Article type: Advanced Review

Article title: Bayesian Models of Child Development

First author: Alison Gopnik*; University of California, Berkeley; gopnik@berkeley.edu
Second author: Elizabeth Bonawitz, Rutgers University, Newark

Abstract

[Bayesian models have been applied to many areas of cognitive science including vision, language, and motor learning. We discuss the implications of this framework for cognitive development. We first present a brief introduction to the Bayesian framework. Bayesian models make assumptions about representation explicit, and provide a detailed account of learning. Furthermore, they can provide an account of developmental transitions and other phenomena in development, such as curiosity and exploration. Drawing on recent work bridging empirical developmental data and modelling, we show that these features of the Bayesian approach provide solutions to problems that elude traditional accounts of learning and raise new areas of investigation.]

Somehow every child solves the baffling and intractable philosophical problem of induction. Children take the plethora of ambiguous information coming in through their senses and turn it into meaningful, abstract, structured representations. Despite centuries of philosophical thought and empirical study, we still don't fully understand how this is possible. How can children's mental representations both support wide-ranging new inferences and change in the light of experience?

This fundamental problem has led to a tension in cognitive science. How can we learn abstract representations from concrete data? The classic nativist response to this question is to say that learning is an illusion. We come equipped with abstract representations. The alternative empiricist response is to say that abstraction is an illusion. Knowledge is just a collection of associations derived from the statistics of our environment.

Nether of these options seems completely right to most empirically-minded developmental psychologists. That's because we see evidence for both abstract representation and

sophisticated learning even in infancy and early childhood. Piaget proposed “constructivism” as a way to bridge the nativist/empiricist divide. Piagetians appealed to an ever-enriching developmental process that employed mechanisms of accommodation and assimilation.

The “theory theory”^[1-4], which was a theoretical offspring of constructivism, was also intended to be an alternative to the nativist and empiricist extremes. It proposed that children’s beliefs are rich, structured and abstract but *defeasible* representations. Even if children are equipped with innate (or rapidly developed) “starting-state” theories, those theories can always be changed in the light of new evidence.

However, both the Piagetian and theory theory accounts essentially restate the problem. They fail to provide a precise account that describes the representational details of children’s beliefs and specifies the mechanisms that support learning.

In the last ten years or so, however, probabilistic models have begun to promise a more precise solution to the nativist-empiricist tension. These approaches are based on ideas from the philosophy of science, machine learning, and artificial intelligence. The probabilistic modelling approach starts by formally describing structured, generative representations of the world. These representations can take many forms including causal graphical models, taxonomies, and logical grammars. But the crucial point is that these models mathematically generate patterns of observable data. Probabilistic models make predictions about the kinds of data you should expect to see, given the generative structure of the model. These generative structures provide a natural characterization of abstract mental representations that allow inferences.

But how could an agent *learn* the structure from the data in the first place? How does a learner who starts with data figure out which structure generated that data? This poses a particularly sticky problem because there are often many possible hypotheses that are compatible with the data – the structure is underdetermined. So how do we know which is the correct hypothesis?

One clever solution is to apply Bayes rule. Bayes rule is a formula for moving backwards from data to structure. Bayes rule tells us how to evaluate the probability of a particular structure, given some observation of data. It does this by combining two pieces of information. First the learner starts with a set of beliefs about which structures are likely, prior to observing any data; this is called the prior. Then the learner considers the probability that each structure would have generated the observed data; this is called the likelihood. Combining the prior and likelihood gives a quantity proportional to the probability of the structure given the data; this is called the posterior. Determining the probability of a particular structure given the data is what learning is all about,

In Bayesian learning, learners can use their existing highly structured knowledge to inform their prior beliefs and likelihoods. But at the same time new data can change that knowledge. In this way, Bayes rule provides a middle ground between classical nativist and empiricist accounts. As in the classical constructivist accounts a learner’s current beliefs will influence her interpretation of the data, but new data can also lead to changes in those beliefs.

The basic idea that cognitive models can generate predictions about data, and that we can invert that process to learn those models from observations, is not new. In fact, it's essential to classic cognitive science accounts of vision and language. The advance in Bayesian models is to integrate probability theory with this basic approach. This leads to a new solution to the under-determination problem. Although we may not ever know for certain which hypothesis is correct, we can still know which hypothesis is most probable given the data. So the learner might consider multiple possible hypotheses, adjusting which hypothesis is more or less likely as new evidence accumulates.

However the general Bayesian framework is just a starting point for understanding cognitive development. To apply the framework to any particular kind of learning much more is necessary—we must specify the generative representations and learning methods in detail. There are numerous kinds of learning that contribute to development, and conceptual development is the one that we focus on here, though the Bayesian approach has been applied to many domains of development including grammar learning^[5] and perception^[6,7].

This approach gives us a way of formalizing a central and productive, though informal, line of research in cognitive development. Quite typically, developmental psychologists explain the behaviour of infants and children by assuming that they have beliefs about plants or animals, words or people – and that those beliefs underpin what they say or do. Similarly, researchers explain developmental changes by assuming that children's beliefs are transformed in the light of evidence. Earlier accounts also informally sketched how prior structured knowledge could constrain children's beliefs and inductions. These inductive constraints, could take the form of framework theories^[1-4], core knowledge^[e.g. 8], or rules for particular domains of learning (e.g. the whole object constraint in word learning^[9]). Bayesian models are a natural way to make these kinds of informal explanations more precise.

Recently, Bayesian models of cognition have been criticised on a few important grounds. This includes concerns about models being underconstrained and failing to make connection to processes^[e.g. 10], concerns that the overall framework is unfalsifiable, and concerns that people do not always behave optimally^[11].

