

Review



Cite this article: Gopnik A. 2020 Childhood as a solution to explore–exploit tensions. *Phil. Trans. R. Soc. B* **375**: 20190502.
<http://dx.doi.org/10.1098/rstb.2019.0502>

Accepted: 24 February 2020

One contribution of 16 to a theme issue ‘Life history and learning: how childhood, caregiving and old age shape cognition and culture in humans and other animals’.

Subject Areas:

cognition, developmental biology, behaviour

Keywords:

development, life history, explore–exploit

Author for correspondence:

Alison Gopnik
e-mail: gopnik@berkeley.edu

Childhood as a solution to explore–exploit tensions

Alison Gopnik

Department of Psychology, University of California, 2121 Berkeley Way, Room 3302, Berkeley, CA 94720-1650, USA

AG, 0000-0002-7183-0867

I argue that the evolution of our life history, with its distinctively long, protected human childhood, allows an early period of broad hypothesis search and exploration, before the demands of goal-directed exploitation set in. This cognitive profile is also found in other animals and is associated with early behaviours such as neophilia and play. I relate this developmental pattern to computational ideas about explore–exploit trade-offs, search and sampling, and to neuroscience findings. I also present several lines of empirical evidence suggesting that young human learners are highly exploratory, both in terms of their search for external information and their search through hypothesis spaces. In fact, they are sometimes more exploratory than older learners and adults.

This article is part of the theme issue ‘Life history and learning: how childhood, caregiving and old age shape cognition and culture in humans and other animals’.

1. Introduction

One of the most distinctive aspects of human life history is our exceptionally extended childhood. Chimpanzees produce as much food as they consume by the time they are around 7 years old; even in forager cultures human children are not self-sufficient until they are at least 15 [1]. Moreover, there is evidence from the fossil record that childhood was extended during the relatively brief period in which modern humans evolved. *Homo sapiens* appear to have had a longer period of immaturity than other hominins, even including Neanderthals [2,3].

Human childhood is expensive. Adults must feed and protect the young into adolescence. Early brain development is also very energetically costly, with more than 60% of 4-year-olds’ calories going to the brain at rest, compared with around 20% for adults [4]. Humans also have shorter inter-birth intervals than their closest primate relatives [5], so they stack up even more of those costly babies.

Humans developed exceptionally extended and varied sources of caregiving investment to deal with this cost, including pair-bonded fathers [6], related and unrelated alloparents [7], and post-menopausal grandmothers [8]. This is in contrast with our closest primate relatives and with most mammals. All of these types of investment are found in other species, but they are relatively rare, and they are not found in the other great apes. No other species appears to have all these types of investment. The role of grandmothers points to an additional life-history feature of humans, our longevity. We live some 20 years longer than chimpanzees.

These changes in life history evolved in concert with well-known changes in human cognitive capacities. No single psychological trait is found in humans but not in any other species, with the possible exception of language. However, the combination of brain and cognitive changes in humans is distinctive, including greatly increased brain size and neuron numbers, and an expansion of the

cortex [9]. These changes also include the emergence of impressive capacities for both physical and social cognition. These include new kinds of physical tool use and physical cognition, particularly causal cognition (e.g. [10–12]), sophisticated capacities for ‘theory of mind’ and social interaction and cooperation [13–15] and impressive capacities for cultural learning and transmission [16,17].

As with caregiving investment, elements of these capacities can be found in other species. Corvids have impressive capacities for ‘folk physics’ (e.g. [18,19]). Social carnivores as well as some other primates have striking capacities for social cognition and cooperation, and recent work suggests that elements of ‘theory of mind’ can be found in great apes (e.g. [20–22]). Cetaceans such as whales and dolphins, as well as birds, show capacities for cultural learning [3,23]. But the full suite of these cognitive capacities and their extent in humans is distinctive.

The rapidity and breadth of the changes in cognition and life history that led to modern humans strongly suggests that many coevolutionary processes emerged in concert rather than that there was a single decisive adaptation. In this paper, I suggest that distinctive cognitive features of childhood, particularly the capacity and motivation for internal and external exploration in a protected period, were one important factor contributing to this coevolutionary cascade.

The fact that both the extended human childhood and the cognitive changes emerged so quickly and in concert, and particularly that the life-history changes are so costly, suggests that there is some adaptive relationship between them. This has, in fact, been suggested in a general way previously [25,26]. But what might this relationship be like in more detail? It is certainly possible that the relationship is epiphenomenal rather than adaptive. That is, a long childhood, or even simply a longer overall lifespan, might be selected for some reason and a larger brain and cortical size might simply be a result of the longer developmental period (e.g. [27]). Or, vice-versa, a larger brain might simply require a longer waiting period of development [28].

Alternatively, there might be some quite specific links between particular adult cognitive capacities and the extended human childhood. For example, Hrdy and Burkart [29], Hawkes [30] and Tomasello [31], in this volume, suggest that social cognition, in particular, was required both to allow infants and children to navigate the demands of engaging a range of carers as well as to allow multiple carers to cooperate. Alternatively, Kaplan *et al.* [1], Boyd & Richerson [16], Henrich [17] and Sterelny [32] suggest specific links between an extended childhood and time for the transmission of cultural skills, both physical and social, that are necessary for adult hunting and foraging.

However, all of these approaches assume a fundamental continuity between the minds of children and adults. I will argue instead that childhood itself, both in humans and other species, involves a distinctive set of cognitive characteristics that actually trade off with adult characteristics. This trade-off is described in the computational literature in terms of an intrinsic tension between exploration and exploitation. The ‘explore’ features of childhood minds include a general capacity for learning and plasticity which allows sensitivity to a wide range of environmental and behavioural possibilities. Moreover, this developmental programme not only involves cognitive, computational and neural differences between children and adults but also differences in

motivation, emotion and action. Juveniles of many species not only show a capacity for learning about new environments, they are actively motivated to do so. This is manifested in the distinctive suite of emotions and activities that are associated with childhood, especially the neophilia and curiosity that lead to active exploration and play and are particularly characteristic of human children.

These capacities are in tension with more adult capacities for skilled action in a particular environment, including cognitive features like attentional focus, inhibition, and executive function and behaviours like long-term, goal-directed, planned action. A long childhood allows for a kind of developmental division of labour, with an early protected period devoted to learning and exploration and a later adult period devoted to directed exploitation based on what has been learned earlier.

Special capacities for learning and exploration, exercised in a protected period of immaturity, would help to enable the whole range of striking human cognitive changes. Many of the specific computations that are involved in particular cognitive capacities, such as intuitive physics or psychology, or social understanding and cooperation, could, in principle, be either built in or restricted to particular domain-specific kinds of learning (this may indeed be the case in other species). A long and wide-ranging period of learning and exploration, in contrast, would help to enable the whole range of human cognitive skills, from physics to psychology, from communication to culture.

