New research shows why you don't need to be perfect to get the job done

June 24 2024

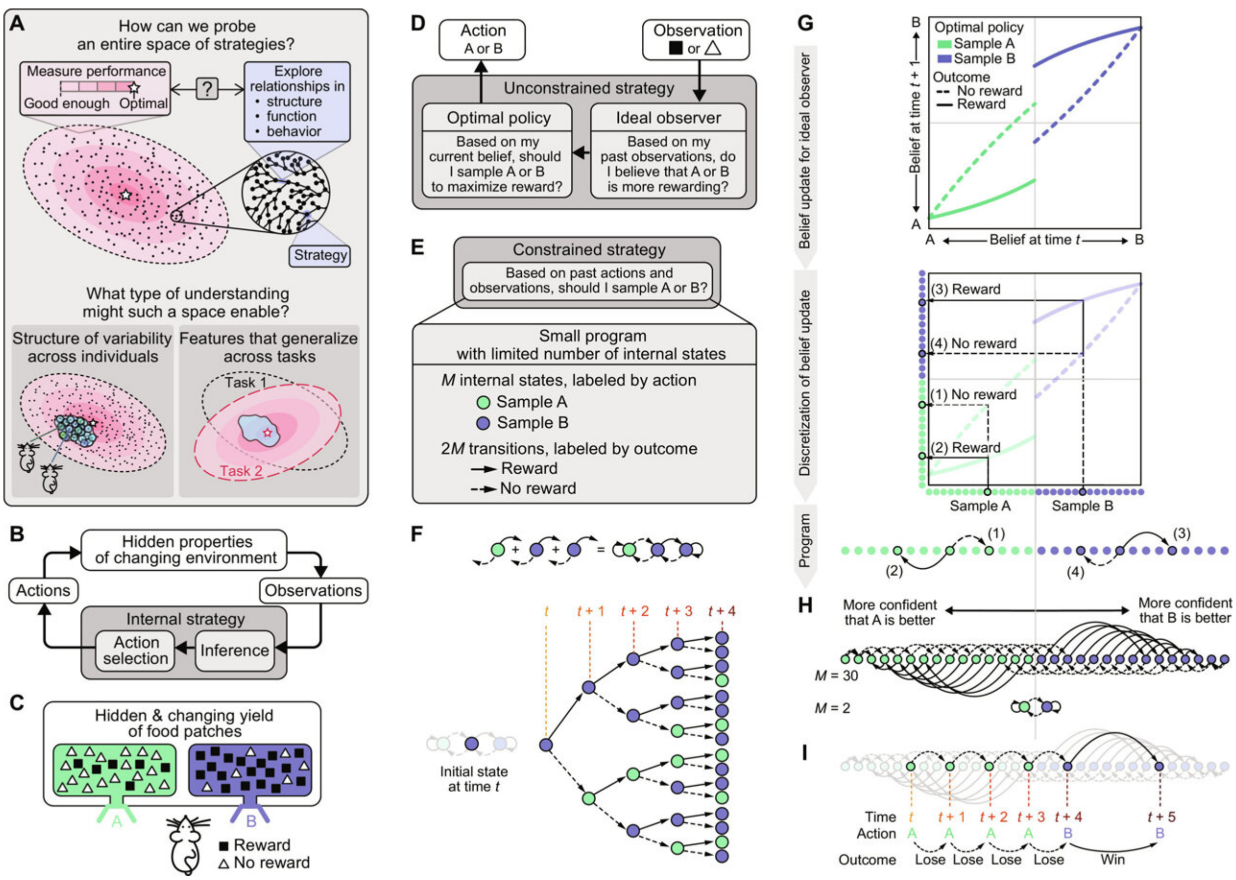

Constructing compact behavioral programs. (A) Top: The space of strategies for solving a task can be large, with many strategies that achieve good enough performance. Bottom: Studying relationships between strategies could provide insight into behavioral variability across animals and tasks. (B) General task setup: An animal makes inferences about hidden properties of the environment to guide actions. (C) Specific task setup: An animal forages from two ports whose reward probabilities change over time. (D) The optimal unconstrained strategy consists of an optimal policy coupled to a Bayesian ideal observer. (E) We formulate a constrained strategy as a small program that uses a limited number of internal states to select actions based on past actions and observations. (F) Each program generates sequences of actions depending on the outcomes of past actions. (G) The optimal unconstrained strategy (D) can be translated into a small program by discretizing the belief update implemented by the ideal Bayesian observer and coupled to the optimal behavioral policy. Top: Optimal belief update. Middle: Belief values can be partitioned into discrete states (filled circles) labeled by the action they specify (blue versus green). The belief update specifies transitions between states, depending on whether a reward was received (solid versus dashed arrows). Bottom: States and transitions represented as a Bayesian program. (H) Top: A 30-state program approximates the Bayesian update in (G) and has two directions of integration that can be interpreted as increasing confidence about either option. Bottom: The two-state Bayesian program, win-stay, lose-go (WSLG), continues taking the same action upon winning (i.e., receiving a reward) and switches actions upon losing (i.e., not receiving a reward). (I) Example behavior produced by the 30-state Bayesian program in (H). Credit: *Science Advances* (2024). DOI: 10.1126/sciadv.adj4064