Many of these concerns reflect problems that come from confounding different levels of analysis. One is the distinction between a *framework* and a *model*. Frameworks are high-level approaches to representing a problem, such as connectionist models, production systems, or generative grammars. Frameworks themselves do not make quantitative predictions and they can typically be used to accommodate many patterns of data – hence the falsifiability concerns. Given any particular pattern of data we could, in principle, construct some connectionist model, production rule, generative grammar or Bayesian model to explain that data.

Instead, frameworks of this kind generate a space of possible specific models – a particular Bayesian, connectionist or grammatical model. The specific assumptions of a model of a particular phenomenon are falsifiable, and critically help inform us about the nature of the

learners representations, biases, and mechanisms for belief revision. If the specific models provide good explanations and predictions about phenomena then the framework is useful.

A second distinction concerns computational or normative versus algorithmic or descriptive models. Probabilistic models provide a computational level account of belief changes and the constraints on learning^[12]. The goal is to describe what the problem is and what the correct solution should be. Thus, Bayesian models provide a story about how powerful childhood learning might be possible^[13,14]. They also give us a normative benchmark, one we can compare to the learning we actually see in children. And this in turn can inform accounts at a process or algorithmic level.

Making Representations Explicit

In order to describe the problem that a learner is solving, the variables of interest must be made explicit. One advantage of the modelling approach is that it forces us to make representations precise and explicit. By itself, Bayes rule is too general to say much. To derive meaningful, quantitative predictions, Bayesian modelling requires explicit details describing the learner's representations and beliefs, the processes by which those representations generate data patterns, and the "priors" – that is, the initial inductive constraints on those representations.

Individual models can thus help us answer specific questions about development. For example, do children follow the statistical assumptions of causality to infer whether one variable was adequately screened-off by another? Does a taxonomic representation of animals help explain children's assumptions about when to extend the meaning of a novel word from one animal to another? Do children attend to the source of information – both how observed events came to be and whether the source of information is reliable?

One way to understand the problem of learning and to ask these questions is to start with the formal tools of modelling, which provide the in principle account of what assumptions are needed in order to solve the problem in a particular way, given statistical learning mechanisms. This is partly what is meant when we say the models operate at the computational level. It is not necessarily a claim that children are carrying-out exact Bayesian inference (a point we turn to in "From computations to algorithms"), but rather a description of the tools a learner might need in order to carry out inference given a particular set of evidence.

Causal Graphical Models

Causal graphical models or "Bayes-nets" were one of the earliest types of representations to be employed in Bayesian models of cognitive development^[15-19]. Causal graphical models formally represent complex causal relationships and generate predictions about the patterns of events those relationships produce. In particular, they allow complex predictions about the correlations among causally related variables. They also allow predictions about "interventions". They allow you to predict what will happen to other variables when you actively alter one variable.

Early studies showed that preschool children accurately inferred causal structures from patterns of correlation and intervention in the way that these models predicted. Children could use these techniques to infer complex structures involving relationships among three variables, discriminating, for example, between common causes, common effects and causal chains.^[19] In some circumstances they could even infer unobserved invisible causes^[18,20-21].

Further studies demonstrated that children's inferences depend on the combined strength of their prior beliefs and the data. For example, children integrate base rate information with new data to make sophisticated causal inferences^[22, 23]. Similarly, Kushnir and Gopnik^[24] and Schulz and Gopnik^[25] showed that children use additional causal factors, such as spatio-temporal and domain-specific prior beliefs, to inform their causal inference following patterns of statistical data. These results demonstrate that, consistent with a general Bayesian framework, children combine prior knowledge and new evidence in sophisticated ways to inform their causal judgements.

Taxonomies

Taxonomies are another representational scheme that has been used in Bayesian models of children's inferences. Earlier research suggested that children have a taxonomic bias in word learning, assuming that kind labels map onto taxonomic categories^[e.g. 26]. Xu and Tenenbaum^[27] modelled children's inferences about the likely meaning of a word with a Bayesian model of hierarchically structured categories. Preschoolers were given a few examples of an item at different levels of a taxonomic hierarchy and were asked which other objects the term applied to. The results were consistent with the Bayesian model's behavior, but not with other models of word learning, such as a purely associative (statistical model) or a purely deductive (constraint-based) model. Bayesian models may thus provide a promising common ground to explain children's fast mapping of words to meanings.

In particular, Xu and Tenenbaum's model demonstrated that a taxonomic representation could be used to produce quantitative predictions about the likelihood of different possible extensions of a concept. Given that the predictions that derived from this structure closely matched children's generalizations, this provides support for the claim that children represent these categories taxonomically. Thus, the Bayesian model both provides a story about how rapid learning may be possible, and also makes explicit the likely representations underlying this learning.

Hierarchical Models

Causal graphical models and taxonomies operate at only one level of abstraction. Griffiths and Tenenbaum^[28] proposed a technique for describing and learning hierarchical Bayesian models. These models include more abstract meta-representations of the structure of possible hypotheses such as, for example, the fact that causal relationships are deterministic or indeterministic, or that they have different logical or relational structures. Hierarchical models enable you to represent what the philosopher Nelson Goodman called "over-hypotheses" – that

is hypotheses about what specific hypotheses will be like^[29]. Such representations are a natural way of representing the kinds of “core knowledge” “framework theories” or “constraints” that have been proposed to constrain children’s inferences.