The idea of childhood as an evolutionary solution to explore–exploit trade-offs also applies widely beyond humans, unlike some of the more specific proposals concerning cooperation and culture. It may seem evolutionarily paradoxical to develop a life history that includes expensive and vulnerable young for a long period. However, across many different species, including birds, both placental and marsupial mammals, and even insects, there is a very general (though not perfect) correlation between the length of immaturity and degree of parental investment and relative brain size, intelligence and a reliance on learning [9,33–36]. Humans are, of course, at the far end of the distribution on all these measures. Explore–exploit trade-offs may help to explain this strikingly general and widespread relationship between an extended childhood and learning.

2. The explore–exploit dilemma

What is the explore–exploit dilemma? Here is a classic example. You can choose whether to go to your old reliable restaurant or try a new place that might be better or might be worse. The first option will give you a good meal but no new information. The second option risks a less good meal but will give you more information—and might lead to something even better than your known choice. Which should you do? It turns out to be remarkably difficult to design a systematic strategy to solve this problem.

The most well-known versions of the problem arise in the context of reinforcement learning algorithms. These algorithms have played a major role in the recent success of machine learning, have important connections to neuroscience and have been applied extensively to foraging behaviour in non-human animals [37]. In reinforcement learning, an agent acts on the world and observes the positive or negative outcomes of

those actions. The agent must choose whether to act to explore unknown properties of the environment (like trying the new restaurant) or to act to exploit known rewards (like sticking with the old favourite). For example, in a classic ‘bandit’ problem, the learner has to decide between pulling two levers. She knows that one lever will produce a reward half the time (the exploit choice), but she can also pull a new lever with an unknown pay-off (the explore choice). When should she go with the modest but sure chance of a win, and when should she risk the mystery lever that might lead to a bigger reward overall?

Solving this explore–exploit trade-off is computationally intractable in realistically complex environments, so agents must use heuristics to find an efficient solution (see [38], for a recent review). With *random* exploration heuristics, the agent might just randomly try new actions. In *directed* exploration algorithms, the agent specifically crafts actions that will be most likely to be informative. Adult humans apply a mix of random and directed exploration heuristics to solve reinforcement learning problems [39].

The classic reinforcement learning problem involves a search for new data, a kind of external exploration. But other areas of computer science, particularly optimization theory, point to a similar set of trade-offs when it comes to finding new hypotheses or solutions—a kind of internal exploration [40]. Consider, for example, a situation in which you are trying to find the best causal explanation for a particular pattern of data. If the space of possibilities is reasonably complex, it will be impossible for any system, human or computer, to consider and compare all the relevant causal hypotheses and search through the whole space. Computer scientists and statisticians often use ‘sampling’ techniques such as Monte Carlo methods to help solve this problem—semi-randomly generating some hypotheses rather than others [41]. We have discovered evidence that people, including young children, do something similar [42–44].

The search and sampling process, however, presents learners with an explore–exploit dilemma. A learner can conduct a narrow search, only revising current hypotheses when the evidence is particularly strong and making small adjustments to current theories to accommodate new evidence. This strategy is most likely to quickly yield a ‘good enough’ solution that will support immediate effective action. But it also means that the learner may fail to imagine a better alternative that is farther from the current hypothesis, such as a hypothesis about an unusual causal relation.

Alternatively, a learner can conduct a more exploratory search, moving to new hypotheses with only a small amount of evidence, and trying out potential hypotheses that are less like the current hypotheses. This strategy is less efficient if the learner’s starting hypothesis is reasonably good and may mean that the learner wastes time imagining unlikely possibilities. But it may also make the learner more likely to discover genuinely new ideas. Drawing on an analogy to statistical physics, computer scientists have described this difference in terms of a contrast between narrower ‘low temperature’ exploit searches and broader ‘high temperature’ exploratory ones [40].

Although there are many different approaches to solving the explore–exploit dilemma, they tend to share a common feature—the learner should begin by exploring and gradually converge on exploiting. There are two reasons for this. One is simply that the learner is likely to start out knowing less and

gradually acquire more knowledge. As the learner knows more, it makes sense to rely more on that prior knowledge and be less motivated to acquire new knowledge. The second reason is that if there is a limited period in which to solve the task, then as that period elapses, there will be fewer opportunities to take advantage of the information that has been acquired through exploration.

Similarly, in the literature on search and sampling, ‘simulated annealing’ [40] is one of the best ways of resolving the tension between high- and low-temperature strategies. Learners, who begin with a broader higher-temperature search and gradually move to a narrower low-temperature search, are most likely to find the optimal solution, just as in metallurgy heating a metal and then cooling it leads to the most robust structure.

Beginning with a high-temperature search and moving to a low-temperature one is particularly important as a way to solve the problem of ‘local optima’. In a low-temperature search, you make small changes to your current state, see if they are improvements, move to the new state if so, and then repeat the process. This is sometimes described as ‘hill-climbing’, gradually ascending to a better state (e.g. [45]). However, if the space is complex, then the current hypothesis may be better than all the local alternatives, but much worse than an alternative that is further away and unlike the current possibilities. Hill-climbing algorithms may take you to the top of a local hill but leave you stuck on a ‘local maximum’, unable to reach an even higher state on a more distant peak.

If you start off by bouncing around the space with a high degree of randomness, you will be more likely to land on this initially unlikely option, along with many others. However, in order for the solution to be stable enough to support action, the learner must eventually settle on a particular option and narrow their search. Simulated annealing, starting off hot and cooling off later, allows the learner both to escape local optima and to settle on stable solutions. In particular, the early high-temperature search allows the learner to explore the high-level structure of the problem space. It gives the learner a sense of the general kinds of solutions that are possible, while the cooler later search allows the learner to hone-in on details.

3. Development and explore–exploit trade-offs

My colleagues and I have argued that this explore-first/exploit-later strategy may be embodied in a developmental division of labour [46–48]. The explore–exploit division of labour might be accomplished in other ways. In the classic reinforcement learning ‘bandit’ tasks, the ‘explore then exploit’ sequence may take place during the course of learning itself, in a single individual for a single problem. In social insects, such as honeybees, different individuals may take on the exploratory versus exploitative roles—scouts contrast with worker bees. Complex human societies may do the same with specialized exploratory roles for shamans or scientists.

But the developmental life-history strategy of an extended exploratory childhood followed by an exploitative adulthood has several advantages. The interests of the child who does the learning and the grown-up who exploits what she has learned are automatically the same, since they are the same individual, obviating the free rider problems in a social

division of labour. Similarly, the caregiving investment that frees up the child for exploration is linked to genetic relatedness—although alloparents may be non-kin, usually the carers are also parents, older siblings or grandparents. Moreover, the temporal sequence of development, being a child first and an adult later, automatically results in the explore-then-exploit sequence. Most of all, protected immaturity allows children to exercise specialized capacities dedicated to learning and exploration, without having to simultaneously exploit.

