your portfolio now.

Such updates are, obviously, nothing more than approximations and bound to fall short. Indeed, there's a name for that shortfall: prediction error. It's the difference between what you think you'll make in the stock market this year, and what you actually do make. The neat thing about learning is that, next year, you can use that new info to change your expectations again and perhaps reduce your prediction error. "When you've completely 'learned' something," notes O'Doherty, "your prediction error goes down to zero."

This sort of trial and error is called model-free reinforcement learning and may explain what goes on when you are starting to form a habit. But there's another strategy: model-based learning.

Explains O'Doherty, "If we're sitting here in my office on the third floor, and I tell you, 'I left $10,000 down in the lobby, and I've told five other people about it,' the only way you can get to that money before they do is to create a map of the building in your brain and compute the value of taking different routes. Chess is the same way; in order to reach your goal, you need to learn how to map out your particular situation."

This, then, is a kind of goal-directed learning. Going after the cash without

a map in mind would be a total bust. You'd just wander down corridor after corridor, aimlessly; those other guys would be out spending the loot before you even found the stairs.

In a Neuron paper last year, O'Doherty and then-postdoc Jan Gläscher described how these two modes of learning interact to help us make critical decisions, and showed that they actually involve different brain areas. Model-free learning was found to involve parts of the striatum including the ventral striatum, while certain aspects of the model-based learning system were found to depend on some areas in the cerebral cortex—the intraparietal sulcus and the dorsolateral prefrontal cortex. "The cortical areas are learning the map that you are going to need in order to perform goal-directed learning," O'Doherty explains. The complete mechanism, of course, is much more complicated. "There is a considerable network of brain areas likely contributing to each of these learning processes."

Learning modes vary over time—with bike riding starting as a goal-directed activity, but later becoming a much less "computationally expensive" habit—and even from moment to moment. In a study published earlier this year, O'Doherty and postdoc Ryan Jessup looked at the gambling strategies used by 31 volunteers playing a roulette-type game (talk about expensive!) while in the fMRI machine. The game involved betting on which of three colors would come up next on a three-colored wheel, but there was a twist—the three regions were of unequal size, and the sizes changed with each spin.

Two radically different approaches quickly emerged. Sometimes players relied on reinforcement learning, a version of model-free learning in which they picked a "lucky" color that had paid off in the past. If the streak turned cold, they'd switch to the "gambler's fallacy"—a more model-based strategy based on the belief that one color was "due," either because it hadn't come up in a while, or because the spins seemed to be following a specific pattern.

The different strategies were reflected in the subjects' brain activity, says Jessup. The dorsal striatum flickered to life when participants were using model-free reinforcement learning—the ones who wagered based on what had worked for them previously—but stayed relatively quiet when bettors were under the sway of the model-based gambler's fallacy.

As it turns out, neither of these strategies is particularly good for this particular game. At the beginning of every experiment, the subjects were explicitly told that the computer was picking winners at random, and that the odds of a color hitting were proportional to the area it occupied. The best bet would thus be the color taking up the biggest piece of the wheel on that spin, regardless of what that color happened to be. Says O'Doherty, "The fact that we repeatedly choose less-promising strategies in the face of a more rational alternative tells us that these learning processes are so deeply ingrained in our brains that they can influence our behavior even in situations where it is actually counterproductive."

THE ACTION NOT TAKEN

But how do you make decisions when you can't bring much to the table in the way of personal experience? How do you assess the roads not taken?

"When you have a set of potential actions—like which movie to see—you learn only about the movie you choose," says behavioral economist Colin Camerer. If the chosen flick was exceptional, the next time there's a movie starring that same actress, you'll be more likely to see it; you'll have learned something from the experience. Similarly, if the movie was awful, you'll know what to do the next time you're considering an offering from that particular director.

As for the movies you passed up? "It would be useful to have a mechanism to learn more about them, too," notes Camerer.

And you do: it's called fictive learning, or learning from the what-ifs of life. In other words, you *can* learn from other peoples' experiences. In the case of the multiplex, you can ask friends who saw the other movies what they thought, and then incorporate those opinions into your worldview for future decision making.

