Kinds of Intelligence

Marta Halina – *Department of History and Philosophy of Science, University of Cambridge*

Table of Contents

Kinds of Intelligence

Marta Halina – *Department of History and Philosophy of Science, University of Cambridge*

Abstract: A common view of intelligent behavior is that it is underpinned by sophisticated mechanisms like imagination, foresight, and casual reasoning, and that such behavior is not a product of innate behavioral programs or associative learning. In this Element, I present this view and argue that it is a mistake. Such a view oversimplifies what we know about innate priors and associative learning—namely, that they scaffold flexible behaviors in complex ways. Rather than attempt to categorize the natural world into organisms that are intelligent on the one hand and rely on rigid, mechanical processes on the other, we should focus on specific mechanisms and behaviors that take an organism's evolutionary, developmental, and ecological context into account. The term "intelligence" should be recognized as an umbrella term that orients researchers towards a broad class of phenomena but is too general for detailed scientific work.

Introduction

In 1770, Wolfgang von Kempelen unveiled to the Viennese court what would become one of the most famous machines in history. The device was "filled with wheels, cylinders, levers, and other pieces of clock-work" (Windisch 1784, p. 22). Kempelen was a person of many skills—he spoke eight languages, studied philosophy and law in Vienna, designed waterworks in Hungary, directed salt mines in Transylvania, and had many other scientific and engineering achievements (Standage 2002). Although it had taken von Kempelen only six months to construct this machine, it would ultimately be displayed around the world, fascinating general audiences, as well as famous figures like Napoleon Bonaparte, Charles Babbage, Benjamin Franklin and Catherine the Great (Standage 2002, p. xii).

The machine that Kempelen built was a mechanical chess player, known as "the Turk". It was famous for having chess-playing abilities that surpassed even the strongest human players, and for the remarkable ability to flexibly adapt to new situations. For example, during the Turk's chess-playing performances, the audience was invited to place the knight on any square on the board. The automaton would then solve the Knight's Tour—visiting every square on the board just once—starting from that position. The Knight's Tour (also known as "Euler's problem" after the mathematician Leonhard Euler) is a difficult problem to solve and was of great interest to mathematicians at the time. To be able to defeat a diverse range of players in chess, as well as solve problems like the Knight's Tour without knowing the starting position in advance, demonstrated a profound ability to flexibly adapt to new situations (Schaffer 1999).

The mechanical Turk fascinated the world by achieving the unachievable: how could a purely mechanical system, one made of cogs and levers, engage in sophisticated and flexible behavior? For many spectators and commentators, this was impossible. For example, the author and traveler, Philip Thicknesse wrote: "That an AUTOMATON may be made to move its hand, its *head*, and its *eyes*, in *certain and regular motions*, is past all doubt; but that an AUTOMATON can be made to move the Chessmen properly, as a *sagacious Player*, in *consequence of the preceding move of a stranger*, who undertakes to play against it, is also UTTERLY IMPOSSIBLE" (1784, p. 5, emphasis original). Given this impossibility, and the fact that the Turk did indeed exhibit such flexible behavior, Thicknesse concluded that this chess-playing automaton must be the result of a trick.

The academic and mechanical engineer, Robert Willis, was similarly skeptical. After viewing the Turk on several occasions, he wrote:

> The phenomena of the Chess Player are inconsistent with the effects of mere mechanism; for, however great and surprising the powers of mechanism may be, the movements which spring from it, are necessarily limited and uniform: it cannot usurp and exercise the faculties of mind; it cannot be made to vary its operations, so as to meet the ever-varying circumstances of a game of chess. This is the province of intellect alone" (1821, p. 11)

Again, a mere mechanism or genuinely autonomous machine can only engage in a limited range of fixed behaviors. Given the Turk's flexibility and adaptability then, its behavior must be controlled or guided by a human (see Standage 2002).

Compare the mechanical Turk with another source of broad public fascination. In the 1970s, it became popular for arcades (like the Chinatown Fair arcade on Mott Street, New York) to include games in which a chicken competed against customers in a game of tic-tac-toe. One of the remarkable things about these chickens was that they seldom lost. Indeed, one Las Vegas casino operator suggested offering a hundred thousand dollars to anyone who succeeded in defeating a chicken in one of these games (an amount later reduced to ten thousand dollars) (Trillin 1999). These human-versus-chicken games were first produced by a company called Animal Behavior Enterprises, founded in Hot Springs, Arkansas, by the psychologists Keller and Marian Breland. The Brelands had trained under the behaviorist B. F. Skinner and were amongst the first psychologists to commercially apply the methods of operant conditioning to animals (Breland & Breland 1951, Bailey & Gillaspy 2005). They had also worked with Skinner to develop the first "smart bombs" guided by trained pigeons (Drumm 2009). Skinner had himself played several games against a tic-tac-toe playing chicken at an annual meeting of the Association for Behavior Analysis and lost (Bihm et al. 2010).

In both the case of the mechanical Turk and the tic-tac-toe playing chicken, it turns out that the seemingly intelligent behaviors exhibited by the automaton and bird were the result of a trick. In the case of the mechanical Turk, Kempelen had created an ingenious set of compartments, functional mechanisms, and decoy mechanisms inside the cabinet on which the Turk played chess. This allowed a human to remain concealed inside the cabinet despite the cabinet doors being opened during the performance for all to see what was inside. The person inside the cabinet could operate the Turk's arm and hand using a sophisticated system of levers. The concealed operator had an internal chessboard that corresponded to the chessboard on top of the cabinet: moving a pointer to a square on the internal chessboard resulted in the Turk's hand being guided to the corresponding square on the external chessboard (Standage 2002).

In the case of the tic-tac-toe playing chicken, the bird was also not the one selecting which move to make against her human competitor. Instead, the chicken had been trained to peck at a switch in response to a flashing cue light (Animal Behavior Enterprises n.d.). When the light flashed, the chicken pecked a single switch which trigged a computer program to select the next move. As the Bird Brain manual states, "Because the chicken plays first, and because her selections are actually made by electronic circuitry, the customer can do no more than tie the game" (Animal Behavior

Enterprises n.d., p. 2). After triggering the last move in the game, the chicken was rewarded with a few grains of food.

The mechanical Turk and tic-tac-toe playing chicken created the illusion of human-like intelligent behavior produced by nonhuman agents—in one case a machine and in the other case a bird. However, probing these two systems reveals that either the intelligent behavior is a direct result of human action (produced by a human secretly operating the device) or an indirect result of human action (a simple computer program designed by a human). More crucially, when stripped of these human guides, we are left with simple mechanical systems, unable to flexibly respond to new situations. Although such systems can be directed to behave intelligently, they are incapable of producing intelligent behavior on their own.

The idea that organisms can be divided into those that are "intelligent" on the one hand, and those that follow simple, mechanical procedures on the other is what I will refer to as the "Standard View" of intelligence. Humans provide one example of biological intelligence: they adapt flexibly to new situations, solve a wide range of problems, engage in complex social situations, and more. Under the Standard View, some organisms, like humans, can flexibility adapt to a highly variable environment by drawing on cognitive capacities such as imagination and causal reasoning. Such organisms are suitably characterized as "intelligent". Other organisms are rigid in their behavior and rely on mechanisms such as innate behavioral programs and associative learning. Such organisms are not intelligent—they work well in some situations but break down when required to function outside of a narrow range of environments. Intelligent agents, in contrast, do not rely on such preprogrammed rules or brute-force training to achieve their goals.

We can see the Standard View operating in comparative psychology. For example, in a review on intelligence in corvids (e.g., rooks, ravens, and crows) and apes (e.g., chimpanzees, bonobos, and orangutans), the comparative psychologists Amanda Seed, Nathan Emery and Nicola Clayton address the question, "what is intelligence"? In answering this question, they write:

> Ever since scholars began discussing animal intelligence it has been a highly divisive issue, and remains so today. At the poles of the debate are two opposite views concerning 'thinking' in animals. The first, the origin of which is attributed to Descartes, is that animals are essentially mindless machines, with their behaviour trigged wholly by external or internal stimuli. The other, most famously articulated by Darwin, is that 'the difference in mind between man and the higher animals… is one of degree and not of kind' (Seed et al. 2009, p. 402)

The comparative psychologist Daniel Hanus also traces current debates on animal intelligence to Descartes' claim that "all nonhuman beings must be best described as complex 'automata'" (Hanus 2016, p. 241; see also Riskin 2016) and the idea that "different cognitive mechanisms are at work in humans and other animals" (2016, p. 242). For our purposes, the crucial question is what role this dichotomy between "mindless machines" on the one hand, and "intelligent, minded creatures" on the other, plays in current scientific and philosophical debates on intelligence.

I argue in section 1 that a key premise of the Standard View of intelligence is that some organisms are best understood as mindless machines. These machines might be ingeniously designed by evolution or shaped by training such that they *appear* intelligent in some circumstances. However, like the mechanical Turk and tic-tac-toe playing chicken, this apparent intelligent behavior is the result of various "tricks". Such tricks can be revealed by modifying the environment and watching the machine break down.

In section 2, I argue that the Standard View of intelligence is mistaken. It is a caricature of the living world that radically distorts what we know about nonhuman animal cognition and behavior. It misleads us into oversimplifying the capacities and behaviors of many biological systems—systems portrayed as simple and rigid when in fact they are complex and exhibit varying degrees of flexibility. Recognizing the true complexity and flexibility of biological systems means rejecting the Standard View that nature can be divided into those animals that are intelligent on the one hand, and those that follow rigid, mechanical procedures on the other.

Once we reject the Standard View, what are the consequences for our use of the term "intelligence"? In section 3, I argue that insofar as the term is meant to broadly apply across species, it cannot be precisely defined. Instead, it serves as an umbrella term that helps vaguely orient researchers towards a class of phenomena that we know little about. Research on this class of phenomena is best undertaken at the level of more specific capacities, like learning and problem solving. Drawing on recent work in philosophy and cognitive science, I advance three strategies that can help us map the space of intelligent systems while moving beyond the Standard View: validating causal relationships within species, thinking of intelligent behavior as multidimensional, and focusing on signatures rather than success.

Why is it important to evaluate nonhuman animal (hereafter "animal") intelligence? Whether or not an animal is intelligent often informs how we think that animal should be practically and ethically treated. For example, some argue that the ability to act intentionally and flexibly is connected to autonomy, and autonomy is sufficient for personhood (Andrews et al. 2018). Under many current laws, such as those in the United States, nonhuman animals such as chimpanzees are legally "things", rather than "persons". As "things", they can be owned and sold like property (Wise 2010). To identify an entity as a "person" however, gives them moral standing and protection under law. As the legal scholar Steven M. Wise writes, "A court confronted with a plaintiff's claim to possess any legal right need only determine the plaintiff's species. If the plaintiff is human, the answer is, 'It is possible that the plaintiff has the legal right she claims.' If the plaintiff is a nonhuman animal, the answer is, 'Impossible'" (Wise 2010, p. 5). Recently, scientists and philosophers have argued that chimpanzees and other animals should be granted the status of "persons" because they have many sophisticated cognitive capacities (Andrews et al. 2018). Thus, our views on animal intelligence can dramatically change how we behave towards other organisms.

1. Intelligence: The Standard View

Consider the following scenario. A raven (*Corvus corax*) encounters a piece of food hanging from a long string tied to a tree branch. The raven would like to eat the food, but it's out of reach. After observing the situation for a minute, she perches above the dangling food, reaches down with her beak, and pulls up the string. This successfully brings the food closer, but it's still out of reach. The raven then takes the string she has pulled up and steps on it with her foot. This frees her beak, which she can now use to reach down and pull up another section of string. This brings the food even closer, but still not close enough. The raven repeats this procedure—pulling up a section of string, stepping on it with her foot while letting go of it with her beak, pulling up the next section of string, etc. until the food is finally within reach. She takes her reward and flies away.

The above situation was described by the biologist Bernd Heinrich (1995) and has since been studied in many different bird and mammal species (Jacob & Osvath 2015). Observing this situation, one might think, "that's clever!" and indeed string-pulling tasks have long been used to determine an animal's capacity for intelligent and insightful problem solving (Shettleworth 2012). In this case, it seems that the raven encounters a puzzle—an out-of-reach piece of food—and rather than knowing how to solve the problem immediately through innate programming or taking a long time to learn the solution through trial and error, she inspects the situation and devises a solution that involves a complex sequence of actions. She then executes this planned solution. The raven's behavior seems intelligent. In order to determine how to evaluate situations such as these, it is useful to define more clearly what one means by "intelligent behavior".

