nature reviews psychology https://doi.org/10.1038/s44159-024-00300-5

Review article

Check for updates

The development of human causal learning and reasoning

Mariel K. Goddu [1,2,3] & Alison Gopnik [4,5]

Abstract

Causal understanding is a defining characteristic of human cognition. Like many animals, human children learn to control their bodily movements and act effectively in the environment. Like a smaller subset of animals, children intervene: they learn to change the environment in targeted ways. Unlike other animals, children grow into adults with the causal reasoning skills to develop abstract theories, invent sophisticated technologies and imagine alternate pasts, distant futures and fictional worlds. In this Review, we explore the development of human-unique causal learning and reasoning from evolutionary and ontogenetic perspectives. We frame our discussion using an 'interventionist' approach. First, we situate causal understanding in relation to cognitive abilities shared with non-human animals. We argue that human causal understanding is distinguished by its depersonalized (objective) and decontextualized (general) representations. Using this framework, we next review empirical findings on early human causal learning and reasoning and consider the naturalistic contexts that support its development. Then we explore connections to related abilities. We conclude with suggestions for ongoing collaboration between developmental, cross-cultural, computational, neural and evolutionary approaches to causal understanding.

Sections

[1]Department of Philosophy, Stanford University, Stanford, CA, USA. [2]Institut für Philosophie, Freie Universität Berlin, Berlin, Germany. [3]Centre for Advanced Study in the Humanities: Human Abilities, Berlin, Germany. [4]Department of Psychology, University of California, Berkeley, Berkeley, CA, USA. [5]Department of Philosophy Affiliate, University of California, Berkeley, Berkeley, CA, USA. e-mail: goddu@stanford.edu; gopnik@berkeley.edu

Introduction

The flexibility and generativity of humans' causal understanding is unparalleled. Humans can develop abstract theories and invent sophisticated technologies, imagine alternate pasts and distant futures, and engage with fictional worlds. Chimpanzees, who share 98.7% of human genetic material[1], do not do any of these things. Human-unique causal reasoning is as defining a characteristic of the species as language.

To ask questions such as "What makes human causal reasoning so different from other species?", or "How and why do these skills develop?" is itself an exercise in causal reasoning. These are questions that researchers in developmental psychology, behavioural ecology, computer science and machine learning all try to answer[2–8]. Like many queries in science and in life, they are 'Why?' questions. A satisfactory answer will not consist of a mere list of facts about what happens – for example, a series of changes between neonatal and adult cognition. Rather, an answer will consist of causal explanations: proposals about the underlying reasons for those changes[9–11].

Understanding the development of human causal learning and reasoning is important for human society. For better or worse, humans change their physical and social environments more than any other species. For example, sophisticated causal reasoning has led to the technologies that produce climate change. But imagining ways to reverse the damage is also a form of causal reasoning. A better understanding of human causal learning and reasoning abilities is important not only for engineering and science education, but also for civic engagement and critical literacy.

There are several approaches to causal understanding in the literature. Some researchers treat it as a simple matter of association[12]. Others propose that it is an elaboration of perceptual skills that enable focus on particular spatiotemporal sequences of events[13–16]. In this paper we take an 'interventionist' perspective on human causal learning and reasoning[2–4,8,9,11]. This perspective defines causal relations as relations between causal variables. Causal variables are discrete factors with values that can change[2,4,17]. For example, wondering what to do to improve your sleep involves identifying causal variables that might influence sleep quality. To posit a causal relation is to posit that changing the value of one variable (the cause) makes a difference to the value of another (the effect). Identifying a causal relation enables you to make predictions about what will happen if specific changes are made ("IF I have less screen time at night, THEN my sleep will improve") and to identify possible causal explanations ("Maybe my sleep improved BECAUSE I stopped watching TV before bed")[2,4]. Identifying causal relations also enables you to design interventions – specific changes made to the values of causal variables[4,9] – to bring about a target outcome ("I will stop using my phone at night IN ORDER TO improve my sleep").

The interventionist view emphasizes that 'difference-making' is the essence of causal knowledge. By isolating a specific variable in a causal network, manipulating that variable, and observing which values of other variables change, you can infer which variables 'make a difference' to others, thus enabling you to imagine the possible outcomes of further causal events. From this perspective, causal reasoning is fundamentally related to counterfactual reasoning. Causal understanding lets you infer what could have happened in the past, what could happen in the future, or what could happen in an alternative fictional reality.

On the interventionist view, mature causal understanding also involves the assumption that the effectiveness of interventions is rooted in objective facts[18]. For example, imagine trying to unlock your front door but finding that it is stuck[19]. Your causal understanding might lead you to posit a problem with either the key or the doorknob. Each candidate causal explanation corresponds to an intervention on a specific causal variable: for example, you might check that you are using the correct key, or you might jiggle the doorknob. These inferences are grounded in your assumption that there are objective conditions that govern the causal relation, independent of your actions[18,20]. That is, even if you do not have detailed knowledge of how locks work[21], you know that there is a particular way that things need to be to enable the door to open. Trying another key or jiggling the knob are relevant causal interventions because they impact this mechanism. Understanding the objective mechanism also enables you to consider other ways to make the desired change, such as inserting a strategically bent paperclip[18,20].

Adult human concepts of causal relations can involve many different kinds of variables and events, across many different contexts and spatiotemporal scales. Understanding the collisions of billiard balls, the 'push' and 'pull' of desires, the movement of the tides, and the cascade of neurotransmitters across a membrane all involve causal relations. There might not be any single objective feature that these relations share – causation might not be a 'natural kind'[18,22]. However, from a psychological perspective, the common feature of these relations is that they support real or theoretical interventions. For example, to say that the moon causes the tides means that an intervention that changed the position of the moon would result in a change in the tides, regardless of whether such an intervention is actually possible.

This notion of causal understanding might seem broad – and it is. But it is not all-encompassing. Understanding the causal dependence of the door's opening on the key is different from understanding the logical dependence between '2 + 2' and '4' or understanding that if something is a triangle, then its interior angles necessarily sum to 180°. Causal reasoning is different from syllogistic reasoning and other forms of logical reasoning. It is also different from reasoning about categorical relations (such as reasoning that someone called a grandmother has grandchildren) or part–whole relations (such as reasoning that a basketball player is part of a team).

In this Review, we examine empirical findings in the development of human-unique causal learning and reasoning through an interventionist frame. We argue that explaining adults' capacities involves identifying the sources of their decontextualized (general) and depersonalized (objective) way of thinking about causation. First, we review the development of causal learning and reasoning over both evolutionary and ontogenetic time. Against this background, we discuss the causal factors that might underlie the development of adult-like causal understanding. Then we explore how human causal learning and reasoning develop in naturalistic contexts and how features of children's everyday experiences shape their abilities. Next, we discuss the relation between causal reasoning and other cognitive abilities. We conclude by discussing open questions and avenues for future research.

Human-unique causal understanding

To explain the development of human-unique causal understanding, we must first identify its key features. Adult humans use causal reasoning effortlessly and automatically every day. Even the most mundane cognitive activities – such as wondering why someone is late for a meeting, wiggling paper trays on a malfunctioning photocopier, or contemplating the consequences of drinking coffee at this hour – involve causal understanding. That is, they involve thinking of causal relations in terms of variables with values that can change, and they exemplify the ability to mentally manipulate these variables without actually changing the world. Crucially, adult human causal understanding is decontextualized

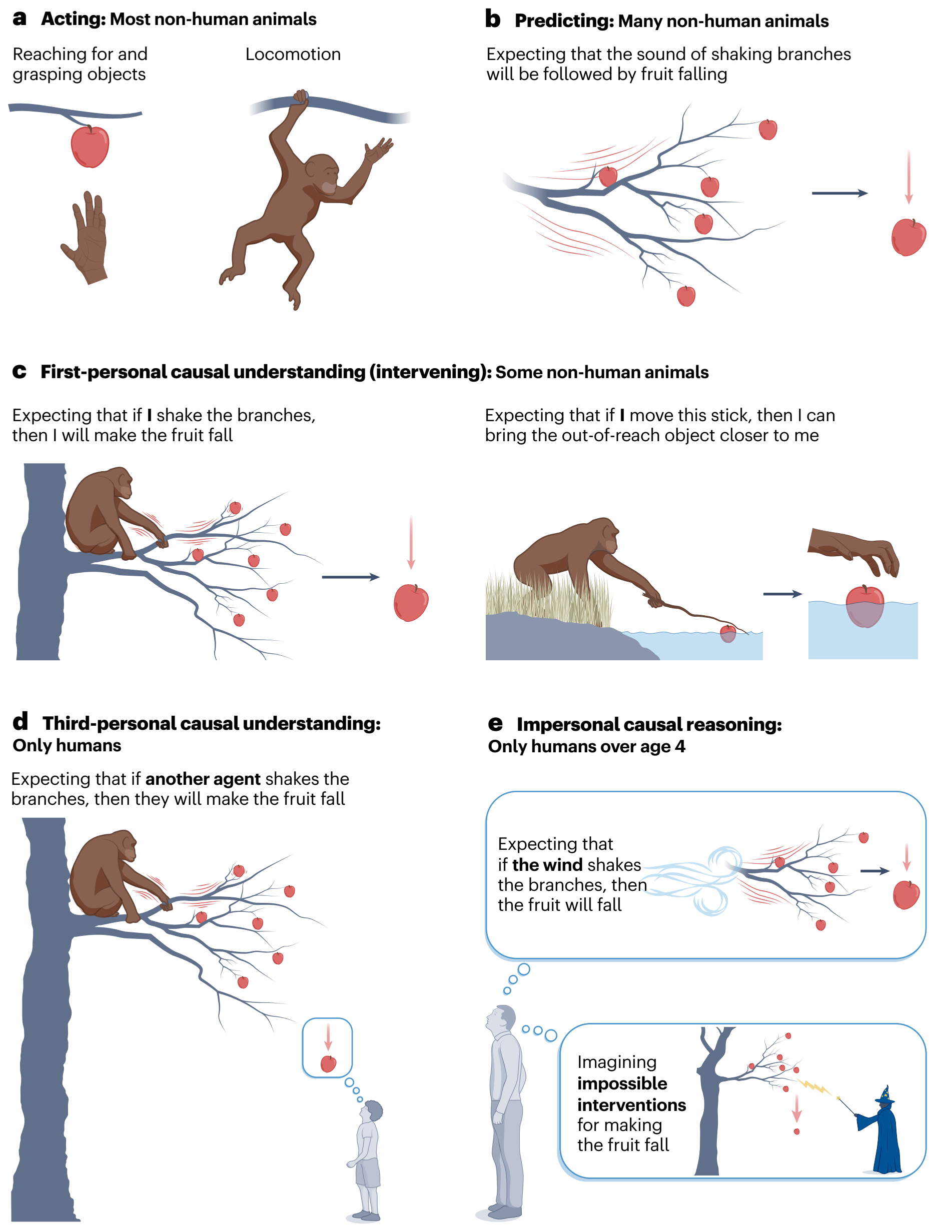

Fig. 1 | Causal reasoning and antecedent abilities in human and non-human animals. a, Acting effectively involves applying knowledge of the differences that one's bodily movements make to one's perceptions and goals. These skills involve interacting with the environment, but not manipulating it to generate an external effect. **b**, Predicting contingencies between external events, such as expecting a positive outcome after hearing a sound, does not involve causal reasoning. **c**, Predicting external events resulting from one's own interventions, such as learning to manipulate one object to change the position of another object, is first-personal causal understanding. **d**, Predicting external events resulting from others' interventions is third-personal causal understanding. **e**, Imagining specific causal events that are independent of agents' actions, and might even be impossible, is impersonal causal understanding. Adult human causal understanding is depersonalized and decontextualized (involves general representations of causal variables). The theoretical distinctions and examples in this figure are inspired by previous accounts[4,9,20,23,41].

and depersonalized: causal relations can be represented in a way that is highly general (not tied to specific contexts) and impersonal (not tied to agents' actions).

Other animals clearly have elements of these abilities (Fig. 1). However, they do not seem to possess a humanlike causal understanding. In this section, we first discuss two dimensions of causal understanding: how decontextualized and depersonalized it is. Then we review the learning and reasoning skills of non-human animals, highlighting the major differences between these skills and the features of adult human causal learning and reasoning.

Dimensions of causal understanding

We suggest that human adult representations of causal relations are decontextualized and depersonalized. These are two orthogonal dimensions on which causal learning and reasoning differ across phylogeny and ontogeny. Decontextualization is a dimension that ranges from knowledge about highly specific causal relations in particular settings to knowledge of more general relations and variables. Depersonalization increases in objectivity along a spectrum of first-personal to third-personal to impersonal.

Decontextualization concerns the specificity of the causal relations that are represented. These can range from very specific sensorimotor phenomena, such as links between internal motor commands and resulting perceptions, to very general relations between variables in scientific theories that apply to entire domains, such as all of physics. Causal understanding often seems to begin (ontogenetically and phylogenetically) by picking out particular, narrowly defined variables in ecologically relevant contexts. Reasoning develops to include a wider range of more general relations between more general variables. An animal who acts in intelligent ways to avoid a particular type of predator or to obtain a particular kind of food without generalizing that knowledge to other predators or types of food has a causal understanding that is restricted to particular contexts. Decontextualized reasoning, by contrast, is fully general and detached from particular events.

The second dimension, depersonalization, involves the relation between causal representations and actions. The interventionist perspective argues that what makes a relation causal is its relation to conceivable interventions. For human adults, causal relations are understood to be rooted in objective mechanisms in the environment, and the actions and interventions can be highly theoretical. Following earlier accounts, we propose that both phylogenetically and ontogenetically earlier forms of causal understanding are grounded in relations between actual actions and outcomes[9,20,23]. In particular, some early forms of causal reasoning, such as reinforcement learning, might only involve cases in which the agent themselves performs an action – a form of causal knowledge that we call 'first-personal' ('I [first person] cause'). Other forms, such as those involved in imitation learning, might include both the agent's own actions and others' actions, a 'third-personal' understanding ('They [third person] cause'). The final, most objective level of causal understanding enables a reasoner to infer causal relations simply by observing events, independent of agents' actions – an 'impersonal' understanding ('It causes'). Fully depersonalized reasoning is detached from agents' actions.

These two dimensions of causal knowledge, decontextualization and depersonalization, are orthogonal. For example, perceptual capacities that enable an agent to predict and track the movements of objects are impersonal, in that they do not rely on the agent's actions, but apply only in very particular contexts. By contrast, reinforcement learning might apply across many contexts involving many objects and relations – and so is relatively decontextualized – but applies only to outcomes of the agent's actions. Thus, it would be first-personal only, rather than impersonal. The literature suggests that phylogenetically and ontogenetically earlier forms of causal understanding tend to be more context-specific, less objective, or both, than later forms.

Sensorimotor learning and causal reasoning

Many evolutionarily ancient sensorimotor capacities can be construed as solutions to causal problems, yet they do not have the same features as adult human causal reasoning. Even very simple organisms (such as earthworms and fish) differentiate between sensory changes caused by themselves and those caused by the environment. Moreover, they can integrate information from different sensory modalities to pinpoint specific causes. For example, a visually perceived shadow and a tactually perceived change in water pressure might both be attributable to an entity passing overhead[24–26].

This kind of context-specific, first-personal causal knowledge also has a role in motor learning and locomotion[27,28]. For example, young humans and chimpanzees who learn to walk and climb must learn which environmental features are difference-making for the success of their actions. In walking, causally relevant variables might include the degree of incline of slopes and the width of gaps; for climbing, these might include the thickness and pliability of branches[29–31]. Similarly, many animals seem to be sensitive to regularities in the dynamics of physical object interactions in the environment – what has been called 'intuitive physics'[32] – in their evolved visual perception (Box 1) and proprioception.

Learning about events

Many non-human animals can also learn about a wide range of events that go beyond simple sensorimotor interactions with the environment. The two major mechanisms that accomplish this feat are associative learning and reinforcement learning. However, in the interventionist framework, only reinforcement learning is seen as a precursor to mature causal understanding.

In associative learning (also called classical or Pavlovian conditioning, or observational learning), two changes become linked such that the organism interprets one event as a cue to predict another[33,34]. For example, consider a dog who learns to anticipate their owner's arrival when they hear a key in the door. Here, an inherently rewarding stimulus (arrival of social attachment figure) elicits a response (emotional excitement). Through associative learning, the sound of the key, which is arbitrary, is correlated with the reward and therefore comes to elicit the response. Associative learning yields correlations and enables predictions, but it does not reflect causal knowledge. Knowing what to expect next – such as anticipating that the sound of the key will be followed by a creaking sound as the door opens – is not the same as understanding that the door's movement generates the creaking. Similarly, predicting that the next letter in the repeated sequence 'ABCABCABCA…' will be 'B' does not involve knowing why or how the sequence is being produced, or what would change the pattern.

From an interventionist perspective, reinforcement learning, unlike associative learning, is a form of causal understanding. In reinforcement learning (also called operant conditioning or instrumental learning), an animal learns to link their own voluntary actions with subsequent changes in the world[35]. However, this understanding is strictly first-personal. For example, a dog who has learned to bark to be let outside is making an intervention: they change the value of one variable (whether they are barking) to make a difference to the value of another (whether the door is open). In model-free reinforcement learning, the animal simply becomes more likely to perform an action when it is followed by a reward (as in canonical operant conditioning)[35,36]. By contrast, in model-based reinforcement learning, the animal learns a wider, more integrated variety of mappings between actions and outcomes, enabling them to flexibly adjust their behaviour according to the context. Model-free learning might be illustrated by a dog who indiscriminately barks at the door, even when a person is not present to open it; model-based learning might be illustrated by a dog who barks only when a person is present. Both model-free and model-based reinforcement learning occur in a range of animals from rats to humans[35,37].