It may seem tautological that childhood precedes adulthood, and more general kinds of learning would precede more specific ones. But it would be possible to have the reverse developmental strategy, in which initial knowledge was highly restricted and oriented to the most significant survival challenges the organism faces, but computationally complex, well developed and highly specified. More generally applicable broad learning mechanisms such as association might be employed later to fine tune responses to a particular environment. This is the strategy implied by ‘core knowledge’ and similar nativist theories of development [49]. Indeed, arguably this is the strategy pursued by highly precocial animals such as chickens and other galliformes. These animals have elaborate but narrow computational capacities in place at birth, and use them competently (e.g. [50]). And, not coincidentally, they mature quickly, and have very brief juvenile periods.

This precocial strategy, with much structure in place early, will be especially valuable in highly constrained and predictable environments. By contrast, an extended period of exploration is particularly valuable in situations in which environments are variable, with a mix of predictability and unpredictability. There are interesting questions and significant formal work about just which kinds of variability and predictability would make extended development adaptive (see [51,52]). In general, though, the value of exploration is directly related to the dimensionality and complexity of both the environment and the potential adaptations to that environment. The broader and more complex the possibilities are, the more exploration will be valuable.

It is plausible that increased environmental variability was associated, in particular, with human evolution. The original trigger may have been increased climactic variability [53], but human behaviours themselves lead to increased environmental variability. Nomadic behaviour means that humans characteristically face novel environments and the human capacity for culture and niche construction creates novel social and physical environments across generations.

The general explore-then-exploit pattern appears to be manifested in a wide range of species with different types of learning capacities, including simpler forms of plasticity, as well as the kinds of sophisticated learning we see in humans. For example, Snell-Rood *et al.* [35] find that in cabbage white butterflies, different developmental strategies are associated with whether the animal relies on learning or is restricted to innate preferences. Individuals who rely more on learning have fewer young and a longer developmental period. Similarly, Frankenhuis & Panchanathan [52] argue that phenotypic plasticity, such as whether and when water fleas grow a protective helmet, is associated with differing developmental trajectories, and that a period of sampling the environment, determining predator density, for example, precedes the commitment to a particular specialized

adaptation. ‘Sensitive periods’ are a similar example of different cognitive adaptations in different developmental periods: the animal in the sensitive period is open to learning in a particular domain, the post-sensitive animal is not. And sensitive periods usually come early in development. In general, a pattern with greater early plasticity leading to a more efficient but inflexible state later on makes sense from the explore-exploit perspective.

But the general principle applies with special force to human evolution. Primates, in general, and humans, in particular, rely on learning particularly heavily, and also have much longer childhoods and much more caregiving investment than chickens, butterflies or fleas. Rather than having a few sensitive periods of plasticity adapted to particular domains, a human mind has to explore the very wide and unpredictable range of human possibilities, both in terms of possible actions and possible hypotheses. You could think of an extended curious childhood with particularly powerful kinds of learning as a kind of turbo-powered super sensitive period—a protected time to extract information from the environment through exploration and to imagine even far-away and unlikely hypotheses.

4. Developmental evidence for the explore–exploit hypothesis

The explore-exploit hypothesis helps to resolve some paradoxes in children’s cognitive development. Human infants and children are remarkably effective learners. In a very short time, they learn about objects, people, animals and plants creating intuitive theories of the physics, biology and psychology of the world around them. They do this well before they go to school, with no explicit teaching [13,54–56].

In the comparative literature, where most of the work on life history has taken place, there has been a tendency to contrast innate hard-wired domain-specific reflexes on the one hand, and domain-general associationist learning mechanisms on the other. But the kinds of learning that human children engage in go well beyond simple association. Even very young human children learn by formulating and testing structured causal hypotheses about the world, updating them in the light of new evidence [56–58]. Cultural learning is particularly important for human children. Children are especially sensitive to information they obtain from others, both in their everyday observation and imitation of actions and in learning from testimony (for reviews, see [59,60]).

Moreover, children do all this learning in spite of the fact that they lack many of the kinds of intelligence that are characteristic of adults. Young children are noisy, variable, unfocused, unpredictable and impulsive. They lack ‘executive function’ abilities and capacities for focused and directed attention and are notoriously bad at long-term planning and deferred gratification [61,62]. In fact, almost by definition, young children cannot even take care of themselves (this is what immaturity means).

In addition to being effective learners in general, human children are especially good at exploratory ‘active learning’ that involves the pursuit of information about the world. A new wave of empirical research shows that toddlers and preschoolers explore the world and collect new information in spontaneous, systematic and rational ways ([63–68]; see

[69] for a review). Even infants systematically explore surprising events [70]. Moreover, children learn from their exploration. Bonawitz *et al.* [67] found that children revised their beliefs based on evidence they generated during exploratory play and Sim & Xu [71] found that even toddlers used the evidence they generated in exploration to formulate new abstract generalizations.

Play is intrinsically on the explore side of the explore–exploit trade-off. By definition, it involves activities that are not designed to accomplish particular goals [72]. The kind of spontaneous active learning described in the studies above is itself a form of play, but other forms of play also have wider exploratory functions. We have argued and presented data that pretend play, in particular, has a close empirical and theoretical connection to counterfactual thinking which is related to the kinds of ‘internal exploration’ that are involved in sampling and hypothesis search [11]. Schulz [69] has suggested that play might also involve a kind of ‘competence exploration’ helping children to learn what kinds of plans might solve a given class of problems.

Burghardt [72], who has extensively studied play behaviour in a wide range of animals, has argued that play can be defined as an activity that requires surplus resources but allows learning and adaptation. This fits well with the computational exploration/exploitation ideas. Childhood is a period in which extra resources are provided to an individual, and the constraints of effective and exploitative action are lifted, in order to allow learning and adaptation to the particular environment. (As Burghardt points out, sleep is another interesting example, as are the social division of labour examples described above.) Across a strikingly wide range of species, play is especially characteristic of young animals.

5. Evidence for an explore–exploit developmental shift

So, human children are highly effective, motivated, active and playful learners, in spite of their weaknesses in executive function. Is there evidence that they are actually more exploratory learners than adults and that these early learning abilities might actually trade off against later exploit abilities?

Comparative evidence suggests that learning, exploration and play are particularly characteristic of juveniles. Younger mice learn to reverse a learned rule more easily than post-pubertal mice [73]. Older monkeys show neural plasticity when they learn an auditory or tactile pattern, but only when the pattern is relevant to their goals—juveniles extract the patterns and demonstrate plasticity independently of goals [74]. Young capuchin monkeys are more prone to create novel behaviours relative to adults [75]. Juvenile rodents are more likely to explore aversive but informative options than adults [76]—in a classic avoidance learning task, young rats actually approach the cue that leads to a shock, preferring it to an uninformative cue, just the opposite of the adult pattern. They are especially likely to do this when the mother is present, a finding recently replicated in human preschoolers [77]. This finding also emphasizes the importance of the other half of the human life-history strategy—the expanded carer investment that allows exploration to take place.