For example, Schulz, Bonawitz, and Griffiths^[30] developed a Bayesian model of children’s cross-domain causal reasoning, such as inferring that psychological anxieties could cause physical illness. Children’s hypotheses were represented using causal graphical models (hypotheses) that captured the various potential causes in the story (e.g. eating cheese, being worried) and the effect (e.g. Bunny getting a tummy ache). The probability of different hypotheses (e.g. a graph where cheese causes tummy aches, but worrying does not versus a graph where worrying causes tummy aches but cheese does not) were given by a framework theory. The framework theory captured general principles about the probability of causes leading to effects within and cross-domains. In this way, the framework theory helped guide the probability of any particular hypothesis being correct, as well as specifying the likelihoods of the data observed in the story, given each particular hypothesis.

Four-year-olds inferences from ambiguous (but informative) statistical evidence corresponded strongly with the model, though younger children failed to learn from the evidence when it conflicted with their strong prior beliefs. A follow-up training study^[31], suggested that both broad prior beliefs and the ability to learn from the statistical evidence in these contexts were responsible for younger children’s failure in the original task.

Hierarchical causal models provide a detailed account of the relationship between theory and evidence in children’s causal reasoning. They also provide an explanation of conflicting findings suggesting that children privilege domain-specific causal knowledge on the one hand or domain-general causal and statistical learning mechanisms on the other.

Logical Grammars

In principle models of learning act as a starting point for age-old developmental questions about what minimal descriptions are necessary for how to get learning off the ground. Thus, they potentially inform our understanding of likely innate constraints, as well as speaking to the learning mechanisms that are required given those constraints. For example, more recently Bayesian models have proposed even more abstract representational structures. Logical grammars have been proposed as a possible broad “language of thought”: other more specific representations can be encoded in this broader language^[32]. For example, Kemp et al^[33] showed how such languages might be used to capture the content in intuitive theories as well as tell a story about how intuitive theories might be learned. Logical grammars have been extended to show how a theory of causality might be learnable^[34]. Causal graphical models, then, would be seen as a specific instance of the more general class of logical grammars.

Empirically, Bonawitz et al^[35, 36] explored preschooler’s solution to a “chicken-and-egg” problem in the domain of magnetism. Causal models assume that we can first specify the causal categories that are being considered and then establish the relations between them. But it is

often hard to say which comes first: learning that objects belong to causal categories or understanding the causal relationships between those categories. Bonawitz et al^[35, 36] extended a logical grammar model from Ullman, Goodman, and Tenenbaum^[37] to solve this problem. Preschoolers were presented with two different simplified magnet learning tasks which required simultaneous inferences about causal laws (e.g. repulsion vs. attraction) and causal categories (e.g. metals vs. magnets). Children were able to solve the problem in a basically rational way – integrating multiple pieces of evidence across different phases of the experiment and abstractly inferring both the correct number of categories and the laws that related those categories.

Two different hierarchical models were tested against children's responses. The first generative model included a bias for producing stick relations among possible logical clauses (which dictated whether categories of objects should stick or repel to each other). The second model did not incorporate this bias. Both models could provide an in principle solution to the chicken-and-egg problem in the domain of magnetism, but the stick bias model captured reflected the pattern of responses generated by the children, suggesting that children might share a similar inductive constraint.

Making Sampling Assumptions Explicit

Another important component of Bayesian models is making explicit assumptions about how the generative model produces the data. In particular, the models can specify whether the data you see is a random or representative sample of the possible data. Imagine visiting a foreign country and trying to learn a new word. The strength of the inferences you can draw about that word's meaning will depend on the context. If you only observe three Dalmatians and your informant tells you that all three are "gavagais", you may be uncertain about whether the label applies only to "Dalmatians" or to "dogs" in general. However, if your informant purposely labels only the three Dalmatians from a broader set of dogs, it will be more likely that the label applies only to Dalmatians. (If the label applied to dogs in general, a helpful teacher would have selected a broader set of examples to label for you.)

Bayesian models make these sampling assumptions explicit. For example, Xu and Tenenbaum^[27] modelled their word learning tasks as following the *strong sampling* assumption: observations were generated by a knowledgeable informant who was assumed to sample randomly from within the space of consistent data. In a different set of studies, Xu and Tenenbaum^[38] showed how these assumptions can be manipulated and how this manipulation influences both model predictions and human behavior. Models (and children) will make different predictions if they believe that the informant is sampling non-randomly.

Models that specify sampling assumptions can also inform inferences about the properties of both objects and people. For example, Gweon, Tenenbaum, and Schulz^[39] showed fifteen-month-old infants' sets of objects and provided cues about whether the objects were drawn randomly or purposely from a box – infants were more likely to assume that a property of the drawn object (squeaking) applied to other objects in the randomly drawn object condition.

Kushnir et al.^[40] showed that 20-month-olds could infer desires and preferences from sampling patterns – they assumed that when people picked objects from a population in a non-random way they preferred those objects. These results suggest that even infants and toddlers are sensitive to details of the generative process that gave rise to the data.

The models and the behavioral data work together to inform our understanding of children's early inferences about others. The models helped clarify the potential variables and sampling assumptions that are otherwise implicit in the problem, and they inform experimental manipulations that demonstrate children are likewise dependent on these assumptions in their own early, sophisticated inferences.

Data are also sometimes generated by a teacher, who intentionally chooses an ideal and representative set of data for a learner. Shafto and Goodman^[41] use a pedagogical Bayesian model to describe how these sampling assumptions can allow learners to make even stronger inferences. These Bayesian pedagogical models are consistent with preschooler's inferences following pedagogical cues^[e.g. see 42]. Other studies suggest that preschooler's exploratory play causal inferences, and imitation^[43,44] are sensitive to this further subtle information about how data were generated.

