

Exploration vs. Exploitation: Adults Are Learning (Once Again) From Children

April 06, 2023



Image credit: A model display of Luxo Jr. in Pixar Fest exhibition, Ocean Terminal, Harbour City, Tsim Sha Tsui, Kowloon, Hong Kong, July 2021. [Achanhk](#), [CC BY-SA 4.0](#), via Wikimedia Commons. Regarded as a breakthrough in computer animation, [Luxo Jr.](#) is a two-minute film produced and released by Pixar in 1986 in which a resilient baby desk lamp and his “parent” play with a ball.

Back in the seventies, my husband, Alvy Ray Smith, was one of the bright young inventors at the famous Xerox PARC, the Palo Alto Research Facility. Alvy was both an artist and an engineer, and he helped design one of the first computers that could make color pictures. Then the bad news came. Xerox management decided they no longer required his services because there was no need for color images. “But color is the future!” he protested. “That may well be,” they said, “but our business is black-and-white copiers.” In fact, a lot of the fundamentals of personal computing, like the mouse and the windows interface, were developed at Xerox Parc but then ignored by central office executives. It took an obscure young guy named Bill Gates at a little company called Microsoft to see their potential.



Alison Gopnik

But when Alvy tells this story, there is a coda. A decade later, he became an executive himself when he cofounded a company called Pixar. And then the young genius engineers began coming to him with exciting new ideas about what computers could do. Rather sadly remembering his young self, he would have to remind them to stick to the company business.

Alvy's story exemplifies a fundamental tension that has recently been the focus of a lot of exciting research in cognitive science and psychology: the trade-off between exploration and exploitation. This tension applies far beyond [entrepreneurship](#), although it's still very relevant there. How do you balance [innovation](#) and implementation, possibility and practicality? How do you resolve the tension between the lure of the crazy new thing and the safe haven of the tried and true?

Computer scientists have formulated this tension mathematically in terms of a search through a high-dimensional space of possibilities. Imagine a very big box of possible actions you can take. You begin at a particular point in that box—that is, you begin with a particular set of techniques for solving familiar problems. When you're faced with a new problem, you move through the box looking for new solutions. You can search in a more narrow and focused way, close to where you started, to quickly find solutions that are just good enough—the exploit option. Or you can search more widely and variably, trying out options that are very different from what you've done before, to try to find the very best solution—the explore option.

A broader, more exploratory search allows for a wider range of possibilities and provides more information. But you may waste a lot of time considering weird solutions that are actually worse than the current one. A narrower, more exploitative search is more likely to lead to an effective solution quickly. But it may leave you stuck in what's called a local optimum. In local optima, all small changes will lead to worse options, although a large shift might lead to a substantially better one.

[See all articles from this issue of the *Observer*.](#)

The Xerox PARC inventions are a nice example. Think about trying to figure out how to write with a computer. You might try small variations on the familiar keyboard —after all, that's what you used for printing and typewriting. Some of us still remember typing ghostly green combinations of letters after

the DOS prompt to try to get the damn machine to work, stuck in the keyboard local optimum. A mouse that you glide around on your desk was very different from that familiar keyboard. It's much further away in the solution space, and it took a leap to imagine how you could use it to cut or copy or edit text.

Exploitation allows you to accumulate resources and succeed in the short run. Quickly narrowing solutions to a single option lets you focus on implementing that option effectively. Exploration, on the other hand, is expensive and requires resources to support you while you explore (the central executive office has to subsidize the R&D division). But exploration may pay off eventually, particularly when the environment is complex and time horizons are long. As Alvy said, color *was* the future.

The bad mathematical news is that there is no simple way to resolve this trade-off—no way to maximize the benefits of exploration and exploitation simultaneously. However, there are different strategies that help balance exploration and exploitation and ensure that exploration takes place despite its short-term cost. Often these strategies involve starting out by exploring and then narrowing in to exploit.

Exploration may pay off eventually, particularly when the environment is complex and time horizons are long. As Alvy said, color was the future.

For example, “simulated annealing” is an algorithmic technique based on the physics of annealing in metallurgy. Heating up a metal and then cooling it makes it more robust. Similarly, machine learning systems begin with a higher “temperature” and a noisier, more random search before gradually “cooling off” to a more detailed and focused search. The learner is like a molecule bouncing around in that big box of potential solutions. If you start by turning the temperature up high, the molecule will move around quickly to many parts of the box, trying a wide variety of options but never settling for long. As the temperature goes down, the molecule moves more slowly until it lands at a particular spot. Simulated annealing plays an essential role in many machine-learning techniques.