Beyond humans, primates and rodents, there are good theoretical reasons to predict that older individuals will be less exploratory and/or neophilic than younger individuals

[78], and the empirical literature, for the most part, seems to agree with this prediction (parrots and corvids: [79], geladas: [80], chimango caracara: [81–83]; great tits: [84,85], hyenas: [86]). In wild spotted hyenas, juveniles were less neophobic, more persistent and exhibited a greater diversity of initial exploratory behaviours, relative to adults, when they were presented with a puzzle box [86]. Aplin [87] found that juvenile great tits were more likely to reverse a learned rule than older individuals. Holzhaider *et al.* [88] found that New Caledonian crows used their exceptionally long fledgling period to explore the possibilities of the objects they would eventually use as tools.

Neuroscience evidence also supports the idea of an explore–exploit shift. Neuroscientists have investigated the origins of both the increased executive control and decreased plasticity that come with age. One set of developments involves synaptic changes. In the early period of development, many more new synaptic connections are made than in adulthood. With age, some of these neural connections are strengthened but others are pruned, transforming a more flexible, sensitive and plastic brain into a more effective and controlled one [89,90].

Increasing executive control in humans is also related to the development of prefrontal areas of the brain and their increasing control over other brain areas. However, neuroscientists have also argued that strong frontal control has costs for exploration and learning [91]. Interference with prefrontal control areas leads to a wider range of responses on a ‘divergent thinking’ task [92], and during learning, there is a characteristic release of frontal control [93]. All this suggests that there is a neural trade-off between increased plasticity and decreased executive function and cognitive control.

The clearest evidence for this hypothesis, however, would be evidence that human children actually perform better than adults on precisely the same tasks when those tasks involve broad search or exploration. This is a challenging empirical agenda. It is difficult to design controlled laboratory experiments to investigate capacities that are spontaneous and uncontrolled by their very nature. Experimental design almost always involves setting a goal for the participants and, by definition, exploit and executive function capacities will make participants better at achieving such goals. Indeed, adults almost always do better on laboratory tasks than children. However, there are interesting and informative exceptions to this pattern. I will outline three such studies from my laboratory in some detail and describe a number of others more briefly.

We have found that younger children show a pattern of broader hypothesis search than adults in their causal learning. Young children are more likely to infer an unusual causal hypothesis than older children and adults [46,47]. For example, the ‘blicket detector’ is a box that lights up and plays music when you put some blocks on it and not others (it is actually controlled remotely by the experimenter). The participant sees a particular pattern of activations and then must infer the causal structure of the machine, determine which blocks are ‘blickets’, and act to activate the machine themselves. Even very young children are surprisingly good at using statistical patterns to determine which particular blocks will activate the machine (see [56] for a review).

We wanted to see whether children could also infer more abstract and general causal features of the machine. Could they figure out whether the machine operated on the

principle that individual blocks made it go or on the principle that combinations of blocks were required? Both children and adults initially assumed that the ‘combination’ principle was less likely than the ‘individual’ principle. Then we showed participants that individual blocks made the machine go, or that only a combination of blocks made it go. Finally, we showed them an ambiguous pattern with new blocks that was compatible with either principle, and asked them to activate the machine. Children and adults were equally good at inferring the more likely ‘individual’ hypothesis, they put the right individual blocks on the machine when that fit the data. But, to our surprise, children were substantially better at inferring the unlikely ‘combination’ hypothesis than adults—they put the right combinations of blocks on the machine when that fit the data, and the adults did not. We then replicated this pattern with low-income children and adults in Peru and low-income children in Headstart programmes in the USA [94].

Similarly, in a social causal task, younger learners were better at accurately using evidence to infer an unlikely causal hypothesis [95]. Adults, particularly in Western cultures, assume that people act the way they do because of their individual personality traits, rather than because they are faced with a particular situation, regardless of the evidence. This is ‘the fundamental attribution error’ or ‘trait bias’ [96]. We showed participants different people acting in different situations. The actions were consistent with either a trait or a situation explanation. For example, Mary might approach a trampoline and a diving board, whereas Josie avoided them (supporting a trait explanation), or Mary and Josie might both approach the trampoline and avoid the diving board (supporting a situation explanation). Then we asked participants to explain the actions—did Mary play on the trampoline because she was brave or because the trampoline was safe? Although the task is quite different from the ‘blicket detector’ task, it also involves inferring a very general and abstract causal schema from the data. And we found a similar result. Preschoolers gave the explanation that was the best fit to the data. But 6-year-olds, like adults, gave trait explanations even when the data suggested a situation explanation.

In a further study [48], we extended the range of ages we studied across the whole period of childhood from preschool through school age and adolescence. Just as children may have cognitive abilities that are qualitatively different from those of adults, different developmental periods may be qualitatively different. Until about 5, children contribute very few resources but make their greatest advances in learning. Around 5–7, there is a qualitative change. Children in contemporary Western society begin school, and in other societies and historical periods, begin more active apprenticeship and involvement in adult work. Another qualitative change comes at adolescence, which is a period of renewed exploration, risk-taking and learning, particularly in the social sphere [97].

We found qualitative changes between the preschool and school-aged children, and between school age and adolescence, but continuity within the school-age period. Interestingly, the changes went in opposite directions for the physical and social causal tasks. There was a qualitative shift at adolescence; like adults, teenagers became much less likely to infer the unlikely physical hypotheses. However, adolescents were actually *more* likely to infer the unlikely social hypothesis than either school-aged children or adults and were as cognitively flexible as the preschoolers.

There is also increasing recent evidence that younger learners prefer to explore rather than exploit and that this enables wider learning. In a very recent study [98], we gave participants a reinforcement learning task in which they could sequentially choose whether to place blocks on a machine. We told them that some blocks would lead to rewards and others to costs but did not reveal what differentiated the two. The actual rule was two-dimensional (e.g. black striped blocks were costly but white striped blocks or black spotted blocks led to rewards). After one negative trial, adults quickly assumed the most obvious rule, that a single feature differentiated the blocks (e.g. all black blocks were costly), and so avoided all the blocks with that feature. But this meant that they never received evidence that showed that the actual rule was more complex and so failed to learn the correct rule. As in other studies, they fell into a ‘learning trap’ [99]. Preschoolers, in contrast, continued to try all the blocks on the machine, and so learned the rule correctly. There is other very recent similar evidence that children explore more than adults in reinforcement learning tasks and learn more as a result, even though they may incur costs to do so [100–102].

In this research, children appear to have the greatest advantage over adults when they must infer hypotheses that have an unusual abstract high-level structure. This makes sense from a computational perspective. High-level abstract schemas typically constrain lower-level hypotheses and shape learner’s interpretation of the data [103]. As learners grow older, these schemes are increasingly well confirmed and become harder to overturn, even though learners may make revisions at the lower level.

There is other evidence for relevant differences between younger and older learners. Younger learners are more able to learn new linguistic distinctions than older learners [104,105] and they are better at imagining new uses for a tool [106]. Younger children also remember information that is outside the focus of goal-directed attention better than adults and older children [107].