1.1 Intelligent Behavior

What does it mean for a sequence of behaviors to be intelligent? Is there a good account of what distinguishes intelligent from unintelligent behavior? When one surveys the cognitive-science literature on intelligence, a common theme emerges. Researchers regularly characterize intelligence as the ability to flexibly respond to a changing environment. The artificial intelligence researchers, Shane Legg and Marcus Hutter, for example, conclude in their review of the psychology literature on intelligence that, "*Intelligence measures an agent's ability to achieve goals in a wide range of environments*" (Legg & Hutter 2007, p. 12, emphasis original). The neurobiologists Gerhard Roth and Ursula Dicke note that researchers have "converged on the view that mental or behavioral flexibility is a good measure of intelligence, resulting in the appearance of novel solutions that are not part of the animal's normal repertoire" (2005, p. 250). Seed and colleagues concur, noting that "many agree that the function of intelligence is to produce flexible adaptive behaviour in the face of environmental complexity and variability" (Seed et al. 2009, p. 403).

The above definitions capture what many would characterize as a form of "general" or "domain-general" intelligence. General intelligence is often distinguished from "narrow" or "domain-specific" skills. One might be highly skilled in a specific domain—for example, play a particular form of chess very well. But such domain-specific skills are compatible with that same agent lacking general intelligence. Computer programs like AlphaGo, for example, vastly outperform humans in games like Go (Silver et al. 2017). They can even be understood as behaving creatively within the domain of a standard game of Go (Halina 2021). However, AlphaGo is unable to achieve goals outside of the context of a standard game of Go (Lake et al. 2017). It's "world model" applies only to this domain, thus we can understand it as lacking general intelligence insofar as we understand general intelligence as the ability to achieve goals in a wide range of domains. As the philosopher Ellen Fridland writes, "paradigmatically intelligent states are not tied to one role or context but can be transferred or applied in multiple roles and contexts" (2015, p. 149).[1] In what follows, I use the term "intelligence" to refer

[1] It is worth flagging that this distinction between domain-general and domain-specific intelligence

to domain-general intelligence of this kind.

Broadly then, we can characterize intelligent behavior as demonstrating the ability to flexibly achieve one's goal. It's worth highlighting two features of this account, which Fridland calls the "success condition" and the "flexibility condition". Fridland (2015) argues that these two conditions are necessary and sufficient for intelligence.[2] The success condition holds that intelligence enables one to successfully achieve one's goal. Thus, simply being able to flexibly respond to the environment is not sufficient for intelligence; instead, this flexibility must be wielded in the service of achieving one's goals, such as the goal of obtaining food or winning a game.

The flexibility condition requires that one achieve one's goals through some change in behavior, representation, or processing (Fridland 2015, p. 145). Achieving one's goal through the same process or action again and again does not capture what many understand to be intelligent behavior. The flexibility condition captures the idea that an intelligent agent must be capable of employing a variety of strategies across a variety of situations to achieve their goals. As Fridland writes, "intelligence requires responding differently, if the situation were different" (2015, p. 154). Flexibility in this sense does not mean changing one's behavior in a random or haphazard fashion, but rather employing strategies and actions in a way that helps one achieve one's goals. Such flexibility requires an agent to "bear the proper systematic and flexible connections to the world" (Fridland 2015, p. 155).

As noted above, intelligence is often described as a domain-general capacity; namely, a capacity that applies to a wide range of situations. This idea is reflected in the expectation that intelligent organisms are capable of "transferability" or the ability to transfer knowledge and skills from one situation to another. Transferability enables agents to achieve their goals in novel situations, despite having never encountered those situations in the past (Fridland 2015, Shanahan et al. 2020). Transferability requires flexibility or the ability to alter one's representations or processing to accommodate new situations. As noted above, although AlphaGo is highly skilled in the domain of a standard game of Go, it cannot transfer the skills it has acquired to a new domain, such as a game of chess or even a non-standard game of Go (Halina 2021). Thus, although AlphaGo is an impressive instance of artificial narrow intelligence, it does not provide an example of artificial general intelligence (Shevlin et al. 2019).

Given the environmental circumstances found on earth, it makes sense that organisms would evolve domain-general intelligence. Many organisms live in a world of unpredictable change. How can an organism achieve its goals—whether locating food or maintaining a social relationship—in such a mercurial world? Intelligence provides an answer. An agent that can flexibly accommodate and overcome new and unexpected challenges will be able to achieve its goals, despite an unpredictable environment. If the raven in our example above had never encountered the situation of food suspended by a piece of string in the past, but it was able to flexibly adapt to this situation and achieve its goal of obtaining the food despite this, then this looks like a hallmark of intelligent behavior, and

depends on providing some account of what is meant by "domain". Providing such an account is tricky (Shetteworth 2013, p. 119) but it is often assumed in the literature that an account of domain can be provided if needed.

[2] More precisely, Fridland (2015) argues that these conditions are required for "learning" which is in turn a criterion of intelligence, but it is consistent with her account to view these as conditions for intelligence (Fridland, personal communication).

provides clear benefits to the organism.

The above provides a broad account of intelligence, which I will adopt for the purposes of this Element. Under this account, intelligent behavior represents the capacity to achieve one's goals in a wide range of environments. When an agent demonstrates the ability to transfer knowledge or skills from one context to new situations, it demonstrates a capacity for intelligence.

1.2 Mechanisms of Intelligent Behavior

How are some animals able to adapt flexibly to new situations in the way described above? A standard answer is that certain cognitive capacities enable intelligent behavior, capacities such as imagination, causal reasoning and future planning. As Emery and Clayton suggest, "complex cognition depends on a 'tool kit' consisting of causal reasoning, flexibility, imagination and prospection" (2004, p. 1903; see also Mikhalevich et al. 2017, Schnell et al. 2021). Under this view, the crow's string-pulling behavior is intelligent insofar as it relies on such a tool kit. The crow uses her causal knowledge to flexibly imagine possible solutions to obtaining the food. She then uses prospection to plan what actions to take to execute the chosen solution.

Broadly, under this view, intelligence involves adapting flexibly to new situations by evaluating possible scenarios in advance of choosing a course of action. Rather than trying out actions in the world to see if they work, such an agent relies on an internal model of the environment and its regularities and can use this model to determine in advance whether a particular action is likely to be successful in the real world. The philosopher Daniel Dennett refers to such organisms as "Popperian creatures" (Dennett 1996). The term "Popperian" comes from the philosopher of science Karl Popper who is well known for his proposal that fruitful science engages in a process of generating and testing hypotheses. Similarly, according to Dennett, Popperian creatures successfully navigate the world by generating and testing possible actions in an "inner environment" before committing to performing a particular action in the world (Dennett 1996, p. 88; see also Godfrey-Smith 2018). Insofar as the inner environment accurately reflects the relevant properties of the world, Popperian creatures can determine the value of an action relative to a particular situation, even if that situation is being encountered for the first time.

The psychologist Alexander Taylor has further suggested that when organisms make use of causal information about the world, they should be recognized as "Pearlian creatures" (after the computer scientist Judea Pearl who advanced a theory of causal and counterfactual inference) (Taylor 2009, Godfrey-Smith 2018). Being a Pearlian creatures requires being a Popperian creature. One must have the capacity to generate and test hypotheses about the world in order to generate and test causal hypotheses about the world. Pearlian creatures then are that subset of Popperian creatures who can utilize causal information about the world.

We characterized intelligent behavior above as the capacity to achieve one's goal in a wide range of situations. An organism with an "inner environment" or world model and the capacity to "try out" candidate actions in this inner environment is well placed to behave intelligently. Such an organism is not dependent on the immediate environment for feedback regarding the value of a potential course of action. Instead, such an agent can rely on their existing knowledge of the world and evaluate actions against this world model. A Pearlian creature can additionally rely on causal regularities in the world. Causal knowledge supports claims about what would happen under intervention across a wide range of causal domains. According to Dennett, Popperian-creature world models must also include some knowledge of oneself. He writes: "A Popperian creature's portable knowledge about the world has to include some modicum of knowledge—know-how—about the omnipresent part of its world that is *itself*. It has to know which limbs are its own, of course, and which mouth to feed, but it also has to know its way around in its own brain, to some extent" (Dennett 1996, p. 142). Thus, Popperian creatures are meant to have not only broad knowledge of the world, but knowledge of themselves, and their place in the world.

According to Dennett and others, humans are Popperian and Pearlian creatures (Dennett 1996; Godfrey-Smith 2018; Dennett 2018). Indeed, humans are a combination of many creatures, drawing on a broad range of learning abilities, as well as cultural and linguistic capacities. This is in contrast to other organisms. As we will see in the following section, some researchers view nonhuman animals are much more limited in their abilities. As Hanus writes, "According to some authors, humans are the only species that can mentally go beyond surface features and are able to cognize the world around them in terms of complex, unobserved predicates and relations (like causality, support, force)" (2016, p. 243). For some of these authors, many nonhuman animals are best understood not as Popperian or Pearlian creatures, but as *merely* Darwinian or Skinnerian ones.

Indeed, it is often flagged that some animals *appear* as if they are engaging in intelligent behavior when in fact they are not. As the psychologist Sara Shettleworth writes, "surprisingly often animal behaviours that appear intelligent at first glance turn out to be the product of remarkably 'stupid' mechanisms" (Shettleworth 2013, p. 10). Dennett also writes, "As researchers regularly discover, the more ingeniously you investigate the competence of nonhuman animals, the more likely you are to discover abrupt gaps in competence" (Dennett 1996, p. 116). Such "gaps in competence" suggest that the mechanisms underlying the observed behaviors do not involve processes like imagination, planning, or causal reasoning. Instead, such creatures might be better understood as "unthinking machines" (Dennett 1996, p. 119). It is to these creatures that we turn to next.

1.3 Unintelligent Behavior

Intelligent behavior is often contrasted with what researchers refer to as "rigid", "simple" or "stupid" behavior. Such behavior might appear intelligent, but upon further probing, one discovers that the behavior is in fact highly inflexible and that it is thus best understood as either a product of innate behavioral programs or associative learning. Behavior such as this is also referred to as "sphexish" after digger wasps (or wasps of the genus *Sphex*) as such wasps have been described as exhibiting such rigid behavior. One study that is often cited as demonstrating this was first reported by the French naturalist Jean Henri Fabre in 1879 (Fabre 1879, Kiejzer 2013). Frabre's study gained broader attention through a popular recounting of it in *The Science of Life* by H. G. Wells, Julian Huxley and G. P. Wells (Keijzer 2013). Wells and colleagues write of the digger wasp:

> The instinctive and machine-like quality of most of their [i.e., wasps] behaviour was clearly shown by some experiments of Fabre on the wasp Sphex, which hunts crickets. When the Sphex has brought a paralyzed cricket to her burrow, she leaves it on the threshold, goes

> inside for a moment, apparently to see that all is well, emerges, and drags the cricket in. While the wasp was inside, Fabre moved the cricket a few inches away. The wasp came out, fetched the cricket back to the threshold, and went inside again—on which Fabre moved again the cricket away. He repeated the procedure forty times, always with the same result; the wasp never thought of pulling the cricket straight in. Drag cricket to the threshold—pop in—pop out—pull cricket in: the sequence of actions seems to be like a set of cog-wheels, each arranged to set the next one going, but permitting of no variation. (Wells et al. 1931, p. 696-697; also quoted in Keijzer 2013, p. 506).

These authors describe the digger wasp's behavior as "instinctive and machine-like". The idea that such rigid behavior must be a product of mechanical processes is later emphasized by the cognitive scientist Douglas Hofstadter. Hofstadter (1982) writes that the Sphex account is a "shocking revelation of the mechanical underpinning in a living creature of what looks like quite reflective behavior" (p. 22). The idea here is that although the digger wasp's behavior looks intelligent, this is effectively a trick. Once one probes the system, as Fabre did, one finds that the behavior is brittle, it easily "breaks" and no longer functions properly. The wasp fails to bring food into its burrow because it cannot adjust to a reality in which the cricket is displaced from the threshold. On a continuum of sphexishness, Hofstadter categorizes the behavior of the digger wasp and other insects as one level above "a stuck record" (1982, p. 22).

Dennett (1984/2015) similarly describes digger wasps as "automata" (p. 11). Commenting on Fabre's study, he writes, "The poor wasp is unmasked; she is not a free agent, but rather at the mercy of brute physical causation, driven inexorably into her states and activities by features of the environment outside her control" (p. 12). He also notes that in this case the "Godlike biologist reaches down and creates a slight dislocation in the wasp's world, revealing her essentially mindless mechanicity" (1984/2015, p. 12). Dennett goes on to consider what it is about the digger wasp's behavior that gives one the strong impression that its actions are determined in a mechanistic way. He suggests that it is not that the wasp's actions are a product of prior causes necessarily that leads to this impression, as human actions are also a product of prior causes. Rather, the key may be that the actions are "so *simply* caused" (1984/2015, p. 13). The digger wasp is regularly cited as an example of rigid, mechanical, unintelligent behavior (Dennett 1998, Carruthers 2004, Raby & Clayton 2009, Fridland 2015; see also citations in Keijzer 2013).