Camerer and colleagues have explored fictive learning in a series of experiments, the most recent involving a game in which 54 volunteers made a series of "investments." At the end of each round, the results of all the investments were revealed to all the players.

The fMRI revealed that processing "rewards not received from actions not taken" happens in the ventral striatum, the same brain area where model-free learning occurs. "We see signals about what you could have done in regions similar to those that encode signals for actual rewards," Camerer remarks. And so, if your friend raves about the movie you didn't see, your brain will tuck that information away in the same place where it normally sticks the information about movies you yourself saw and enjoyed.

There is, however, a slight catch: some of Camerer's earlier studies have found that the weight you give to a fictive account of something's value is about half what it would have been if you'd learned about it yourself.

Which means that the next time you check your local cinema's listings, you're still likely to snub *The Hangover Part II* if you hated Part I—no matter how many times Uncle Joe insists you absolutely *must* check it out.

And that, my friend, is what learning is all about. E&S

PICTURE CREDITS:
14–15, 20, 21 — Keiko Satoh; 16, 17, 19 — Lance Hayashida; 17 —Eugene Lubenov and Thanos Siapas

Athanassios (Thanos) Siapas is a professor of computation and neural systems. His work is funded by the Bren Foundation, the McKnight Foundation, the James S. McDonnell Foundation, the Whitehall Foundation, and the National Institutes of Health.

John O'Doherty is a professor of psychology. His work is funded by the NSF, the NIDA, and the Gordon and Betty Moore Foundation.

Colin Camerer is the Kirby Professor of Behavioral Economics. His work is funded by the NSF, the Gordon and Betty Moore Foundation, the Lipper Family Foundation, a Tamagawa University Global Center of Excellence grant, and the Trilience Foundation.

Mary Kennedy is the Davis Professor of Biology. Her work is funded by the National Institute of Mental Health, the National Institute of Neurological Disorders and Stroke, the National Institute on Drug Abuse (NIDA), the Gordon and Betty Moore Center for Integrative Study of Cell Regulation, the Della Martin Foundation, Allen and Lenabelle Davis, the Hereditary Disease Foundation, the CHDI Foundation, the Hicks Foundation, the John Douglas French Foundation for Alzheimer's Research, the Joyce Foundation, the National Science Foundation (NSF), and the Howard Hughes Medical Institute.

LEARNING GONE WRONG

Learning isn't all cupcakes and bike rides. Somewhere amid the synapses and theta waves, things can go awry: memories melt away, habits turn obsessive, positive rewards turn into addictions.

"So often, mental illnesses are derangements of the brain's regulatory behavior, defects in this machinery," notes Mary Kennedy. "If we want better and more specific drugs to treat these illnesses, we really need to understand what is going on at every level."

Thanos Siapas is particularly concerned about vulnerabilities in the ever-so-delicate memory-making circuitry of the hippocampus. "That's a huge part of the impetus to study these systems in such detail," he says. "We want to know exactly how learning and memory work, so that we can build machines that can learn as much as humans, or repair memory problems like Alzheimer's."

O'Doherty, for his part, is focused on how just a little bit of "overexuberance" on the part of a person's habit system could lead to obsessive-compulsive disorder or addiction. O'Doherty notes, "For instance, the habitual learning system could get hijacked by drugs of abuse. If a smoker has a cigarette every time she has a cup of coffee, drinking coffee will become a cue that will signal the response of lighting up a cigarette."

And yet, addiction isn't a given for every single person who picks up a cigarette, nor does every Las Vegas visitor wind up a compulsive gambler. Which is why it's important to tease apart the differences in behavior and brain wiring between folks who become addicted to gambling or nicotine or recreational drugs and those who don't. Says O'Doherty, "We need not only to investigate how learning goes wrong, but to look at the people in which it goes wrong," as he, Professor of Economics and Neuroscience Antonio Rangel (BS '93), and others are beginning to do. E&S