Humans are a cultural species, with variable and dynamic environments. For such a species, it may be particularly important for each new generation to quickly discover novel abstract features of the environment. Cultural learning risks a particular kind of local optima, the result of the intrinsic tension between imitation and innovation. Adults who are faced with a changing social or physical environment may only be willing to make small local changes to the representations they have already learned and that support their actions and plans. The young children of the next generation, in contrast, may be more willing to consider a variety of high-level schemes to explain the data they see, allowing them to eventually make broader and more accurate predictions. To take two contemporary examples, my own computer schema involves a keyboard interface, with exceptions for touch and voice, and my marriage schema involves heterosexual couples with the exceptional addition of gay marriage. By contrast, my preschool grandchildren assume that computers generally work by touch and voice and that gender is irrelevant to marriage and make different predictions and produce different actions as a result. This ability to discover novel high-level features of the environment at an early age, and so to shape further learning, could provide considerable advantages later on, in the exploit phase (see [48] for further discussion).

We still have not discovered the full extent and nature of children's exploratory behaviour and there may well be contexts where adults are able to explore more widely than children. However, many adult types of exploration, such as those involved in formal and informal science and technology, actually also require exploit abilities. In particular, there is evidence that both human and primate juveniles are less able to actually produce useful new tools than adults [108,109], although younger children are better able to imagine new uses for an existing artefact than older ones [106]. This task and other classic adult 'insight' tasks require both explore and exploit abilities. Actually creating an effective new tool requires both the ability to imagine alternatives to the current options and the ability to select which of those options is most likely to be effective, as well as to realize that option. These are classic exploit abilities that characteristically require executive function.

'Over-imitation' is another example of a behaviour that appears to suggest that children explore more narrowly than adults. When children imitate complex tool-use behaviours, they often include even unnecessary details, initially suggesting that they only consider a narrow range of options for action [110]. More recent studies, however, have shown that over-imitation actually reflects more sophisticated inferences about the physical and social world. Over-imitation varies depending on how much children know about the physical causal relations in the task [111], and on whether the demonstrator's action is accidental, intentional or pedagogical [112] and it varies depending on the social context, for example, children are more likely to over-imitate when the demonstrators are present and when actions are arbitrary and might be interpreted as rituals [113,114] (see review in [60]). Over-imitation does not appear to emerge because children have an automatic impulse to narrowly reproduce the actions of others. Instead, it appears to stem from the fact that children believe that the demonstrator is a knowledgeable expert from their own group attempting to teach them about a tool or instructing them in a social ritual. Insofar as children learn from a range of demonstrators

and combine that information with their own experiences and inferences, imitation may actually help them consider a broader rather than narrower range of possibilities, though further studies are necessary to determine if this is true.

6. Conclusion

The explore/exploit trade-off helps to make sense of the general and widespread relationship between life history and learning. Across many different kinds of organisms, the length of immaturity and the complexity and flexibility of behaviour and cognition are correlated. The idea is particularly applicable, however, to human evolution, which involved particularly dramatic changes on both these dimensions. To return to the introduction, it is highly plausible that human evolution involved coevolutionary cascades between multiple changes in life history, behaviour and cognition. For example, a single change in the unfolding of the entire lifespan could have simultaneously allowed a long period of childhood exploration, larger brain size, more possibilities for cultural learning and a wider range of carers, particularly elders such as grandmothers, all interacting in a positive evolutionary feedback loop. Language, in turn, may have magnified all these changes. But the exploratory character of childhood cognition may have played an important role. Certainly, more empirical evolutionary and computational investigations of the links between childhood, learning, exploration and play would be illuminating.

Data accessibility. This article has no additional data.

Competing interests. I declare I have no competing interests.

Funding. This paper was supported by grants from NSF, the Bezos Foundation and the Templeton Foundation.

Acknowledgements. I am grateful to Eliza Kosoy for help in the preparation of the manuscript, to Susan Perry for invaluable comparative advice and information, to Tom Griffiths and Chris Lucas for computational insights and to the participants and sponsors of the Life History and Learning Workshop at Berkeley.

References

- Kaplan H, Hill K, Lancaster J, Hurtado AM. 2000 A theory of human life history evolution: diet, intelligence, and longevity. *Evol. Anthropol.* **9**, 156–185. (doi:10.1002/1520-6505(2000)9:4<156::AID-EVAN5>3.0.CO;2-7)
- Hublin J, Neubauer S, Gunz P. 2015 Brain ontogeny and life history in *Pleistocene hominins*. *Phil. Trans. R. Soc. B* **370**, 20140062. (doi:10.1098/rstb.2014.0062)
- Smith TM *et al.* 2010 Dental evidence for ontogenetic differences between modern humans and Neanderthals. *Proc. Natl Acad. Sci. USA* **107**, 20 923–20 928. (doi:10.1073/pnas.1010906107)
- Kuzawa C *et al.* 2014 Metabolic costs of human brain development. *Proc. Natl Acad. Sci. USA* **111**, 13 010–13 015. (doi:10.1073/pnas.1323099111)
- Gurven M, Walker R. 2005 Energetic demand of multiple dependents and the evolution of slow human growth. *Proc. R. Soc. B* **273**, 835–841. (doi:10.1098/rspb.2005.3380)
- Chapais B. 2009 *Primeval kinship: how pair-bonding gave birth to human society*. Cambridge, MA: Harvard University Press.
- Hrdy SB. 2011 *Mothers and others*. Cambridge, MA: Harvard University Press.
- Hawkes K, O'Connell JF, Jones NB, Alvarez H, Charnov EL. 1998 Grandmothering, menopause, and the evolution of human life histories. *Proc. Natl Acad. Sci. USA* **95**, 1336–1339. (doi:10.1073/pnas.95.3.1336)
- Herculano-Houzel S. 2012 The remarkable, yet not extraordinary, human brain as a scaled-up primate brain and its associated cost. *Proc. Natl Acad. Sci. USA* **109**(Suppl. 1), 10 661–10 668. (doi:10.1073/pnas.1201895109)
- Byrne RW. 1997 The technical intelligence hypothesis: an alternative evolutionary stimulus to intelligence. In *Machiavellian intelligence II: extensions and evaluations* (eds A Whiten, R Byrne), 289–311. Cambridge, UK: Cambridge University Press.
- Buchsbaum DB, Weisberg S, Gopnik AD. 2012 The power of possibility: causal learning, counterfactual reasoning, and pretend play. *Phil. Trans. R. Soc. B* **367**, 2202–2212. (doi:10.1098/rstb.2012.0122)
- Rosati AG. 2017 Foraging cognition: reviving the ecological intelligence hypothesis. *Trends Cogn. Sci.* **21**, 691–702. (doi:10.1016/j.tics.2017.05.011)
- Gopnik A, Wellman HM. 1994 The theory theory. In *Mapping the mind: domain specificity in cognition and culture* (eds L Hirschfeld, S Gelman), p. 257. Cambridge, UK: Cambridge University Press.
- Byrne RW, Whiten A. 1990 *Machiavellian intelligence: social expertise and the evolution of intellect in monkeys, apes, and humans*. Oxford, UK: Oxford University Press.