Integrating multiple sources of information

One of the other tensions in developmental psychology stems from the fact that children seem to use many different sources of information in making inferences. For example, there is evidence that children use perceptual, statistical and socio-cultural information in their inferences about word meaning. This has led to theoretical battles about which kinds of information are most important. One of the advantages of probabilistic models is that they provide a natural way to integrate multiple kinds of data, and also predict the contributing role of these information sources in different contexts. For example, both statistical and social information might independently lend probabilistic weight to one hypothesis rather than another. Moreover, joint inferences can be described in which multiple hypotheses and their interactions can be considered simultaneously. In particular, more recent Bayesian models incorporate rich information about the social world and provide a better account of children's inferences in social contexts.

For example, to explain how children bridge social and causal inferences, Shafto et al.^[45] developed a model of epistemic trust as an explanatory account for three- and four-year-olds' behavior on trust tasks. Unlike previous accounts of preschooler's trust reasoning, this model simultaneously infers an informant's knowledge, intent, and the true state of the world and it does a better job of capturing preschoolers' behavior. Furthermore, the model predicted that developmental changes between three- and four-years of age stemmed from changing beliefs about helpfulness. Shafto et al.'s^[45] computational model explained how a learner might be able to simultaneously make inferences about the informant's trustworthiness and the true state of the world. (See also^[46].)

Bayesian models can also integrate other information. For example, an actor's choice of objects can be combined with information about the properties of those objects. Models with these components make predictions about an actor's future novel object preference^[e.g., 47]. Lucas et al's^[47] joint inference model captures developmental data demonstrating that children are sensitive to the sampling population when they determine preferences^[40] as well as to the degree of property overlap between objects when they decide whether to extend preference generalization^[48, 49]. A single model can integrate the results of what appear to be very different developmental investigations of preference learning.

Models that integrate different kinds of inferences have also informed our understanding of social reasoning in infants. The Bayesian inverse planning framework explored by Hamlin et al^[50] helps explain how a rational observer might reason about the mental states of an actor. The model makes a few assumptions: agents rationally plan actions given their goals and beliefs and there are different classes of agents (helpers or hinderers) whose goals are either complementary or contradictory to another agent's goals. The model explains how the same actions can lead to different judgments about the agent because inferences depend on combining information about the agent's beliefs, knowledge, and goals. Hamlin et al^[50] found that infants' behavior fit the rational inverse planning model. Furthermore, the model provided a better account of infants' performance than accounts suggesting that infants rely solely on perceptual cues.

Transitions in Development

So far we have discussed how Bayesian models inform our understanding of the representational frameworks, sampling assumptions, and rich mutual inferences that children are able to make. We have shown that when children are given a particular kind of data, they draw conclusions that are consistent with those representations, assumptions and inferences.

So we can use specific models to characterize the content of children's knowledge at different stages of development. But, importantly, modelling can also explain the mechanisms that are responsible for transitions between stages. For example, Bayesian models naturally capture the trade-off between simplicity and goodness-of-fit that often drives cognitive change^[51]. The likelihood term in the Bayesian model will always prefer the more specific hypothesis; that is the data that fit the hypothesis best. However, less specific hypotheses are more likely in general¹,

¹ Less specific hypotheses may be more likely for a number of reasons. For example, models with more variables will be more complex and there will be a larger space of possible variants in this space of more complex models. There could be costs associated with a framework theory generating a model with a greater number of variables. These costs need not be arbitrary, but may fall out naturally, as the probability of any particular model must decrease to account for a growing number of these variants.

and so will be weighted more heavily in the prior, and will be preferred if data are relatively sparse. Scientific theorizing invokes a preference for simplicity, as in the use of Occam's razor. Recent developmental data suggests that children do too. For example, preschoolers prefer explanations with fewer causal variables^[52].

This Bayesian Occam's razor can capture classic developmental transitions in several domains. For example, Goodman et al^[53] showed how preschooler's false-belief understanding can be described as a transition between two causal models. The earlier "naive realist model" is simpler than the "knower model" and so is initially preferred. The knower model is more complex and thus has greater explanatory power and comes to be more probable as more data accumulates. Goodman et al^[53] predicted a transitional asymmetry following surprising evidence and this corresponded with preschooler's false belief performance².

In other work, Lucas et al^[55] developed a Bayesian model that captures another developmental transition. Empirical studies suggest that toddlers initially believe that people all share common preferences but eventually learn that individuals can have different preferences^[56]. The Lucas et al. model^[55] depends on the fact that the shared-preference belief is more parsimonious than the different-preference belief and so is initially favored. However, with experience, data eventually favor the more complex model.

Gerken^[57, 58] has suggested that infants also show this trade-off between simplicity and evidence when they learn linguistic rules. Kemp and Tenenbaum^[59] have demonstrated how other radical developmental shifts (e.g. a shift from a simpler cluster model of animal organization to a more complex tree model) can naturally fall out of a Bayesian model as data are acquired.

Hierarchical Bayesian models and the blessing of abstraction

Like classic "constraint theories" hierarchical Bayesian models put limits on what children will infer from data and so help to solve the under-determination problem. Unlike the constraints in such theories however, these higher-order constraints can themselves be learned from data in a Bayesian way. In fact, several recent empirical studies show that even infants can infer "overhypotheses" as well as more specific causal and taxonomic hypotheses^[60-63]. Learning such abstract overhypotheses could help account for some of the large qualitative conceptual changes we see in development.