Another example of behavior that seems intelligent but appears sphexish upon closer inspection is dead reckoning or path integration in the desert ant. A desert ant will search for food hundreds of meters from its nest. Once it finds food, it runs straight back to its nest, even from great distances, as if it knows exactly where the nest is. However, if one moves the ant and its food to another location before it starts heading home, then it will travel the same direction and distance as would have been appropriate given its prior location and begin searching for its nest there (Shettleworth 2013, p. 10). The ant's ability to locate its nest works well given a particular set of circumstances. But if the situation changes unexpectedly, it seems unable to recognize this and adjust its behavior appropriately. Instead, its behavior is brittle. It is inflexible and domain-specific: that is, it seems to only works in particularly situations and cannot be flexibly adapted to new situations.

1.4 Mechanisms of Unintelligent Behavior

According to the Standard View, unintelligent behavior like that described above is best understood as resulting from simple or "stupid" mechanisms. The two classes of simple mechanisms most often referred to are innate behavioral programs and associative learning. The idea is that these two kinds of mechanisms produce behaviors that are appropriate in some circumstances, but which fail to adapt outside of the contexts in which they were designed or trained to operate. We will review these two kinds of mechanisms before turning to a critique of the Standard View in section 2.

1.4.1 Innate Behavioral Programs

The general idea behind the claim that some behaviors are a result of innate programming is that they are a product of mechanisms put in place by evolution and are not susceptible to change through learning. The idea is that the organism is like the hardware of a computer and the computer has some pre-installed programs that automatically run when triggered by the right internal or external input. The programs however are not flexible in the sense that it can adapt to new situations (Raby & Clayton 2009). Such programs fail when an agent faces a situation that is different from the situation in which the behavior normally operates.

The philosopher Matteo Mameli argues that a minimal condition for an account of innateness is that it is incompatible with learning: "if a trait is innate then it is not learned and if it is learned then it is not innate" (2008, p. 721). This minimal condition emerges from the literature on innateness. As Mameli observes, "arguments for the view that a specific trait is innate are very often arguments aimed at showing that the trait in question could not possibly be the result of learning" (p. 721). Dennett refers to organisms that behave according to innate programming, "Darwinian creatures" (Dennett 1996, pp. 83-85). These creatures are a product of Darwinian evolution by natural selection. They are well designed insofar as they are shaped by millions of years of selective forces, but their innate behavioral programs cannot be changed or updated by learning or other cognitive mechanisms like imagination and foresight.

Behaviors that are characterized as innate behavioral programs include reflexes and fixed action patterns. Reflexes are simple, automatic responses to a stimulus. For example, when a particle causes one to sneeze or a puff of air causes one to blink. These behaviors are involuntarily elicited by a stimulus (Powell et al. 2017, p. 94-95). Fixed action patterns are also elicited in this way but consist of a more complex sequence of behaviors. In this case, a stimulus will cause the sequence of behaviors to be involuntarily set into motion. A reflex is often understood as "hardwired" in the sense that there is a fixed neural structure linking sensory receptors to a motor response. Thus, there is little to no flexibility or control in how one responds once a stimulus has triggered the reflex. Fixed action patterns are also often conceived as hardwired in this way. In their review of cephalopod intelligence, for example, the behavioral ecologist Alexandra Schnell and colleagues define a "hardwired predisposition" as '[f]ixed action patterns triggered by a cue" (Schnell et al. 2021, p. 164). In a review of the term "hardwired" in the popular and scientific literature, the neuroscientist Giordana Grossi shows that the term has seen a dramatically increase in usage over the past 50 years (Grossi 2017). Grossi argues that the term "hardwired" and "innate" are used synonymously, and within the psychology and neuroscience literature, the term "hardwired" has several core meanings, including evolved, genetically determined, not learned, automatic, inflexible, and indicating a fixed or invariant relationship between a stimulus and motor response (Grossi 2017).

Raby and Clayton (2009) provide the digger wasp's behavior in Fabre's experiment as an example of a fixed action pattern (p. 315). Other examples commonly cited as fixed action patterns include nest building by birds, web construction by spiders, and mound building by termites. These activities are also described as resulting from simple "genetically determined rules" (Bluff et al. 2007; Hansell and Ruxton 2007). Carruthers (2004) proposes a general architecture for innate behavioral programs, which he calls an "unminded" behavioral architecture. He writes that the "innate behaviors of many mammals, birds, reptiles, and insects can be explained in such terms" (p. 207). Under this model, an organism's innately coded action patterns are typically triggered when that organism receives the appropriate percept as input, given a particular bodily state. In other words, the releasing factors include the organism's bodily state (e.g., a state of hunger) and perceptual information (e.g., nearby prey) and these trigger an action schemata or innate behavioral program that determines the motor response (Carruthers 2004, p. 207). According to Carruthers, examples of rigid behavior, like that of the digger wasp, are best understood as operating with this kind of underlying architecture (2004, p. 211-212). Not only are such organisms unintelligent, but "[i]f this were the full extent of the flexibility of insect behaviors, then there would be no warrant for believing that insects have minds at all" (Carruthers 2004, p. 212).

1.4.2 Associative Learning

Innate behavioral programs are one kind of simple mechanism generally believed to give rise to unintelligent behavior. A second major type of mechanism believed to give rise to unintelligent behavior is associative learning. Associative learning is the capacity to associate responses or outcomes with a particular stimulus, such as a rat learning to associate the pressing of a lever with food delivery. In contrast to intelligent behaviors, behaviors learned through associative learning are typically thought to be inflexible and acquired through trial and error.

The mechanisms responsible for associative learning are generally characterized as simple (Dacey 2016, Hanus 2016). As the philosopher Mike Dacey writes, associations are often described as "*links between representations that are sequentially activated in the process*" (Dacey 2016, p. 3765, emphasis original). Associative learning provides an account of how such simple links between representations are formed. Under this account, such links between mental representations are "formed passively and automatically as a direct consequence of contiguous (with some restrictions) pairings of those physical stimuli" (Mitchell et al. 2009, p. 184). Once such links are formed, then the perception of one stimulus activates the other linked representations. Clayton and colleagues highlight the mechanistic nature of associative processes (Clayton et al. 2006). Like the above accounts, they characterize associative learning as the process of nodes being activated by events (like stimuli), which cause excitatory or inhibitory connections between nodes to be formed. This chain of connections then ultimately controls associatively learned behavior. Clayton and colleagues hold that associative processes are "mechanistic" in the sense that they "gain their explanatory power by analogy to physical processes" (Clayton et al. 2006, p. 198). As they write, "[t]he processes by which associative structures control behaviour are constrained only by their mechanistic-like properties (Clayton et al. 2006, pp. 198-199).

Associative learning is also characterized as simple on the level of neural implementation. For example, Hanus (2016) writes that many "highlight the apparent similarity between the proposed automatic link-formation mechanism and the neurophysiological hardware in which it is implemented" (p. 242; see also Buckner 2011). Similarly, Dennett notes that the competences resulting from associative learning, "tend to be anchored in the specific tissues that are modified by training" (Dennett 1996, p. 132). He describes the training process involved in associative learning as taking place in "a network of nerve cells" (Dennett 1996 p. 87). Dacey (2016) writes that, "[t]hrough much of the history of the concept [of association], it has been considered a strength that associations operate mechanically, such that they are better candidates for neural realization than more nebulous concepts like the will" (p. 3769). Such a mechanical understanding of associations is illustrated by behaviorist accounts of the early twentieth century, which described neurons as linking stimuli and responses in the manner of a switchboard relay. Such a mechanical understanding of association continues today but is instead grounded in accounts of neural networks. Under this view, associative learning is taken to involve the adjusting of connections between neurons, according to various rules (Dacey 2016, p. 3769; Dennett, 1996, p. 87).

Dennett refers to organisms that rely on associative learning as "Skinnerian creatures" (Dennett 1996; see also Godfrey-Smith 2018). Skinnerian creatures are more flexible than Darwinian creatures because they can adapt to new situations during their lifetime, rather than over generations. However, this flexibility falls short of the kind of flexibility needed for intelligence. This is because, unlike Popperian creatures, Skinnerian creatures must use trial-and-error learning to identify appropriate actions in new situations. Thus, they cannot behave successfully in new situations without first failing for some time. If such agents have the systematic connections to the world needed to make good choices, then this has come at the end of a long process of training. As Dennett writes, "Skinnerian creatures ask themselves, "What do I do next?" and haven't a clue how to answer until they have taken some hard knocks" (Dennett 1996, p. 100). Skinnerian creatures survive insofar as they make "lucky first moves" while Popperian creatures survive because "the truly stupid moves are weeded out before they're hazarded in 'real life'" (Dennett 1996, p. 88). Skinnerian creatures are incapable of "one-shot learning" or learning about the world from one or a few examples but must instead "endure the arduous process of trial-and-error in the harsh world" (Dennett 1996, p. 88).

Although Skinnerian creatures are more flexible than Darwinian creatures, they are still often characterized as simple and mechanical. Like Darwinian creatures, their behavior is believed to be rigid and easy to break. One example of such purportedly brittle behavior is imprinting or what Dennett calls "Mamataxis". Here a young organism learns the behavioral response of moving towards it's mother by being exposed to the mother as a stimulus (see Bolhuis et al. 1990 for a review of imprinting as a form of associative learning). Dennett argues that Mamataxis is brittle: that if chicks, for example, do not imprint on the mother soon after birth, they will imprint on any large moving object and follow that object instead, as if it were their mother (Dennett 1996, p. 104). The idea that animals can imprint on objects other than their mothers was famously illustrated by the Austrian ethologist, Konrad Lorenz, who had greylag goslings imprinted on him and argued that such imprinted behavior was very rigid (Lorenz 1937). As Dennett writes, "reliable Mamataxis can be achieved with a bag of simple tricks. The talent is normally robust in simple environments, but a creature armed with such a simple system is easily 'fooled,' and when it is fooled, it trundles to its misfortune without appreciation of its folly" (Dennett 1996, p. 105).

Behaviors that result from associative learning are typically understood as unintelligent in the sense of being rigid, brittle or sphexish. We should expect Skinnerian creatures to be mechanical or "mindless" in this way. However, which creatures, if any, are Skinnerian in nature? According to Dennett, "if there are any purely Skinnerian creatures, capable only of blind trial-and-error learning, they are to be found among the simple invertebrates" (Dennett 1996, pp. 92-93). Dennett holds that most animals are capable of associative learning but remains agnostic regarding which creatures are limited to only this form of learning. The categories of Darwinian, Skinnerian, and Popperian creatures are also best understood as nested rather than exclusive: although a Darwinian creature is only Darwinian, a Skinnerian creature is Skinnerian *and* Darwinian, and a Popperian creature is Popperian, Skinnerian and Darwinian (Godfrey-Smith 2018, Dennett 2018).

Although it is unclear which creatures are merely Darwinian or Darwinian and Skinnerian, the Standard View holds that we must eliminate the possibility that a creature is Darwinian or Skinnerian before concluding that it is exhibiting intelligent behavior. Like the digger wasp bringing a cricket into her burrow or a chick engaging in Mamataxis, a behavior might look intelligent, but upon further inspection be a product of simple innate behavioral programs or associative learning. Indeed, when considering behavior that looks intelligent in animals like the piping plover, hare and gazelle (Dennett 1996, pp. 122-125), Dennett writes that the needs of these animals to engage in what looks like sophisticated behavior "can probably be provided by networks designed almost entirely by Darwinian mechanisms, abetted here and there by Skinnerian mechanisms" (Dennett 1996, p. 130). Although determining whether these organisms are in fact merely Darwinian or Skinnerian in nature is an empirical question, it is one we must answer, and eliminate as a possible explanation of their behavior, before concluding that the observed behavior is intelligent.

The idea that we must eliminate the possibility that an organism is relying on innate behavioral programs or associative learning before concluding that it is behaving intelligently is central to the Standard View. Indeed, although there is some variation in views regarding the exact mechanisms underlying intelligent behavior, there is broad consensus regarding the mechanisms underlying unintelligent behavior—namely, innate behavioral programs and associative learning. As Seed and colleagues write, animal intelligence is, "usually defined by exclusion, rather than by some positive assessment of the mechanisms underpinning it" (Seed et al. 2009, p. 402) and this "principle of exclusion" defines intelligence as complex and flexible behavior that "cannot easily be explained in terms of simple conditioning, or hardwired action patterns" (Seed et at. 2019, p. 410). In other words, to determine whether an organism is behaving intelligently, we need to first exclude the possibility that they are Darwinian or Skinnerian creatures like the digger wasp who has "tricked" us into thinking it is intelligent, when in fact it is not.

1.5 Conclusion

The above provides a sketch of the Standard View of intelligent behavior. According to this view, intelligent behavior involves successfully achieving one's goals in a wide range of circumstances and is often underpinned by mechanisms such as imagination, future planning, and causal reasoning. Unintelligent behavior, on the other hand, is rigid and brittle and a product of simple mechanisms like innate behavioral programs and associative learning. One method for distinguishing intelligent behavior from unintelligent behavior involves probing the behavior of an organism and seeing whether that behavior fails when small changes are made to the environment. Once a behavior has been uncovered as unintelligent in this way, then we know that we are observing the products of Darwinian or Skinnerian, rather than Popperian, processes. Such mechanisms do not enable organisms to act flexibility in such a way that allows them to achieve their goals in new situations. Instead, successful performance is limited to a narrow range of environments.