15. Tomasello M. 1999 *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
16. Boyd R, Richerson PJ. 1988 *Culture and the evolutionary process*. Chicago, IL: University of Chicago Press.
17. Henrich J. 2015 *The secret of our success*. Princeton, NJ: Princeton University Press.
18. Hunt G, Gray R. 2003 Diversification and cumulative evolution in New Caledonian crow tool manufacture. *Proc. R. Soc. Lond. B* **270**, 867–874. (doi:10.1098/rspb.2002.2302)
19. Kacelnik A. 2009 Tools for thought or thoughts for tools? *Proc. Natl Acad. Sci. USA* **106**, 10 071–10 072. (doi:10.1073/pnas.0904735106)
20. Hare B. 2007 From nonhuman to human mind: what changed and why? *Curr. Dir. Psychol. Sci.* **16**, 60–64. (doi:10.1111/j.1467-8721.2007.00476.x)
21. Call J, Tomasello M. 2008 Does the chimpanzee have a theory of mind? 30 years later. *Trends Cogn. Sci.* **12**, 187–192. (doi:10.1016/j.tics.2008.02.010)
22. Krupenye C, Kano F, Hirata S, Call J, Tomasello M. 2016 Great apes anticipate that other individuals will act according to false beliefs. *Science* **354**, 110–114. (doi:10.1126/science.aaf8110)
23. Rendell L, Whitehead H. 2001 Culture in whales and dolphins. *Behav. Brain Sci.* **24**, 309–324. (doi:10.1017/S0140525X0100396X)
24. Aplin LM, Farine DR, Morand-Ferron J, Cockburn A, Thornton A, Sheldon BC. 2015 Experimentally induced innovations lead to persistent culture via conformity in wild birds. *Nature* **518**, 538–541. (doi:10.1038/nature13998)
25. Bruner JS. 1972 Nature and uses of immaturity. *Am. Psychol.* **27**, 687–708. (doi:10.1037/h0033144)
26. Bjorklund DF, Green BL. 1992 The adaptive nature of cognitive immaturity. *Am. Psychol.* **47**, 46. (doi:10.1037/0003-066X.47.1.46)
27. Finlay BL, Uchiyama R. 2015 Developmental mechanisms channeling cortical evolution. *Trends Neurosci.* **38**, 69–76. (doi:10.1016/j.tins.2014.11.004)
28. Bogin B. 1997 Evolutionary hypotheses for human childhood. *Am. J. Phys. Anthropol.* **104**, 63–89. (doi:10.1002/(SICI)1096-8644(1997)25+<63::AID-AJPA3>3.0.CO;2-8)
29. Hrdy SB, Burkart JM. 2020 The emergence of emotionally modern humans: implications for language and learning. *Phil. Trans. R. Soc. B* **375**, 20190499. (doi:10.1098/rstb.2019.0499)
30. Hawkes K. 2020 Cognitive consequences of our grandmothering life history: cultural learning begins in infancy. *Phil. Trans. R. Soc. B* **375**, 20190501. (doi:10.1098/rstb.2019.0501)
31. Tomasello M. 2020 The adaptive origins of uniquely human sociality. *Phil. Trans. R. Soc. B* **375**, 20190493. (doi:10.1098/rstb.2019.0493)
32. Sterelny K. 2012 *The evolved apprentice*. Cambridge, MA: MIT Press.
33. Bennett PM, Harvey PH. 1985 Brain size, development and metabolism in birds and mammals. *J. Zool.* **207**, 491–509. (doi:10.1111/j.1469-7998.1985.tb04946.x)
34. Weisbecker V, Goswami A. 2010 Brain size, life history, and metabolism at the marsupial/placental dichotomy. *Proc. Natl Acad. Sci. USA* **107**, 16 216–16 221. (doi:10.1073/pnas.0906486107)
35. Snell-Rood EC, Davidowitz G, Papaj DR. 2011 Reproductive tradeoffs of learning in a butterfly. *Behav. Ecol.* **22**, 291–302. (doi:10.1093/beheco/arq169)
36. Snell-Rood EC. 2013 An overview of the evolutionary causes and consequences of behavioural plasticity. *Anim. Behav.* **85**, 1004–1011. (doi:10.1016/j.anbehav.2012.12.031)
37. Sutton R, Barto A. 1998 *Introduction to reinforcement learning*. Cambridge, MA: MIT Press.
38. Schulz E, Gershman SJ. 2019 The algorithmic architecture of exploration in the human brain. *Curr. Opin Neurobiol.* **55**, 7–14. (doi:10.1016/j.conb.2018.11.003)
39. Wilson RC, Geana A, White JM, Ludvig EA, Cohen JD. 2014 Humans use directed and random exploration to solve the explore–exploit dilemma. *J. Exp. Psychol. Gen.* **143**, 2074. (doi:10.1037/a0038199)
40. Kirpatrick S, Gelatt M, Vecchi M. 1983 Optimization by simulated annealing. *Science* **220**, 671–680. (doi:10.1126/science.220.4598.671)
41. Roberts C, Casella G. 1999 *Monte Carlo statistical methods*. Berlin, Germany: Springer.
42. Denison S, Bonawitz E, Gopnik A, Griffiths T. 2013 Rational variability in children’s causal inferences: the Sampling Hypothesis. *Cognition* **126**, 2, 285–300. (doi:10.1016/j.cognition.2012.10.010)
43. Bonawitz E, Denison S, Gopnik A, Griffiths TL. 2014 Win-stay, lose-sample: a simple sequential algorithm for approximating Bayesian inference. *Cognit. Psychol.* **74**, 35–65. (doi:10.1016/j.cogpsych.2014.06.003)
44. Bonawitz E, Denison S, Griffiths TL, Gopnik A. 2014 Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends Cogn. Sci.* **18**, 497–500. (doi:10.1016/j.tics.2014.06.006)
45. Davis L. 1991 *Handbook of genetic algorithms*. New York, NY: Van Nostrand Reinhold.
46. Lucas CG, Bridgers S, Griffiths TL, 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**, 284–299. (doi:10.1016/j.cognition.2013.12.010)
47. Gopnik A, Griffiths TL, Lucas CG. 2015 When younger learners can be better (or at least more open-minded) than older ones. *Curr. Dir. Psychol. Sci.* **24**, 87–92. (doi:10.1177/0963721414556653)
48. Gopnik A, O’Grady S, Lucas CG, Griffiths TL, Wente A, Bridgers S, Aboody R, Fung H, Dahl RE. 2017 Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proc. Natl Acad. Sci. USA* **114**, 7892–7899. (doi:10.1073/pnas.1700811114)
49. Spelke ES. 2000 Core knowledge. *Am. Psychol.* **55**, 1233. (doi:10.1037/0003-066X.55.11.1233)
50. Wood JN, Prasad A, Goldman JG, Wood SM. 2016 Enhanced learning of natural visual sequences in newborn chicks. *Anim. Cogn.* **19**, 835–845. (doi:10.1007/s10071-016-0982-5)
51. Fawcett TW, Frankenhuus WE. 2015 Adaptive explanations for sensitive windows in development. *Front. Zool.* **12**(S1), 53. (doi:10.1186/1742-9994-12-S1-53)
52. Frankenhuus W, Panchanathan K. 2011 Balancing sampling and specialization: an adaptationist model of incremental development. *Proc. R. Soc. B* **278**, 3558–3565. (doi:10.1098/rspb.2011.0055)
53. Potts R. 1996 Evolution and climate variability. *Science* **273**, 922–923. (doi:10.1126/science.273.5277.922)
54. Gelman SA, Wellman HM. 1991 Insides and essences: early understandings of the non-obvious. *Cognition* **38**, 213–244. (doi:10.1016/0010-0277(91)90007-Q)
55. Carey S. 2009 *The origin of concepts*. Oxford, UK: Oxford University Press.
56. Gopnik A, Wellman HM. 2012 Reconstructing constructivism: causal models, Bayesian learning mechanisms, and the theory theory. *Psychol. Bull.* **138**, 1085. (doi:10.1037/a0028044)
57. Gopnik A, Glymour C, Sobel D, Schulz L, Kushnir T, Danks D. 2004 A theory of causal learning in children: causal maps and Bayes nets. *Psychol. Rev.* **111**, 1–31. (doi:10.1037/0033-295X.111.1.3)
58. Kushnir T, Xu F (eds). 2012 *Advances in child development and behavior: rational constructivism in cognitive development*. Waltham, MA: Elsevier.
59. Harris P. 2012 *Trusting what you’re told*. Cambridge, MA: Harvard University Press.
60. Gopnik A. 2016 *The gardener and the carpenter: what the new science of child development tells us about the relationship between parents and children*. New York, NY: Macmillan.
61. Munakata Y, Casey BJ, Diamond A. 2004 Developmental cognitive neuroscience: progress and potential. *Trends Cogn. Sci.* **8**, 122–128. (doi:10.1016/j.tics.2004.01.005)
62. Carlson SM. 2005 Developmentally sensitive measures of executive function in preschool children. *Dev. Neuropsychol.* **28**, 595–616. (doi:10.1207/s15326942dn2802_3)
63. Schulz LE, Gopnik A, Glymour C. 2007 Preschool children learn about causal structure from conditional interventions. *Dev. Sci.* **10**, 322–332. (doi:10.1111/j.1467-7687.2007.00587.x)
64. Schulz LE, Bonawitz EB. 2007 Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Dev. Psychol.* **43**, 1045. (doi:10.1037/0012-1649.43.4.1045)
65. Cook C, Goodman ND, Schulz LE. 2011 Where science starts: spontaneous experiments in preschoolers’ exploratory play. *Cognition* **120**, 341–349. (doi:10.1016/j.cognition.2011.03.003)
66. Legare CH. 2012 Exploring explanation: explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Dev.* **83**, 173–185. (doi:10.1111/j.1467-8624.2011.01691.x)

