{neuroscience}

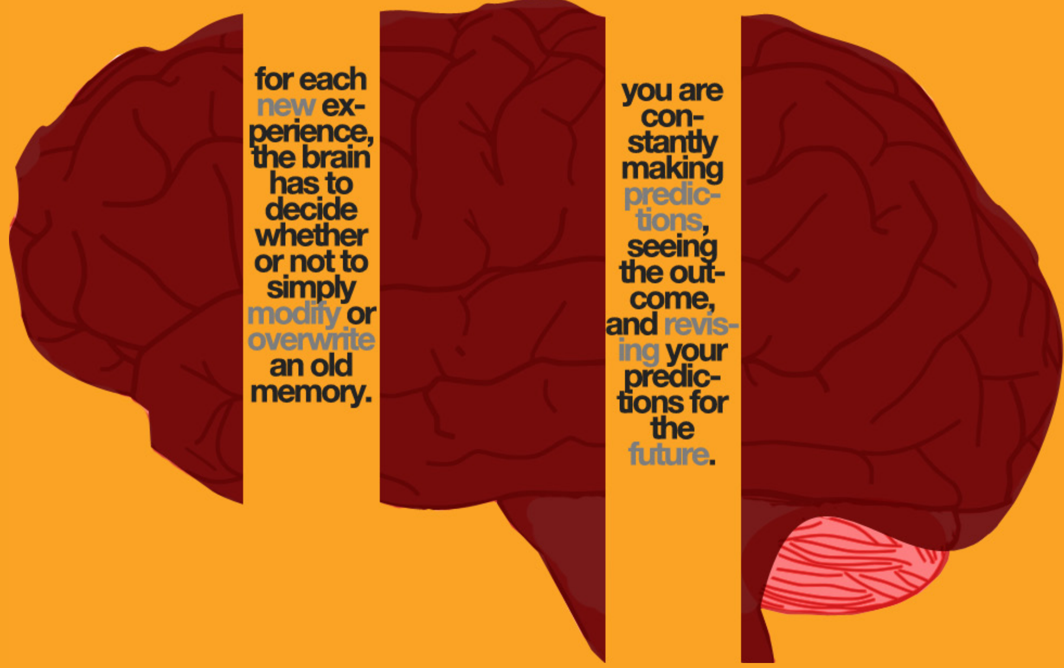

learning what to learn

by **ANNE MASLAN** *with* **Prof. YAEL NIV (NEU/PSY)**

"Even if I love pasta, I will not order it at a Chinese restaurant," says Professor Niv. The simple task of choosing an entrée on a restaurant's menu requires your brain to take a myriad of considerations into account. Did I like this dish last time? Would this food be good at this type of restaurant? What would be the freshest option? Yet at the same time, your brain must filter out all other stimuli that don't affect that given decision, like the color of the tablecloth. How does your brain determine what's important?

In the lab, learning is usually studied in very specified situations with one stimulus, for example: a study in which a rat receives food for pressing on a lever. This type of learning from feedback is called "reinforcement learning." The main idea is that people are constantly making predictions (for instance, about how good the ordered entrée will be, or, for rats, whether pressing the lever will result in food falling into the food cup). Then, when the person experiences what actually happens in the world, they can compute a "prediction error" - the difference between the prediction and the result. This prediction error enables people to revise their predictions (and actions) in the future according to this error.

Professor Niv's lab focuses on this constant process of sorting the relevant from the spurious – a distinction that is central to real world learning. Professor Niv notes, "In the real world, the 'stimulus' is not a simple light or a specific action. It is everything around you. How do learn efficiently about so many things at once? Luckily, you don't have to, as most of what's around you is not relevant to your current task… we're trying to understand how we learn to attend only to the aspects of the environment that are relevant, in order to efficiently learn a task." If a person focused on what table they were sitting at while deciding which entrée to order, that would not be efficient, as they would have to learn anew, at each table in the restaurant, what dishes they like best. In Professor Niv's words, her lab is studying the central question: "How do we learn what to learn about?"

In order to study real-world learning, Professor Niv's lab tests how people solve simple decision-making tasks. These tasks are equivalent to small puzzles, such as

"YOU ARE CONSTANTLY MAKING PREDICTIONS, SEEING THE OUTCOME, AND REVISING YOUR PREDICTIONS FOR THE FUTURE."