Tenenbaum et al^[64] present a computational story of how one might be able to 'grow a mind', by virtue of this hierarchical machinery in Bayesian models. They describe several case studies in which abstract knowledge can be rapidly inferred from remarkably little data. In fact, Goodman

² The Bayesian account of the false belief transition has inspired additional developmental studies. For example, the additional variable in Goodman et al's^[53] more complex knower model can be made salient to children at the threshold of false belief understanding, thus improving their performance on false-belief tasks^[54].

et al^[34], present cases in which inference at these higher levels of abstraction may *precede* inferences at the lower level. They call this the “Blessing of Abstraction”. The Blessing of Abstraction naturally falls out because each additional degree of freedom at higher levels of abstraction receives evidence from all the variables at each level below.

This is in contrast to both traditional nativist and empiricist accounts which assume that learning more abstract knowledge depends on first learning more concrete kinds of knowledge. In fact, nativist arguments often rest on the fact that abstract generalizations are in place very early. But the blessing of abstraction is consistent with developmental data suggesting that children may sometimes learn abstract rules earlier than more concrete ones.

Other studies describe how hierarchical inference can lead to developmental leaps. Piantadosi, Tenenbaum, and Goodman^[65] developed a hierarchical Bayesian model to account for numerical development. Their model can account for the inductive leap young children make when they transition from understanding a few number words, to the rich system that affords adult-like numerical understanding. A similar approach has been used to explain children’s acquisition of quantifier semantics^[66]. Lucas and Griffiths^[67] have developed hierarchical models that describe how learners might infer abstract causal principals, and they have applied this model to explain developmental differences in learning these forms^[62-63]. Seiver, Gopnik & Goodman showed how children could infer abstract social concepts such as concepts of traits, from specific behavioural data^[68].

Considering additional questions in developmental psychology

Another benefit of the general Bayesian framework is that it provides a possible story about when young learners should be interested in exploring or attending to a particular variable. Bonawitz et al^[69], suggest that children may be curious when two (or more) competing explanations for the data are approximately equal, either because the evidence fails to distinguish the plausible hypotheses or because the evidence is strongly consistent with a weakly held belief and simultaneously inconsistent with a strongly held prior belief^[see also 70]. Work by Schulz and colleagues support this idea, showing that children choose to explore in cases where the prior beliefs and evidence interact in a way that leads competing hypotheses to be roughly equivalent^[69, 71-73].

For example, research by Bonawitz, van Schijndel, Friel, and Schulz^[69] examined children’s exploratory play, explanation, and learning in the domain of balance understanding. Children were first given a test that characterized their stage of balance understanding. Younger children typically are “center theorists” believing that objects always balance at their geometrical center, while older children are “mass theorists” recognizing that the distribution of weight of the object has to be considered. Then children were shown a block that balanced in a way that was either consistent or inconsistent with their prior theory. Children were given the opportunity to explore the block. They preferred to play with the block when the evidence contradicted their beliefs. Because evidence that was surprising to children with one theory was consistent to

children with the other theory, these results demonstrated the importance of considering both the effects of evidence and of prior beliefs. After they had played, children were asked to explain the balance event. Again, and consistent with general predictions of the Bayesian framework, children's explanations depended on both their prior beliefs and the evidence that they observed.

On a similar theme, Bayesian analysis has been used to help explain infant looking time results^[74]. A Bayesian ideal observer model would predict that optimal learning occurs for material that is not too simple (already learned) or too complex (unknowable). Kidd et al's^[74] results suggest that infants prefer stimuli that are moderately complex (predictable) given a set of probabilistic expectations. Additional Bayesian analyses revealed that these results hold for individual infants^[75], and ongoing work suggests extension in infant auditory cognition^[76]. This application of a Bayesian model may help resolve long standing concerns regarding the interpretation of results from habituation and preferential looking paradigms.

From computations to algorithms

Most of the Bayesian models of cognitive development have functioned at the computational level^[12] of analysis. They characterize representations, and they demonstrate how those representations may change with learning.

However, these models don't describe how the mind carries out these inferences in detail nor do they require that the learner need necessarily compute the Bayesian analytical solution. Indeed, a major drawback of the Bayesian approach is the vast space of possible hypotheses to be considered – how could a learner actually enumerate and evaluate each one in real time? Answering this question is one of the key problems for future work.

There are *approximation* algorithms that can produce behavior that looks Bayesian on aggregate, but that does not require that learner is actually carrying out Bayesian inference over the (potentially) vast space of possible hypotheses. In machine learning the solutions to this problem involve sampling just a few hypotheses at a time, randomly selected from a probability distribution, and then testing those hypotheses against the data. (These sampling *algorithms* should not be confused with the previously discussed sampling *assumptions* built into Bayesian models). These algorithms can be shown to converge to ideal Bayesian computational solutions in the long run, but they are much more computationally tractable. They can also provide a rational account of adult behavior^[e.g. 77, 78].

There is evidence that children also “sample” hypotheses from a probability distribution in this way^[79]. Denison et al. showed children a box full of red and blue chips in different proportions, and asked them to guess the color of a chip invisibly selected at random, often several times. Children's responses were variable: the same child would sometimes say red and sometimes say blue. However, the proportion of “red” or “blue” responses closely tracked the probability of the

relevant hypotheses, children said “red” more often when that was more likely to be the correct answer. This is a signature of sampling.

There are lots of ways in which a learner could sample hypotheses. The simplest idea is that each time a learner observes new data, she recomputes the updated posterior and samples a guess from that updated distribution. This kind of approach to updating predicts that subsequent guesses from a single learner will be independent. That is, knowing that a learner prefers a specific hypothesis at a particular time tells you nothing about what hypothesis he is likely to have after the next observation of data. Another possibility is that a learner tends to maintain a hypothesis that makes a successful prediction and only tries a new hypothesis when the data weigh against the original choice – a “win-stay: lose-shift” strategy. This means that an individual will tend towards “stickiness” – she will be more likely to keep the current hypothesis, and this will lead to dependency between responses. A learner following either algorithm would appear to randomly vary from one hypothesis to the next, but importantly both algorithms share the property that behavior on aggregate produces responses consistent with an optimal Bayesian model.