With the Standard View of intelligence in hand, in the next section, we will consider contemporary work on nonhuman animal behavior and cognition. Based on this work, I argue that we should reject the Standard View, as it fails to accommodate the empirical and theoretical state of the art on animal minds. In particular, the idea that animals can be sorted into Darwinian, Skinnerian and Popperian creatures is a fiction. It distorts the true capacities of organisms to such a degree that our best option is to dissolve these categories altogether.

2. Against the Standard View

Our understanding of innate behavioral programs has changed dramatically over the past few decades (Versace et al. 2018). Associate learning is also now recognized as a powerful tool for adapting to new circumstances (Heyes 2012). In this section, I show how innate behavioral programs and associative learning both result in flexible and domain-general abilities. Thus, these hallmarks of intelligence are not limited to organisms with capacities like imagination and causal reasoning. Moreover, the fact that an organism relies on innate behavioral programs and associative learning does not mean its resulting behavior will be unintelligent.

Crucially my point here is not that we should understand flexible behavior in nonhuman animals as “merely” the product of innate behavioral programs and associative learning. My aim instead is to show that Darwinian and Skinnerian creatures, as they are traditionally conceived, do not exist. Those organisms that rely on innate behavioral programs and associative learning (which includes humans) are not limited to rigid and brittle behavior. Indeed, given what we know about the form these mechanisms take across the animal kingdom, we should not expect to encounter organisms that are sphexish in the way that the digger wasp has been characterized. Indeed, as we will see, it appears that even this often-repeated account of the digger wasp is a myth. Intelligent behavior is not a product of so-called “sophisticated” cognition alone but arises from innate behavioral programs and associate learning as well.

2.1 Innateness and Intelligence

Do innate behavioral programs result in rigid, brittle behavior? For those organisms that do rely on innate behavioral programs, are they best understood as Darwinian creatures? The answer to both questions I argue is “no.” Crucially, innate behavioral programs often work in concert with experience to produce appropriately flexible responses. Such “programs” are not typically “triggered” by a specific stimulus causing a rigid sequence of behaviors to unfold. Instead, organisms enter the world with adaptive priors. And these adaptive priors scaffold learning, making it faster and more effective.

A common model organism for investigating innate behavioral programs is the newborn chicken (*Gallus gallus*). Chicks are precocial, meaning they have relatively mature sensory-motor systems, and can act independently to a high degree, from birth. Soon after birth, chicks seem to have some understanding of occlusion and solidity (that a solid object cannot pass through another solid object), numerosity, the ordinal value of numbers, basic athematic, geometrical relationships, as well as other abilities (Vallortigara 2012, Versace & Vallortigara 2015).

Chicks often do not require long periods of trial-and-error learning to behave successfully in the world. However, chicks are also not equipped with a set of innate behavioral programs that predetermine behavior. Instead, they are born with mechanisms that guide learning. Crucially, these mechanisms guide learning in a way that leaves room for environmental variation. For example, in the case of imprinting, there are mechanisms that orient the chick towards objects that visually resemble a mother hen. However, given there is much variation in the form that a mother hen can take, these orienting mechanisms are not overly specific, as this would result in too many false negatives (Versace et al. 2018). Instead, the mechanisms bias the chick to orient towards objects exhibiting properties like biological motion and face-like features. As the psychologist Elisabetta Versace and colleagues write, "[o]ptimal learning mechanisms must trade being sufficiently open to allow a wide range of stimuli to be stored as imprinted memories, against being sufficiently specific to avoid imprinting on inappropriate objects that in natural environments coexist with the chick's mother and siblings" (Versace et al. 2018, pp. 963-964). It is findings like these that have led researchers working in this area to conclude that innate predispositions and learning mechanisms are dynamically interdependent with experiences driving the emergence of predispositions and predispositions guiding learning. As Orsola Rosa-Slava and colleagues write: "Predispositions are not fixed and immutable mechanisms" (p. 9). Instead, such mechanisms respond to and guide learning throughout the ontogeny of an organism.

Adaptive priors such as these also allow organisms like the chick to successfully generalize from a limited number of cases (Versace et al. 2018, p. 964). Chicks recognize a mother hen from many points of views and against many different backgrounds. They also recognize their siblings despite these siblings transforming during development (Versace et al. 2018, p. 963). This is possible because the mechanisms involved in imprinting allow an agent to generalize from an imprinted object to novel objects in remarkable ways. For example, in one study, researchers showed that newborn ducklings have the capacity to identify logical relations among objects and generalize these abstract relations to novel stimuli. Domesticated mallard ducklings were exposed to pairs of objects one hour after hatching. These objects were either the same in their color and shape or different in one of these properties. Later, when presented with a new pair of objects, ducklings preferred those exhibiting the relation (same or different) to which they had been initially exposed, suggesting they had imprinted on the original abstract relation and were able to identify it in the new case (Martinho & Kacelnik 2016). This and related studies suggest that "simple" mechanisms like imprinting are sensitive to high-level patterns in the world like abstract relations between objects and can facilitate the transfer of knowledge of these patterns to new situations.

Thus, the role of innate behavioral programs in producing behavior are more complex than portrayed by the Standard View. Even canonical examples of purported inflexible behavior, determined by innate behavioral programs, have found to be more complicated than traditionally described. For example, the animal behavior scientist Alexis Breen and colleagues note that, "public and scholarly consensus [holds] that bird nests are achieved by instinct alone" (Breen et al. 2016, p. 83). However, after reviewing almost 150 years of data on avian nest building, they conclude that a variety of forms of learning are involved in nest building, including social learning. Breen (2021) builds on this finding, arguing that nest building would make a promising (but hitherto neglected) framework for investigating material culture in animals. Similarly, studies suggest that spiders adjust their web design based on both long-term web-building experience and recent experiences capturing prey (Heiling & Herberstein 1999). Finally, other “hardwired” or “innate” behaviors, such as mating and fighting, are now known to be highly variable. For example, attack behavior in mice can be understood as “fixed” in the sense that there is a motor circuit dedicated to executing the behavior. However, there is large individual variability regarding when and whether an individual attacks another, based on the individual’s experience, as well as the potential rewards and costs (such as metabolic and opportunity costs) associated with making an attack. Similarly, mating behavior, partner preferences, social attraction, and social avoidance in vertebrates like rodents are affected by experience in nuanced ways. For example, if a mouse is defeated by another individual, it will avoid that individual for a few days; if it is defeated many times, however, it will sometimes exhibit general social avoidance and “depression-like behavior” such as a decreased preference for previously rewarding stimuli (Wei 2021, pp.1609-1611). As researchers reviewing “innate” vertebrate social behavior write: “The combination of learning-dependent and -independent modulatory mechanisms makes seemingly stereotypical social behaviors incredibly flexible and adaptive” (Wei et al. 2021, p. 1614; see also Gorostiza 2018).

It is worth noting that humans likely share many innate priors with other animals. It is challenging to investigate such predispositions in humans, given the difficulty of controlling for experience after birth. However, there is evidence that soon after birth, humans have some understanding of solidity, occlusion, number, geometry, social attractiveness, physical danger, the difference between animate and inanimate objects, causality, and other things (Spelke & Kinzler 2007, Platt & Spelke 2009, Vallortigara 2012, Versace & Vallortigara 2015). One recent study, for example, suggests that three-month-old human infants already have some intuitions regarding causal agency. When observing others acting causally in the world (for example, reaching for and touching a ball, causing it to illuminate and emit a sound), they look longer when the agent’s reach is unnecessarily circuitous, rather than direct. This contrasts with observing non-causal actions (actions and effects lacking spatiotemporal continuity—for example, the ball activating not on contact). In this case, infants do not look longer at the circuitous versus direct reach. This suggests that three-month-old infants have some understanding of the goals of agents acting causally in the world and the cost of such actions. At three months, infants do not yet reach for or grasp objects, so this causal knowledge cannot be acquired from experiencing their own actions. The authors conclude that, “before infants can reach for objects themselves, they represent other people’s reaching actions in accord with the abstract concept of ‘cause,’ a concept that may function together with the associated concepts of ‘cost’ and ‘goal’” (Liu et al. 2019, p. 5). They further suggest that such concepts are present at birth and guide learning and note that this is consistent with what we know about the abilities of precocial animals—for example, newly hatched chicks prefer self-propelled objects over objects that have been caused to move by an external force (Liu et al. 2019, p. 5; Mascalzoni et al. 2010; see also Mascalzoni et al. 2013).

The above discussion shows that innate behavioral programs work in concert with learning and plasticity to provide organisms, such as newly hatched chicks, with tools for navigating an unpredictable world. When one examines contemporary research on other purportedly rigid behaviors that are a product of innate behavioral programs, one finds that the picture is more complex than that advanced by the Standard View. As Shettleworth writes,

> attempting to classify behavior as *learned* as opposed to *innate* is meaningless. Trial-and-error learning likely perfects the crows' skill, but it operates on appropriate motor patterns which they are predisposed to engage in. By the same token, tool use is not innate either, if by *innate* we mean performed without any relevant experience. And if we mean by *innate* not modifiable by experience once it is performed, that cannot be correct either. Every moment of an organism's development from the very beginning results from a seamless interplay of the learned and the innate, or genes and environment." (2013, pp. 13-14, emphasis original)

Considerations such as these have led some philosophers to reject the concept of innateness altogether. For example, Mateo Mameli (2008) argues that the concept of innateness, as it is typically used in the literature, conflates numerous properties that are in fact distinct—it treats properties like "not learned," "genetically encoded" and "inflexible" as a natural cluster, when in fact they constitute a clutter (see also Bateson & Mameli 2007). Recognizing this is important for avoiding bad inferences, such as using the concept of innateness to infer that a trait with high heritability will also be difficult to modify through environmental intervention. Unless there is good evidence that these properties are causally connected in some way, such inferences are dubious. Mameli writes: "Even if it turns out that INNATENESS is not theoretically useful, it does not follow that the Nativist Debates are misguided or pointless. It follows instead that there are better ways of conducting (at least some of) these debates, ways that do not make any use of INNATENESS" (2008, p. 735).

Similarly, we have seen that the properties used to characterize Skinnerian creatures (simple, innate mechanical structures that lead to inflexible, domain-specific responses) do not in fact cluster in the natural world, according to contemporary research. Creatures that depend on innate concepts and priors do so in a way that is integrated with learning rather than distinct from it. Such "innate" structures also lead to flexible and domain-general behavior, allowing organisms to transfer knowledge to new situations. Innate behavioral programs underpin intelligent behavior.

2.2 Associative Learning and Intelligence

To begin, it is worth noting that the term "associative" does not refer to a single, simple cognitive mechanism. As the philosopher Mike Dacey points out, terms like "association" are "so abstract that they are merely *filler terms* that could be realized by many different mechanisms" (2016, p. 3764, emphasis original). Dacey argues that associative models are often mistakenly interpreted as representing simple psychological or neural mechanisms, such as those held by the Standard View. But when one looks closely at particular associative models, one finds that they typically do not represent simple mechanisms, but instead abstract away from the details of mechanisms. Thus, associative models are simple insofar as they leave out mechanistic detail, but the mechanisms they represent are not simple in this way. As Dacey writes, associative models "are useful when we don't know enough about the mechanistic detail. They can be an early step in a top-down characterization of the process (that is, moving from and [sic] abstract, partial characterization of the process to one including more causal and mechanistic detail)" (2016, p. 3779). Thus, holding that an organism relies on "association" does not imply that it employs a simple psychological or neural mechanism in the sense maintained by the Standard View.

Psychologists have made similar observations. The comparative psychologist Daniel Hanus, for example, has noted that contemporary associative models are diverse (Hanus 2016). Indeed, he writes that the models can be so diverse that "nearly any empirical finding could potentially be simulated by an associative model" (Hanus 2016, p. 243). Given this diversity, Hanus argues that associative models do not form a coherent category of models. Like Dacey, Hanus is also skeptical that associative learning can be understood as instantiated in simple neural processes (Hanus 2016, p. 242). He notes that, first, it is unclear that the theoretical entities posited by associative models correspond to anything physiological. Second, if associative learning were simple in the sense of being implemented by large look-up tables, for example, then this would be computationally very expensive. As the psychologist C. Randy Gallistel writes with respect to neural network models of dead reckoning construed in terms of a look-up-table architecture: "Such an architecture is prodigally wasteful of material resources. It is nakedly exposed to combinatorial explosions that lurk behind every tree in the computational forest" (2008, p. 240).