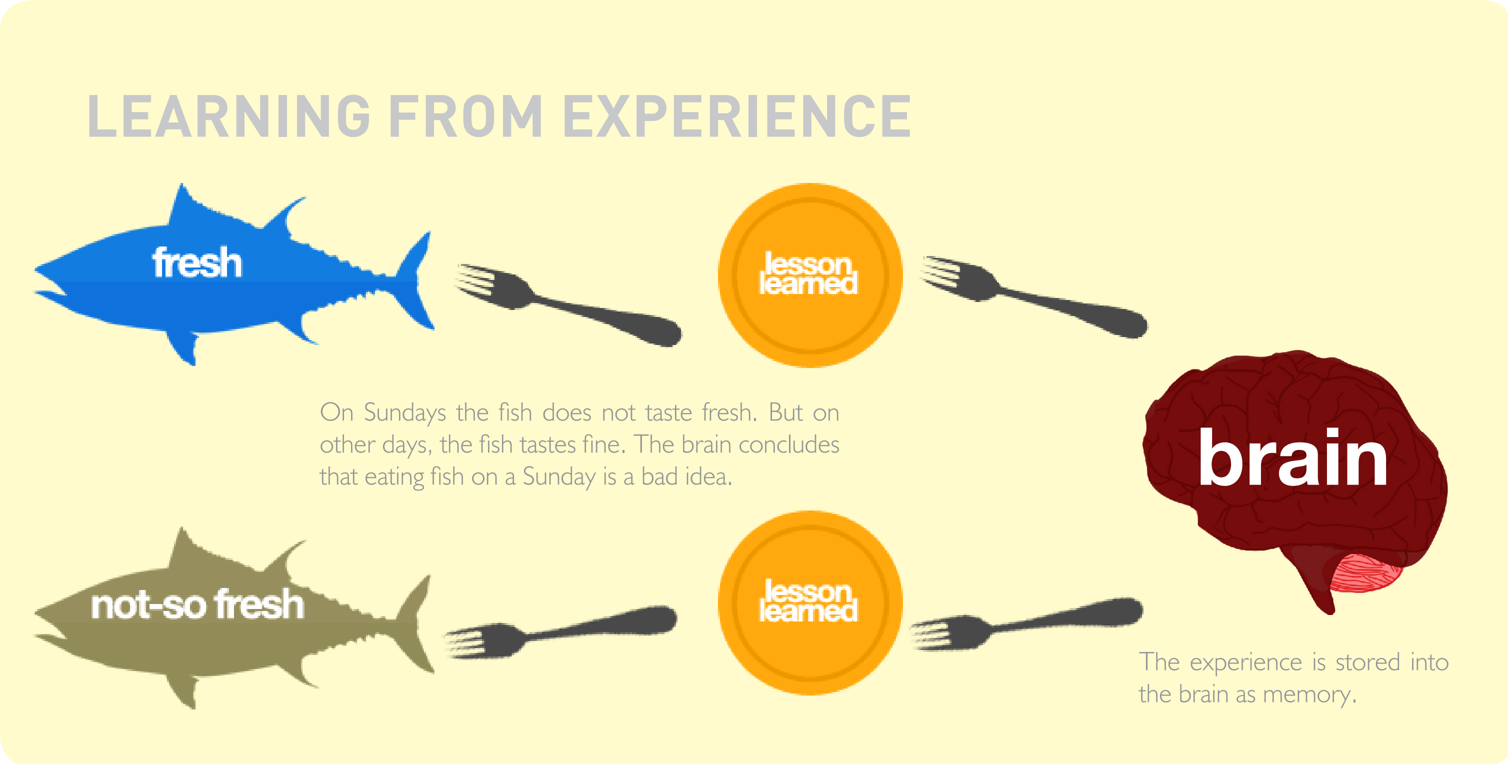

estimating the number of circles on a computer screen. After each trial, the participant is shown the correct answer. By measuring how the feedback from one trial affects what the participant pays attention to on the next trial, Professor Niv is able to track learning curves and begin to unravel how the brain learns what matters.

While the brain must be able to reduce the situation at hand to important information, it must also be able to add hidden dimensions that determine reward. Let's return to the restaurant example. Suppose you order fish at a certain restaurant you attend regularly. You start to notice that some days the fish is delicious and fresh, and others do not taste as superior. Your brain may begin to detect a pattern. It's common knowledge among food connoisseurs that ordering fish on a Sunday night is a bad idea because there is no delivery on Sundays. The cause of the fishiness was likely not immediately apparent to you, but was rather a hidden cause.

To understand how hidden causes affect our learning and decision-making, Professor Niv seeks to understand what future actions each bit of learning will affect and how people decide to separate or combine experiences. In Professor Niv's words: "We carve our experiences into groups, and we learn about each group separately; we don't just lump everything together. You can think about each group as a hidden cause."

"YOUR BRAIN IS CONSTANTLY TRYING TO FIND PATTERNS AND DEDUCE CAUSE AND EFFECT."

This grouping by similarity has tremendous implications when considering the distinction between updating an old memory and creating a new one. People must decide whether what they learn from a specific experience should modify an old memory of a similar experience or create a new memory of an entirely distinct experience. For example, if your friend has an immense fear of spiders, you may try to reassure your friend that most spiders are in fact harmless. However, this will not help your friend update an old memory; your friend will instead create a new memory in the context of a conversation with you, and the phobia will remain. Professor Niv's research suggests that if you really want to help your friend get over the phobia, you should create a new experience that is as similar as possible to those that made your friend afraid in the first place, so this experience won't seem new but rather will be "lumped together" with the old memory. This process will help the initial memory to be gradually rewritten.

The intuition behind this type of memory re-wiring relates to the story of the frog in hot water: if a frog is placed in hot water, it will immediately jump out. However, if the frog is placed in cold water that is slowly heated, the frog will not sense the danger and will remain in the pot. Rewriting memories is similar. The new experience must be introduced gradually and in a way that is similar enough to the original experience to be recognized and grouped with the old memory. Therefore, in order to learn from past experiences, the learned information must be grouped with the other relevant information, so that the learning can be translated to future similar situations.

By seeking to understand how the brain learns what to learn about, the potential for tailoring teaching for individual learning styles is a possible next step. Directing people's attention to relevant things and avoiding the extraneous can help people learn better and ultimately improve their decision-making process.