To explore these algorithms, Bonawitz, Denison, Gopnik, and Griffiths^[80] provided children and adults with a “mini-microgenetic” experiment, in which learners were presented with new data gradually and trial by trial asked about their beliefs after each new presentation of evidence. By comparing the subsequent hypotheses generated by each participant trial-by-trial to the predictions of these two algorithms, Bonawitz et al^[80] were able to show that children and adults produce dependencies that are the signature of the win-stay, lose-shift algorithm. (See also^[81] for a review.)

Once again this approach not only shows how it is possible for children to solve inductive problems, it also illuminates a classic developmental problem -- the variability that is characteristic of children’s belief revision. There is substantial variability in children’s responses, and children often entertain multiple hypotheses and strategies at once^[e.g. 82]. This variability was one of the factors that originally led Piaget to describe young children’s behavior as irrational. Indeed, such findings have led some researchers to suggest that children’s behavior is always intrinsically variable and context-dependent^[e.g. 83-85]. Other researchers assume that this variability is simply the result of extraneous factors like noise or information-processing limitations.

The sampling hypothesis provides a rational way of explaining this variability. If children implicitly sample from a distribution of hypotheses then we would expect their responses to be variable, and yet also to reflect the probability of different beliefs.

The idea that children sample from hypotheses and search through possibilities also suggests some interesting developmental hypotheses. Lucas et al^[62] suggest that this search may take place in different ways at different developmental periods. Younger children may search in a

more exploratory way while older children and adults search in a more constrained way, and this may explain developmental differences.

Conclusion

Bayesian modelling can provide new insights into old problems in cognitive development. It is important to emphasize, however, that Bayesianism is a very general approach to development that must be instantiated in particular models in order to provide testable hypotheses --it is a broad framework theory which allows researchers to construct much more specific particular theories of children's beliefs and learning. A principal advantage of the Bayesian framework is that it allows those particular theories to be phrased in precise and transparent ways.

Bayesian models can precisely characterize the representations that children use in different domains at different stages of development. They can also provide precise accounts of learning and of developmental transitions. Nonetheless, these models have only begun to scratch the surface. Much work must still be done to specify the content of children's representations in any particular domain, and the changes in those representations.

As we've seen Bayesian models can be made more complex to solve ever more complex inference problems, like inferring abstract over-hypotheses from concrete data, or simultaneously integrating hypotheses about teachers, objects and word meanings. But as the models grow the size of the potential hypothesis spaces grows too. Algorithmic models have only just begun to address the question of how a learner might be able to search through such spaces in real time. Furthermore, connecting these algorithms to the brain remain an important challenge, although growing evidence suggests that Bayesian approaches may be relevant at the neural level^[e.g. 86-89].

Bayesian models are not appropriate for every question in development. But they are particularly well-designed to address the questions of induction that are at the heart of many issues in cognitive development. By applying this approach to children's cognitive development, we may better understand how even very young children develop rich, abstract knowledge about the world.

References

- 1 Carey, S. (1985). *Conceptual change in childhood*. Cambridge, MA: MIT Press/Bradford Books.
- 2 Gopnik, A., & Meltzoff, A. (1997). *Words, thoughts and theories*. Cambridge, MA: MIT Press.
- 3 Murphy, G., & Medin, D. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.

- 4 Wellman, H., & Gelman, S. (1992) Cognitive development: Foundational theories of core domains. *Annual Review of Psychology*, 43, 337-375.
- 5 Perfors, A., Tenenbaum, J. B., & Regier, T. (2011). The learnability of abstract syntactic principles. *Cognition*, 118(3), 306-338.
- 6 Gori M, Del Viva M, Sandini G, Burr D.C. (2008) Young children do not integrate visual and haptic form information. *Curr. Biol.* **18**, 694-698. [doi:10.1016/j.cub.2008.04.036](https://doi.org/10.1016/j.cub.2008.04.036).
- 7 Nardini M, Jones P, Bedford R, Braddick O (2008) Development of cue integration in human navigation. *Curr. Biol.* **18**, 689–693. [doi:10.1016/j.cub.2008.04.021](https://doi.org/10.1016/j.cub.2008.04.021)
- 8 Carey, S., & Spelke, E. (1996). Science and core knowledge. *Philosophy of science*, 515-533.
- 9 Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, 14(1), 57-77.
- 10 Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34(04), 169-188.
- 11 Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological bulletin*, 138(3), 389.
- 12 Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*, Henry Holt and Co. Inc., New York, NY.
- 13 Gopnik, A. (2012). Scientific thinking in young children: Theoretical advances, empirical research, and policy implications. *Science*, 337(6102), 1623-1627.
- 14 Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, 138(6), 1085.
- 15 Pearl, J. (2000). *Causality: models, reasoning and inference* (Vol. 29). Cambridge: MIT press.
- 16 Glymour, C. (2003). Learning, prediction and causal Bayes nets. *Trends in cognitive sciences*, 7(1), 43-48.
- 17 Gopnik, A., Sobel, D. M., Schulz, L. E., & Glymour, C. (2001). Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental psychology*, 37(5), 620.
- 18 Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: causal maps and Bayes nets. *Psychological review*, 111(1), 3.
- 19 Schulz, L. E., Gopnik, A., & Glymour, C. (2007). Preschool children learn about causal structure from conditional interventions. *Developmental Science*, 10(3), 322-332.