It's always tempting to take a good dichotomy—implicit versus explicit, nature versus nurture, introvert versus extrovert—and apply it everywhere, especially when it comes with a tasty rhyme or a nice alliteration. As someone once said, there are two kinds of people: the ones who divide everything into dichotomies and the ones who don't. But the explore–exploit contrast really does help us understand a remarkably wide range of psychological phenomena.

Reinforcement learning is a classic psychological idea, but it also plays an important role in modern artificial intelligence and neuroscience. An agent takes a particular action and notes the result—if it leads to a reward they should repeat that action; if it doesn't, they shouldn't. But if the agent only tries actions that were successful before (the exploit strategy) they might never discover an even more effective new possibility. That takes exploration. So sometimes a machine or an animal or a person should try something new even if they don't know how it will work out. Computer scientists, cognitive scientists, and neuroscientists have worked out formal models of the best way to do this. But in general, it's best to do more exploration early, and then narrow in on the most rewarding actions.

I've argued that biological organisms may employ a similar technique over their lifespan. They begin with a protected exploratory period—[childhood](#), when resources are provided by caregivers—and gradually shift to a more competent but constrained period of exploitation. My slogan is that childhood

is evolution's way of resolving the explore–exploit trade-off and performing simulated annealing. In fact, many characteristics of children that look like bugs from the exploit perspective may be features from the explore perspective. Children are notoriously unfocused, deficient in executive function and long-term planning, and both literally and metaphorically noisy. This makes children bad at acting effectively but good at [learning](#), exploration, and discovery.

[View past presidential columns](#)

Of course, children can only do this because they have caregivers who look after them when they're in the R&D phase. Children really only need one exploit strategy: Be as cute as you possibly can be and get people to love you and take care of you. Fortunately, they are extremely effective at that.

You can see this explore–exploit sequence in many different areas of development. In my lab, we've shown that young children can learn about unusual causal systems better than adults can. In language learning, infants begin with the ability to learn all the phonetic contrasts in the world's languages. As they grow older and gain more experience, they narrow the range of possibilities they consider, and as adults they have a hard time even hearing distinctions in other languages, let alone learning to use them. In both these cases, an early broader search allows more exploration of the potential space, but a more finely tuned system is more efficient later.

The explore–exploit sequence shows up in studies of attention and memory, too. Adults are better than children at paying attention to events that are relevant to their goals, and they remember relevant events better too. But new studies show that children are better at noticing and remembering incidental events—information that isn't relevant now but might let you learn something that will be useful later. We say that children have trouble paying attention, but they really have trouble *not* paying attention. For them, something interesting always beats out something important.

You can see the same explore–exploit sequence in [brain development](#). In infancy and early childhood, many new neural connections are formed and our brain is more “plastic”—it changes more easily. As we grow older, the connections we use a lot become stronger and more effective, and the other possibilities are pruned away. Psychologists like APS William James Award Fellows Janet Werker and Patricia Kuhl have shown how these brain developments are related to the early changes in language development. In the same way, the executive prefrontal cortex is undeveloped at first but gradually comes to exert more control over the rest of the brain. Prefrontal control is essential for effective action, but psychologists like Sharon Thompson-Schill, another APS Fellow, have shown that it may limit creativity and exploration.

We say that children have trouble paying attention, but they really have trouble not paying attention. For them, something interesting always beats out something important.

This may all sound rather discouraging for grown-ups who seem stuck with the dull exploit role while

kids have all the fun. But exploration and exploitation are complementary too. Children couldn't explore without adults who exploit, and all of us reap the rewards of focused attention and effective action. Perhaps it's comforting that some of the most effective computer-science techniques use repeated cycles of exploration and exploitation rather than a single shift. Human adults have also found ways to periodically return to a state of more child-like exploration, plasticity, and openness. Places like Xerox PARC in the seventies or Bell Labs in the sixties provided protected spaces for engineers to play. Meditation, mystical experiences, even psychedelics may restore a wider focus and more exploratory attitude. Or you might try my favorite technique and spend some time playing with a 2-year-old.

***Feedback on this article? Email apsobserver@psychologicalscience.org or login to comment.
Interested in writing for us? Read our [contributor guidelines](#).***