Model-based learning in particular might prefigure more depersonalized forms of causal reasoning. However, reinforcement learning has several major limitations. First, although the interventions and outcomes and their relations can be varied and general, the causal relations that can be learned are often quite constrained. Reinforcement learning is motivated by utility: the effect is a reward or punishment. Typically, in animals, the effect must occur close in space and time to the cause[23]. Second, as noted, the interventions are first-personal: the causal relation is tied to the animal's own goal-directed action[23]. Even clever non-human animals (such as New Caledonian crows and chimpanzees) have difficulty learning causal relations from events other than their own actions. Learning by observing others' goal-directed actions (third-personal causal learning), or even by generalizing from the consequences of one's own accidental movements, will not typically enable a non-human animal to learn a causal intervention[20,23,38–43].

There is a debate in the comparative literature about the scope and limits of third-personal causal understanding in non-human animals[41,44,45] (for a theoretical overview, see ref. 46). Although there is some evidence for imitation learning, particularly in non-human primates, it is not always clear whether these results reflect genuine imitation or simpler processes such as emulation or stimulus enhancement. For example, watching another animal drop a stone on a nut might lead the first animal to attend to the stone and then independently discover its causally relevant properties through trial and error or reinforcement learning[46]. This type of learning could occur without the first animal's viewing the second animal's actions as 'dropping the stone in order to crack the nut'. Overall, it appears that some genuine imitative learning occurs in non-human primates. However, this third-personal causal learning is much less extensive than it is in humans, where imitation and attention to others' actions is pervasive even in infancy[47] (Box 2).

On the surface, tool use by non-human animals seems to involve a sensitivity to difference-making variables: agents must recognize that certain features (such as shape) are causally relevant, whereas others (such as colour) are not. However, this sensitivity is often highly specific to particular tools and contexts. Associative and reinforcement learning, plus some understanding of particular physical regularities, perhaps as part of perception (Box 1), might explain much apparent causal understanding in this domain[39]. Although many tool-using animals (including corvids and non-human primates) seem to attend to particular physical properties when they use, select and modify tools[48–55], tool use often requires lengthy trial-and-error learning, fails to generalize to novel circumstances, and breaks down when the animal's own goal-directed action is not involved[38,40,44,56].

A form of causal reasoning is also manifest in studies designed to test 'diagnostic inference'[57,58]. In these tasks, animals must reason about regularities of the physical environment to make a reward-relevant decision, such as where to search for food. Some species of birds and non-human primates can infer the presence of food from cues to volume and solidity, contact causality and weight[59–67]. For instance, some animals seem to understand that food under a wooden board causes it to slant, food inside a cup causes rattling, and food inside a container makes it heavier. However, this apparent causal understanding might be more parsimoniously explained by the animals' understanding of a particular subset of perceptual regularities – either innately or due to associative learning. This understanding contrasts with more decontextualized representations of causal variables and their difference-making relations. A dramatic illustration comes from 'floppy stick' experiments with non-human primates. When an experimenter waves two sticks to show a subject that one stick is rigid and the other sticklike object is floppy, chimpanzees do not reliably select the rigid object to bring a distant reward into reach. Instead, they often select and make many

Box 1

Infants' visual expectations about object dynamics

A literature on 'intuitive physics' and 'core cognition' has investigated infants' expectations about events that they observe, but do not themselves produce[13,32,85,88]. Studies of infants' visual attention to external events have produced two main sets of findings. First, within the first year of life, infants become sensitive to visual stimuli corresponding to events that adults also judge to be causal. Second, by the age of nine months, infants begin to discriminate between visual stimuli that depict physically possible versus impossible events[13,32,85,88].

There are several interpretations of these findings. The 'core cognition view' is that infants' visual perception is sensitive to causal relations in the same way as it is to more basic features such as shape, colour or motion[13]. This work builds on a theoretical suggestion from vision science that causality may be a high-level component of visual perception — as automatic and irresistible as three-dimensional depth perception[317]. In this view, when infants witness a 'launching' event (displays in which one object approaches and contacts another object, which then begins to move at a similar speed, as in a collision of billiard balls), they construe the motion of the first ball as making a difference to the position of the second ball. Core cognition also holds that infants possess 'intuitive physics': highly general knowledge of the laws governing physical object interactions[13,32]. The label 'intuitive physics' suggests an understanding of the impersonal causal relations involved in external events.

An alternative view is that infants are highly skilled at tracking statistical contingencies between typical events and are therefore surprised by abnormal contingencies. However, infants might nonetheless lack an understanding of general causal principles[86]. Evidence for this view includes the fact that infants' expectations about objects seem to be acquired in relatively piecemeal, context-specific ways. For example, although four-month-old children recognize that a tall object will not be occluded by a short one, they do not recognize that the tall object will not be contained by the shorter one until seven months of age[87].

Other findings suggest that visual expectations might not always guide behaviour. For example, four-month-old infants look longer if an object that was dropped above a temporarily occluded shelf is shown below the shelf when the occluder is lifted, relative to infants who see the physically possible outcome[85]. However, 24-month-old toddlers who witness the same event will search for the object underneath the shelf rather than on top of it[318]. The same phenomenon has been documented in non-human primates[319]. Core cognition theorists have dismissed these failures as a form of perseveration, in which the toddlers or monkeys temporarily forget their general conviction that solid objects cannot pass through other solid objects[320,321]. The alternative explanation is that children's early visual expectations about the regularities of object interactions do not reflect the kind of highly general, impersonal causal understanding that determines interventions.

It is unclear how infants' early emerging expectations are related to an adult-like (decontextualized and depersonalized) understanding of difference-making causation. It is plausible that the two types of knowledge co-develop and reinforce each other. For example, correlational knowledge of observable regularities, coupled with observations of agent-involving events (Box 2), might guide early goal-directed behaviour, which in turn gives rise to causal understanding.

Box 2

Third-personal causal knowledge

Human adults, young children and some non-human animals have third-personal causal knowledge[23]: merely observing another agent's causal intervention (such as banging a nut with a rock) is sufficient for them to infer that performing the same intervention themselves would make the same difference (that the nut would crack). This ability might depend on seeing others as agents who are intentionally trying to bring about particular effects.

There is converging evidence that even infants understand that actions are directed at goals. For instance, nine-month-old infants will imitate others' actions on objects in intelligent ways[115,116]. Moreover, in looking-time studies, six-month-old infants expect that a human hand will consistently reach for the same object but not that an inanimate stick or claw will do so[322].

One possibility is that as infants become more competent at moving their bodies to accomplish their goals, they learn to map their own goal-directed actions onto the actions of others[323,324]. In this view, the development of third-personal causal understanding would be the result of mapping 'my' goal/effect onto 'their' goal/effect — perhaps facilitated by a general innate cross-modal mapping between one's own interventions and those of others[94,325]. Support for this view comes from experiments in which pre-reaching three-month-old infants are outfitted with Velcro 'sticky mittens' that enhance their reaching success. Following this experience, the three-month-old infants behave more like (reaching) six-month-old infants in looking-time experiments that test their attention to others' goal-directed actions[326–328].

An opposing perspective from the 'core cognition' view holds that human infants have innately specified concepts of agent, goal, cost and cause[13,307–311,329]. This theoretical position proposes that humans are born with the ability to view agents as difference-makers. However, infants' sensitivity to the distinctive visual features of agent-caused events might be more parsimoniously explained by a general sensitivity to the features of reliable statistical relationships. Salient state changes, close cause–effect spatiotemporal proximity, and invariant contingencies (see ref. 329 for a mega-analysis of experiments and methods) might enable infants to make predictions about the visual features of agent-involving events, although they do not yet understand such sequences as causal interventions. To show genuine causal understanding, researchers would need to show that very young infants try to reproduce other agents' interventions to achieve similar goals. We suggest several such experiments in the conclusion below.

attempts with the floppy stick, suggesting that they do not understand that the property of rigidity is causally relevant to achieving the goal[39].

A handful of studies provide evidence that non-human animals do possess some kinds of more decontextualized and depersonalized causal knowledge. For instance, apes make different interpretations of the same observation depending on causal cues: they infer the presence of yogurt based on seeing a trail of yogurt drips near a hiding location only if the traces appeared after an occluded yogurt-hiding event and not if they were present beforehand[68]. Similarly, chimpanzees seem to discount confounded causes, particularly when their own actions lead to an effect[43]. Several experiments suggest that rats can learn about causal structure from observation, which in turn seems to enable them to make effective interventions[69,70]. However, evidence for adult human-like, fully decontextualized and depersonalized causal understanding in non-human animals is very limited overall.

Development of human causal understanding

Adult human causal understanding assumes that there are objective causal relations that exist independently of one's actions, but that still support interventions. As recently as the mid-twentieth century, psychologists thought that this ability emerged only gradually in the school-age years[71,72]. However, there is now diverse evidence for causal learning and reasoning in very young children. The major developments occur between birth and five years of age.

Most of the research we describe here comes from laboratory experiments conducted with predominantly 'WEIRD' (Western, educated, industrialized, rich and democratic) populations[73]. It is likely that there is interesting and important variation to be discovered across human cultures, and cross-cultural studies would illuminate which aspects of this trajectory reflect particular patterns of experience (see 'Causal learning in naturalistic contexts').

Causal learning and reasoning in infancy

As with non-human animals, humans' ability to generate and imagine causal interventions is probably prefigured in the ability to control bodily movements and to perceptually track environmental regularities. Humans are a highly visual species born without developed locomotion, and eye movements are among infants' first controlled actions[74–76]. Visual behaviour develops through effortful practice: achieving control over where one is looking (and therefore what one sees) depends on implicit learning about which muscle movements make a difference to perception[75,77–79]. Thus eye movements might be understood as a precursor to difference-making interventions.

This learning seems to occur in parallel with the development of early expectations about external events, such as the ability to predict regularities of the physical world in visual perception. Within the first year of life, infants discriminate visual stimuli corresponding to events that adults judge to be causal[14]. For example, between six and seven months of age, infants become sensitive to 'launching' events[15,16,80–83]. Between seven and nine months of age, infants discriminate between launching and 'triggering' events (in which one object approaches and contacts another object, which then begins to move faster than the first object)[84].

Between the ages of four and nine months — and in one study, as young as 2.5 months[85] — infants begin to discriminate between visual stimuli that depict physically possible versus impossible events. For example, infants look longer at displays in which objects appear to float unsupported in midair, pass through other objects, or travel along discontinuous paths than they do at closely matched displays in which object interactions accord with the laws of physics[13,32,85–88]. Some researchers suggest that these findings reflect infants' understanding of particular types of physical causal event, but the interpretation of these findings is subject to debate (Box 1). Another possibility is that early abilities to make predictions in visual perception later combine with a more action-based causal understanding as infants gain experience intervening in the physical world.

In addition to their abilities to predict and detect anomalies in observed events, infants also show anticipatory looking behaviour to visual stimuli that exemplify difference-making patterns of statistical dependence. Two studies showed that eight-month-old infants[89], but not five-month-old infants[90], learned to predict patterns in the sequential appearance of static images (without animated collisions or agents) in ways that were consistent with a causal interpretation. For example, infants might see that two pictures A and B that occurred together were always followed by a third picture C, and that C followed A by itself, but did not follow B by itself. They would then expect C when A occurred, but not when B did. However, these experiments use a dependent measure based on prediction rather than intervention. It is therefore an open question whether this should be interpreted as causal learning. In the interventionist perspective, only an experiment demonstrating that infants learned from observational evidence to make an intervention would provide evidence for genuine causal understanding.

However, even before these developments, within the first weeks of life infants exert control over the social environment. Interventions that fulfill social desires – such as making noises, faces and gestures that bring a caregiver closer – are among the most meaningful and salient differences that infants learn to make. One of their earliest causal interventions might be crying[91]. Several studies suggest that by four to eight weeks of age, crying shifts from being merely 'expressive' to being targeted at goals, such as gaining a caregiver's attention[92,93]. By six weeks of age, infants also initiate imitative games with their caregivers – interactions in which social partners contingently mimic each other's facial movements[94]. These behaviours reflect an early understanding that certain actions make a difference to the social environment.

By at least three months of age[95], but perhaps as early as one month of age[96,97], infants are capable of basic first-personal reinforcement learning with respect to the physical environment. For instance, when experimenters tie one end of a ribbon to a mobile and the other to an infant's foot, participants rapidly learn which movements are causally relevant for generating the visually rewarding effect[95,98]. However, similar to non-human animals, infants often struggle to generalize: if the ribbon is tied to the other foot or the child is moved to a slightly different context, they need to relearn the causal relation[99]. These very young infants also seem to seek out contingencies between actions and outcomes even when the outcome is not intrinsically rewarding[98,99].

Between four and five months of age, infants' automatic grasping and 'manual babbling' (spontaneous and relatively random finger, hand and wrist movements) develops into targeted reaching, grasping and mastery of the human 'precision grip'[100,101]. These motor accomplishments facilitate increasingly precise manipulations of the physical environment, thereby expanding opportunities for discovering difference-making relations[101–103]. At the same time, a fundamental interest in executing goals feeds back into the physiological developments themselves. For example, in studies where infants are outfitted with 'sticky mittens' (Velcro mittens to which experimental stimuli can easily attach), subjects can pick up objects before their grasping skills are fully mastered. These infants subsequently engage in more reaching and grasping compared to infants without sticky mitten experience[104,105].

Like non-human animals, infants struggle to reason about impersonal causal relations independent of their own actions. For example, learning to recognize difference-making variables that are relevant for successful spoon use (such as where to grip and how to orient the bowl of the spoon) requires many attempts – approximately ten months' worth of practice according to one study[106]. Causal interventions using tools are also highly personalized: rates of success decrease dramatically if the interventions are directed at other agents (such as spoon-feeding a teddy bear) or at maintaining external conditions (such as figuring out how to hold a shovel to avoid water sloshing out)[107,108]. Before approximately nine months of age, infants also fail in 'means-to-ends reasoning' tasks that require them to consider causal relations between external physical objects (such as using a rake to retrieve an out-of-reach toy). Although an infant might understand that spatiotemporal contact is required, they frequently fail to recognize causally relevant variables, such as which end of the rake will be more helpful for achieving the goal[109–111].

Although many studies have focused on infants' abilities to understand causal relations in the physical environment, difference-making is perhaps even more evident in social contexts. Researchers have suggested that infants' ability to make social interventions undergoes a 'revolution' at nine months of age[112]. For instance, at this age infants learn to manipulate others' attention by pointing. They also start to produce basic communicative signs, such as raising their arms above their head in the 'pick me up' gesture[113].

The ability to share joint attention to objects and events in the environment facilitates language acquisition, an even more powerful means of manipulating other agents[112,113]. A large proportion of infants' early utterances encode causal events[111,114]. Examples include 'more' (referring to diverse events involving recurrence) and 'allgone' (referring to diverse events of disappearance). Initially, these utterances are restricted to events the children cause themselves (such as repeating an earlier action or hiding an object) and only gradually become more depersonalized and decontextualized[114]. These early words reflect both infants' interest in causal events and adults' propensity to point them out[109,111].

Moreover, unlike non-human animals, infants appear to be capable of something like third-personal causal learning from an early age. By nine months of age, infants will imitate an experimenter's novel actions on objects that bring about a novel effect[115,116]. At 14 months of age, infants who witness an experimenter place their head on a box that subsequently lights up will produce the action themselves to bring about the effect – even if they have never done so before, and even a week after seeing the demonstration[116]. Moreover, this imitation seems to reflect an understanding of the demonstrators' intentions rather than just their physical actions: if infants see that the experimenter is wrapped in a blanket that prevents them from using their hands, and then uses their head to activate the machine, the infants will use their own hands rather than their head[117]. By 24 months of age, infants can clearly act to bring about a particular causal effect: they will differentially imitate an action that leads to a particular goal[118,119].