- 20 Schulz, L. E., Goodman, N. D., Tenenbaum, J. B., & Jenkins, A. C. (2008). Going beyond the evidence: Abstract laws and preschoolers' responses to anomalous data. *Cognition*, 109(2), 211-223.
- 21 Schulz, L. E., & Sommerville, J. (2006). God does not play dice: Causal determinism and preschoolers' causal inferences. *Child development*, 77(2), 427-442.
- 22 Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive science*, 28(3), 303-333.
- 23 Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and blickets: Effects of knowledge on causal induction in children and adults. *Cognitive science*, 35(8), 1407-1455.
- 24 Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, 43(1), 186.
- 25 Schulz, L. E., & Gopnik, A. (2004). Causal learning across domains. *Developmental psychology*, 40(2), 162.
- 26 Markman, E. M. (1991). The whole-object, taxonomic, and mutual exclusivity assumptions as initial constraints on word meanings. *Perspectives on language and thought: Interrelations in development*, 72-106.
- 27 Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological review*, 114(2), 245.
- 28 Tenenbaum, J. B., Griffiths, T. L., & Niyogi, S. (2007). Intuitive theories as grammars for causal inference. *Causal learning: Psychology, philosophy, and computation*, 301-322.
- 29 Goodman, N. (1955). *Fact, Fiction and Forecast*. Bobbs-Merrill. Indianapolis, IN.
- 30 Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared cause tummy aches? Naive theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental psychology*, 43(5), 1124.
- 31 Bonawitz, E., Fischer, A., & Schulz, L. (2012). Teaching 3.5-year-olds to revise their beliefs given ambiguous evidence. *Journal of Cognition and Development*, 13(2), 266-280.
- 32 Kemp, C., Goodman, N. D., & Tenenbaum, J. B. (2010). Learning to learn causal models. *Cognitive Science*, 34(7), 1185-1243.
- 33 Kemp, C., Tenenbaum, J. B., Niyogi, S., & Griffiths, T. L. (2010). A probabilistic model of theory formation. *Cognition*, 114(2), 165-196.

- 34 Goodman, N. D., Ullman, T. D., & Tenenbaum, J. B. (2011). Learning a theory of causality. *Psychological review*, 118(1), 110.
- 35 Bonawitz, E., Ullman, T., Gopnik, A., & Tenenbaum, J. (2012, November). Sticking to the Evidence? A computational and behavioral case study of micro-theory change in the domain of magnetism. In *Development and Learning and Epigenetic Robotics (ICDL), 2012 IEEE International Conference on* (pp. 1-6). IEEE.
- 36 Bonawitz, E., Ullman, T., Gopnik, A., & Tenenbaum, J.B. (*in review*) Sticking to the evidence? A computational and behavioral case study of micro-theory change in the domain of magnetism.
- 37 Ullman, T., Goodman, N. D., & Tenenbaum, J.B. (2012), Theory learning as stochastic search in the language of thought. *Cognitive Development*.
- 38 Xu, F., & Tenenbaum, J. B. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental science*, 10(3), 288-297.
- 39 Gweon, H., Tenenbaum, J. B., & Schulz, L. E. (2010). Infants consider both the sample and the sampling process in inductive generalization. *Proceedings of the National Academy of Sciences*, 107(20), 9066-9071.
- 40 Kushnir, T., Xu, F., & Wellman, H. M. (2010). Young children use statistical sampling to infer the preferences of other people. *Psychological Science*, 21(8), 1134-1140.
- 41 Shafto, P., & Goodman, N. (2008). Teaching games: Statistical sampling assumptions for learning in pedagogical situations. In *Proceedings of the thirtieth annual conference of the cognitive science society* (pp. 1632-1637). Austin, TX: Cognitive Science Society.
- 42 Csibra, G., & Gergely, G. (2006). Social learning and social cognition: The case for pedagogy. *Processes of change in brain and cognitive development. Attention and performance XXI*, 249-274.
- 43 Bonawitz, E., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Instruction limits spontaneous exploration and discovery. *Cognition*, 120(3), 322-330.
- 44 Buchsbaum, D., Gopnik, A., Griffiths, T. L., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, 120(3), 331-340.
- 45 Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental science*, 15(3), 436-447.
- 46 Sobel, D.M., & Kushnir, T. (2013) Knowledge matters: How children evaluate the reliability of testimony as a process of rational inference. *Psychological Review*, Vol 120(4), 779-797.

- 47 Lucas, C., Gopnik, A., & Griffiths, T. L. (2010). Developmental differences in learning the forms of causal relationships. In *Proceedings of the 32nd annual conference of the Cognitive Science Society* (pp. 28-52).
- 48 Fawcett, C. A., & Markson, L. (2010). Children reason about shared preferences. *Developmental psychology*, 46(2), 299.
- 49 Fawcett, C. A., & Markson, L. (2010). Similarity predicts liking in 3-year-old children. *Journal of experimental child psychology*, 105(4), 345-358.
- 50 Kiley Hamlin, J., Ullman, T., Tenenbaum, J., Goodman, N., & Baker, C. (2013). The mentalistic basis of core social cognition: experiments in preverbal infants and a computational model. *Developmental science*, 16(2), 209-226.
- 51 Jefferys, W. H., & Berger, J. O. (1992). Ockham's razor and Bayesian analysis. *American Scientist*, 80(1), 64-72.
- 52 Bonawitz, E. B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental psychology*, 48(4), 1156.
- 53 Goodman, N. D., Baker, C. L., Bonawitz, E. B., Mansinghka, V. K., Gopnik, A., Wellman, H., ... & Tenenbaum, J. B. (2006, July). Intuitive theories of mind: A rational approach to false belief. In *Proceedings of the twenty-eighth annual conference of the cognitive science society* (pp. 1382-1387).
- 54 Pham, K., Bonawitz, E., & Gopnik, A. (2012) Seeing who sees: Contrastive access helps children reason about other minds. *Proceedings of the Thirty-fourth Cognitive Science Society*.
- 55 Lucas, C. G., Griffiths, T. L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., Hu, J., & Markson, L. (under revision). The child as econometrician: A rational model of preference understanding in children. *PLoS ONE*.
- 56 Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: evidence from 14- and 18-month-olds. *Developmental psychology*, 33(1), 12.
- 57 Gerken, L. (2006). Decisions, decisions: Infant language learning when multiple generalizations are possible. *Cognition*, 98(3), B67-B74.
- 58 Gerken, L. (2010). Infants use rational decision criteria for choosing among models of their input. *Cognition*, 115(2), 362-366.
- 59 Kemp, C., & Tenenbaum, J. B. (2008). The discovery of structural form. *Proceedings of the National Academy of Sciences*, 105(31), 10687-10692.