67. Bonawitz EB, van Schijndel TJ, Friel D, Schulz L. 2012 Children balance theories and evidence in exploration, explanation, and learning. *Cognit. Psychol.* **64**, 215–234. (doi:10.1016/j.cogpsych.2011.12.002)
68. Ruggeri A, Swaboda N, Sim ZL, Gopnik A. 2019 Shake it baby, but only when needed: preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition* **193**, 104013. (doi:10.1016/j.cognition.2019.104013)
69. Schulz L. 2012 The origins of inquiry: inductive inference and exploration in early childhood. *Trends Cogn. Sci.* **16**, 382–389. (doi:10.1016/j.tics.2012.06.004)
70. Stahl AE, Feigenson L. 2015 Observing the unexpected enhances infants' learning and exploration. *Science* **348**, 91–94. (doi:10.1126/science.aaa3799)
71. Sim ZL, Xu F. 2017 Learning higher-order generalizations through free play: evidence from 2- and 3-year-old children. *Dev. Psychol.* **53**, 642. (doi:10.1037/dev0000278)
72. Burghardt GM. 2005 *The genesis of animal play: testing the limits*. Cambridge, MA: MIT Press.
73. Johnson C, Wilbrecht L. 2011 Juvenile mice show greater flexibility in multiple choice reversal learning than adults. *Dev. Cogn. Neurosci.* **1**, 540–551. (doi:10.1016/j.dcn.2011.05.008)
74. Buonomano DV, Merzenich MM. 1998 Cortical plasticity: from synapses to maps. *Annu. Rev. Neurosci.* **21**, 149–186. (doi:10.1146/annurev.neuro.21.1.149)
75. Perry SE, Barrett BJ, Godoy I. 2017 Older, sociable capuchins (*Cebus capucinus*) invent more social behaviors, but younger monkeys innovate more in other contexts. *Proc. Natl Acad. Sci. USA* **114**, 7806–7813. (doi:10.1073/pnas.1620739114)
76. Moriceau S, Sullivan RM. 2006 Maternal presence serves as a switch between learning fear and attraction in infancy. *Nat. Neurosci.* **9**, 1004–1006. (doi:10.1038/nn1733)
77. Tottenham N, Shapiro M, Flannery J, Caldera C, Sullivan RM. 2019 Parental presence switches avoidance to attraction learning in children. *Nat. Hum. Behav.* **3**, 1070–1077. (doi:10.1038/s41562-019-0656-9)
78. Sherratt TN, Morand-Ferron J. 2018 The adaptive significance of age-dependent changes in the tendency of individuals to explore. *Anim. Behav.* **138**, 59–67. (doi:10.1016/j.anbehav.2018.01.025)
79. O'Hara M, Auersperg AM. 2017 Object play in parrots and corvids. *Curr. Opin. Behav. Sci.* **16**, 119–125. (doi:10.1016/j.cobeha.2017.05.008)
80. Bergman TJ, Kitchen DM. 2009 Comparing responses to novel objects in wild baboons (*Papio ursinus*) and geladas (*Theropithecus gelada*). *Anim. Cogn.* **12**, 63. (doi:10.1007/s10071-008-0171-2)
81. Biondi LM, Bó MS, Vassallo AI. 2010 Inter-individual and age differences in exploration, neophobia and problem-solving ability in a Neotropical raptor (*Milvago chimango*). *Anim. Cogn.* **13**, 701–710. (doi:10.1007/s10071-010-0319-8)
82. Biondi LM, Guido J, Madrid E, Bó MS, Vassallo AI. 2013 The effect of age and sex on object exploration and manipulative behavior in a Neotropical raptor, the Chimango Caracara, *Milvago chimango*. *Ethology* **119**, 221–232. (doi:10.1111/eth.12056)
83. Biondi LM, Guido JM, Bó MS, Muzio RN, Vassallo AI. 2015 The role of stimulus complexity, age and experience in the expression of exploratory behaviour in the Chimango Caracara, *Milvago chimango*. *Anim. Cogn.* **18**, 139–150. (doi:10.1007/s10071-014-0785-5)
84. Morand-Ferron J, Cole EF, Rawles JE, Quinn JL. 2011 Who are the innovators? A field experiment with 2 passerine species. *Behav. Ecol.* **22**, 1241–1248. (doi:10.1093/beheco/arr120)
85. Morand-Ferron J, Hamblin S, Cole EF, Aplin LM, Quinn JL. 2015 Taking the operant paradigm into the field: associative learning in wild great tits. *PLoS ONE* **10**, e0133821. (doi:10.1371/journal.pone.0133821)
86. Benson-Amram S, Holekamp KE. 2012 Innovative problem solving by wild spotted hyenas. *Proc. R. Soc. B* **279**, 4087–4095. (doi:10.1098/rspb.2012.1450)
87. Aplin LM. 2019 Culture and cultural evolution in birds: a review of the evidence. *Anim. Behav.* **147**, 179–187. (doi:10.1016/j.anbehav.2018.05.001)
88. Holzhaider JC, Hunt GR, Gray RD. 2010 Social learning in new Caledonian crows. *Learn. Behav.* **38**, 206–211. (doi:10.3758/LB.38.3.206)
89. Huttenlocher PR. 1990 Morphometric study of human cerebral cortex development. *Neuropsychologia*. **28**, 517–527. (doi:10.1016/0028-3932(90)90031-1)
90. Huttenlocher PR. 2009 *Neural plasticity*. Cambridge, MA: Harvard University Press.
91. Thompson-Schill SL, Ramscar M, Chrysikou EG. 2009 Cognition without control: when a little frontal lobe goes a long way. *Curr. Dir. Psychol. Sci.* **18**, 259–263. (doi:10.1111/j.1467-8721.2009.01648.x)