Recent work on the evolution of cognition also suggests that a wide range of organisms engage in forms of associative learning that support flexible behavior. One form of such learning is what researchers call "unlimited associative learning" (Bronfman et al. 2016a, Bronfman et al. 2016b, Ginsburg & Jablonka 2019, 2021, Birch et al. 2020). Unlimited associative learning (UAL) is a form of associative learning in that it involves learning associations between objects, events, and actions. However, it is effectively unlimited in the range of associations that can be formed. An animal with UAL can associate novel stimuli, as well as compound, multimodal stimuli composed of different elements or action-patterns. Such an animal can learn to associate stimuli even if there is a temporal gap between them. Finally, an animal with UAL can engage in second-order conditioning or form associations between new stimuli and actions and prior associations (Bronfman et al. 2016, Bronfman et al. 2018, Ginsburg & Jablonka 2021). According to Bronfman and colleagues, "a system enabling UAL entails the integration of information, leads to a massive increase of discrimination, and allows the generation of flexible goal-directed behavior" (Bronfman et al. 2016b, p. 12). An organism with UAL then has the capacity to flexibly adapt to a wide range of situations.

Researchers studying UAL also believe it is widely distributed across the animal kingdom. As Ginsburg and Jablonka write, "A survey of the learning literature suggests that these [unlimited associative] learning capacities are present in three phyla: in almost all vertebrates, some arthropods (including honeybees and cockroaches) and some cephalopod molluscs (the colloid cephalopods: octopods, squid and cuttlefish)" (Ginsburg & Jablonka, 2021, p. 6; see also Ginsburg and Jablonka 2019). If these researchers are correct, then we should expect flexible behavior to be found in numerous organisms, underpinned by a form of associative learning.

One might object that UAL is not what most researchers have in mind when they appeal to associative learning. Instead, they have simple associative processes in mind like classical and operant conditioning. However, here we can return to the point made by Hanus, that the category of associative models is incredibly diverse. It is not clear what, if anything, gives this category of models theoretical or mechanistic coherence. It is also not clear that traditional associative models are simple, rather than simply abstract, as argued by Dacey. Thus, excluding UAL from the category of associative models on the grounds that it is too different or too complex compared to other associative models would be difficult to justify without first clarifying what is meant by "different" and "complex."

Thus, associative models of cognition and behavior do not always represent simple mechanisms, and some forms of associative learning produce complex, flexible behavior. These findings are consistent with research on the role of associative learning in human behavior. For example, associative learning has been shown to contribute not just to basic human functions, but a range of complex behaviors. Indeed, a study conducted by the psychologist Scott Barry Kaufman and colleagues suggests that associative learning is one of the mechanisms underlying general intelligence (*g*) understood as a "positive manifold" or the tendency for a battery of cognitive tests to be positively correlated with one another (Kaufman et al. 2009). Kaufman and colleagues found that individual differences in associative learning predicts *g*. This is the case even when one controls for other mechanisms thought to underly *g*, such as working memory or processing speed. As Kaufman et al. (2009) write, according to their analysis, "associative learning, working memory, and processing speed all made statistically independent contributions to *g*. This finding suggests that each of these elementary cognitive processes may represent a mechanism that contributes differentially to general intelligence" (p. 379).

There is evidence that associative learning contributes to a range of other human cognitive abilities and behaviors, including a sense of agency, imitation, social learning, flexible planning, learning higher-order relationships between multiple environmental outcomes, and others (Heyes 2012, Lind 2018). The fact that a behavior is a product of associative learning then does not mean it is rigid or unintelligent. These and other considerations have led researchers like the psychologist Konstantinos Voudouris to conclude that associative and cognitive processes are often not mutually exclusive (Voudouris 2020).

Even if some abilities, like causal reasoning and associative learning, are best understood as distinct, it is not clear that relying on causal reasoning is "smarter" than relying on procedural rules learned through association. For example, the trap-tube task is a common paradigm used to probe an organism's causal reasoning abilities (Seed et al. 2006). In this task, participants are presented with a transparent tube baited with a reward (usually food). To obtain the reward, participants must use their body (e.g., finger or beak) or a tool (such as a stick or rake) to retrieve (by pushing or pulling) the reward from the tube. The tube, however, contains various traps that must be avoided if the reward is to be successfully extracted. If the reward falls into a trap, it can no longer be retrieved. Various animals such as chimpanzees, rooks, and New Caledonian crows have been tested using this paradigm. However, often the results have been inconclusive, as there is some simple procedural rule that participants could have acquired through associative learning that might explain their performance (Seed et al. 2011). However, when adult humans are tested on the same task, their pattern of performance is also more consistent with the implementation of a procedural rule rather than the application of abstract causal principles (Silva & Silva 2006). Presumably, adult humans have some understanding of the causal principles involved in this task (e.g., that unsupported objects fall). If they do, however, they do not always make use of that causal knowledge. Perhaps it is somehow "smarter" to rely on procedural rules instead (we will return to the topic of trade-offs in section 3).

The above suggests that the Standard View is mistaken. Associative learning is not a simple mechanism that gives rise to unintelligent behavior. Quite the contrary, associative models are enormously diverse and often leave mechanistic details unspecified. Associative learning also gives rise to a wide range of flexible behaviors from unlimited associative learning in molluscs to general intelligence in humans. Associative learning is also implicated in cognitive processes like flexible planning, which are viewed by some psychologists as key processes involved in producing intelligent behavior (see section 1.2).

2.3 Conclusion

Contemporary work on innate behavioral programs and associative learning suggests that these mechanisms work in concert with other cognitive mechanisms and give rise to flexible behavior. The idea that innate behavioral programs and associative learning are best understood as simple mechanisms that lead to rigid behavior is mistaken. There is no neat mapping of intelligent behaviors onto one set of mechanisms and unintelligent behaviors onto another. There is also no neat divide between Darwinian, Skinnerian, Popperian and Pearlian creatures. Instead, empirical research suggests that both human and nonhuman animals rely on a tangled web of innate priors, learning mechanisms, and other abilities to flexibly engage with the world. These mechanisms typically do not operate in isolation but depend on each other.

Given the above, how do we make sense of sphexish behavior? Are at least some animals best understood as simple automata? As we have seen, many behaviors are more flexible than characterized by the Standard View, such as nest construction and web building. Indeed, even the often-repeated case of the digger wasp appears to be a myth. The philosopher Fred Keijzer (2013) notes that already in the original digger-wasp study, Fabre noted that there was variation among individuals of the same species, with some individuals repeating the behavior of checking the den after replacing the displaced cricket, and others not. Later attempts to reproduce the behavioral repetition in Fabre's study also failed. The entomologists George and Elizabeth Peckham, for example, found that the wasps adapted to displaced prey after a few trials; other replications showed that longer repetitions could be elicited, but that wasps adapted eventually. Finally, more recent studies suggest that the digger wasps' behavioral repetitions might not be superfluous but rather an adaptive response (see Keijzer 2011, 2013).

Recall that Hofstadter categorized the behavior of the digger wasp as one level above "a stuck record". Keijzer's analysis suggests that the stuck record here might instead be the human tendency to retell the digger wasp story and others like it. He writes:

> While the *Sphex* story does not teach us a lot about insect behavior, the real interest of the *Sphex* story might lie somewhere else: the tenacity of the story might teach us something important about *human* thought and behavior. The interesting fact is not so much the presumed endless repetition made by the wasp, but the endless repetition of humans retelling the story as a matter of significance despite all the available counterevidence. (Keijzer 2013, p. 515-516, emphasis original)

If there is poor evidence for sphexish behavior, why do researchers continue to appeal to this example? Keijzer suggests that the story reinforces the human intuition that an organism can be interpreted as agential and minded on the one hand, and mechanical and mindless on the other. He writes, "what initially seems like a sign of mind in insects is suddenly shown to be a mere mechanism" (Kiejzer 2013, p. 516). We can extend this analysis to innate behavioral programs and associative learning. As we have seen, the Standard View attributes sphexish behavior to these processes and this serves as a contrast class for identifying intelligent behavior. In other words, the Standard View draws on the idea that innate behavioral programs and associative learning give rise to rigid behavior to reinforce the idea that we have a handle on what constitutes intelligence. Identifying intelligent organisms requires excluding these alternatives.

This point can be made more precise by drawing on contrastivism in philosophy. Contrastivism is the view that reasons are always relative to contrast classes. Within the context of epistemology, the idea is that reasons are required for knowledge and justified belief. Given that a reason is relative to a contrast class, what is included in that contrast class is important for determining whether one has knowledge. For example, if you see an animal, and the appearance of that animal leads you to conclude that it is a zebra, as opposed to a rhinoceros or horse or any other large mammal with which you are familiar, then you are justified in believing it is a zebra (Dreske 1970, Sinnott-Armstrong 2008). However, you are not justified in believing it is a zebra as opposed to a perfect simulation of a zebra. Your visual experience rules out that the animal is a rhinoceros or horse but does not rule out that what you see is a perfect simulation of a zebra. Contrast classes are important for determining whether there is sufficient evidence to rule out the alternatives, and thus whether someone is justified in their belief. As Sinnott-Armstrong writes, "Someone, *S*, is justified out of a contrast class, *C*, in believing a proposition, *P*, when and only when *S* is able to rule out all other members of *C* but is not able to rule out *P*" (2008, p. 259).

For our purposes, what matters is that the Standard View regularly treats innate behavioral programs and associative learning as the relevant contrast class for believing that a system is intelligent. In doing so, it seeks evidence to eliminate these alternatives. However, the picture it paints of these processes, that they are simple, mechanical and lead to inflexible behaviors is incorrect. As we have seen, innate priors and associative learning work in concert with other mechanisms and give rise to a wide range of flexible behaviors. Supplying evidence that these processes are not at work does not give us insight into whether a system is intelligent.

Even if sphexish behavior does exist in the form of rigid or "stupid" responses to certain situations, such responses are unlikely to comprise the whole behavioral repertoire of an organism. Also, as we have seen, humans rely on a combination of associative learning, innate mechanisms, and other cognitive strategies for a wide range of sophisticated behaviors. Although such constraints might lead to behavior that appears suboptimal, irrational or unintelligent in some situations, it might be that the behavior is optimal against a background of various trade-offs over the lifetime of an organism (see Shettleworth 2013, p. 68). We explore such trade-off in the following section, as well as consider how best to understand and investigate intelligent systems once the Standard View has been cast aside.

3. Kinds of Intelligence

The cognitive scientist Aaron Sloman advanced the idea of a space of possible minds in the 1980s. In doing so, he hoped to move away from dichotomous thinking concerning minds and behavior. Sloman writes:

> A common approach to this space of possible 'behaving systems', to coin a neutral phase, is to seek a single sharp division, between those with minds, consciousness, souls, thoughts, or

> whatever, and those without. Where to draw the line then becomes a major problem, with protagonists of the uniqueness of man [sic], or of living things, or champions of machine mentality, all disputing the location of the boundary, all offering different criteria for allocating things to one side or another. (Sloman, 1984, para. 3).

Sloman argues that we should reject such dichotomies and instead "acknowledge that there are *many* discontinuities, or divisions within the space of possible systems: the space is not a continuum, nor is it a dichotomy" (1984, para. 11, emphasis original). Sloman encourages researchers to focus instead on exploring the space of possible minds or behaving systems by, first, surveying and classifying the abilities of different systems and, second, seeking explanations for the abilities (and inabilities) of these classified systems. I believe this is the approach we should take for studying intelligent systems. Instead of attempting to neatly divide behaving systems into intelligent and unintelligent, or rank them on a scale of intelligence, we must recognize that intelligent systems are diverse: domain-general, flexible behavior takes a variety of forms and is underpinned by a variety of mechanisms (including innate priors and associative learning).

Much comparative psychology and cognitive ethology today is dedicated to this project of describing, classifying, and explaining behavior across a wide range of cultures, species, individuals (biological and artificial), developmental periods, etc. However, as we have seen, when it comes to intelligence, there is a tendency towards dichotomous thinking. The Standard View aims to divide the natural world according to those organisms that are rigid in their behavior and depend on simple, mechanical processes, and those organisms that can behave flexibly and depend on more sophisticated processes like causal reasoning and planning. The central aim of this Element has been to show that this dichotomy does not reflect what we currently know about cognition and behavior. Eliminating this dichotomy however leads to the question, how should we structure the space of intelligent systems? Also, if there is no compelling distinction between intelligent and unintelligent systems, why not eliminate the concept of intelligence altogether?

In this section, I address these questions. I first revisit the question, "what is intelligence?" to see whether we can pin down a precise definition that does not rely on the Standard View, but that can also be applied across species or taxa. I argue that the unique trade-offs each organism faces precludes such an approach and that we should also avoid relying on humans as a standard for intelligence. I then draw on three recent suggestions in philosophy and cognitive science for research strategies that better capture the diversity of mechanisms and behavioral abilities found across the animal kingdom. These strategies provide a good starting point and guide regarding how to construct a space of intelligent systems. I conclude by considering whether we should avoid or eliminate the term "intelligence" altogether, given our inability to define it precisely. I follow other philosophers in holding that umbrella terms such as these can help orient a research community towards a class of phenomena that we do not yet understand, and that doing so may facilitate research. Crucially, however, we must reject the Standard View and not contrast intelligence with innate behavioral programs and associative learning but instead recognize that intelligence emerges from a heterogenous suite of mechanisms and takes a variety of forms.