- 60 Dewar, K. M., & Xu, F. (2010). Induction, Overhypothesis, and the Origin of Abstract Knowledge Evidence From 9-Month-Old Infants. *Psychological Science*, 21(12), 1871-1877.
- 61 Walker, C. M., & Gopnik, A. (2014). Toddlers Infer Higher-Order Relational Principles in Causal Learning. *Psychological science*, 0956797613502983.
- 62 Lucas, C. G., Gopnik, A., Griffiths, T., & Bridgers, S. (2014). When children are better(or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*.
- 63 Gopnik, A., Griffiths, T., & Lucas, C. (in press) When young learners can be better (or at least more open-minded) than older ones. *Current Directions in Psychological Science*.
- 64 Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022), 1279-1285.
- 65 Piantadosi, S. T., Tenenbaum, J. B., & Goodman, N. D. (2012). Bootstrapping in a language of thought: A formal model of numerical concept learning. *Cognition*, 123(2), 199-217.
- 66 Piantadosi, S., Goodman, N. & Tenenbaum, J.B. (in revision) Modeling the acquisition of quantifier semantics: a case study in function word learnability.
- 67 Lucas, C. G., & Griffiths, T. L. (2010). Learning the form of causal relationships using hierarchical Bayesian models. *Cognitive Science*, 34, 113-147.
- 68 Seiver, E., Gopnik, A., & Goodman, N. D. (2013). Did she jump because she was the big sister or because the trampoline was safe? Causal inference and the development of social attribution. *Child development*, 84(2), 443-454.
- 69 Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive psychology*, 64(4), 215-234.
- 70 Bonawitz, E. B., & Schulz, L. (2007). Children's Rational Exploration. In *Proc. AAAI*.
- 71 Cook, C., Goodman, N. D., & Schulz, L. E. (2011). Where science starts: Spontaneous experiments in preschoolers' exploratory play. *Cognition*, 120(3), 341-349.
- 72 Gweon, H., & Schulz, L. (2008). Stretching to learn: Ambiguous evidence and variability in preschoolers' exploratory play. In *Proceedings of the 30th annual meeting of the Cognitive Science Society* (pp. 570-574).
- 73 Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental psychology*, 43(4), 1045.

- 74 Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PLoS One*, 7(5), e36399.
- 75 Piantadosi, S.T., Kidd, C. & Aslin, R.N. (*in press*) Rich analysis and rational models: Inferring individual behavior from infant looking data. *Developmental Science*.
- 76 Kidd, C, Piantadosi, S.T., & Aslin, R.N. (*in review*). The Goldilocks effect in infant auditory cognition.
- 77 Vul, E., & Pashler, H. (2008). Measuring the crowd within probabilistic representations within individuals. *Psychological Science*, 19(7), 645-647.
- 78 Bonawitz, E. B., & Griffiths, T. L. (2010). Deconfounding hypothesis generation and evaluation in Bayesian models. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2260-2265).
- 79 Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2012). Rational variability in children's causal inferences: The Sampling Hypothesis. *Cognition*.
- 80 Bonawitz, E., Denison, S., Gopnik, A., & Griffiths, T. L. (2014). Win-Stay, Lose-Sample: A simple sequential algorithm for approximating Bayesian inference. *Cognitive psychology*, 74, 35-65.
- 81 Bonawitz, E., Denison, S., Griffiths, T. L., & Gopnik, A. (2014). Probabilistic models, learning algorithms, and response variability: sampling in cognitive development. *Trends in cognitive sciences*, 18, 497-500.
- 82 Siegler, R. S. (1996). *Emerging minds: The process of change in children's thinking*. New York: Oxford University Press
- 83 Greeno, J. G. (1998). The situativity of knowing, learning, and research. *American psychologist*, 53(1), 5.
- 84 Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge university press.
- 85 Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA. MIT Press.
- 86 Doya, K. (Ed.). (2007). *Bayesian brain: Probabilistic approaches to neural coding*. MIT Press.
- 87 Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature neuroscience*, 9(11), 1432-1438.

88 Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual review of neuroscience*, 26(1), 381-410.

89 Rao, R. P. (2004). Bayesian computation in recurrent neural circuits. *Neural computation*, 16(1), 1-38.

Further Reading/Resources

Perfors, A., Tenenbaum, J.B., Griffiths, T.L., & Xu, F. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 120, 302-321.

Related Articles

Article ID	Article title
COGSCI-220	Rational (Bayesian) Models of Cognition
COGSCI-013	Bayesian and Computational Learning Theory
COGSCI-313	Cognitive Development