Edge

To arrive at the edge of the world's knowledge, seek out the most complex and sophisticated minds, put them in a room together, and have them ask each other the questions they are asking themselves.

https://www.edge.org/response-detail/25530
Printed On Wed January 11th 2023

2014 : WHAT SCIENTIFIC IDEA IS READY FOR RETIREMENT?

In the News [60]
|
Contributors [177] | View All Responses [184]

Oliver Scott Curry
Senior Researcher, Director, The Oxford Morals Project, Institute of Cognitive and Evolutionary Anthropology, University of Oxford

Associationism

How do birds fly? How do they stay up in the air? Suppose a textbook told you that the answer was 'levitation', and proceeded to catalogue the different types of levitation (Stationary, Mobile), its laws ("What goes up must come down", "Lighter things levitate longer") and constraints (Quadrupedalism). You'd rapidly realise that flying was not well understood, and also that the belief in levitation was obscuring the need for, and holding back, a proper scientific account of aerodynamics.

Unfortunately, a similar situation applies to the question 'How do animals learn?'. Textbooks will tell you that the answer is 'association', and will proceed to catalogue the various types (Classical, Operant), its laws (Rescorla-Wagner), and constraints (Autoshaping, Differential Conditionality, Blocking). You will be told that association is the ability of organisms to make connections between any given stimulus and any given outcome or response—the sound of a bell with the arrival of food, or the left-branch of a maze with the administration of pain—merely through (repeated) exposure to their pairing. And you will be told that, because association treats all stimuli equally, it can in principle enable an organism to learn anything.

The problem is that, as with levitation, no-one has ever set out a mechanism that could perform such a feat. And no-one ever will, because such a mechanism is not possible in theory, and hence not possible in practice. At any given time, an organism is confronted by an infinite number of potential stimuli, and subsequently, an infinite number of potential outcomes. A day in the life of a rat, for example, might include waking up, blinking, walking east, twitching its nose, being trampled on, eating a berry, hearing a rumbling noise, sniffing a mate, experiencing a temperature of 5°C, being chased, watching the sun go down, defecating, feeling nauseous, finding its way home, having a fight, going to sleep, and so on. How does the rat discern that, of all the possible combinations of stimuli and outcomes, it was the berry alone that made him feel sick? Just as answers presuppose a question, data presuppose a theory. In the absence of a prior theory that specifies what to look for, and which relationships to test, there is no way of sorting through this chaos to identify useful patterns. And yet what is the defining feature of associative learning? It is the absence of a prior theory. So, like levitation, association is hollow—a misleading redescription of the very phenomenon that is in need of explanation.

Critics have, for centuries, pointed out this problem with associationism (sometimes called the problem of induction, or the frame problem). And, in recent decades, there have been countless empirical demonstrations that animals—ants learning their way home, birds learning song, or rats learning to avoid food—do not learn in the way that associationism suggests. And yet, associationism (whether as empiricism, behaviourism, conditioning, connectionism, or plasticity) refuses to die, and keeps rising again, albeit encrusted by ever more ad hoc exceptions, anomalies and constraints. Its proponents refuse to abandon it, perhaps because they believe there is no alternative.

But there is. In communication theory, information is the reduction of prior uncertainty. Organisms are 'uncertain' because they are composed of conditional adaptations that adopt different states under different conditions. These mechanisms can be described in terms of the decision rules that they embody—'if A, then B', or 'If you detect light, then move towards it'. Uncertainty about which state to adopt (to B or not to B), is resolved by attending to the specified conditions (A). The reduction of uncertainty by one half constitutes one 'bit' of information; and so a single decision rule is a one-bit processor. By favouring adaptations with more branching decision rules, natural selection can design more sophisticated organisms that engage in more sophisticated information processing, asking more questions of the world before coming to a decision. This framework explains how animals acquire information and learn from their environments. For the rat, a rule is, "If you ate something and subsequently felt sick, then avoid that food in future"; it has no such rule fingering sunsets, nose twitching, or fighting, which is why it never makes those connections. Similarly, this account explains why organisms facing different ecological problems, composed of different clusters of such mechanisms, are able to learn different things.

So much for rats. What about humans, who obviously can learn things that natural selection never prepared them for? Surely we must be able to levitate? Not at all; the same logic of uncertainty and information processing must apply. If humans are able to learn novel things, then this must be because they are able to generate novel uncertainty—to invent, imagine, create new theories, hypotheses and predictions, and hence to ask new questions of the world. How? The most likely answer is that humans have a range of innate ideas about the world (to do with colour, shape, forces, objects, motion, agents and minds), which they are able to recombine (almost at random) in an endless variety of ways (as when we dream), and then test these novel conjectures against reality (by means of the senses). And successful conjectures are themselves recombined, and revised, to build ever more elaborate theoretical systems. So, far from constraining learning, our biology makes it possible: providing the raw materials, guiding the process to a greater or lesser degree, liberating us to think altogether unprecedented thoughts, and fostering the growth of knowledge. This is how we learn from experience—and all without a whiff of association.

Look, nobody disputes that birds fly; the only question is how. Similarly, nobody disputes that humans and other animals learn; the only question is how. Working out the alternative account of learning will involve identifying which innate ideas humans posses, what rules are used to combine them, and how they are revised. But for this to happen, we must first accept not only that association is not the answer, but that association is not even an answer. Only then will the science of learning stop levitating, and take off for real.

Return to Table of Contents