One study suggests that this sort of third-personal causal learning integrates expectations about 'intuitive physics'[120] (Box 1). In these experiments, eleven-month-old infants witnessed actions that resulted in physically impossible events: for example, an experimenter's hand pushed a ball down a ramp and then the ball appeared to pass through a solid wall, or an experimenter's hand pushed a toy car to the edge of a

Table 1 | Milestones in the development of causal reasoning

Domain	Key milestone	Example behaviour	Age range observed	Selected references
Predictive visual scanning and oculomotor control	Infants learn to control eye movements and predictively track changes in environment	Tracking the trajectory of a sequence of lights	0–3 months	74,75, 77–79
Reinforcement and contingency learning	Infants learn to generate desired states of the world via social and physical interventions Earliest FIRST-PERSONAL causal understanding	Crying to attract attention Kicking to move a mobile	0–3 months	92–97, 99,305
Visual expectations about regularities of the physical world	Infants visually differentiate between typical versus atypical dynamics of physical objects	Looking longer at a visual display in which an object appears to float unsupported in the air than one in which it rests on top of another object	2.5–12 months	85–88
Grasping and reaching	Infants gain enhanced manipulative control over the physical environment	Grabbing a blanket and pulling it closer	4–5 months	101–105
Visual expectations about external events involving agents	Infants visually differentiate between events involving non-agents versus agents seeking goals	Looking longer at a display in which an agent appears to take an inefficient path in approaching another object	6–12 months	306–311
Visual expectations about object motion and spatiotemporal contact in causal events	Infants visually differentiate between typical versus atypical dynamics of physical object interactions	Dishabituating to the reversal of a 'launching' event Looking longer at a 'triggering' event after being habituated to launching events	6–12 months	15,16,80–84
Locomotion	Infants learn which perceivable features of the environment are difference-making with respect to effective locomotion	Re-learning which inclines are safe for descent in walking following transition from crawling	6–12 months	29,30,250
Statistical learning and predictive looking behaviour consistent with causal understanding	Infants learn to visually predict statistical patterns that are consistent with causality in visually observed displays that do not involve agents or collisions	Looking longer at the place where a picture will next appear, given previously observed patterns of statistical dependence in the sequential appearance of other images	8 months	89,90
Joint attention and early language development	Infants learn to make a difference to social partners' attention, facilitating attention to external events	Pointing and learning to attend to social partners' points	9 months	112,113
Imitating others' causal interventions	Infants learn causal interventions by observing other agents' actions Earliest THIRD-PERSONAL causal understanding	Imitating an experimenter's novel causal intervention (such as using one's head to turn on a light)	9–14 months	115–117
Exploring or imitating surprising physical events caused by others' interventions	Infants notice unusual causal outcomes of other agents' interventions	Preferentially dropping a toy car that another agent pushed off a surface, which then floated in midair	11 months	120,186
Causal language	Infants and toddlers learn and use language to refer to difference-making interventions and events	Using 'allgone' to refer to the outcomes of diverse interventions that cause disappearance (such as finishing a cup of juice or popping a bubble)	12–24 months	110,114
Causal learning and reasoning from others' interventions	Toddlers identify difference-making variables and generalize causal rules after observing evidence from another agent's causal interventions	Placing a (new) green triangle block on a machine after observing that a green rectangle and a green cylinder, but not a red cube, resulted in a blicket detector playing music	16 months to 3 years	8,118,119, 123,125,130, 158,208
Pretending	Toddlers understand and generate fictional identities and events for themselves, other people and objects	Using a wooden block as soap to 'wash' a teddy bear	24 months	202,203
Causal learning and reasoning about external events that do not involve agents	Four-year-old children learn causal relations by observing evidence that does not involve their own or another agent's intervention Earliest IMPERSONAL causal understanding	Observing that a toy car's movement (without an agent) generates a light stimulus, then moving the car when prompted to reproduce the phenomenon ('Can you make the light turn on?')	4 years	118,158
Rationally revising beliefs in light of evidence	Preschoolers and early school-aged children learn novel and counterintuitive causal relations and revise pre-existing beliefs in light of evidence	Learning that being scared can cause a tummyache Learning that spatiotemporal contact is not required to cause a blicket detector to play	4–5 years	128,137,138
Forming 'overhypotheses' about causal relations	Preschoolers and early-school-aged children identify and generalize causally relevant variables and the abstract forms of cause–effect relations	Predicting that a wizard who caused a tree and a cow to shrink will also cause a doughnut to shrink	4–6 years	134,139–143

Table 1 (continued) | Milestones in the development of causal reasoning

Domain	Key milestone	Example behaviour	Age range observed	Selected references
Spontaneously exploring causal relations in free play	Preschoolers and school-aged children design experiments, identify causally relevant variables, and disambiguate causal relations in free play	Spontaneously flipping a pair of glued-together pop beads around to test which bead makes the blicket detector go	4–8 years	187–190,194
Causal conceptual understanding of natural kind categories	School-aged children appeal to causal 'essences' when reasoning about biological concepts	Believing that a raccoon cannot be turned into a skunk by changing its appearance	5–10 years	146–148
Causal conceptual understanding of complexity and mechanism	School-aged children show an appreciation of the complexity of causal mechanisms and processes	Believing that a machine with more complex insides will be able to perform a greater diversity of functions	5–10 years	149–154
Past counterfactual reasoning	Children reason about the possibility of past events having gone differently than they actually did	Reasoning about where a ball would have gone if an obstacle had not been present Reasoning about where a character would have found the candy if another character had not stolen it	5–12 years	268,279, 283,312

solid surface and then the car appeared to float in midair. Next, infants were provided with the opportunity to interact with the ball or car. Infants who saw the ball pass through the wall opted to bang more often than drop it — as if attempting to reproduce the intervention that caused it to pass through another object — and children who saw the car float opted to drop more often than bang it — as if attempting to reproduce the intervention that caused floating. Along with other work[117–119], this result suggests that children were not simply imitating the experimenter's actions, but rather understood those actions in a causal way. Furthermore, infants seemed to interpret the actions in the light of broader expectations about physical events. Those expectations might be derived either from innate 'core knowledge' or from previous statistical learning (Box 1). However, given that the experimental stimuli in this study involved agents — a human hand always initiated the sequence of unusual events — it is plausible that infants' behaviour does not reflect a decontextualized and depersonalized understanding of general physical causal relations. Rather, the infants' behaviour might merely reflect third-personal causal understanding, an understanding of cause and effect that is directly linked to agents' actions. Only displays in which another object caused the ball or car's movement, or was not shown at all, would directly test for impersonal causal understanding.

Causal learning in early childhood

By around the age of four years, children demonstrate more sophisticated capacities for depersonalized and decontextualized causal learning and reasoning (Table 1).

One influential line of research brings together developmental studies and formal models of causal inference developed in philosophy of science and computer science[4,6–9]. Causal hypotheses can be represented as probabilistic models of the world that generate patterns of evidence, which enable inverse Bayesian inference[2,8]. From a particular pattern of evidence, one can infer the probability that different possible causal structures produced that pattern[2,121,122]. In the causal Bayes net formalism, causal relations are described using directed graphs that systematically generate patterns of conditional probability among the variables and the outcomes of interventions on them. If a reasoner observes a pattern of conditional probabilities and interventions, they can infer which causal graphs might have produced that pattern. (For helpful tutorials, see refs. 2,121,122).

The Bayes net formalism inspired a research programme exploring whether young children can make similar inferences[2,3,8]. The basic method behind this research has been to show children novel causal systems, often new toys in which entities light up, move up and down, or spin[8,123–126]. In other paradigms, a novel system is presented as part of a story[127,128].

As an example, imagine someone places a red and a blue block on a machine, which plays music when certain objects are placed on top of it (a 'blicket detector'[125]; Fig. 2). It is possible that only the red block caused the effect, that only the blue block did, or that they both did. Each of these causal hypotheses can be represented by a different causal graph and will systematically generate different predictions. For example, if the blue block is the only cause, then the red block by itself will not make the machine go, but the blue block by itself will. By looking at further evidence, such as what happens when single blocks are placed on the machine, these alternative possibilities can be differentiated.

In these experiments, children observe the conditional probabilities among the values of variables in the system and the outcomes of interventions on them. For example, they might see that a red block makes a blicket detector activate three out of four times, whereas a blue block is effective only one out of eight times, or that a blue block and red block together activate the machine, but the blue block by itself does not. Experimenters then prompt the children either to produce causal interventions themselves (for example, asking them to make the machine go) or to answer causal questions that require them to make predictions, generate explanations or reason counterfactually (such as asking if the machine will go if a certain block is placed on it). The answers to these questions demonstrate whether children have inferred the correct causal hypotheses from the data. To a striking degree, children make inferences that are consistent with the predictions of probabilistic generative models in general, and causal Bayes nets, in particular[2,8]. Moreover, children's inferences go well beyond what would be expected from simple association, reinforcement learning or intuitive physics knowledge (for further discussion, see refs. 4,8,9).

Children as young as 16 months can infer causal structure from patterns of covariation between cause and effect[19], and toddlers can learn from probabilistic data to infer relative causal strength or power[119,129]. In addition, children at this age not only infer that simple properties of objects (such as the shape or colour of blocks placed on the detector)

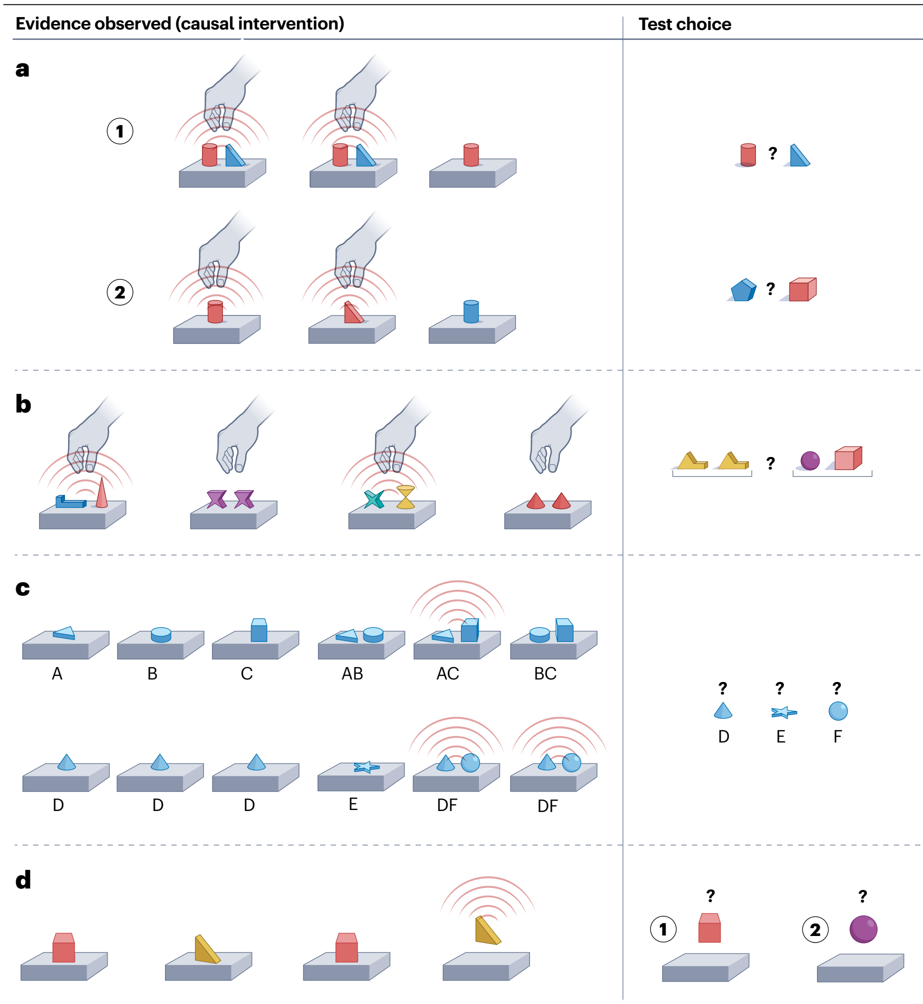

Fig. 2 | The blicket detector paradigm. **a**, Participants observe an experimenter's causal interventions on a machine that lights up or plays music when certain objects, called blickets, are placed on it (a blicket detector). Then, they are prompted to identify the blickets and to 'make the machine play music'[8,123,125]. (1) Toddlers (19–24-month-olds) learn which block is difference-making. Here, they choose to place the blue triangle on the machine. (2) Toddlers also generalize causal rules. Here, they choose to place the novel red cube on the machine. **b**, One variation shows that toddlers learn and generalize 'relational' causal rules[130]. Here, they choose the novel 'not-same' pair (right). **c**, Another variation shows that preschoolers learn and generalize 'conjunctive' causal rules, such as 'A and C together cause music'[141]. Here, preschoolers are more likely to endorse the conjunctive hypothesis – that both D and F are required – than adults, who tend to choose F alone. **d**, Another variation shows that preschoolers learn and generalize unusual causal rules[137]. (1) Preschoolers intervene to make the machine play music by hovering a previously inactive block above the machine. (2) Preschoolers generalize the unusual rule to a new block. Part **b** adapted with permission from ref. 130, Sage. Part **c** adapted with permission from ref. 141, Elsevier. Part **d** adapted with permission from ref. 137, APA.

are causally relevant, but also infer more abstract and general causal rules that are grounded in relations between objects. For example, 18- to 30-month-old toddlers can learn that two identical objects placed on top of a blicket detector will make it go, whereas two objects that are not the same will not[130]. Toddlers also make causal inferences spontaneously in more naturalistic settings. For example, 16-month-old infants learn from evidence to infer whether the failure of a toy's activation was more likely to be due to the toy's malfunction or to their own incompetence[19].

By four years of age, children's causal understanding is quite sophisticated. They can learn complex causal structure (including causal chains, common effect structures and common cause structures) from conditional probabilities[126]. They can infer unobserved hidden causes as well as observable ones[8,131,132]. From this age, children also reason about causal relations in the biological, psychological and social domains[133–136]. For example, they can learn from evidence to infer whether someone acts a certain way because of their situation (such as the presence of a dangerous object) or because of a personality trait (such as because they are timid)[134].

Preschoolers can also integrate and override prior causal knowledge in the face of new evidence in a Bayesian way[137,138]. For example, three-year-old and four-year-old children initially assume that a block must physically contact the blicket detector to make it go (in Bayesian terms, they have a strong 'prior' for contact causation), but they can rapidly learn from evidence to infer that a distant block can also be effective[137]. Similarly, four-year-old children, though not three-year-old children, can learn to infer counterintuitive cross-domain causality (such as that being scared, a psychological cause, can cause a tummy ache, a physical effect)[128].

Preschoolers also seem to have deterministic expectations about physical causal relations: for example, they seem to assume one-to-one mappings between distinct causes and their effects[139]. Perhaps consequently, they infer that hidden causes are responsible for apparently stochastic causal patterns[132]. Preschoolers also learn overhypotheses, more abstract causal rules about the form of the relation between causes and their effects[134,140–143]. For example, three-year-old and four-year-old children who observe that a novel cause makes certain objects larger will infer that the same cause will make other, perceptually dissimilar objects larger, too[142]. Preschoolers can also infer the logical form of causal relations (such as whether they are conjunctive or disjunctive)[140,141] and can sometimes do so better than adults can[141,144].

This suite of causal reasoning abilities is clearly in place by the age of four, and some are present as young as two years of age. All these examples involve learning in laboratory experiments. However, the ability to infer abstract causal forms also seems to guide naturalistic learning. For example, children often acquire larger framework theories of a domain (such as biology) before they learn specific facts about objects and events[145,146]. These 'intuitive theories' are often associated with ideas about internal essences that cause observable features (such as the idea that there is something inside skunks that generates their skunk-ness). The theories then seem to guide and constrain children's acquisition of novel concepts within the category[146–148].

By the age of five, children are also capable of reasoning about causal mechanisms. For example, they might have expectations about the way that internal causal complexity will relate to the diversity of the functions that a system can perform[149–154] – even though, like adults, they often do not have detailed knowledge of how the mechanisms actually work[155].

Thus, by the age of five years, human children have extensive decontextualized and depersonalized causal understanding. But, of course, causal understanding continues to develop. For learners at any age, acquiring specific causal knowledge (for example, about the inner workings of a laptop or the human body) often depends on a combination of factors, such as personal motivation and the availability of cultural knowledge. As causal expertise in specific domains increases, so too does general causal understanding. For example, expert scientists are better than novices at identifying types of causal pattern, such as positive feedback loops, across diverse scenarios[156,157].

Limitations of early causal understanding

There are some limitations on children's early causal learning and reasoning. These limitations parallel some of the limitations of non-human animals. As we have said, some of the capacities of infants and toddlers seem to be limited to specific expectations in particular situations rather than applying more generally: their causal understanding is not yet decontextualized. Before the age of four years, young children, like non-human animals, can also learn more wide-ranging and general causal relations. However, they seem to focus on those that are directly linked to agents' actions, rather than seeing causal relations as objective features of the world that happen to be manipulable by agents: their causal understanding is not yet depersonalized.

On the other hand – unlike non-human animals – human children from infancy on seem to readily infer causal relations from both their own actions (first-personal causal understanding) and the actions of others (third-personal causal understanding; Box 2). However, unlike preschoolers, toddlers and infants do not yet infer causality simply from observing correlations. That is, seeing statistical contingencies between two external events will not prompt them to try to intervene on the first variable to generate a change in the second. They have not yet developed adult-like, impersonal causal understanding.

The blicket detector studies involve causal effects that are the outcome of human actions. The demonstrator activates the machine, and the children must infer its causal structure. However, two studies using separate experimental paradigms suggest that toddlers make genuinely causal inferences such as these in only three cases: when causal events are the product of another agent's intervention (as in the blicket detector studies); when causal events involve spatiotemporal contact between objects; or when causal events are described using causal language[118,158].

In one experiment, children repeatedly saw a toy car collide with a block, followed by an attractive effect in which an egg-shaped object lit up. They also saw that the effect did not occur when the car collided with a different block. Crucially, the experimenters manipulated whether someone pushed the car towards the two blocks or whether the car simply moved by itself. In both conditions – whether or not the events involved agents – toddlers learned to predict the effect: they looked at the egg-shaped object only when the car contacted the causally relevant block, indicating that they expected the light-up effect to follow this event[118]. (This finding is consistent with previous work suggesting that infants learn to anticipate visual sequences that exemplify difference-making patterns of statistical dependence[89,90]).

The difference between conditions (agent-caused versus non-agent-caused) became apparent when the experimenter asked the children to make a causal intervention ('Can you make the egg light up?'). If another person had pushed the car, two-year-old children and three-year-old children readily pushed the car towards the correct block, indicating that they had inferred the causal relation from the other person's action. However, they did not do so if they had only seen the car move by itself without human intervention. In a similar experiment, two-year old and three-year-old children also made the correct intervention if the car moved by itself but the experimenter described the events using causal language ("Look! The car makes it go")[158]. By contrast, four-year-old children spontaneously made the correct intervention in both the agent-caused and non-agent-caused conditions.

Before the age of four years, children's understanding of causality appears to be tightly linked to their understanding of their own and others' actions. These experiments suggest that young children see agents as 'making things happen', but they seem to perceive the rest of reality differently. They are limited to first- and third-personal causal understanding. External events are viewed as predictable sequences that can be anticipated, but not as objective causal relationships that can be generated, controlled and explained. This is consistent with a large empirical literature showing that human infants focus on agents. It also complements theoretical proposals that emphasize the importance of social cognition to humans' general intelligence and evolutionary success[159–161].

Limitations of the causal Bayes approach

There are at least two major challenges for the interventionist approach, formalized in terms of causal Bayes nets. Interestingly, empirical evidence suggests that children solve these computationally intractable problems from an early age.