3.1 What is Intelligence?

In section 1, we broadly defined intelligent behavior as the ability to flexibly achieve one's goals in a wide range of situations. Although the Standard View couples such behavioral flexibility with capacities like imagination, planning and causal reasoning, and behavioral inflexibility with innate behavioral programs and associative learning, we have seen that cognition and behavior does not lend itself to such a neat divide. The latter mechanisms also give rise to behavioral flexibility. If we think the concept of intelligence is worth retaining, then we need to provide an approach that serves as an alternative to the Standard View.

Recall that in section 1 we made the distinction between narrow or domain-specific intelligence on the one hand, and general intelligence on the other. This concept of general intelligence is what Shettleworth (2013) seems to have in mind when discussing the capacities of organisms like Clark's nutcrackers (*Nucifraga columbiana*). She notes that Clark's nutcrackers (a corvid species native to North America) have excellent spatial memory when it comes to retrieving the thousands of seeds that they store for the winter. However, when one examines their other memory abilities, such as remembering the colors of items or remembering where items stored by other birds are located, then they are either no better or worse than other corvids. Shettleworth concludes that animals like the Clark's nutcracker are "not exceptionally 'intelligent' in general, just especially good at specific skills that are important for their survival and reproduction" (Shettleworth, 2013, p. 11). According to Shettleworth, the nutcracker's spatial memory is an example of an adaptive specialization of cognition and not general intelligence.

This distinction between domain-specific cognitive adaptations and domain-general intelligence suggests that perhaps we could define intelligence along these lines. That is, rather than commit to any specific mechanisms underlying intelligent behavior, we could instead categorize organisms according to behavior alone—namely, according to the degree of domain-general flexibility they exhibit. Setting the problem of how to delineate a domain aside, this approach immediately runs into what me might call the "more is not better" problem. What does it mean for a particular species of corvid to be more intelligent with respect to a capacity like memory, for example? If species A remembers more properties than species B, is it more intelligent? What about if it remembers more properties for a greater length of time? The problem with this approach is that it fails to consider the necessary trade-offs that occur with the exercise of any cognitive capacity. Agents who remember too much are often impaired. In the words of patient "AJ", the first person diagnosed with highly superior autobiographical memory (HSAM): "Whenever I see a date flash on the television (or anywhere else for that matter) I automatically go back to that day and remember where I was, what I was doing, what day it fell on and on and on and on and on. It is non-stop, uncontrollable and totally exhausting" (Parker et al. 2006, p. 35). AJ has detailed autobiographical memories of every day of her life stretching back decades. Patients with HSAM, however, suffer from impairments such as obsessive thinking and difficulty attending to the present. Moreover, neuroscientists like Blake Richards argue that forgetting has various critical functions, such as allowing agents to generalize and make predictions: remembering too much may lead to overfitting to past experience (Richards et al. 2017). Indeed, people with severely deficient autobiographical memory (SDAM) seem to excel at abstract thinking and problem solving (Gravitz 2019).

Similar trade-offs occur in the context of flexible problem solving. Research led by the developmental psychologist Alison Gopnik and colleagues suggests that humans become less flexible as they grow older: adults are less likely than children, for example, to adopt an unfamiliar hypothesis, despite the hypothesis being consistent with the evidence (Gopnik et al. 2017). Instead, adults prefer familiar hypotheses even if those hypotheses are less consistent with the evidence. This means that children outperform adults on problem-solving tasks when the solutions to those problems require adopting an unfamiliar hypothesis. Gopnik and colleagues interpret this and similar findings as demonstrating a trade-off between exploration and exploitation. Early in life, learners maximize exploration at the expense of exploitation. Exploring new, unfamiliar solutions to problems, however, means taking a risk of getting it wrong. Later in life, learners favor exploitation at the expense of exploration. Here there is less risk of getting it wrong insofar as the learner has experienced success applying this hypothesis. Instead, the risk is that the solution is suboptimal. When one ceases exploration too early, one might mistake a locally optimal solution for a globally optimal one. One way out of this problem is to begin with a broad search (more exploration) followed by a narrower search (more exploitation), which is what we find in developing humans (Lucas et al. 2014, Gopnik et al. 2015).

It thus seems that behaving intelligently does not mean simply having more of one ability or another, but rather striking an appropriate balance amongst a wide range of trade-offs. Furthermore, the nature of the trade-offs will vary depending on the environment, as well as the evolutionary and ecological history of an organism. Intelligence under this view is not a one-dimensional capacity of which more is better, but rather a multi-dimensional capacity that requires a context-dependent balance between trade-offs.

Considerations such as the above suggest that there is no precise definition or metric for evaluating the intelligence of animals. It might be tempting at this point to appeal to humans as a paradigm example of an intelligent organism. As the psychologist Robert Sternberg notes: "Most evolutionary approaches place humans at the top of some kind of scale of intelligence. They view humans as supremely intelligence" (Sternberg 2010, p. 251). If this is the case, why not use humans as a standard by which to measure the intelligence of other species?

There are several problems with this approach. First, like other animals, humans have their own unique specializations. Philosophers of mind have long been concerned with overcoming human chauvinism when evaluating the mental properties of other organisms (Block 1978). A chauvinist approach or theory is one that tends to deny that a system has mental properties when in fact it does (that is, the approach produces many false negatives). If we construct a benchmark for intelligence based on theories of human intelligence, then we will likely overlook many if not all cases of intelligence that diverge from the human form. This is a problem, given that the brains and bodies of organisms like the kea, octopus and bee differ dramatically from the primate brain. While humans and chimpanzees evolutionarily diverged 5-7 million years ago, humans and cephalopods, like the octopus, diverged around 600 million years ago. As the philosopher Peter Godfrey-Smith writes, "cephalopods are an independent experiment in the evolution of large brains and complex behavior" (Godfrey-Smith 2017, p. 9; see Halina 2018 for discussion). If a desideratum of comparative cognition is to understand the minds and behaviors of creatures like the octopus, then we cannot rely on humans as our standard (Birch 2020, Shevlin 2021).

In addition to the concern that intelligence is realized in dramatically different ways, there is the worry that our understanding of human intelligence is biased. The philosopher Cameron Buckner argues that we have good reason to think that we have an inflated conception of human abilities (Buckner 2013). Psychology studies suggest that humans consistently overestimate their own cognitive sophistication: attributing actions that are a result of situational factors, for instance, to rational choice. This tendency to overestimate our cognitive capacities might be exacerbated in the case of intelligence, as this concept has a long history of being used by humans to legitimize power hierarchies (with the more intelligent having the right to command or control the less intelligent). Thus, identifying humans as "intelligent" likely conveys not just information about our cognitive or behavioral abilities, but also where we stand on a social or dominance hierarchy (Cave 2020). Given this, we should treat claims of human intelligence and "supreme intelligence" with caution and work carefully to control for cognitive biases in our investigations of diverse intelligences (see Dacey 2017, Shevlin & Halina 2019).

We can think of the above concerns in terms of the observational bias known as the streetlight effect, illustrated by the following anecdote: Imagine you come across someone searching for their keys under a streetlight. You offer to help the person search for their keys, and when after some time you cannot find them, you ask, "are you sure you lost them here?" They reply, no, that they lost their keys elsewhere, but that they're searching here because "this is where the light is". Similarly, searching for intelligence under the streetlight of our conception of human intelligence severely limits our understanding of the diverse forms that intelligence might take across the animal kingdom. It also risks adopting an inflated view of intelligence and using this to evaluate the capacities of other organisms. If we would like to genuinely understand the kinds of intelligence that are found on earth (and possibly elsewhere), then we must adopt a different approach towards mapping this space.

3.2 Towards a New Approach

Developing a compelling account of what renders a behavior "intelligent" is challenging. It appears there is currently no underlying process or principle uniting those behaviors we might pick out as "intelligent". Although we can fall back on our original definition of intelligence as the ability to solve problems or achieve one's goals in a wide range of situations, we must leave behind the idea that such behavioral flexibility tracks a particular suite of mechanisms (and excludes others). We must also recognize that organisms face many trade-offs: in the face of limited resources, more will often not be better.

In developing an approach to intelligence, I suggest we begin by following Sternberg (2003) and think of intelligence as an interaction, rather than a property of an individual. He writes: "Intelligence has meaning by virtue of the kinds of problems one needs to solve in some environment. If there are no meaningful problems to solve, there is no meaningful intelligence or intelligent problem solving. Similarly, competencies and expertise exist only with respect to some kind of environmental niche" (pp. 252-253). Sternberg's point is that the concept of intelligence is meaningful only when understood as involving an interaction with an environment. Organisms are not intelligent in themselves; it is the interaction between an organism and its environment that is meaningfully described as "intelligent". Moving from individuals to interactions transforms the project of mapping the space of intelligent systems. Rather than mapping where particular categories of organisms (such as humans, honey bees, and jellyfish) fall with respect to their intelligence, we might instead focus on the interaction between the systems that make up an organism and the systems that make up the environment. For example, rather than ask, are honey bees intelligent? Or, how intelligent are honey bees? We can ask, given 1) the evolutionary and developmental history of the organism under study,

2) this organism's immediate environment, and 3) a particular goal or task, what kind of behavioral flexibility should we expect? What do we find? How is this flexibility achieved? How do these behavioral patterns and processes compare with those of other organisms? Crucially, with respect to this last question, we are not asking "which is more intelligent?" but simply "what similarities and differences do we find?" Understanding these similarities and differences provides a clearer picture of the many ways cognitive mechanisms work together to achieve adaptive behavior, and the role of the environment in this achievement.

Similar proposals regarding how to approach animal cognition and intelligence have been made in the literature. For example, the psychologists Michael Colombo and Damian Scarf write that researchers should shift their focus from "evaluating how animals differ in 'intelligence'" to "specific, definable, and measurable capacities that allow direct comparisons to be made between species" (Colombo & Scarf 2020, p. 2). Similarly, in an analysis on "rationality" in nonhuman animals, the philosopher David Papineau and psychologist Cecilia Heyes argue that, "research should refocus on specific explanations of how animals do specific things, rather than on the presence or absence of some general or ideal form of rationality that contrasts with associative mechanisms" (Papineau & Heyes 2006, p. 187; see also Allen 2014). I suggest that this applies to intelligence research as well. Although the coarse-grained contrast between rigid, mechanically produced behavior on the one hand, and flexible, intelligent behavior on the other, might have served useful functions in the past, it is now an obstacle to research.

How then should we proceed? Three recent suggestions provide useful starting points. First, the zoologist Corina Logan and colleagues outline an approach for understanding the relationship between brain morphologies and behavioral capacities, which goes "beyond brain size" and involves "deemphasizing coarse-grained notions of 'intelligence'" (Logan et al. 2018, p. 56). After reviewing the limitations of research seeking to draw a meaningful link between intelligence and absolute or relative brain size, they argue that a more nuanced approach is needed. One of their key suggestions is that researchers should take a "bottom-up" approach towards validating the relationship between behavior and neural mechanisms by focusing on individuals within a species. This would help ensure the validity of the posited causal relationships and avoid potential reifications of hypothetical mechanisms. It would also help place cognitive abilities within an evolutionary, developmental, and ecological context (Logan et al. 2018, p. 72). Once such intraspecies accounts are in place, they could then be scaled across taxa. As Logan et al. write, "rather than positing or assuming a coarse-grained, cross-taxa category and applying it across a range of cases (thus losing ecological relevance and increasing the potential for *post hoc* explanations and reification), the bottom-up approach makes scaling a much more piecemeal, empirically tractable matter" (Logan et al. 2018, p. 74; see also Brown 2018). Such a bottom-up approach to the study of intelligence would have similar advantages. Rather than starting with a coarse-grained conception of intelligence (or one reflecting what we take to be human intelligence) and trying to apply this across taxa, we should start by validating intraspecies accounts of behavioral flexibility and the mechanisms responsible for such flexibility.

A second recent suggestion comes from the philosopher Tobias Starzak and evolutionary biologist and psychologist Russell Gray. Focusing on causal cognition, they advance a framework aimed at helping researchers move away from viewing behavior as a product of sophisticated and human-like causal understanding on the one hand, and simple and nonhuman-like mechanisms on the other. They write: "Over and over again the familiar refrain is, 'do animals have complex human-like cognitive abilities or can their behavior be explained in terms of simpler processes such as associative learning?'" (Starzak & Gray 2021, p. 2). Like Logan and colleagues, they propose a more fine-grained approach and think a critical step in the right direction is to get a better handle on what is meant by "causal understanding". They propose three parameters of causal cognition: sources of causal information, integration, and explicitness. Focusing on the first parameter, an organism can gain causal information from a wide range of sources, such as experience of their own behavior, experience of the behavior of others, innate priors (as we saw in section 2.1), and other sources. One can thus compare individuals and species according to which sources of causal information they exploit (and when and how). Starzak and Gray also note that some parameters are conceptually dissociable and whether they dissociate in biological organisms is an empirical question. Such an approach highlights the many dimensions along which a general cognitive ability may vary, moving us away from an all-or-none or complex-versus-simple way of thinking.