When neuroscientists think about the strategy an animal might use to carry out a task—like finding food, hunting prey, or navigating a maze—they often propose a single model that lays out the best way for the animal to accomplish the job.

But in the real world, animals—and humans—may not use the optimal way, which can be resource-intensive. Instead, they use a strategy that's good enough to do the job but takes a lot less brain power.

In new research appearing in *Science Advances*, Janelia scientists set out to better understand the possible ways an animal could successfully solve a problem, beyond just the best strategy.

The work shows there is a huge number of ways an animal can accomplish a simple foraging task. It also lays out a theoretical framework for understanding these different strategies, how they relate to each other, and how they solve the same problem differently.

Some of these less-than-perfect options for accomplishing a task work nearly as well as the optimal strategy but with a lot less effort, the researchers found, freeing up animals to use precious resources to handle multiple tasks.

"As soon as you release yourself from being perfect, you would be surprised just how many ways there are to solve a problem," says Tzuhsuan Ma, a postdoc in the Hermundstad Lab, who led the research.

The new framework could help researchers start examining these "good enough" strategies, including why different individuals might adapt different strategies, how these strategies might work together, and how generalizable the strategies are to other tasks. That could help explain how the brain enables behavior in the real world.

"Many of these strategies are ones we would have never dreamed up as possible ways of solving this task, but they do work well, so it's entirely possible that animals could also be using them," says Janelia Group Leader Ann Hermundstad. "They give us a new vocabulary for understanding behavior."

Looking beyond perfection

The research began three years ago when Ma started wondering about the different strategies an animal could possibly use to accomplish a simple but common task: choosing between two options where the chance of being rewarded changes over time.

The researchers were interested in examining a group of strategies that fall between optimal and completely random solutions: "small programs" that are resource-limited but still get the job done. Each program specifies a different algorithm for guiding an animal's actions based on past observations, allowing it to serve as a model of animal behavior.

As it turns out, there are many such programs—about a quarter of a million. To make sense of these strategies, the researchers first looked at a handful of the top-performing ones. Surprisingly, they found they were essentially doing the same thing as the optimal strategy, despite using fewer resources.

"We were a little disappointed," Ma says. "We spent all this time searching for these small programs, and they all follow the same computation that the field already knew how to mathematically derive without all this effort."

But the researchers were motivated to keep looking—they had a strong intuition that there had to be programs out there that were good but different from the optimal strategy. Once they looked beyond the very best programs, they found what they were looking for: about 4,000 programs that fall into this "good enough" category. And more importantly, more than 90% of them did something new.

They could have stopped there, but a question from a fellow Janelian spurred them on: How could they figure out which strategy an animal was using?

The question prompted the team to dive deep into the behavior of individual programs and develop a systematic approach to thinking about the entire collection of strategies. They first developed a mathematical way to describe the programs' relationships to each other through a network that connected the different programs. Next, they looked at the behavior described by the strategies, devising an algorithm to reveal how one of these "good enough" programs could evolve from another.

They found that small changes to the optimal program can lead to big changes in behavior while still preserving performance. If some of these new behaviors are also useful in other tasks, it suggests that the same program could be good enough for solving a range of different problems.

"If you are thinking about an animal not being a specialist who is optimized to solve just one problem, but rather a generalist who solves many problems, this really is a new way to study that," Ma says.

The new work provides a framework for researchers to start thinking beyond single, optimal programs for animal behavior. Now, the team is focused on examining how generalizable the small programs are to other tasks, and designing new experiments to determine which program an animal might be using to carry out a task in real time. They are also working with other researchers at Janelia to test their theoretical framework.

"Ultimately, getting a strong grasp on an animal's behavior is an essential prerequisite to understanding how the brain solves different types of problems, including some that our best artificial systems only solve inefficiently, if at all," Hermundstad says. "The key challenge is that animals might be using very different strategies than we might initially assume, and this work is helping us uncover that space of possibilities."

More information: Tzuhsuan Ma et al, A vast space of compact strategies for effective decisions, *Science Advances* (2024). DOI: 10.1126/sciadv.adj4064

Provided by Howard Hughes Medical Institute