First, formal approaches usually specify the relevant variables and relations in advance. But one unanswered question is how agents choose the right variables to test, consider, or otherwise represent in their mental models in the first place. Philosophers have called this the 'problem of causal selection'[162] or the 'problem of variable choice'[163]; it is also related to the problem of 'feature extraction' in machine learning[164]. One study showed that children as young as three can solve this problem by considering which general kinds of variables – rather than which specific values of variables – made a difference in the past. For example, if changing the colour of a watering can (pink versus yellow) previously made a difference to whether a seed grew, but changing the pattern on a flower pot (striped versus polka dotted) did not, children will choose the watering can as the relevant causal variable to change in a new task in which they are prompted to cause a new kind of seed to grow. Crucially, this pattern occurs even when they are faced with values of the variables they have never seen before (such as a purple can and a checkered pot)[139].

Another major question concerns how agents come to consider a limited set of hypotheses and how they search through this space of possibilities – the 'search problem'[143]. A formally ideal learner would consider the entire set of possible causal relations, causes or outcomes, but this presents a computationally intractable challenge. Even if this problem were solved, there is still an implementation issue: searching through such a vast space is not possible for biological agents with limited time and resources. Several studies suggest that children as young as four years of age might solve this problem by using rational sampling strategies to generate and test out new hypotheses serially[165–167]. Other findings with adults and children suggest that the possibilities they entertain might be constrained by action-relevant considerations: those that are probable, normative (such as morally permissible or conventional), physically possible and good for the agent are more likely to be represented[168–170].

Causal learning in naturalistic contexts

Human causal reasoning develops in diverse naturalistic contexts. Three contexts are particularly important: exploration, play and sociocultural practices (Fig. 3). Here, we emphasize the relations between the learning opportunities that these contexts afford and how they might support the development of depersonalized and decontextualized causal understanding. We also highlight specific cross-cultural factors that might influence this development and suggest avenues for future research.

Exploration

Early exploratory behaviour seems to be driven by two major, distinct motivations. Exploration can be directed at unexpected or otherwise unusual events (to discover causal explanations) and at controllable events (to discover invariant causal relations that hold across many contexts). Both forms of exploration help children to learn specific causal relations and generalize from them.

Several interrelated literatures investigate children's curiosity and spontaneous experimentation in response to unexpected events. The 'active learning' framework portrays children as self-directed learners who seek and produce evidence for themselves[171,172]. The related 'child-as-scientist' (or 'rational constructivist') approach emphasizes that children are interested in surprising results and perform hypothesis-testing behaviours[111,146,147,173–178]. Empirical research guided by these frameworks has demonstrated that children's exploration is directed: as early as 7–8 months of age, they preferentially attend to events that will let them learn the most[179–185]. By 11 months of age, infants who observe a surprising event show targeted exploration[120]. However, they do not explore when they receive subsequent evidence that the apparent irregularity was actually typical[186]. By the age of four years, children recognize ways to disambiguate multiple possible causal relations and even design spontaneous experiments to test them[124,183,187–194].

Other approaches suggest that much of children's exploration is motivated by a desire to pursue their goals and gain control. The 'search for invariance' proposal suggests that preschoolers might be motivated to discover causal relations that generalize across contexts[189,195]. This motivation might help to explain 'positive testing' behaviour, which persists into the school-aged years. In positive testing, children opt to repeat the same or very similar causal interventions rather than testing alternative hypotheses[196,197]. Moreover, from toddlerhood, children seem to explore and learn about causal relations by setting novel goals and exploring the space of potential actions, consequences and solutions, rather than by directly exploring the features of objects[198]. One possibility is that children are seeking 'empowerment' (see below) — contingency between their actions and outcomes — in addition to seeking information or novelty.

Play

Play — defined broadly as non-functional, often repetitive behaviour that occurs when animals are healthy, rested and safe — occurs in many forms across animal species[199]. For humans, playing presents opportunities for children to learn and generalize causal relations.

One proposal suggests that human play fundamentally involves setting and solving arbitrary problems[200]. Turning a stack of cushions into a fort or a pile of blocks into a tower often requires learning about interventions and causally relevant features that are useful for achieving goals more generally. In addition to learning about physical causal relations, playing with siblings and friends can promote social causal learning as children experiment with ways to persuade, manipulate and cooperate with others. Playing with others also provides opportunities for observational causal learning[201].

Pretend play emerges in toddlerhood[202,203]. Pretence involves a kind of counterfactual reasoning, or assuming possible but non-actual values for causal variables and playing out their consequences[204,205]. Through childhood, pretence behaviour becomes more elaborate[206]. There is empirical evidence that pretence, causal reasoning and counterfactual reasoning about the past are positively related: children who are better at pretence are also better at counterfactual reasoning[207,208]. Thus, pretence might both reflect and promote children's thinking about causal possibilities[204,205].

Sociocultural contexts

Children's interactions with adults and peers provide opportunities for observational causal learning, narrative and explanatory reasoning, and question-asking. These activities might enhance children's causal learning and reasoning by directing their attention to specific causal relations. They might also aid in the development of decontextualized and depersonalized causal understanding.

In observational causal learning, children learn by watching others and reasoning about their goal-directed actions[118,158,209]. Children's observational causal learning might be especially enhanced for interventions involving artefacts, such as tools[210,211]. By the age of two, toddlers rapidly learn artefact functions from observation and approach artefacts in ways that differ from the approaches of tool-using non-human primates[212,213]. One theoretical proposal has credited humans' early emerging 'teleological-intentional stance' for their enhanced tool-use abilities and causal learning more generally, relative to other animals[214]. One study demonstrates that preschoolers spontaneously inferred an unusual causal rule from artefact design alone: when a blicket detector had openings on both sides for placing blocks, children were more likely to infer that two blocks were required to make it activate[215].

Language is also important. Adults directly and indirectly highlight causal relations in their verbal interactions with children, and toddlers are more likely to make causal inferences when events are described with causal language[158]. Adults' questions and prompts help children to frame predictions and interpret causal events, and stimulate children's exploration and explanation[216–223]. Reading stories provides opportunities for children to infer novel causal relations[127] and form expectations about event categories[224]. Prompting children to recount past events or imagine future ones, such as prompting them to think about what to pack for school tomorrow, might expand children's understanding of temporal dependence[225–227].

Children can also extract causal knowledge from adults by asking questions[228,229]. As with self-directed exploration, children often ask questions that will optimally enhance their current knowledge

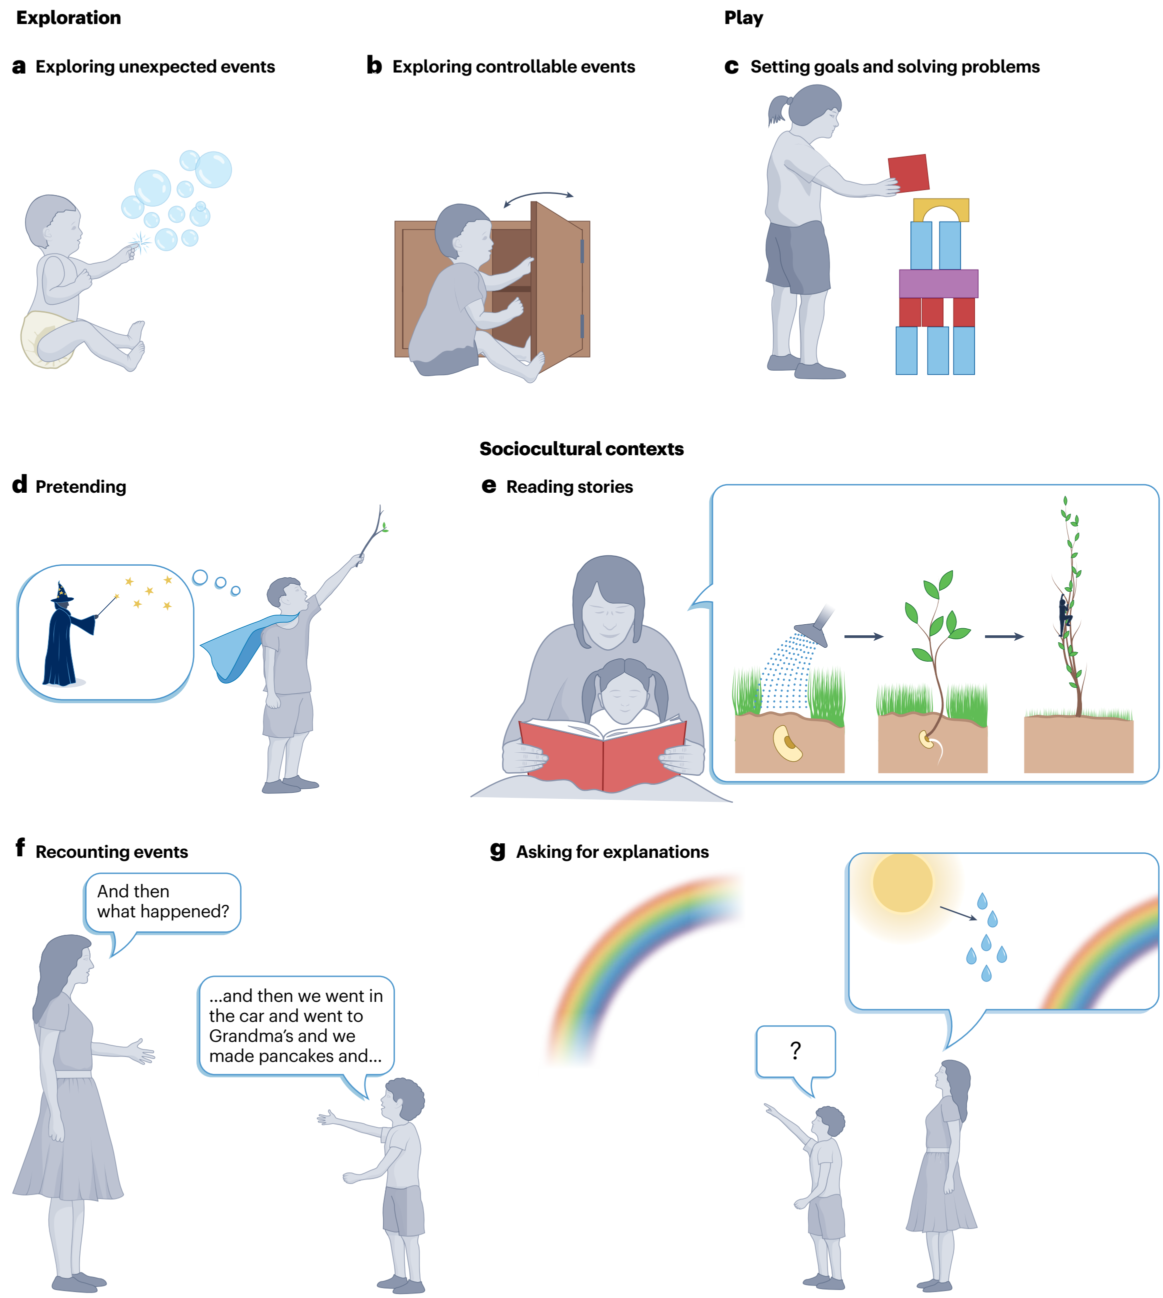

Fig. 3 | Naturalistic contexts that support the development of human causal reasoning. a, Children preferentially attend to and explore unexpected or otherwise unusual objects and events, such as transparent spherical entities floating in midair[120,173,174,176,177,180–182,185,186]. **b**, Children explore invariant causal relations generated by their own goal-directed actions, such as the difference that opening and closing a door makes to what they can see[187,189,195]. **c**, Children explore specific causal variables and relations that are relevant for accomplishing their goals in play, such as the difference that the shapes of blocks makes to success in building block towers[200]. **d**, Pretend play, such as pretending that a stick is a magic wand, involves mentally manipulating the values of causal variables and their relations[202–208]. **e**, Stories help children attend to novel causal variables and their relations[127,224]. **f**, Responding to prompts to recount or explain promotes children's understanding of causal relations and temporal dependence between events[217–220,225,228,232–238,313]. **g**, Question-asking can result in receiving causal explanations[223,228–231], which help children to attend to particular causal variables and their relations[216–220].

state[230,231]. Events that are surprising or inconsistent might be especially likely to trigger children's own production of causal explanations[232,233], which might both enhance their learning[234–238] and prompt discussion with knowledgeable adults[217,219,220].

It might not be a coincidence that exploration, play, curiosity and endless 'why' questions are especially characteristic of young children. In addition to having a uniquely powerful causal understanding, humans also have a uniquely long protected childhood[239]. This lengthy childhood might have evolved to enable wide-ranging causal learning through exploration, play and social interactions.

Only a few studies have looked at causal learning across cultures. Low-income 4–5-year-old children in the USA and Peru were as adept at causal learning as middle-class children in the USA[240]. However, one study found that three-year-old children from cultures that valued interdependence interpreted a causal event as the product of relational causation (involving multiple objects), but children from cultures that valued individuality inferred individual causation (involving single objects)[241]. Other studies have demonstrated that children from urban and rural indigenous cultures reason differently about biological causes[242] and that children who have experience with pets are better at causal reasoning about animals[243].

Because most previous experiments have been conducted with WEIRD populations, there are a variety of sociocultural factors in early causal understanding that have yet to be systematically tested cross-culturally. These might include differences in the age of formal schooling and direct pedagogy[244–247], norms of question-asking and explanation[248,249], and practices surrounding parental supervision and involvement in children's locomotion, exploration and other activities[250]. Investigating the influence of these factors on the development of causal understanding will open up exciting and important avenues for future research.

Development of related abilities

We have highlighted the ways in which exploration, play and social interaction might influence the development of uniquely human causal understanding. Here, we explore connections to related areas of development.

Social and emotional development

Causal reasoning and learning are often thought of as capacities that fundamentally concern the physical world. However, there are many ways in which developing causal understanding intersects with social and emotional development.

Interpersonal skills and communicative abilities are fundamentally ways of exerting control and 'making things happen' in the social environment. Infants' social causal interventions include contingent imitation, crying to summon a caregiver, and pointing[92–94,112,113]. These early skills become more sophisticated across the preschool and school-aged years. The development of persuasive abilities — such as the ability to cite evidence to change others' beliefs[251] — can be viewed as another form of social causal reasoning. Social contexts might also lend special salience to causal outcomes and structures[168,169,252,253]. For instance, the concepts of preventing and allowing are often especially relevant in sociomoral judgement[168,169].

Emotional development can also influence children's causal learning. According to attachment theory, a child's primary caregiver serves as the child's literal or figurative 'secure base for exploration'[254]. Thus, children who are securely attached might spend more time exploring and discovering causal relations. Indeed, one recent study showed that 3–5-year-old children were more likely to explore when a caregiver was present[255]. More generally, some work argues that an early nurturing environment enables more exploration[239] and that early adverse experiences might cause a premature shift from exploration to exploitation, curtailing causal learning[256]. Future work should explore the effects of attachment style and caregiving environments on early causal understanding across cultures[257].

Individual differences in 'epistemic emotions', or feelings related to knowing and learning, might also serve to highlight specific causal relations. One proposal suggests that humans might have a particularly strong explanation-seeking drive, relieved by the characteristic 'aha!' feeling of insight[258]. Causal learning might also be related to emotions that trigger information search, such as wonder, awe, surprise, and the feelings that accompany 'cognitive closure'[259–262].

Reasoning about possibilities

The interventionist and Bayes net approaches to causal understanding assume that young learners are — at least implicitly — capable of reasoning about possibilities. First, causal interventions seem to reflect a belief in a non-actual but possible way the world might become, conditional on one's actions ("If I kick my foot, then the mobile will move"). Second, causal learning often depends on considering multiple hypothesized causal relations. As noted above, even 16-month-old infants can apparently consider multiple possible causes for an outcome: when a toy fails to activate, infants' actions (for example, trying a different toy versus looking to a caregiver for help) suggest that they have inferred whether the failure was more likely to be due to the toy or to their own incompetence[19]. Furthermore, when two possible causes are equally likely to generate an observed effect, 18–30-month-old toddlers flexibly switch between them in their causal interventions[263].

These results suggest that even toddlers consider multiple causal possibilities, especially in the context of their own actions. This parallels recent theoretical proposals from cognitive science, neuroscience and behavioural ecology, which suggest that the ability to consider action-relevant possibilities might have driven the general evolution of cognition[170,264–267]. For instance, early locomoting organisms needed to consider possible threats to avoid, rewards to approach, or paths to pursue towards a goal.

Children's early competence stands in puzzling contrast to research that suggests that they struggle to understand and reason about possibilities in other domains and tasks[268–273] (Fig. 4). Future work might explore how developmentally early reasoning (both in evolution and ontogeny) about first-personal causal, action-relevant possibilities might develop into more depersonalized and decontextualized reasoning about what is objectively possible[274,275].

In humans, one such ontogenetically later-emerging ability is past counterfactual reasoning. This is the type of thinking that adults often perform spontaneously in light of 'near misses' (such as thinking, "If only I had left 5 minutes earlier!" when late to a lecture)[276]. Adults can also engage in counterfactual reasoning about even more decontextualized and depersonalized, and even impossible, events (such as thinking, "What if aliens had landed in 300 BCE and implemented a global matriarchal society?"). In past counterfactual reasoning, the reasoner must set aside what they know about the actual world, imagine that the values of causal variables that have already been determined are mutable, and then play out the consequences of hypothetical interventions on these variables[277–279].

There is considerable debate in the literature about what constitutes counterfactual reasoning, in general, and past counterfactual reasoning in particular. One recent theoretical proposal suggests that elements of counterfactual reasoning are present even in visual event perception in adults[280,281]. This would suggest that the capacity to reason counterfactually might develop in visual perception prior to other measures, but this has not yet been directly tested. However, other proposals hold that genuine counterfactual reasoning does not emerge until as late as 12 years of age[277] (for reviews see refs. 278,282).