We might similarly break down intelligent behavior according to its purported properties. For example, we might decide that problem-solving speed, multi-modal sensory information integration, and inhibition are all important dimensions of behavioral flexibility (see Buckner 2014). This would then provide us with the more tractable task of evaluating how these specific abilities are instantiated in a given organism and how they vary across taxa. There is, however, the further question of whether these dimensions are in fact dimensions of one thing. That is, whether they form a natural cluster or kind that captures what researchers mean when they use the term "behavioral flexibility" or "intelligent behavior". Buckner (2014) provides several proposals for thinking that properties such as those listed above do naturally cluster. However, the empirical evidence is still mixed. For example, although it is often assumed that behavioral flexibility requires inhibition: that one must inhibit a previously learned response to successfully adapt to a new situation, studies in humans and nonhuman animals have failed to find such a correlation (Logan et al. 2020). In studies of great-tailed grackles, behavioral flexibility has also failed to correlate with problem solving abilities and problem solving speed (Logan 2016). Logan writes that this result, "reveals how little we know about behavioral flexibility, and provides an immense opportunity for future research to explore how individuals and species can use behavior to react to changing environments" (2016, p. 25).

Finally, Taylor and colleagues have recently suggested that when investigating intelligence comparative psychologists should move from "success testing" to "signature testing" (Taylor 2014, Bastos & Taylor 2020, Taylor et al. 2021). Success testing focuses on whether an organism has succeeded or failed to solve a problem. Signature testing, in contrast, examines patterns in information processing, including errors, biases, and limitations. The problem with success testing is that simply knowing whether an organism has passed a test provides little information regarding the cognitive processes underlying that performance. Signature testing provides more information in this respect. As Taylor and colleagues write, "errors and biases can be strongly diagnostic, because while there is often only one way to solve a problem, there are many ways to make an error at a task or have a bias in how information is processed" (2021, p. 6). If two organisms rely on the same cognitive mechanisms to solve a problem, they should exhibit the same suite of signatures or patterns of errors, limitations, etc. Given that the signature-testing approach provides additional constraints on the cognitive hypothesis space, it can help researchers move away from interpreting task performance as either the result of "rich" human-like cognition or "lean" cognitive processes like associative learning, and instead home in on forms of cognition that fall between these extremes (Taylor et al. 2021, pp. 12-13).

The above strategies (validating causal relationships within species, thinking of intelligent behavior as multidimensional, and focusing on signatures rather than success) should help us map the space of intelligent systems in a way that moves beyond the Standard View. Rather than categorizing organisms as intelligent or not, these strategies seek to provide a fine-grained account that is sensitive to the unique contextual factors in which an organism or system is embedded. Once we move in this direction, however, one might ask, "what role is left for the concept of intelligence?" Intelligence and behavioral flexibility are vaguely defined. In attempting to map the space of intelligent systems, shouldn't our first task be to provide a precise definition of intelligence so that we know what to include (and exclude)?

Here I agree with the philosopher Colin Allen that some terms, like "cognition" and "life", are umbrella terms that need not be precisely defined. The same I believe is true of "intelligence." Such terms can be given a working definition, like "the capacity to achieve one's goals in a wide range of environments.", but such working definitions are often as loose as the original concept. Rather than precisely defining intelligence, such definitions might instead help orient novices to the phenomenon of interest. As Allen writes:

> To insist on precisely defining the terms that make up a working definition is to put the cart before the horse inasmuch orienting towards phenomena worthy of further investigation does not depend upon the kind of precisification that follows from detailed study of those phenomenon—just as fruitful investigation of samples of soft, yellow metal did not depend on providing, in advance, precise definitions of the terms making up that working definition of 'gold'. (Allen 2017, p. 4239)

The purpose of the term "intelligence" then might be to orient researchers towards a general class of phenomena before we know the precise nature of that class of phenomena. Given the unknown nature of the target, it might be a "superficial kind" in the sense that the members of the kind have little but the mark by which we sort them in common (Hacking 2005). This, however, need not serve as an obstacle to research. As Allen notes, research in cognitive science tends to operate on the level of specific capacities, like memory, learning, and problem solving, rather than "cognition" or "intelligence". Thus, this research can proceed despite the umbrella terms being vaguely defined. However, we should expect research on specific capacities to feed back into our understanding of the umbrella concept.

Although we need not provide a precise definition of intelligence, we should prevent the concept from distorting what we have learned about animal cognition and behavior thus far. As we have seen, the Standard View of intelligence distorts the natural world by attempting to divide it into those creatures that are intelligent and those that rely on simple, mechanical processes. Given our current knowledge of capacities like innate priors and associative learning, however, the Standard View of intelligence is clearly in need of an update. The investigative strategies proposed here can help us better understand the specific mechanisms and abilities that fall under the umbrella term "intelligence". However, we must also be explicit that this term is too vaguely defined to do any real scientific or philosophical work. We currently do not know exactly what class of mechanisms and behaviors it picks out. Rather than prejudging the situation (based on sweeping theories and intuitions), we should limit the use of the term to the role of "vague guide", but let the real guidance come from research on specific mechanisms and behaviors.

3.3 Conclusion

We have seen that the line between intelligent and unintelligent organisms is not as clear as assumed by the Standard View. First, innate behavioral programs and associative learning can give rise to flexible behavior, and there is good evidence to think that they do in a wide range of organisms. Second, there is reason to think that the rigid, mechanical behavior traditionally associated with Darwinian and Skinnerian creatures is a myth. Mapping out the space of intelligent systems is best accomplished using a bottom-up approach that focuses on the specific capacities and behaviors of organisms within their evolutionary, developmental, and ecological context. The term "intelligence" can still be used to loosely orient scientific research, but it should not be mistaken for a precise account of mechanisms and behavior.

We began this Element with the cases of the mechanical Turk and tic-tac-toe playing chicken. The skeptical response to such systems was that they must result from some trick: intelligent, flexible behavior could not possibly be the product of simple mechanisms; thus, there must be a human behind the scenes pulling the strings. A less skeptical response to these systems was to stand in wonder and curiosity: how could nonhuman mechanisms give rise to such intelligent behavior? Although the skeptics turned out to be correct in the case of the mechanical Turk and tic-tac-toe playing chicken, contemporary research suggests that this skeptical heuristic fails when applied to animals more broadly. The world does not consist of only simple or human-like mechanisms, but instead a multitude of processes that fit neither category, and which many of us have trouble grasping in an intuitive way. Rather than trying to fit these processes into the more familiar categories of "simple" or "human-like", I suggest we embrace the feelings of wonder and curiosity and recognize that we are facing something that will take some time to intuitively understand.

References

Allen, C. (2014). Models, Mechanisms, and Animal Minds: Models, Mechanisms, and Animal Minds. *The Southern Journal of Philosophy*, **52**, 75–97.

Allen, C. (2017). On (not) defining cognition. *Synthese*, **194**(11), 4233–4249.

Andrews, K., Comstock, G., Crozier, G.K.D., Donaldson, S., Fenton, A., John, T.M., Johnson, L.S.M., Jones, R.C., Kymlicka, W., Meynell, L. and Nobis, N. (2018). *Chimpanzee Rights: The Philosophers' Brief*, London: Routledge.

Animal Behavior Enterprises. (n.d.). *Bird Brain Manual*.

Bailey, R. E., & Gillaspy, J. A. (2005). Operant psychology goes to the fair: Marian and Keller Breland in the popular press, 1947–1966. *The Behavior Analyst*, **28**(2), 143–159.

Bastos, A. P. M., & Taylor, A. H. (2020). Kea show three signatures of domain-general statistical inference. *Nature Communications*, **11**(1), 828.

Bateson, P., & Mameli, M. (2007). The innate and the acquired: Useful clusters or a residual distinction from folk biology? *Developmental Psychobiology*, **49**(8), 818–831.

Bihm, E. M., Gillaspy, J. A., Abbott, H. J., & Lammers, W. J. (2010). More Misbehavior of Organisms: A PSI Chi Lecture by Marian and Robert Bailey. *The Psychological Record*, **60**(3), 505–521.

Birch, J. (2020). The Search for Invertebrate Consciousness. *Noûs*, 1–21.

Birch, J., Ginsburg, S., & Jablonka, E. (2020). Unlimited Associative Learning and the origins of consciousness: a primer and some predictions. *Biology & Philosophy*, **35**(6), 56.

Block, N. (1978). Troubles with Functionalism. *Minnesota Studies in the Philosophy of Science*, **9**, 261–325.

Bluff, L. A. (2006). Tool-Related Cognition in New Caledonian Crows. *Comparative Cognition & Behavior Reviews*, **1**, 1–25.

Bolhuis, J. J., de Vos, G. J., & Kruijt, J. P. (1990). Filial imprinting and associative learning. *The Quarterly Journal of Experimental Psychology*, **42**(3), 313–329.

Breen, A. J. (2021). Animal culture research should include avian nest construction. *Biology Letters*, **17**(7), 20210327.

Breen, A. J., Guillette, L. M., & Healy, S. D. (2016). What Can Nest-Building Birds Teach Us? *Comparative Cognition & Behavior Reviews*, **11**, 83–102.

Breland, K., & Breland, M. (1951). A field of applied animal psychology. *American Psychologist*, **6**(6), 202–204.

Bronfman, Z. Z., Ginsburg, S., & Jablonka, E. (2016a). The Evolutionary Origins of Consciousness: Suggesting a Transition Marker. *Journal of Consciousness Studies*, **23**(9–10), 7–34.

Bronfman, Z. Z., Ginsburg, S., & Jablonka, E. (2016b). The Transition to Minimal Consciousness through the Evolution of Associative Learning. *Frontiers in Psychology*, **7**, 1–16.

Brown, R. L. (2019). Infer with care: A critique of the argument from animals. *Mind & Language*, **34**(1), 21–36.

Buckner, C. (2011). Two Approaches to the Distinction between Cognition and 'Mere Association.' *International Jounral of Comparative Psychology*, **24**(4), 314–348.

Buckner, C. (2013). Morgan's Canon, meet Hume's Dictum: avoiding anthropofabulation in cross-species comparisons. *Biology & Philosophy*, **28**(5), 853–871.

Buckner, C. (2015). A property cluster theory of cognition. *Philosophical Psychology*, **28**(3), 307–336.

Cave, S. (2020). The Problem with Intelligence: Its Value-Laden History and the Future of AI. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, New York NY USA: ACM, pp. 29–35.

Carruthers, P. (2004). On Being Simple Minded. *American Philosophical Quarterly*, **41**(3), 205–220.
Dennett, D. C. (1998). *Brainchildren: Essays on Designing Minds*, Cambridge, Mass: MIT Press.

Clayton, N., Emery, N., & Dickinson, A. (2006). The rationality of animal memory: Complex caching strategies of western scrub jays. In S. Hurley & M. Nudds, eds., *Rational Animals?*, Oxford

University Press, pp. 197–216.

Colombo, M., & Scarf, D. (2020). Are There Differences in "Intelligence" Between Nonhuman Species? The Role of Contextual Variables. *Frontiers in Psychology*, **11**, 2072.

Dacey, M. (2016). Rethinking associations in psychology. *Synthese*, **193**(12), 3763–3786.

Dacey, M. (2017). Anthropomorphism as Cognitive Bias. *Philosophy of Science*, **84**(5), 1152–1164.

Dennett, D. C. (1996). *Kinds of Minds: Towards an Understanding of Consciousness*, New York, NY: Basic Books.

Dennett, D. (2018). Reflections on Peter Godfrey-Smith. In B. Huebner, ed., *The Philosophy of Daniel Dennett*, Vol. 1, Oxford: Oxford University Press, pp. 250–254.

Dretske, F. I. (1970). Epistemic Operators. *The Journal of Philosophy*, **67**(24), 1007–1023.

Drumm, P. (2009). Applied Animal Psychology at an American Roadside Attraction: Animal Behavior Enterprises and The IQ Zoo of Hot Springs, Arkansas. *The American Journal of Psychology*, **122**(4), 537–545.

Emery, N. J., & Clayton, N. S. (2004). The Mentality of Crows: Convergent Evolution of Intelligence in Corvids and Apes. *Science*, **306**(5703), 1903–1907.

Fridland, E. (2015). Learning Our Way to Intelligence: Reflections on Dennett and Appropriateness. In C. Muñoz-Suárez and F. De Brigard, eds., *Content and Consciousness Revisited*, Cham: Springer International Publishing, pp. 143–161.

Gallistel, C. R. (2008). Learning and Representation. In J. H. Byrne, ed., *Learning and Memory: A Comprehensive Reference*, New York, NY: Elsevier Ltd., pp. 227–242.