92. Chrysikou EG, Hamilton RH, Coslett HB, Datta A, Bikson M, Thompson-Schill SL. 2013 Noninvasive transcranial direct current stimulation over the left prefrontal cortex facilitates cognitive flexibility in tool use. *Cogn. Neurosci.* **4**, 81–89. (doi:10.1080/17588928.2013.768221)
93. Bassett DS, Wymbs NF, Porter MA, Mucha PJ, Carlson JM, Grafton ST. 2011 Dynamic reconfiguration of human brain networks during learning. *Proc. Natl Acad. Sci. USA* **108**, 7641–7646. (doi:10.1073/pnas.1018985108)
94. Wente A, Kimura K, Walker CM, Banerjee N, Flecha M, MacDonald B, Lucas C, Gopnik A. 2019 Causal learning across culture and SES. *Child Dev.* **90**, 859–875. (doi:10.1111/cdev.12943)
95. Seiver E, Gopnik A, Goodman ND. 2013 Did she jump because she was the big sister or because the trampoline was safe? Causal inference and the development of social attribution. *Child Dev.* **84**, 443–454. (doi:10.1111/j.1467-8624.2012.01865.x)
96. Kelley HH. 1973 The processes of causal attribution. *Am. Psychol.* **28**, 107–128. (doi:10.1037/h0034225)
97. Crone EA, Dahl RE. 2012 Understanding adolescence as a period of social–affective engagement and goal flexibility. *Nat. Rev. Neurosci.* **13**, 636–650. (doi:10.1038/nrn3313)
98. Liquin E, Gopnik A. 2019 Learning from approach-avoid decisions: children explore and learn more than adults. In *Paper presented at the Society for Research in Child Development, Baltimore, MD*.
99. Rich AS, Gureckis TM. 2018 The limits of learning: exploration, generalization, and the development of learning traps. *J. Exp. Psychol. Gen.* **147**, 1553–1570. (doi:10.1037/xge0000466)
100. Schulz E, Wu C, Ruggieri A, Meder B. 2019 Searching for rewards like a child means less generalization and more directed exploration. *Psychol. Sci.* **30**, 1561–1572. (doi:10.1177/0956797619863663)
101. Sumner E, Li A, Perfors A, Hayes B, Navarro D, Sarnecka B. 2019 The exploration advantage: children's instinct to explore allows them to find information that adults miss. *PsyArxiv*. (doi:10.31234/osf.io/h437v)
102. Nussenbaum K, Hartley CA. 2019 Reinforcement learning across development: what insights can we draw from a decade of research? *Dev. Cogn. Neurosci.* **40**, 100733. (DOI:10.1016/j.dcn.2019.100733)
103. Kemp C, Tenenbaum JB, Niyogi S, Griffiths TL. 2010 A probabilistic model of theory formation. *Cognition* **114**, 165–196. (doi:10.1016/j.cognition.2009.09.003)
104. Werker JF, Hensch TK. 2015 Critical periods in speech perception: new directions. *Annu. Rev. Psychol.* **66**, 173–196. (doi:10.1146/annurev-psych-010814-015104)
105. Kuhl PK. 2004 Early language acquisition: cracking the speech code. *Nat. Rev. Neurosci.* **5**, 831–843. (doi:10.1038/nrn1533)
106. German TP, Defeyter MA. 2000 Immunity to functional fixedness in young children. *Psychon. Bull. Rev.* **7**, 707–712. (doi:10.3758/bf03213010)
107. Plebanek DJ, Sloutsky VM. 2017 Costs of selective attention: when children notice what adults miss. *Psychol. Sci.* **28**, 723–732. (doi:10.1177/0956797617693005)
108. Kendal RL, Coe RL, Laland KN. 2005 Age differences in neophilia, exploration, and innovation in family groups of callitrichid monkeys. *Am. J. Primatol.* **66**, 167–188. (doi:10.1002/ajp.20136)
109. Neldner K, Redshaw J, Murphy S, Tomaselli K, Davis J, Dixon B, Nielsen M. 2019 Creation across culture: children's tool innovation is influenced by cultural and developmental factors. *Dev. Psychol.* **55**, 877–889. (doi:10.1037/dev0000672)
110. Horner V, Whiten A. 2005 Causal knowledge and imitation/emulation switching in chimpanzees (*Pan troglodytes*) and children (*Homo sapiens*). *Anim. Cogn.* **8**, 164–181. (doi:10.1007/s10071-004-0239-6)

111. Williamson RA, Meltzoff AN. 2011 Own and others' prior experiences influence children's imitation of causal acts. *Cognit. Dev.* **26**, 260–268. (doi:10.1016/j.cogdev.2011.04.002)
112. Buchsbaum D, Gopnik A, Griffiths TL, Shafto P. 2011 Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition* **120**, 331–340. (doi:10.1016/j.cognition.2010.12.001)
113. Nielsen M, Blank C. 2011 Imitation in young children: when who gets copied is more important than what gets copied. *Dev. Psychol.* **47**, 1050–1053. (doi:10.1037/a0023866)
114. Legare CH, Wen NJ, Herrmann PA, Whitehouse H. 2015 Imitative flexibility and the development of cultural learning. *Cognition* **142**, 351–361. (doi:10.1016/j.cognition.2015.05.020)