Recent studies suggest that counterfactual reasoning about concrete causal systems, such as physical collision events[279] and blicket detectors[207,208,283], can emerge as young as four years of age – at about the same time as measures of impersonal causal understanding. Perhaps the capacity to represent causal variables in a more depersonalized way is related to the capacity to mentally manipulate them. Future research is needed to investigate possible connections between these capacities.

Summary and future directions

We have identified causal factors that help to explain the emergence of humans' causal understanding from simpler ontogenetic

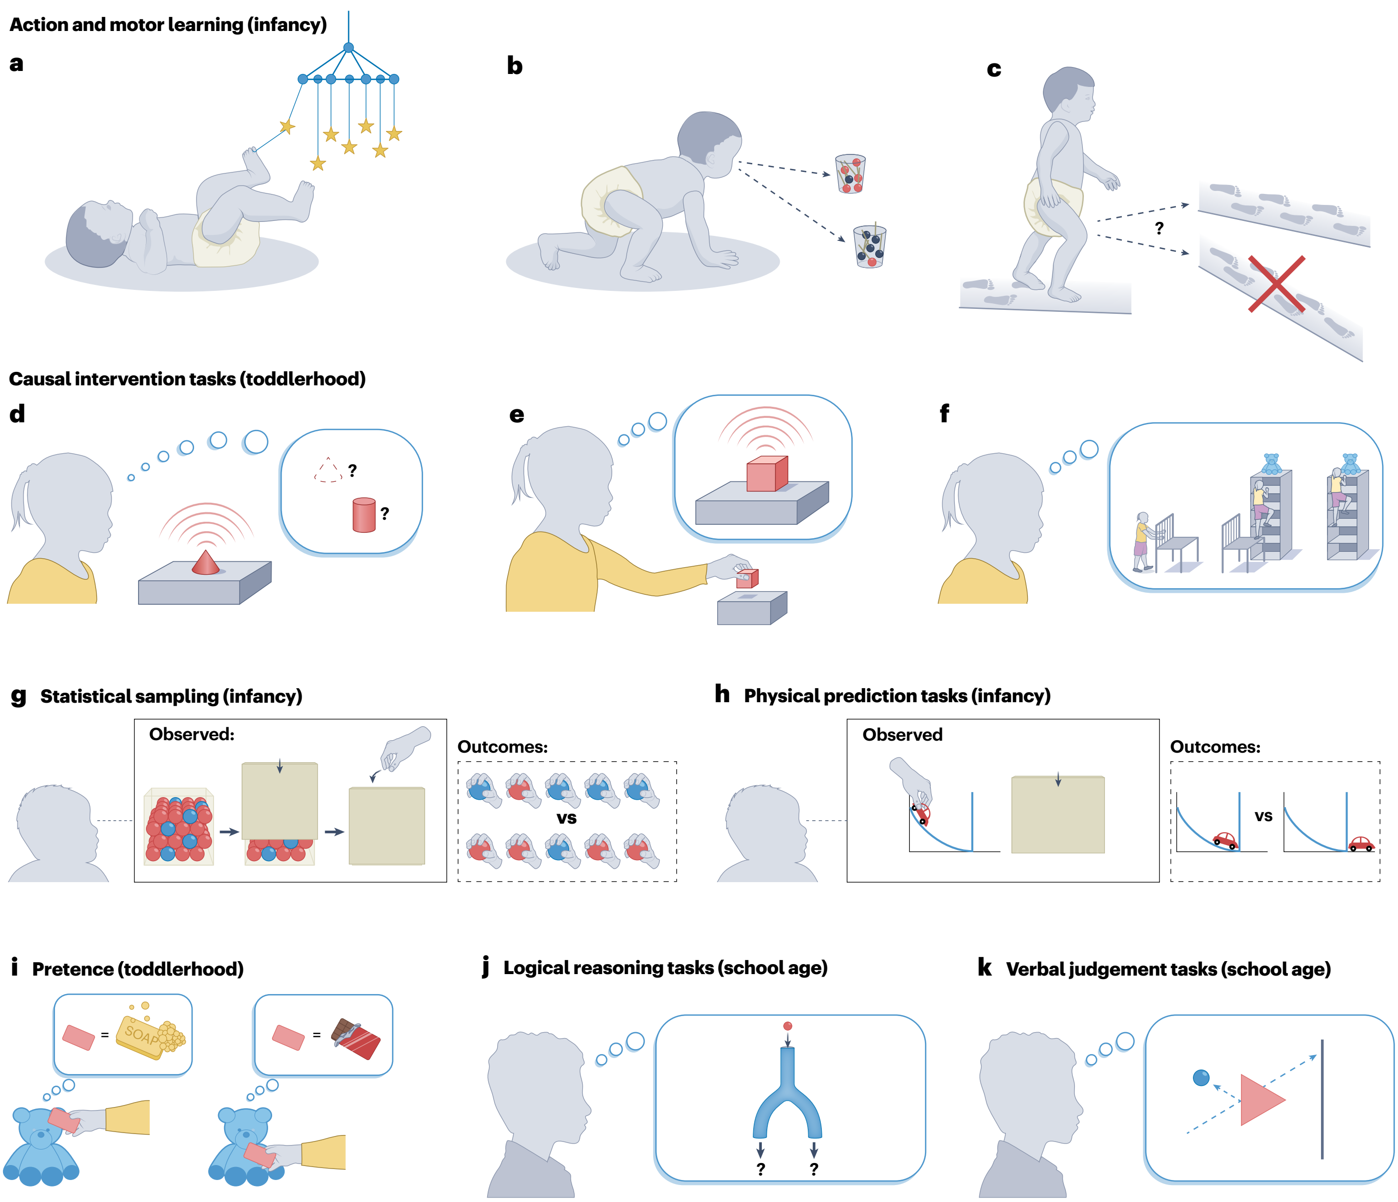

Fig. 4 | Reasoning about possibilities across contexts. a, Reinforcement learning involves reasoning about possible causes, such as whether one's foot movement generated a mobile's swaying[92–97,99,305]. **b**, Decision-making involves considering possible actions, such as which of two containers of treats to approach[314]. **c**, Learning to walk involves reasoning about possible causes, such as whether stumbling was caused by an error in action or by the environment (such as a steep slope)[29,30]. **d**, Learning a causal rule involves reasoning about possible causes, such as whether colour or shape made a machine activate[8,123,125]. **e**, Making causal interventions involves reasoning about possible action outcomes, such as, 'Placing a block on the blicket detector might make music play'[8,123,125]. **f**, Means–ends reasoning involves considering possible courses of actions to solve a problem[106]. **g**, Predicting a sampling event involves sensitivity to possible outcomes, such as the relative likelihood of selecting a red versus blue ball[315,316]. **h**, Exploring objects involved in unexpected events involves differentiating between typical and atypical possible outcomes[120,186]. **i**, Pretending involves appreciating alternative possible functions and identities[202–208]. **j**, Reasoning about mutual exclusivity involves reasoning about contradicting possibilities, such as at which arm of a Y-shaped tube a reward will fall out of[268–271,273]. **k**, Responding to questions about counterfactuals involves making judgements about possibilities, such as what would have happened if the triangle were not there[207,208,277–279,282,283].

and phylogenetic beginnings. Adult capacities depart from those of young children and non-human animals in two key ways. Most animals — including human infants — seem to have a very limited, context-specific understanding of causal events. Non human animals' more general causal understanding is tightly linked to their own causal interventions: they have a first-personal causal understanding. Very young children (and, to a limited extent, some other species) can also reason about other agents' interventions: they have a third-personal causal understanding. By contrast, adult humans' causal understanding is depersonalized (objective; independent of their own or other agents' actions) and decontextualized (applicable across many situations).

Modern adult human causal understanding is probably the result of many coinciding phylogenetic and ontogenetic developments. Physical adaptations for manual dexterity might have enabled more precise manipulations of the environment and, correspondingly, made causal relations more immediate and salient. Social cognitive adaptations that support joint attention, social learning and attention to others' actions might also have contributed. The opportunity to explore and play during a long, protected childhood might enable particularly extensive causal learning[239]. Finally, sociocultural adaptations (such as language and tools) and cooperative social institutions support causal understanding across the lifespan. Investigating potential cross-cultural differences in developing causal understanding is a crucial area for ongoing research.

Another promising direction for future research is to explore potential relations between children's developing notion of an objective reality that exists independently of what people think and a notion of causality that exists independently of what people do. The literature on theory of mind suggests that children begin to understand that the world can be viewed differently from different perspectives from around the age of four[284–292]. Children around this age also begin to reason about causality independent of agents' actions[118,158]. Another line of research investigates children's propensity to provide teleological explanations for natural phenomena (such as 'The sun rises to wake us up'). Impersonal causal understanding might result from a gradual generalization of cause from agent-involving to non-agent-involving events[22,211,293–295].

Experimental paradigms could directly compare infants' learning and generalization from observing external events versus from themselves intervening to make events happen. Experiments like this would clarify the relation between infants' visual expectations about physical events and their expectations about their own goal-directed actions (Box 1). One hypothesis might be that infants use violations of visual expectations as cues to explore. Consequently, they enrich their causal understanding of the environment by actually interacting with it[120,184–186].

Researchers should also investigate early learning from others' actions (Box 2). One important direction might be to test whether children generalize observed causal interventions to new contexts — especially in the social domain. As we have emphasized, many early causal interventions change the social rather than the physical world: infants use their faces and voices to 'reach' into their social environments long before they can physically reach with their hands[91–94]. For example, imagine an infant who observed that an agent's novel vocalization caused a desirable toy to come closer. An infant who reproduced the vocalization to get the toy themselves — that is, an infant who attempted to intervene after merely observing this evidence — would show that they had learned a causal relation. Experiments like this would provide evidence for genuine third-personal causal understanding in infants who are too young to move their limbs.

Another relevant research direction relates cognitive development to machine learning[296–299]. Designing computers that can solve certain problems might illuminate how humans solve those problems, and vice versa. Recent deep learning technologies extract statistical patterns from very large training sets of texts and images and then use those patterns to generate new text and images. These models have made striking advances: they can generate grammatical sentences and realistic-looking images in response to prompts. In some specific causal vignettes, machines can generate predictions, such as predicting that a glass that falls on a hard surface will break. Yet, at present, preschoolers outperform machines on causal reasoning tasks. For example, machines do not infer novel causal uses of a familiar object[296] or learn causal structure in blicket detector tasks[297]. These findings are consistent with the general point that associations, no matter how extensive and complex, can allow predictions to be made, but are not like the causal knowledge that supports interventions and counterfactuals[4,8,9].

Other techniques might be more promising. Reinforcement learning systems are also highly effective and, as we have noted, the structure of this learning is closer to causal understanding. One interesting parallel to children's causal learning may come from approaches that use intrinsic ('epistemic') rewards rather than external rewards[300]. These 'curiosity-based' systems optimize for novelty or information gain rather than, say, a high score. An approach that might be especially related to causal learning uses 'empowerment' as a reward. Empowerment is the mutual information between actions and their outcomes, capturing how well an agent can predict the effects of its actions[301–303]. Maximizing empowerment gives an agent more wide-ranging control of their environment, independent of external rewards. Empowerment is closely related to the interventionist view of causation — if X causes Y, then intervening to change X should lead to a change in Y — in that causation enables control. Therefore, increasing empowerment should increase causal understanding, and vice versa[304].

Like the development of causal reasoning itself, the continued evolution of causal reasoning research will be the joint product of coinciding advances across many fields of study. We hope that this Review will be difference-making for the recognition of new opportunities and continued discoveries.

Published online: 26 April 2024

References

1. Prüfer, K. et al. The bonobo genome compared with the chimpanzee and human genomes. *Nature* **486**, 527–531 (2012).
2. Gopnik, A. & Wellman, H. M. Reconstructing constructivism: causal models, Bayesian learning mechanisms, and the theory theory. *Psychol. Bull.* **138**, 1085–1108 (2012).
This paper provides an introduction to the interventionist approach formalized in terms of causal Bayes nets.
3. Gopnik, A. & Schulz, L. (eds.) *Causal Learning: Psychology, Philosophy and Computation* (Oxford Univ. Press, 2007).
This edited volume provides an interdisciplinary overview of the interventionist approach.
4. Pearl, J. *Causality* (Cambridge Univ. Press, 2009).
5. Penn, D. C. & Povinelli, D. J. Causal cognition in human and nonhuman animals: a comparative, critical review. *Annu. Rev. Psychol.* **58**, 97–118 (2007).
6. Schölkopf, B. et al. Toward causal representation learning. *Proc. IEEE* **109**, 612–634 (2021).
7. Spirtes, P., Glymour, C. & Scheines, R. in *Causation, Prediction, and Search* Lecture Notes in Statistics Vol 81, 103–162 (Springer, 1993).
8. Gopnik, A. et al. A theory of causal learning in children: causal maps and Bayes nets. *Psychol. Rev.* **111**, 3–32 (2004).
This seminal paper reviews the causal Bayes nets approach to causal reasoning and suggests that children construct mental models that help them understand and predict causal relations.
9. Pearl, J. & Mackenzie, D. *The Book of Why: The New Science of Cause and Effect* (Basic Books, 2018).
This book provides a high-level overview of approaches to causal inference from theoretical computer science for a general audience.