Ginsburg, S., & Jablonka, E. (2019). *The Evolution of The Sensitive Soul: Learning and The Origins of Consciousness*, Cambridge, MA: MIT Press.

Ginsburg, S., & Jablonka, E. (2021). Evolutionary transitions in learning and cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **376**(1821), 20190766.

Godfrey-Smith, P. (2018). Towers and Trees in Cognitive Evolution. In B. Huebner, ed., *The Philosophy of Daniel Dennett*, Vol. 1, Oxford: Oxford University Press, pp. 225–249.

Gopnik, A., Griffiths, T. L., & Lucas, C. G. (2015). When Younger Learners Can Be Better (or at Least More Open-Minded) Than Older Ones. *Current Directions in Psychological Science*, **24**(2), 87–92.

Gopnik, A., O'Grady, S., Lucas, C. G., … Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, **114**(30), 7892–7899.

Gorostiza, E. A. (2018). Does Cognition Have a Role in Plasticity of "Innate Behavior"? A Perspective From Drosophila. *Frontiers in Psychology*, **9**, 1502.

Gravitz, L. (2019). The Importance of Forgetting. *Nature*, **571**, S12–S14.

Grossi, G. (2017). Hardwiring: innateness in the age of the brain. *Biology & Philosophy*, **32**(6), 1047–1082.

Hacking, I. (2005). Why race still matters. *Daedalus*, **134**(1), 102–116.

Halina, M. (2018). Octopuses as conscious exotica. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, **67**, 28–31.

Halina, M. (2021). Insightful artificial intelligence. *Mind & Language*, **36**(2), 315–329.

Hansell, M., & Ruxton, G. (2008). Setting tool use within the context of animal construction behaviour. *Trends in Ecology & Evolution*, **23**(2), 73–78.

Hanus, D. (2016). Causal reasoning versus associative learning: A useful dichotomy or a strawman battle in comparative psychology? *Journal of Comparative Psychology*, **130**(3), 241–248.

Heiling, A. M., & Herberstein, M. E. (1999). The role of experience in web-building spiders (Araneidae). *Animal Cognition*, **2**(3), 171–177.

Heinrich, B. (1995). An Experimental Investigation of Insight in Common Ravens (Corvus corax). *The Auk*, **112**(4), 994–1003.

Heyes, C. (2012). Simple minds: a qualified defence of associative learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **367**(1603), 2695–2703.

Hofstadter, D. R. (1982). Metamagical Themas: Can Inspiration be Mechanized? *Scientific American*, **247**(3), 18–34.

Jacobs, I. F., & Osvath, M. (2015). The string-pulling paradigm in comparative psychology. *Journal of Comparative Psychology*, **129**(2), 89–120.

Kaufman, S. B., DeYoung, C. G., Gray, J. R., Brown, J., & Mackintosh, N. (2009). Associative learning predicts intelligence above and beyond working memory and processing speed. *Intelligence*, **37**(4), 374–382.

Keijzer, F. A. (2011). *Representation and Behavior*, Cambridge, MA: MIT Press.

Keijzer, F. (2013). The *Sphex* story: How the cognitive sciences kept repeating an old and questionable anecdote. *Philosophical Psychology*, **26**(4), 502–519.

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, **40**, e253.

Legg, S., & Hutter, M. (2007). Universal Intelligence: A Definition of Machine Intelligence. *ArXiv:0712.3329 [Cs]*. Retrieved from http://arxiv.org/abs/0712.3329

Lind, J. (2018). What can associative learning do for planning? *Royal Society Open Science*, **5**, 1–14.

Liu, S., Brooks, N. B., & Spelke, E. S. (2019). Origins of the concepts cause, cost, and goal in prereaching infants. *Proceedings of the National Academy of Sciences*, **116**(36), 17747–17752.

Logan, C. J. (2016). Behavioral flexibility and problem solving in an invasive bird. *PeerJ*, **4**, e1975.

Logan, C. J., Avin, S., Boogert, N., … Montgomery, S. H. (2018). Beyond brain size: Uncovering the neural correlates of behavioral and cognitive specialization. *Comparative Cognition & Behavior Reviews*, **13**, 55–89.

Logan, C. J., McCune, K., MacPherson, M., … Wascher, C. (2020). *Are the more flexible great-tailed grackles also better at behavioral inhibition?* (preprint), PsyArXiv. Retrieved from https://osf.io/vpc39

Lorenz, K. Z. (1937). The Companion in the Bird's World. *The Auk*, **54**(3), 245–273.

Lucas, C. G., Bridgers, S., Griffiths, T. L., & Gopnik, A. (2014). When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*, **131**(2), 284–299.

Mameli, M. & Journal of Philosophy, Inc. (2008). On Innateness: The Clutter Hypothesis and the Cluster Hypothesis. *Journal of Philosophy*, **105**(12), 719–736.

Martinho, A., & Kacelnik, A. (2016). Ducklings imprint on the relational concept of "same or different." *Science*, **353**(6296), 286–288.

Mascalzoni, E., Regolin, L., & Vallortigara, G. (2010). Innate sensitivity for self-propelled causal agency in newly hatched chicks. *Proceedings of the National Academy of Sciences*, **107**(9), 4483–4485.

Mascalzoni, E., Regolin, L., Vallortigara, G., & Simion, F. (2013). The cradle of causal reasoning: newborns' preference for physical causality. *Developmental Science*, **16**(3), 327–335.

Mikhalevich, I., Powell, R., & Logan, C. (2017). Is behavioural flexibility evidence of cognitive complexity? How evolution can inform comparative cognition. *Interface Focus*, **7**(3), 20160121.

Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, **32**(2), 183–198.

Papineau, D., & Heyes, C. (2006). Rational or associative? Imitation in Japanese quail. In S. Hurley & M. Nudds, eds., *Rational Animals?*, Oxford University Press, pp. 187–196.

Parker, E. S., Cahill, L., & McGaugh, J. L. (2006). A Case of Unusual Autobiographical Remembering. *Neurocase*, **12**(1), 35–49.

Platt, M. L., & Spelke, E. S. (2009). What can developmental and comparative cognitive neuroscience tell us about the adult human brain? *Current Opinion in Neurobiology*, **19**(1), 1–5.

Powell, R. A., Honey, P. L., & Symbaluk, D. G. (2017). *Introduction to Learning and Behavior, Fifth Edition*, Mexico: Cengage Learning.

Raby, C. R., & Clayton, N. S. (2009). Prospective cognition in animals. *Behavioural Processes*, **80**(3), 314–324.

Richards, B. A., & Frankland, P. W. (2017). The Persistence and Transience of Memory. *Neuron*, **94**(6), 1071–1084.

Riskin, J. (2016). *The Restless Clock: A History of the Centuries-Long Argument over What Makes Living Things Tick*, Chicago: The University of Chicago Press.

Rosa-Salva, O., Mayer, U., Versace, E., Hébert, M., Lemaire, B. S., & Vallortigara, G. (2021). Sensitive periods for social development: Interactions between predisposed and learned mechanisms. *Cognition*, **213**, 1–13.

Roth, G., & Dicke, U. (2005). Evolution of the brain and intelligence. *Trends in Cognitive Sciences*,

9(5), 250–257.

Rowe, C., & Healy, S. D. (2014). Measuring variation in cognition. *Behavioral Ecology*, **25**(6), 1287–1292.

Schaffer, S. (1999). Enlightened Automata. In C. William, J. Golinski and S. Schaffer, eds., *The Sciences in Enlightened Europe*. London: The University of Chicago Press, pp. 126–165.

Schnell, A. K., Amodio, P., Boeckle, M., & Clayton, N. S. (2021). How intelligent is a cephalopod? Lessons from comparative cognition. *Biological Reviews*, **96**(1), 162–178.

Seed, A. M., Emery, N., & Clayton, N. (2009). Intelligence in Corvids and Apes: A Case of Convergent Evolution? *Ethology*, **115**(5), 401–420.

Seed, A., Hanus, D., & Call, J. (2011). Causal Knowledge in Corvids, Primates, and Children. In T. McCormack, C. Hoerl, & S. Butterfill, eds., *Tool Use and Causal Cognition*, Oxford University Press, pp. 89–110.

Seed, A. M., Tebbich, S., Emery, N. J., & Clayton, N. S. (2006). Investigating Physical Cognition in Rooks, Corvus frugilegus. *Current Biology*, **16**(7), 697–701.

Shanahan, M., Crosby, M., Beyret, B., & Cheke, L. (2020). Artificial Intelligence and the Common Sense of Animals. *Trends in Cognitive Sciences*, **24**(11), 862–872.

Shettleworth, S. J. (2012). Do animals have insight, and what is insight anyway? *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, **66**(4), 217–226.

Shettleworth, S. J. (2013). *Fundamentals of Comparative Cognition*, Oxford: Oxford University Press.

Shevlin, H. (2021). Non-human Consciousness and The Specificity Problem: A Modest Theoretical Proposal. *Mind & Language*, **36**(2), 297–314.

Shevlin, H., & Halina, M. (2019). Apply rich psychological terms in AI with care. *Nature Machine Intelligence*, **1**(4), 165–167.

Shevlin, H., Vold, K., Crosby, M., & Halina, M. (2019). The limits of machine intelligence: Despite progress in machine intelligence, artificial general intelligence is still a major challenge. *EMBO Reports*, **20**(10), 1–5.

Silva, F. J., & Silva, K. M. (2006). Humans' folk physics is not enough to explain variations in their tool-using behavior. *Psychonomic Bulletin & Review*, **13**(4), 689–693.

Silver, D., Schrittwieser, J., Simonyan, K., … Hassabis, D. (2017). Mastering the game of Go without human knowledge. *Nature*, **550**(7676), 354–359.

Sinnott-Armstrong, W. (2002). What's in a Contrast Class? *Analysis*, **61**(1), 75–84.

Sinnott-Armstrong, W. (2008). A Contrastivist Manifesto. *Social Epistemology*, **22**(3), 257–270.

Sloman, A. (1984). The Structure of The Space of Possible Minds. In S. Torrance and E. Horwood, eds., The Mind and The Machine: Philosophical Aspects of Artificial Intelligence, John Wiley & Sons, pp. 35–42.

Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, **10**(1), 89–96.

Standage, T. (2002). *The Turk: The Life and Times of the Famous Eighteenth-Century Chess-Playing Machine*, New York: Berkley Books.

Starzak, T. B., & Gray, R. D. (2021). Towards ending the animal cognition war: a three-dimensional model of causal cognition. *Biology & Philosophy*, **36**(2), 9.

Sternberg, R. J. (2003). Biological Intelligence. In R. J. Sternberg & E. L. Grigorenko, eds., *The Psychology of Abilities, Competencies, and Expertise*, Cambridge: Cambridge University Press, pp. 240–262.

Taylor, A. H. (2009). *Complex Cognition in the New Caledonian Crow (Corvus moneduloides)*, Thesis (PhD), The University of Auckland.

Taylor, A. H. (2014). Corvid cognition: Corvid cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, **5**(3), 361–372.

Taylor, A. H., Bastos, A., Brown, R., & Allen, C. (2021). The signature-testing approach to mapping biological and artificial intelligences. Manuscript submitted for publication.

Ticknesse, P. (1784). *The Speaking Figure, and the Automaton Chess-Player, Exposed and Detected*, London: John Stockdale.

Trillin, C. (1999). The Chicken Vanishes: But who Taught it to Play Ticktacktoe? Finally, the Mystery is Solved. *The New Yorker*, 8 February pp. 38–41.

Vallortigara, G. (2012). Core knowledge of object, number, and geometry: A comparative and neural approach. *Cognitive Neuropsychology*, **29**(1–2), 213–236.

Versace, E., Martinho-Truswell, A., Kacelnik, A., & Vallortigara, G. (2018). Priors in Animal and Artificial Intelligence: Where Does Learning Begin? *Trends in Cognitive Sciences*, **22**(11), 963–965.

Versace, E., & Vallortigara, G. (2015). Origins of Knowledge: Insights from Precocial Species. *Frontiers in Behavioral Neuroscience*, **9**, 1–10.

Voudouris, K. (2020). *An Interdisciplinary Approach to the Associative-Cognitive Distinction*, Thesis (Master of Philosophy), University of Cambridge.

Wei, D., Talwar, V., & Lin, D. (2021). Neural circuits of social behaviors: Innate yet flexible. *Neuron*, **109**(10), 1600–1620.

Wells, H. G., Huxley, J., & Wells, G. P. (1931). *The Science of Life*, London: Cassell and Company Limited.

Willis, R. (1821). *An Attempt to Analyse the Automaton Chess Player, of Mr de Kempelen. With an Easy Method of Imitating the Movements of That Celebrated Figure*, London: J. Booth.

Windisch, M. C. G. (1784). *Inanimate Reason; or a Circumstantial Account of That Astonishing Piece of Mechanism, M. de Kempelen's Chess-Player*, London: S. Bladon.

Wise, S. M. (2010). Legal Personhood and The Nonhuman Rights Project. *Animal Law*, **17**, 1–12.