10. Lombrozo, T. & Vasil, N. in *Oxford Handbook of Causal Reasoning* (ed. Waldmann, M.) 415–432 (Oxford Univ. Press, 2017).
11. Woodward, J. *Making Things Happen: A Theory of Causal Explanation* (Oxford Univ. Press, 2005).
This book outlines the 'interventionist' approach to causation and causal explanation in philosophy.
12. Shanks, D. R. & Dickinson, A. Associative accounts of causality judgment. *Psychol. Learn. Motiv.* **21**, 229–261 (Elsevier, 1988).
13. Carey, S. *The Origin of Concepts* (Oxford Univ. Press, 2009).
14. Michotte, A. *The Perception of Causality* (Routledge, 2017).
15. Leslie, A. M. The perception of causality in infants. *Perception* **11**, 173–186 (1982).
This canonical study shows that 4.5-month-old and 8-month-old infants are sensitive to spatiotemporal event configurations and contingencies that adults also construe as causal.
16. Leslie, A. M. Spatiotemporal continuity and the perception of causality in infants. *Perception* **13**, 287–305 (1984).
17. Danks, D. *Unifying the Mind: Cognitive Representations as Graphical Models* (MIT Press, 2014).
18. Godfrey-Smith, P. in *The Oxford Handbook of Causation* (eds Beebee, H., Hitchcock, C. & Menzies, P.) 326–338 (Oxford Academic, 2009).
19. Gweon, H. & Schulz, L. 16-month-olds rationally infer causes of failed actions. *Science* **332**, 1524–1524 (2011).
20. Woodward, J. *Causation with a Human Face: Normative Theory and Descriptive Psychology* (Oxford Univ. Press, 2021).
This book integrates philosophical and psychological literatures on causal reasoning to provide both normative and descriptive accounts of causal reasoning.
21. Rozenblit, L. & Keil, F. The misunderstood limits of folk science: an illusion of explanatory depth. *Cogn. Sci.* **26**, 521–562 (2002).
22. Ismael, J. *How Physics Makes us Free* (Oxford Univ. Press, 2016).
23. Woodward, J. in *Causal Learning: Psychology, Philosophy, and Computation* (eds Gopnik, A. & Schulz, L.) Oxford Series in Cognitive Development 19–36 (Oxford Academic, 2007).
24. Godfrey-Smith, P. *Other Minds: The Octopus, the Sea, and the Deep Origins of Consciousness* (Farrar, Straus and Giroux, 2016).
25. Körding, K. P. et al. Causal inference in multisensory perception. *PLoS One* **2**, e943 (2007).
26. Wei, K. & Körding, K. P. in *Sensory Cue Integration* (eds Trommershäuser, J., Kording, K. & Landy, M. S.) Computational Neuroscience Series 30–45 (Oxford Academic, 2011).
27. Haith, A. M. & Krakauer, J. W. in *Progress in Motor Control: Neural, Computational and Dynamic Approaches* (eds Richardson, M. J., Riley, M. A. & Shockley, K.) 1–21 (Springer, 2013).
28. Krakauer, J. W. & Mazzoni, P. Human sensorimotor learning: adaptation, skill, and beyond. *Curr. Opin. Neurobiol.* **21**, 636–644 (2011).
29. Adolph, K. E. An ecological approach to learning in (not and) development. *Hum. Dev.* **63**, 180–201 (2020).
30. Adolph, K. E., Hoch, J. E. & Ossmy, O. in *Perception as Information Detection* (eds Wagman, J. B. & Blau, J. J. C.) 222–236 (Routledge, 2019).
31. Riesen, A. H. The development of visual perception in man and chimpanzee. *Science* **106**, 107–108 (1947).
32. Spelke, E. S. *What Babies Know: Core Knowledge and Composition* Vol. 1 (Oxford Univ. Press, 2022).
33. Pearce, J. M. & Bouton, M. E. Theories of associative learning in animals. *Annu. Rev. Psychol.* **52**, 111–139 (2001).
34. Wasserman, E. A. & Miller, R. R. What's elementary about associative learning? *Annu. Rev. Psychol.* **48**, 573–607 (1997).
35. Sutton, R. S. & Barto, A. G. Reinforcement learning: an introduction. *Robotica* **17**, 229–235 (1999).
36. Gershman, S. J., Markman, A. B. & Otto, A. R. Retrospective revaluation in sequential decision making: a tale of two systems. *J. Exp. Psychol. Gen.* **143**, 182–194 (2014).
37. Gershman, S. J. Reinforcement learning and causal models. In *Oxford Handbook of Causal Reasoning* (ed. Waldmann, M.) 295 (Oxford Univ. Press, 2017).
38. Taylor, A. H. et al. Of babies and birds: complex tool behaviours are not sufficient for the evolution of the ability to create a novel causal intervention. *Proc. R. Soc. B* **281**, 20140837 (2014).
This study shows that complex tool use does not entail the ability to understand and create novel causal interventions: crows do not learn causal interventions from observing the effects of their own accidental behaviours.
39. Povinelli, D. J. & Penn, D. C. in *Tool Use and Causal Cognition* (eds McCormack, T., Hoerl, C. & Butterfill, S.) 69–88 (Oxford Univ. Press, 2011).
40. Povinelli, D. J. & Henley, T. More rope tricks reveal why more task variants will never lead to strong inferences about higher-order causal reasoning in chimpanzees. *Anim. Behav. Cogn.* **7**, 392–418 (2020).
41. Tomasello, M. & Call, J. *Primate Cognition* (Oxford Univ. Press, 1997).
42. Visalberghi, E. & Tomasello, M. Primate causal understanding in the physical and psychological domains. *Behav. Process.* **42**, 189–203 (1998).
43. Völter, C. J., Sentís, I. & Call, J. Great apes and children infer causal relations from patterns of variation and covariation. *Cognition* **155**, 30–43 (2016).
44. Tennie, C., Call, J. & Tomasello, M. Untrained chimpanzees (*Pan troglodytes schweinfurthii*) fail to imitate novel actions. *Plos One* **7**, e41548 (2012).
45. Whiten, A., Horner, V., Litchfield, C. A. & Marshall-Pescini, S. How do apes ape? *Anim. Learn. Behav.* **32**, 36–52 (2004).
46. Moore, R. Imitation and conventional communication. *Biol. Phil.* **28**, 481–500 (2013).
47. Meltzoff, A. N. & Marshall, P. J. Human infant imitation as a social survival circuit. *Curr. Opin. Behav. Sci.* **24**, 130–136 (2018).
48. Boesch, C. & Boesch, H. Optimisation of nut-cracking with natural hammers by wild chimpanzees. *Behaviour* **83**, 265–286 (1983).
49. Chappell, J. & Kacelnik, A. Tool selectivity in a non-primate, the New Caledonian crow (*Corvus moneduloides*). *Anim. Cogn.* **5**, 71–78 (2002).
50. Weir, A. A., Chappell, J. & Kacelnik, A. Shaping of hooks in New Caledonian crows. *Science* **297**, 981–981 (2002).
51. Wimpenny, J. H., Weir, A. A., Clayton, L., Rutz, C. & Kacelnik, A. Cognitive processes associated with sequential tool use in New Caledonian crows. *PLoS One* **4**, e6471 (2009).
52. Manrique, H. M., Gross, A. N.-M. & Call, J. Great apes select tools on the basis of their rigidity. *J. Exp. Psychol. Anim. Behav. Process.* **36**, 409–422 (2010).
53. Mulcahy, N. J., Call, J. & Dunbar, R. I. Gorillas (*Gorilla gorilla*) and orangutans (*Pongo pygmaeus*) encode relevant problem features in a tool-using task. *J. Comp. Psychol.* **119**, 23–32 (2005).
54. Sanz, C., Call, J. & Morgan, D. Design complexity in termite-fishing tools of chimpanzees (*Pan troglodytes*). *Biol. Lett.* **5**, 293–296 (2009).
55. Visalberghi, E. et al. Selection of effective stone tools by wild bearded capuchin monkeys. *Curr. Biol.* **19**, 213–217 (2009).
56. Seed, A., Hanus, D. & Call, J. in *Tool Use and Causal Cognition* (eds McCormack, T., Hoerl, C. & Butterfill, S.) 89–110 (Oxford Univ. Press, 2011).
57. Völter, C. J. & Call, J. in *APA Handbook of Comparative Psychology: Perception, Learning, and Cognition* (eds Call, J. et al.) 643–671 (American Psychological Association, 2017).
This chapter provides a comprehensive overview of causal and inferential reasoning in non-human animals, highlighting (1) the difference between prediction versus causal knowledge and (2) the organization of non-human animals' knowledge (stimulus- and/or context-specific versus general and structured).
58. Völter, C. J., Lambert, M. L. & Huber, L. Do nonhumans seek explanations? *Anim. Behav. Cogn.* **7**, 445–451 (2020).
59. Call, J. Inferences about the location of food in the great apes (*Pan paniscus*, *Pan troglodytes*, *Gorilla gorilla*, and *Pongo pygmaeus*). *J. Comp. Psychol.* **118**, 232–241 (2004).
60. Call, J. Apes know that hidden objects can affect the orientation of other objects. *Cognition* **105**, 1–25 (2007).
61. Hanus, D. & Call, J. Chimpanzees infer the location of a reward on the basis of the effect of its weight. *Curr. Biol.* **18**, R370–R372 (2008).
62. Hanus, D. & Call, J. Chimpanzee problem-solving: contrasting the use of causal and arbitrary cues. *Anim. Cogn.* **14**, 871–878 (2011).
63. Petit, O. et al. Inferences about food location in three cercopithecine species: an insight into the socioecological cognition of primates. *Anim. Cogn.* **18**, 821–830 (2015).
64. Heimbauer, L. A., Antworth, R. L. & Owren, M. J. Capuchin monkeys (*Cebus apella*) use positive, but not negative, auditory cues to infer food location. *Anim. Cogn.* **15**, 45–55 (2012).
65. Schloegl, C., Schmidt, J., Boeckle, M., Weiß, B. M. & Kotrschal, K. Grey parrots use inferential reasoning based on acoustic cues alone. *Proc. R. Soc. B* **279**, 4135–4142 (2012).
66. Schloegl, C., Waldmann, M. R. & Fischer, J. Understanding of and reasoning about object–object relationships in long-tailed macaques? *Anim. Cogn.* **16**, 493–507 (2013).
67. Schmitt, V., Pankau, B. & Fischer, J. Old world monkeys compare to apes in the primate cognition test battery. *PLoS One* **7**, e32024 (2012).
68. Völter, C. J. & Call, J. Great apes (*Pan paniscus*, *Pan troglodytes*, *Gorilla gorilla*, *Pongo abelii*) follow visual trails to locate hidden food. *J. Comp. Psychol.* **128**, 199–208 (2014).
69. Blaisdell, A. P., Sawa, K., Leising, K. J. & Waldmann, M. R. Causal reasoning in rats. *Science* **311**, 1020–1022 (2006).
70. Leising, K. J., Wong, J., Waldmann, M. R. & Blaisdell, A. P. The special status of actions in causal reasoning in rats. *J. Exp. Psychol. Gen.* **137**, 514–527 (2008).
71. Flavell, J. H. *The Developmental Psychology Of Jean Piaget* (Van Nostrand, 1963).
72. Piaget, J. *The Construction of Reality in the Child* (Routledge, 2013).
73. Henrich, J., Heine, S. J. & Norenzayan, A. Most people are not WEIRD. *Nature* **466**, 29–29 (2010).
74. Ayzenberg, V. & Behrmann, M. Development of visual object recognition. *Nat. Rev. Psychol.* **3**, 73–90 (2023).
75. Bronson, G. W. Changes in infants' visual scanning across the 2- to 14-week age period. *J. Exp. Child. Psychol.* **49**, 101–125 (1990).
76. Aslin, R. N. in *Eye Movements. Cognition and Visual Perception* (eds Fisher, D. F., Monty, R. A. & Senders, J. W.) 31–51 (Routledge, 2017).
77. Miranda, S. B. Visual abilities and pattern preferences of premature infants and full-term neonates. *J. Exp. Child. Psychol.* **10**, 189–205 (1970).
78. Haith, M. M., Hazan, C. & Goodman, G. S. Expectation and anticipation of dynamic visual events by 3.5-month-old babies. *Child Dev.* **59**, 467–479 (1988).
79. Harris, P. & MacFarlane, A. The growth of the effective visual field from birth to seven weeks. *J. Exp. Child. Psychol.* **18**, 340–348 (1974).
80. Cohen, L. B. & Amsel, G. Precursors to infants' perception of the causality of a simple event. *Infant. Behav. Dev.* **21**, 713–731 (1998).
81. Leslie, A. M. & Keeble, S. Do six-month-old infants perceive causality? *Cognition* **25**, 265–288 (1987).
82. Oakes, L. M. & Cohen, L. B. Infant perception of a causal event. *Cogn. Dev.* **5**, 193–207 (1990).
This canonical study shows that 10-month-olds, but not 6-month-olds, discriminate between causal versus non-causal events.

83. Kotovsky, L. & Baillargeon, R. Calibration-based reasoning about collision events in 11-month-old infants. *Cognition* **51**, 107–129 (1994).
84. Kominsky, J. F. et al. Categories and constraints in causal perception. *Psychol. Sci.* **28**, 1649–1662 (2017).
85. Spelke, E. S., Breinlinger, K., Macomber, J. & Jacobson, K. Origins of knowledge. *Psychol. Rev.* **99**, 605 (1992).
86. Baillargeon, R. in *Language, Brain, and Cognitive Development: Essays in Honor of Jacques Mehler* (ed. Dupoux, E.) 341–361 (MIT Press, 2001).
87. Hespos, S. J. & Baillargeon, R. Infants' knowledge about occlusion and containment events: a surprising discrepancy. *Psychol. Sci.* **12**, 141–147 (2001).
88. Spelke, E. S. Principles of object perception. *Cogn. Sci.* **14**, 29–56 (1990).
89. Sobel, D. M. & Kirkham, N. Z. Blickets and babies: the development of causal reasoning in toddlers and infants. *Dev. Psychol.* **42**, 1103–1115 (2006).
90. Sobel, D. M. & Kirkham, N. Z. Bayes nets and babies: infants' developing statistical reasoning abilities and their representation of causal knowledge. *Dev. Sci.* **10**, 298–306 (2007).
91. Bell, S. M. & Ainsworth, M. D. S. Infant crying and maternal responsiveness. *Child Dev.* **43**, 1171–1190 (1972).
92. Jordan, G. J., Arbeau, K., McFarland, D., Ireland, K. & Richardson, A. Elimination communication contributes to a reduction in unexplained infant crying. *Med. Hypotheses* **142**, 109811 (2020).
93. Nakayama, H. Emergence of amae crying in early infancy as a possible social communication tool between infants and mothers. *Infant. Behav. Dev.* **40**, 122–130 (2015).
94. Meltzoff, A. N. & Moore, M. K. in *The Body and the Self* (eds Bermúdez, J. L., Marcel, A. J. & Eilan, N.) 3–69 (MIT Press, 1995).
95. Rovee, C. K. & Rovee, D. T. Conjugate reinforcement of infant exploratory behavior. *J. Exp. Child. Psychol.* **8**, 33–39 (1969).
96. Hillman, D. & Bruner, J. S. Infant sucking in response to variations in schedules of feeding reinforcement. *J. Exp. Child. Psychol.* **13**, 240–247 (1972).
97. DeCasper, A. J. & Spence, M. J. Prenatal maternal speech influences newborns' perception of speech sounds. *Infant. Behav. Dev.* **9**, 133–150 (1986).
98. Watson, J. S. & Ramey, C. T. Reactions to response-contingent stimulation in early infancy. *Merrill-Palmer Q. Behav. Dev.* **18**, 219–227 (1972).
99. Rovee-Collier, C. in *Handbook of Infant Development* 2nd edn (ed. Osofsky, J. D.) 98–148 (John Wiley & Sons, 1987).
100. Twitchell, T. E. The automatic grasping responses of infants. *Neuropsychologia* **3**, 247–259 (1965).
101. Wallace, P. S. & Whishaw, I. Q. Independent digit movements and precision grip patterns in 1–5-month-old human infants: hand-babbling, including vacuous then self-directed hand and digit movements, precedes targeted reaching. *Neuropsychologia* **41**, 1912–1918 (2003).
102. Von Hofsten, C. Mastering reaching and grasping: the development of manual skills in infancy. *Adv. Psychol.* **61**, 223–258 (1989).
103. Witherington, D. C. The development of prospective grasping control between 5 and 7 months: a longitudinal study. *Infancy* **7**, 143–161 (2005).
104. Needham, A., Barrett, T. & Peterman, K. A pick-me-up for infants' exploratory skills: early simulated experiences reaching for objects using 'sticky mittens' enhances young infants' object exploration skills. *Infant. Behav. Dev.* **25**, 279–295 (2002).
105. van den Berg, L. & Gredebäck, G. The sticky mittens paradigm: a critical appraisal of current results and explanations. *Dev. Sci.* **24**, e13036 (2021).
106. Keen, R. The development of problem solving in young children: a critical cognitive skill. *Annu. Rev. Psychol.* **62**, 1–21 (2011).
This paper provides an overview of the developmental trajectory of 'problem-solving'skills in young children, integrating findings from perception and motor development studies with cognitive problem-solving studies.
107. Claxton, L. J., McCarty, M. E. & Keen, R. Self-directed action affects planning in tool-use tasks with toddlers. *Infant. Behav. Dev.* **32**, 230–233 (2009).
108. McCarty, M. E., Clifton, R. K. & Collard, R. R. The beginnings of tool use by infants and toddlers. *Infancy* **2**, 233–256 (2001).
109. Gopnik, A. & Meltzoff, A. N. Semantic and cognitive development in 15- to 21-month-old children. *J. Child. Lang.* **11**, 495–513 (1984).
110. Gopnik, A. & Meltzoff, A. N. in *The Development of Word Meaning: Progress in Cognitive Development Research* (eds Kuczaj, S. A. & Barrett, M. D.) 199–223 (Springer, 1986).
111. Gopnik, A. & Meltzoff, A. N. *Words, Thoughts, and Theories* (Mit Press, 1997).
112. Tomasello, M. in *Early Social Cognition: Understanding Others in the First Months of Life* (ed. Rochat, P.) 301–314 (Lawrence Erlbaum Associates, 1999).
113. Tomasello, M. & Farrar, M. J. Joint attention and early language. *Child Dev.* **57**, 1454–1463 (1986).
114. Gopnik, A. Words and plans: early language and the development of intelligent action. *J. Child. Lang.* **9**, 303–318 (1982).
This paper proposes that language acquisition tracks with conceptual developments in infants' and toddlers' abilities in goal-directed action and planning.
115. Meltzoff, A. N. Infant imitation and memory: nine-month-olds in immediate and deferred tests. *Child. Dev.* **59**, 217–225 (1988).
116. Meltzoff, A. N. Infant imitation after a 1-week delay: long-term memory for novel acts and multiple stimuli. *Dev. Psychol.* **24**, 470–476 (1988).
117. Gergely, G., Bekkering, H. & Király, I. Rational imitation in preverbal infants. *Nature* **415**, 755–755 (2002).
118. Meltzoff, A. N., Waismeyer, A. & Gopnik, A. Learning about causes from people: observational causal learning in 24-month-old infants. *Dev. Psychol.* **48**, 1215–1228 (2012).
This study demonstrates that 2-year-old and 3-year-old children learn novel causal relations from observing other agents' interventions (observational causal learning).
119. Waismeyer, A., Meltzoff, A. N. & Gopnik, A. Causal learning from probabilistic events in 24-month-olds: an action measure. *Dev. Sci.* **18**, 175–182 (2015).
120. Stahl, A. E. & Feigenson, L. Observing the unexpected enhances infants' learning and exploration. *Science* **348**, 91–94 (2015).
This study demonstrates that 11-month-old children pay special visual and exploratory attention to objects that appear to violate the laws of physics as the result of an agent's intervention.
121. Perfors, A., Tenenbaum, J. B., Griffiths, T. L. & Xu, F. A tutorial introduction to Bayesian models of cognitive development. *Cognition* **120**, 302–321 (2011).
122. Gopnik, A. & Bonawitz, E. Bayesian models of child development. *Wiley Interdiscip. Rev. Cognit. Sci.* **6**, 75–86 (2015).
This paper is a technical introduction and tutorial in the Bayesian framework.
123. Gopnik, A., Sobel, D. M., Schulz, L. E. & Glymour, C. Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Dev. Psychol.* **37**, 620–629 (2001).
124. Schulz, L. E. & Bonawitz, E. B. Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Dev. Psychol.* **43**, 1045–1050 (2007).
125. Gopnik, A. & Sobel, D. M. Detecting blickets: how young children use information about novel causal powers in categorization and induction. *Child. Dev.* **71**, 1205–1222 (2000).
126. Schulz, L. E., Gopnik, A. & Glymour, C. Preschool children learn about causal structure from conditional interventions. *Dev. Sci.* **10**, 322–332 (2007).
127. Walker, C. M., Gopnik, A. & Ganea, P. A. Learning to learn from stories: children's developing sensitivity to the causal structure of fictional worlds. *Child. Dev.* **86**, 310–318 (2015).
128. Schulz, L. E., Bonawitz, E. B. & Griffiths, T. L. Can being scared cause tummy aches? Naive theories, ambiguous evidence, and preschoolers' causal inferences. *Dev. Psychol.* **43**, 1124–1139 (2007).
129. Kushnir, T. & Gopnik, A. Young children infer causal strength from probabilities and interventions. *Psychol. Sci.* **16**, 678–683 (2005).
130. Walker, C. M. & Gopnik, A. Toddlers infer higher-order relational principles in causal learning. *Psychol. Sci.* **25**, 161–169 (2014).
This paper shows that 18–30-month-old infants can learn relational causal rules and generalize them to novel stimuli.
131. Sobel, D. M., Yoachim, C. M., Gopnik, A., Meltzoff, A. N. & Blumenthal, E. J. The blicket within: preschoolers' inferences about insides and causes. *J. Cogn. Dev.* **8**, 159–182 (2007).
132. Schulz, L. E. & Sommerville, J. God does not play dice: causal determinism and preschoolers' causal inferences. *Child. Dev.* **77**, 427–442 (2006).
133. Schulz, L. E. & Gopnik, A. Causal learning across domains. *Dev. Psychol.* **40**, 162–176 (2004).
134. Seiver, E., Gopnik, A. & Goodman, N. D. 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 (2013).
135. Vasilyeva, N., Gopnik, A. & Lombrozo, T. The development of structural thinking about social categories. *Dev. Psychol.* **54**, 1735–1744 (2018).
136. Kushnir, T., Xu, F. & Wellman, H. M. Young children use statistical sampling to infer the preferences of other people. *Psychol. Sci.* **21**, 1134–1140 (2010).
137. Kushnir, T. & Gopnik, A. Conditional probability versus spatial contiguity in causal learning: preschoolers use new contingency evidence to overcome prior spatial assumptions. *Dev. Psychol.* **43**, 186–196 (2007).
138. Kimura, K. & Gopnik, A. Rational higher-order belief revision in young children. *Child. Dev.* **90**, 91–97 (2019).
139. Goddu, M. K. & Gopnik, A. Learning what to change: young children use "difference-making" to identify causally relevant variables. *Dev. Psychol.* **56**, 275–284 (2020).
140. Gopnik, A. et al. 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 (2017).
141. Lucas, C. G., Bridgers, S., Griffiths, T. L. & Gopnik, A. 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 (2014).
This paper shows that young children learn and generalize unusual causal relationships more readily than adults do.
142. Goddu, M. K., Lombrozo, T. & Gopnik, A. Transformations and transfer: preschool children understand abstract relations and reason analogically in a causal task. *Child. Dev.* **91**, 1898–1915 (2020).
143. Magid, R. W., Sheskin, M. & Schulz, L. E. Imagination and the generation of new ideas. *Cogn. Dev.* **34**, 99–110 (2015).
144. Liquin, E. G. & Gopnik, A. Children are more exploratory and learn more than adults in an approach–avoid task. *Cognition* **218**, 104940 (2022).
145. Erickson, J. E., Keil, F. C. & Lockhart, K. L. Sensing the coherence of biology in contrast to psychology: young children's use of causal relations to distinguish two foundational domains. *Child Dev.* **81**, 390–409 (2010).
146. Keil, F. C. *Concepts, Kinds, and Cognitive Development* (MIT Press, 1992).
147. Carey, S. *Conceptual Change in Childhood* (MIT Press, 1987).

148. Gelman, S. A. *The Essential Child: Origins of Essentialism in Everyday Thought* (Oxford Univ. Press, 2003).
149. Ahl, R. E., DeAngelis, E. & Keil, F. C. "I know it's complicated": children detect relevant information about object complexity. *J. Exp. Child. Psychol.* **222**, 105465 (2022).
150. Chuey, A. et al. No guts, no glory: underestimating the benefits of providing children with mechanistic details. *npj Sci. Learn.* **6**, 30 (2021).
151. Keil, F. C. & Lockhart, K. L. Beyond cause: the development of clockwork cognition. *Curr. Dir. Psychol. Sci.* **30**, 167–173 (2021).
152. Chuey, A., Lockhart, K., Sheskin, M. & Keil, F. Children and adults selectively generalize mechanistic knowledge. *Cognition* **199**, 104231 (2020).
153. Lockhart, K. L., Chuey, A., Kerr, S. & Keil, F. C. The privileged status of knowing mechanistic information: an early epistemic bias. *Child Dev.* **90**, 1772–1788 (2019).
154. Kominsky, J. F., Zamm, A. P. & Keil, F. C. Knowing when help is needed: a developing sense of causal complexity. *Cogn. Sci.* **42**, 491–523 (2018).
155. Mills, C. M. & Keil, F. C. Knowing the limits of one's understanding: the development of an awareness of an illusion of explanatory depth. *J. Exp. Child Psychol.* **87**, 1–32 (2004).
156. Goldwater, M. B. & Gentner, D. On the acquisition of abstract knowledge: structural alignment and explication in learning causal system categories. *Cognition* **137**, 137–153 (2015).
157. Rottman, B. M., Gentner, D. & Goldwater, M. B. Causal systems categories: differences in novice and expert categorization of causal phenomena. *Cogn. Sci.* **36**, 919–932 (2012).
158. Bonawitz, E. B. et al. Just do it? Investigating the gap between prediction and action in toddlers' causal inferences. *Cognition* **115**, 104–117 (2010).
This study demonstrates that the ability to infer causal relations from observations of correlational information without an agent's involvement or the use of causal language develops at around the age of four years.
159. Herrmann, E., Call, J., Hernández-Lloreda, M. V., Hare, B. & Tomasello, M. Humans have evolved specialized skills of social cognition: the cultural intelligence hypothesis. *Science* **317**, 1360–1366 (2007).
160. Tomasello, M. *Becoming Human: A Theory of Ontogeny* (Harvard Univ. Press, 2019).
161. Henrich, J. *The Secret of our Success* (Princeton Univ. Press, 2015).
162. Hesslow, G. in *Contemporary Science and Natural Explanation: Commonsense Conceptions of Causality* (ed. Hilton, D. J.) 11–32 (New York Univ. Press, 1988).
163. Woodward, J. The problem of variable choice. *Synthese* **193**, 1047–1072 (2016).
164. Khalid, S., Khalil, T. & Nasreen, S. A survey of feature selection and feature extraction techniques in machine learning. In *2014 Science and Information Conf*. 372–378 (IEEE, 1988).
165. Bonawitz, E., Denison, S., Griffiths, T. L. & Gopnik, A. Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends Cogn. Sci.* **18**, 497–500 (2014).
166. Bonawitz, E., Denison, S., Gopnik, A. & Griffiths, T. L. Win–stay, lose–sample: a simple sequential algorithm for approximating Bayesian inference. *Cogn. Psychol.* **74**, 35–65 (2014).
167. Denison, S., Bonawitz, E., Gopnik, A. & Griffiths, T. L. Rational variability in children's causal inferences: the sampling hypothesis. *Cognition* **126**, 285–300 (2013).
168. Samland, J., Josephs, M., Waldmann, M. R. & Rakoczy, H. The role of prescriptive norms and knowledge in children's and adults' causal selection. *J. Exp. Psychol. Gen.* **145**, 125–130 (2016).
169. Samland, J. & Waldmann, M. R. How prescriptive norms influence causal inferences. *Cognition* **156**, 164–176 (2016).
170. Phillips, J., Morris, A. & Cushman, F. How we know what not to think. *Trends Cogn. Sci.* **23**, 1026–1040 (2019).
171. Gureckis, T. M. & Markant, D. B. Self-directed learning: a cognitive and computational perspective. *Perspect. Psychol. Sci.* **7**, 464–481 (2012).
172. Saylor, M. & Ganea, P. *Active Learning from Infancy to Childhood* (Springer, 2018).
173. Goddu, M. K. & Gopnik, A. in *The Cambridge Handbook of Cognitive Development* (eds Houdé, O. & Borst, G.) 299–317 (Cambridge Univ. Press, 2022).
174. Gopnik, A. Scientific thinking in young children: theoretical advances, empirical research, and policy implications. *Science* **337**, 1623–1627 (2012).
175. Weisberg, D. S. & Sobel, D. M. *Constructing Science: Connecting Causal Reasoning to Scientific Thinking in Young Children* (MIT Press, 2022).
176. Xu, F. Towards a rational constructivist theory of cognitive development. *Psychol. Rev.* **126**, 841 (2019).
177. Xu, F. & Kushnir, T. Infants are rational constructivist learners. *Curr. Dir. Psychol. Sci.* **22**, 28–32 (2013).
178. Lapidow, E. & Bonawitz, E. What's in the box? Preschoolers consider ambiguity, expected value, and information for future decisions in explore-exploit tasks. *Open. Mind* **7**, 855–878 (2023).
179. Kidd, C., Piantadosi, S. T. & Aslin, R. N. The Goldilocks effect: human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS One* **7**, e36399 (2012).
180. Ruggeri, A., Swaboda, N., Sim, Z. L. & Gopnik, A. Shake it baby, but only when needed: preschoolers adapt their exploratory strategies to the information structure of the task. *Cognition* **193**, 104013 (2019).
181. Sim, Z. L. & Xu, F. Another look at looking time: surprise as rational statistical inference. *Top. Cogn. Sci.* **11**, 154–163 (2019).
182. Sim, Z. L. & Xu, F. Infants preferentially approach and explore the unexpected. *Br. J. Dev. Psychol.* **35**, 596–608 (2017).
183. Siegel, M. H., Magid, R. W., Pelz, M., Tenenbaum, J. B. & Schulz, L. E. Children's exploratory play tracks the discriminability of hypotheses. *Nat. Commun.* **12**, 3598 (2021).
184. Schulz, E., Wu, C. M., Ruggeri, A. & Meder, B. Searching for rewards like a child means less generalization and more directed exploration. *Psychol. Sci.* **30**, 1561–1572 (2019).
185. Schulz, L. Infants explore the unexpected. *Science* **348**, 42–43 (2015).
186. Perez, J. & Feigenson, L. Violations of expectation trigger infants to search for explanations. *Cognition* **218**, 104942 (2022).
187. Cook, C., Goodman, N. D. & Schulz, L. E. Where science starts: spontaneous experiments in preschoolers' exploratory play. *Cognition* **120**, 341–349 (2011).
This study demonstrates that preschoolers spontaneously perform causal interventions that are relevant to disambiguating multiple possible causal structures in their free play.
188. Lapidow, E. & Walker, C. M. Learners' causal intuitions explain behavior in control of variables tasks. *Dev. Psychol.* (in the press).
189. Lapidow, E. & Walker, C. M. Rethinking the "gap": self-directed learning in cognitive development and scientific reasoning. *Wiley Interdiscip. Rev. Cogn. Sci.* **13**, e1580 (2022).
This theory paper provides a complementary viewpoint to 'child-as-scientist', or Bayesian 'rational constructivist', account, arguing that children seek to identify and generate evidence for causal relations that are robust across contexts (and thus will be reliable for causal intervention).
190. Lapidow, E. & Walker, C. M. Informative experimentation in intuitive science: children select and learn from their own causal interventions. *Cognition* **201**, 104315 (2020).
191. Moeller, A., Sodian, B. & Sobel, D. M. Developmental trajectories in diagnostic reasoning: understanding data are confounded develops independently of choosing informative interventions to resolve confounded data. *Front. Psychol.* **13**, 800226 (2022).
192. Fernbach, P. M., Macris, D. M. & Sobel, D. M. Which one made it go? The emergence of diagnostic reasoning in preschoolers. *Cogn. Dev.* **27**, 39–53 (2012).
193. Buchanan, D. W. & Sobel, D. M. Mechanism-based causal reasoning in young children. *Child. Dev.* **82**, 2053–2066 (2011).
194. Sobel, D. M., Benton, D., Finiasz, Z., Taylor, Y. & Weisberg, D. S. The influence of children's first action when learning causal structure from exploratory play. *Cogn. Dev.* **63**, 101194 (2022).
195. Lapidow, E. & Walker, C. M. *The Search for Invariance: Repeated Positive Testing Serves the Goals of Causal Learning* (Springer, 2020).
196. Klayman, J. Varieties of confirmation bias. *Psychol. Learn. Motiv.* **32**, 385–418 (1995).
197. Zimmerman, C. The development of scientific thinking skills in elementary and middle school. *Dev. Rev.* **27**, 172–223 (2007).
198. Rule, J. S., Tenenbaum, J. B. & Piantadosi, S. T. The child as hacker. *Trends Cogn. Sci.* **24**, 900–915 (2020).
199. Burghardt, G. M. *The Genesis of Animal Play: Testing the Limits* (MIT Press, 2005).
200. Chu, J. & Schulz, L. E. Play, curiosity, and cognition. *Annu. Rev. Dev. Psychol.* **2**, 317–343 (2020).
201. Schulz, L. The origins of inquiry: inductive inference and exploration in early childhood. *Trends Cogn. Sci.* **16**, 382–389 (2012).
202. Harris, P. L., Kavanaugh, R. D., Wellman, H. M. & Hickling, A. K. in *Monographs of the Society for Research in Child Development* https://doi.org/10.2307/1166074 (Society for Research in Child Development, 1993).
203. Weisberg, D. S. in *The Oxford Handbook of the Development of Imagination* (ed. Taylor, M.) 75–93 (Oxford Univ. Press, 2013).
204. Gopnik, A. & Walker, C. M. Considering counterfactuals: the relationship between causal learning and pretend play. *Am. J. Play.* **6**, 15–28 (2013).
205. Weisberg, D. S. & Gopnik, A. Pretense, counterfactuals, and Bayesian causal models: why what is not real really matters. *Cogn. Sci.* **37**, 1368–1381 (2013).
206. Root-Bernstein, M. M. in *The Oxford Handbook of the Development of Imagination* (ed. Taylor, M.) 417–437 (Oxford Univ. Press, 2013).
207. Buchsbaum, D., Bridgers, S., Skolnick Weisberg, D. & Gopnik, A. The power of possibility: causal learning, counterfactual reasoning, and pretend play. *Phil. Trans. R. Soc. B* **367**, 2202–2212 (2012).
208. Wente, A., Gopnik, A., Fernández Flecha, M., Garcia, T. & Buchsbaum, D. Causal learning, counterfactual reasoning and pretend play: a cross-cultural comparison of Peruvian, mixed-and low-socioeconomic status US children. *Phil. Trans. R. Soc. B* **377**, 20210345 (2022).
209. Buchsbaum, D., Gopnik, A., Griffiths, T. L. & Shafto, P. Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition* **120**, 331–340 (2011).
210. Csibra, G. & Gergely, G. 'Obsessed with goals': functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychol.* **124**, 60–78 (2007).
211. Kelemen, D. The scope of teleological thinking in preschool children. *Cognition* **70**, 241–272 (1999).
212. Casler, K. & Kelemen, D. Young children's rapid learning about artifacts. *Dev. Sci.* **8**, 472–480 (2005).
213. Casler, K. & Kelemen, D. Reasoning about artifacts at 24 months: the developing teleo-functional stance. *Cognition* **103**, 120–130 (2007).
214. Ruiz, A. M. & Santos, L. R. 6. in *Tool Use in Animals: Cognition and Ecology* 119–133 (Cambridge Univ. Press, 2013).
215. Walker, C. M., Rett, A. & Bonawitz, E. Design drives discovery in causal learning. *Psychol. Sci.* **31**, 129–138 (2020).
216. Butler, L. P. & Markman, E. M. Finding the cause: verbal framing helps children extract causal evidence embedded in a complex scene. *J. Cogn. Dev.* **13**, 38–66 (2012).

217. Callanan, M. A. et al. Exploration, explanation, and parent–child interaction in museums. *Monogr. Soc. Res. Child. Dev.* **85**, 7–137 (2020).
218. McHugh, S. R., Callanan, M., Jaeger, G., Legare, C. H. & Sobel, D. M. Explaining and exploring the dynamics of parent–child interactions and children's causal reasoning at a children's museum exhibit. *Child Dev*. https://doi.org/10.1111/cdev.14035 (2023).
219. Sobel, D. M., Letourneau, S. M., Legare, C. H. & Callanan, M. Relations between parent–child interaction and children's engagement and learning at a museum exhibit about electric circuits. *Dev. Sci.* **24**, e13057 (2021).
220. Willard, A. K. et al. Explain this, explore that: a study of parent–child interaction in a children's museum. *Child Dev.* **90**, e598–e617 (2019).
221. Daubert, E. N., Yu, Y., Grados, M., Shafto, P. & Bonawitz, E. Pedagogical questions promote causal learning in preschoolers. *Sci. Rep.* **10**, 20700 (2020).
222. Yu, Y., Landrum, A. R., Bonawitz, E. & Shafto, P. Questioning supports effective transmission of knowledge and increased exploratory learning in pre-kindergarten children. *Dev. Sci.* **21**, e12696 (2018).
223. Walker, C. M. & Nyhout, A. in *The Questioning Child: Insights From Psychology and Education* (eds Butler, L. P., Ronfard, S. & Corriveau, K. H.) 252–280 (Cambridge Univ. Press, 2020).
224. Weisberg, D. S. & Hopkins, E. J. Preschoolers' extension and export of information from realistic and fantastical stories. *Infant. Child. Dev.* **29**, e2182 (2020).
225. Tillman, K. A. & Walker, C. M. You can't change the past: children's recognition of the causal asymmetry between past and future events. *Child. Dev.* **93**, 1270–1283 (2022).
226. Rottman, B. M., Kominsky, J. F. & Keil, F. C. Children use temporal cues to learn causal directionality. *Cogn. Sci.* **38**, 489–513 (2014).
227. Tecwyn, E. C., Mazumder, P. & Buchsbaum, D. One- and two-year-olds grasp that causes must precede their effects. *Dev. Psychol.* **59**, 1519–1531 (2023).
228. Liquin, E. G. & Lombrozo, T. Explanation-seeking curiosity in childhood. *Curr. Opin. Behav. Sci.* **35**, 14–20 (2020).
229. Mills, C. M., Legare, C. H., Bills, M. & Mejias, C. Preschoolers use questions as a tool to acquire knowledge from different sources. *J. Cogn. Dev.* **11**, 533–560 (2010).
230. Ruggeri, A., Sim, Z. L. & Xu, F. "Why is Toma late to school again?" Preschoolers identify the most informative questions. *Dev. Psychol.* **53**, 1620–1632 (2017).
231. Ruggeri, A. & Lombrozo, T. Children adapt their questions to achieve efficient search. *Cognition* **143**, 203–216 (2015).
232. Legare, C. H. & Lombrozo, T. Selective effects of explanation on learning during early childhood. *J. Exp. Child. Psychol.* **126**, 198–212 (2014).
233. Legare, C. H. Exploring explanation: explaining inconsistent evidence informs exploratory, hypothesis-testing behavior in young children. *Child Dev.* **83**, 173–185 (2012).
234. Walker, C. M., Lombrozo, T., Legare, C. H. & Gopnik, A. Explaining prompts children to privilege inductively rich properties. *Cognition* **133**, 343–357 (2014).
235. Vasil, N., Ruggeri, A. & Lombrozo, T. When and how children use explanations to guide generalizations. *Cogn. Dev.* **61**, 101144 (2022).
236. Walker, C. M., Bonawitz, E. & Lombrozo, T. Effects of explaining on children's preference for simpler hypotheses. *Psychon. Bull. Rev.* **24**, 1538–1547 (2017).
237. Walker, C. M. & Lombrozo, T. Explaining the moral of the story. *Cognition* **167**, 266–281 (2017).
238. Walker, C. M., Lombrozo, T., Williams, J. J., Rafferty, A. N. & Gopnik, A. Explaining constrains causal learning in childhood. *Child. Dev.* **88**, 229–246 (2017).
239. Gopnik, A. Childhood as a solution to explore–exploit tensions. *Phil. Trans. R. Soc. B* **375**, 20190502 (2020).
240. Wente, A. O. et al. Causal learning across culture and socioeconomic status. *Child. Dev.* **90**, 859–875 (2019).
241. Carstensen, A. et al. Context shapes early diversity in abstract thought. *Proc. Natl Acad. Sci. USA* **116**, 13891–13896 (2019).
This study provides evidence for developmentally early emerging cross-cultural differences in learning 'individual' versus 'relational' causal rules in children from individualist versus collectivist societies.
242. Ross, N., Medin, D., Coley, J. D. & Atran, S. Cultural and experiential differences in the development of folkbiological induction. *Cogn. Dev.* **18**, 25–47 (2003).
243. Inagaki, K. The effects of raising animals on children's biological knowledge. *Br. J. Dev. Psychol.* **8**, 119–129 (1990).
244. Cole, M. & Bruner, J. S. Cultural differences and inferences about psychological processes. *Am. Psychol.* **26**, 867–876 (1971).
245. Rogoff, B. & Morelli, G. Perspectives on children's development from cultural psychology. *Am. Psychol.* **44**, 343–348 (1989).
246. Rogoff, B. Adults and peers as agents of socialization: a highland Guatemalan profile. *Ethos* **9**, 18–36 (1981).
247. Shneidman, L., Gaskins, S. & Woodward, A. Child-directed teaching and social learning at 18 months of age: evidence from Yucatec Mayan and US infants. *Dev. Sci.* **19**, 372–381 (2016).
248. Callanan, M., Solis, G., Castañeda, C. & Jipson, J. in *The Questioning Child: Insights From Psychology and Education* (eds Butler, L. P., Ronfard, S. & Corriveau, K. H.) 73–88 (Cambridge Univ. Press, 2020).
249. Gauvain, M., Munroe, R. L. & Beebe, H. Children's questions in cross-cultural perspective: a four-culture study. *J. Cross-Cult. Psychol.* **44**, 1148–1165 (2013).
250. Adolph, K. E. & Hoch, J. E. Motor development: embodied, embedded, enculturated, and enabling. *Annu. Rev. Psychol.* **70**, 141–164 (2019).
251. Schleihauf, H., Herrmann, E., Fischer, J. & Engelmann, J. M. How children revise their beliefs in light of reasons. *Child. Dev.* **93**, 1072–1089 (2022).
252. Vasil, N. et al. Structural explanations lead young children and adults to rectify resource inequalities. *J. Exp. Child Psychol.* **242**, 105896 (2024).
253. Koskuba, K., Gerstenberg, T., Gordon, H., Lagnado, D. & Schlottmann, A. What's fair? How children assign reward to members of teams with differing causal structures. *Cognition* **177**, 234–248 (2018).
254. Bowlby, J. The Bowlby–Ainsworth attachment theory. *Behav. Brain Sci.* **2**, 637–638 (1979).
255. Tottenham, N., Shapiro, M., Flannery, J., Caldera, C. & Sullivan, R. M. Parental presence switches avoidance to attraction learning in children. *Nat. Hum. Behav.* **3**, 1070–1077 (2019).
256. Frankenhuis, W. E. & Gopnik, A. Early adversity and the development of explore–exploit tradeoffs. *Trends Cogn. Sci.* **27**, 616–630 (2023).
257. Van IJzendoorn, M. H. & Kroonenberg, P. M. Cross-cultural patterns of attachment: a meta-analysis of the strange situation. *Child Dev.* **59**, 147–156 (1988).
258. Gopnik, A. Explanation as orgasm. *Minds Mach.* **8**, 101–118 (1998).
259. Gottlieb, S., Keltner, D. & Lombrozo, T. Awe as a scientific emotion. *Cogn. Sci.* **42**, 2081–2094 (2018).
260. Valdesolo, P., Shtulman, A. & Baron, A. S. Science is awe-some: the emotional antecedents of science learning. *Emot. Rev.* **9**, 215–221 (2017).
261. Keil, F. C. *Wonder: Childhood and the Lifelong Love of Science* (MIT Press, 2022).
262. Perez, J. & Feigenson, L. Stable individual differences in infants' responses to violations of intuitive physics. *Proc. Natl Acad. Sci. USA* **118**, e2103805118 (2021).
263. Goddu, M. K., Sullivan, J. N. & Walker, C. M. Toddlers learn and flexibly apply multiple possibilities. *Child. Dev.* **92**, 2244–2251 (2021).
264. Cisek, P. Resynthesizing behavior through phylogenetic refinement. *Atten. Percept. Psychophys.* **81**, 2265–2287 (2019).
265. Pezzulo, G. & Cisek, P. Navigating the affordance landscape: feedback control as a process model of behavior and cognition. *Trends Cogn. Sci.* **20**, 414–424 (2016).
266. Cisek, P. Cortical mechanisms of action selection: the affordance competition hypothesis. *Phil. Trans. R. Soc. B* **362**, 1585–1599 (2007).
267. Tomasello, M. *The Evolution of Agency: Behavioral Organization From Lizards to Humans* (MIT Press, 2022).
268. Beck, S. R., Robinson, E. J., Carroll, D. J. & Apperly, I. A. Children's thinking about counterfactuals and future hypotheticals as possibilities. *Child. Dev.* **77**, 413–426 (2006).
269. Robinson, E. J., Rowley, M. G., Beck, S. R., Carroll, D. J. & Apperly, I. A. Children's sensitivity to their own relative ignorance: handling of possibilities under epistemic and physical uncertainty. *Child. Dev.* **77**, 1642–1655 (2006).
270. Leahy, B. P. & Carey, S. E. The acquisition of modal concepts. *Trends Cogn. Sci.* **24**, 65–78 (2020).
271. Mody, S. & Carey, S. The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition* **154**, 40–48 (2016).
272. Shtulman, A. & Carey, S. Improbable or impossible? How children reason about the possibility of extraordinary events. *Child Dev.* **78**, 1015–1032 (2007).
273. Redshaw, J. & Suddendorf, T. Children's and apes' preparatory responses to two mutually exclusive possibilities. *Curr. Biol.* **26**, 1758–1762 (2016).
274. Phillips, J. S. & Kratzer, A. Decomposing modal thought. *Psychol. Rev.* (in the press).
275. Vetter, B. Abilities and the epistemology of ordinary modality. *Mind* (in the press).
276. Kahneman, D. & Tversky, A. Variants of uncertainty. *Cognition* **11**, 143–157 (1982).
277. Rafetseder, E., Schwitalla, M. & Perner, J. Counterfactual reasoning: from childhood to adulthood. *J. Exp. Child Psychol.* **114**, 389–404 (2013).
278. Beck, S. R. & Riggs, K. J. Developing thoughts about what might have been. *Child Dev. Perspect.* **8**, 175–179 (2014).
279. Kominsky, J. F. et al. The trajectory of counterfactual simulation in development. *Dev. Psychol.* **57**, 253 (2021).
This study uses a physical collision paradigm to demonstrate that the content of children's counterfactual judgements changes over development.
280. Gerstenberg, T. What would have happened? Counterfactuals, hypotheticals and causal judgements. *Phil. Trans. R. Soc. B* **377**, 20210339 (2022).
281. Gerstenberg, T., Goodman, N. D., Lagnado, D. A. & Tenenbaum, J. B. A counterfactual simulation model of causal judgments for physical events. *Psychol. Rev.* **128**, 936 (2021).
282. Nyhout, A. & Ganea, P. A. The development of the counterfactual imagination. *Child Dev. Perspect.* **13**, 254–259 (2019).
283. Nyhout, A. & Ganea, P. A. Mature counterfactual reasoning in 4- and 5-year-olds. *Cognition* **183**, 57–66 (2019).
284. Moll, H., Meltzoff, A. N., Merzsch, K. & Tomasello, M. Taking versus confronting visual perspectives in preschool children. *Dev. Psychol.* **49**, 646–654 (2013).
285. Moll, H. & Tomasello, M. Three-year-olds understand appearance and reality — just not about the same object at the same time. *Dev. Psychol.* **48**, 1124–1132 (2012).
286. Moll, H. & Meltzoff, A. N. How does it look? Level 2 perspective-taking at 36 months of age. *Child. Dev.* **82**, 661–673 (2011).
287. Gopnik, A., Slaughter, V. & Meltzoff, A. in *Children's Early Understanding of Mind: Origins and Development* (eds Lewis, C. & Mitchell, P.) 157–181 (Routledge, 1994).
288. Gopnik, A. & Astington, J. W. Children's understanding of representational change and its relation to the understanding of false belief and the appearance–reality distinction. *Child Dev.* **59**, 26–37 (1988).
289. Doherty, M. & Perner, J. Metalinguistic awareness and theory of mind: just two words for the same thing? *Cogn. Dev.* **13**, 279–305 (1998).
290. Wimmer, H. & Perner, J. Beliefs about beliefs: representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* **13**, 103–128 (1983).

291. Flavell, J. H., Speer, J. R., Green, F. L., August, D. L. & Whitehurst, G. J. in *Monographs of the Society for Research in Child Development* 1–65 (Society for Research in Child Development, 1981).
292. Wellman, H. M., Cross, D. & Watson, J. Meta-analysis of theory-of-mind development: the truth about false belief. *Child. Dev.* **72**, 655–684 (2001).
293. Kelemen, D. & DiYanni, C. Intuitions about origins: purpose and intelligent design in children's reasoning about nature. *J. Cogn. Dev.* **6**, 3–31 (2005).
294. Kelemen, D. Why are rocks pointy? Children's preference for teleological explanations of th natural world. *Dev. Psychol.* **35**, 1440 (1999).
295. Vihvelin, K. *Causes, Laws, and Free Will: Why Determinism Doesn't Matter* (Oxford Univ. Press, 2013).
296. Yiu, E., Kosoy, E. & Gopnik, A. Transmission versus truth, imitation versus innovation: what children can do that large language and language-and-vision models cannot (yet). *Persp. Psychol. Sci*. https://doi.org/10.1177/17456916231201401 (2023).
297. Kosoy, E. et al. Towards understanding how machines can learn causal overhypotheses. Preprint at *arXiv* https://doi.org/10.48550/arXiv.2206.08353 (2022).
298. Frank, M. C. Baby steps in evaluating the capacities of large language models. *Nat. Rev. Psychol.* **2**, 451–452 (2023).
299. Frank, M. C. Bridging the data gap between children and large language models. *Trends Cogn. Sci*. **27**, https://doi.org/10.1016/j.tics.2023.08.007 (2023).
300. Schmidhuber, J. A possibility for implementing curiosity and boredom in model-building neural controllers. In *Proc. Int. Conf. on Simulation of Adaptive Behavior: From Animals to Animats* 222–227 (MIT Press, 1991).
301. Volpi, N. C. & Polani, D. Goal-directed empowerment: combining intrinsic motivation and task-oriented behaviour. *IEEE Trans. Cogn. Dev. Syst.* **15**, 361–372 (2020).
302. Salge, C., Glackin, C. & Polani, D. in *Guided Self-Organization: Inception. Emergence, Complexity and Computation* Vol. 9 (ed. Prokopenko, M.) 67–114 (Springer, 2014).
303. Klyubin, A. S., Polani, D. & Nehaniv, C. L. Empowerment: a universal agent-centric measure of control. In *2005 IEEE Congr. on Evolutionary Computation* 128–135 (IEEE, 2005).
304. Gopnik, A. Empowerment as causal learning, causal learning as empowerment: a bridge between Bayesian causal hypothesis testing and reinforcement learning. Preprint at https://philsci-archive.pitt.edu/id/eprint/23268 (2024).
305. Rovee-Collier, C. K., Sullivan, M. W., Enright, M., Lucas, D. & Fagen, J. W. Reactivation of infant memory. *Science* **208**, 1159–1161 (1980).
306. Kominsky, J. F., Li, Y. & Carey, S. Infants' attributions of insides and animacy in causal interactions. *Cogn. Sci.* **46**, e13087 (2022).
307. Lakusta, L. & Carey, S. Twelve-month-old infants' encoding of goal and source paths in agentive and non-agentive motion events. *Lang. Learn. Dev.* **11**, 152–175 (2015).
308. Saxe, R., Tzelnic, T. & Carey, S. Knowing who dunnit: infants identify the causal agent in an unseen causal interaction. *Dev. Psychol.* **43**, 149–158 (2007).
309. Saxe, R., Tenenbaum, J. & Carey, S. Secret agents: inferences about hidden causes by 10- and 12-month-old infants. *Psychol. Sci.* **16**, 995–1001 (2005).
310. Liu, S., Brooks, N. B. & Spelke, E. S. Origins of the concepts cause, cost, and goal in prereaching infants. *Proc. Natl Acad. Sci. USA* **116**, 17747–17752 (2019).
311. Liu, S. & Spelke, E. S. Six-month-old infants expect agents to minimize the cost of their actions. *Cognition* **160**, 35–42 (2017).
312. Nyhout, A. & Ganea, P. A. What is and what never should have been: children's causal and counterfactual judgments about the same events. *J. Exp. Child. Psychol.* **192**, 104773 (2020).
313. Legare, C. H. The contributions of explanation and exploration to children's scientific reasoning. *Child. Dev. Perspect.* **8**, 101–106 (2014).
314. Denison, S. & Xu, F. Twelve-to 14-month-old infants can predict single-event probability with large set sizes. *Dev. Sci.* **13**, 798–803 (2010).
315. Denison, S., Reed, C. & Xu, F. The emergence of probabilistic reasoning in very young infants: evidence from 4.5-and 6-month-olds. *Dev. Psychol.* **49**, 243 (2013).
316. Alderete, S. & Xu, F. Three-year-old children's reasoning about possibilities. *Cognition* **237**, 105472 (2023).
317. Scholl, B. J. & Tremoulet, P. D. Perceptual causality and animacy. *Trends Cogn. Sci.* **4**, 299–309 (2000).
318. Hood, B., Carey, S. & Prasada, S. Predicting the outcomes of physical events: two-year-olds fail to reveal knowledge of solidity and support. *Child. Dev.* **71**, 1540–1554 (2000).
319. Hood, B. M., Hauser, M. D., Anderson, L. & Santos, L. Gravity biases in a non-human primate? *Dev. Sci.* **2**, 35–41 (1999).
320. Hood, B. M. Gravity does rule for falling events. *Dev. Sci.* **1**, 59–63 (1998).
321. Hood, B. M. Gravity rules for 2-to 4-year olds? *Cognit. Dev.* **10**, 577–598 (1995).
322. Woodward, A. L. Infants selectively encode the goal object of an actor's reach. *Cognition* **69**, 1–34 (1998).
323. Woodward, A. L. Infants' grasp of others' intentions. *Curr. Dir. Psychol. Sci.* **18**, 53–57 (2009).
324. Woodward, A. L., Sommerville, J. A., Gerson, S., Henderson, A. M. & Buresh, J. The emergence of intention attribution in infancy. *Psychol. Learn. Motiv.* **51**, 187–222 (2009).
325. Meltzoff, A. N. 'Like me': a foundation for social cognition. *Dev. Sci.* **10**, 126–134 (2007).
326. Gerson, S. A. & Woodward, A. L. The joint role of trained, untrained, and observed actions at the origins of goal recognition. *Infant. Behav. Dev.* **37**, 94–104 (2014).
327. Gerson, S. A. & Woodward, A. L. Learning from their own actions: the unique effect of producing actions on infants' action understanding. *Child. Dev.* **85**, 264–277 (2014).
328. Sommerville, J. A., Woodward, A. L. & Needham, A. Action experience alters 3-month-old infants' perception of others' actions. *Cognition* **96**, B1–B11 (2005).
329. Liu, S. & Almeida, M. Knowing before doing: review and mega-analysis of action understanding in prereaching infants. *Psychol. Bull.* **149**, 294–310 (2023).

Acknowledgements
The authors acknowledge their funding sources: the Alexander von Humboldt Foundation, the Templeton World Charity Foundation (0434), the John Templeton Foundation (6145), the Defense Advanced Research Projects Agency (047498-002 Machine Common Sense), the Department of Defense Multidisciplinary University Initiative (Self-Learning perception Through Real World Interaction), and the Canadian Institute for Advanced Research Catalyst Award. For helpful discussions, comments, and other forms of support, the authors thank: S. Boardman, E. Bonawitz, B. Brast-McKie, D. Buchsbaum, M. Deigan, J. Engelmann, T. Friend, T. Gerstenberg, S. Kikkert, A. Kratzer, E. Lapidow, B. Leahy, T. Lombrozo, J. Phillips, H. Rakoczy, L. Schulz, D. Sobel, E. Spelke, H. Steward, B. Vetter, M. Waldmann and E. Yiu.

Author contributions
M.K.G. wrote the article and created the figures. A.G. contributed substantially to discussion of the content and to multiple revisions. A.G. and M.K.G. together reviewed and edited the manuscript before submission.

Competing interests
The authors declare no competing interests.

Additional information
Peer review information *Nature Reviews Psychology* thanks Jonathan Kominsky and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.