In H. Geffner, R. Dechter, and J.Y. Halpern (Eds.), Probabilistic and Causal Inference: The Works of Judea Pearl, Association for Computing Machinery, 2021.

30

Causal Models and Cognitive Development

Alison Gopnik (University of California at Berkeley)

Pearl's work has had an important influence on the field of cognitive development. In particular, in hundreds of empirical studies, causal models, combining ideas about probability, intervention, and counterfactuals, have turned out to play an essential role in children's everyday knowledge. Even very young children learn such models from data in the ways that Pearl suggested. New frontiers in the project of understanding children's causal learning include sampling, active search and experimentation, and combining causal models with deep learning and deep reinforcement learning techniques.

In the year 2000, more than 20 years ago, my graduate students and I made a weekly trek across the campus and up the hill to the computer science department. We were there as part of a reading group discussing a brand-new book, *Causality* by Judea Pearl. Those students went on to become distinguished faculty, and 20 years later, they and *their* students, and many other psychologists, are still working on problems that were inspired by that book and those conversations. So am I.

Why would developmental psychologists, usually found sitting in tiny chairs opposite toddlers in preschools, immerse themselves in a volume of statistics and equations? The book, and Pearl's work, in general, speaks to a foundational problem that is at the core of the study of cognitive development. Cognitive development and machine learning belong to the same natural category, along with the philosophy of science, epistemology, and vision science, even if they live in opposite corners of the campus. (And all these disciplines are in a different natural category than sociologically closer ones like adult cognitive psychology, cognitive neuroscience, and philosophy of mind.)

Developmental psychology, machine learning, and philosophy of science might seem like strange bedfellows, but they are all trying to solve the same problem—sometimes called the problem of induction. How can we know anything about the world around us? After all, the information that reaches us from that world is just a stream of photons at our retinas and disturbances of air at our ear drums. And yet we come to know about people and poodles, tables and toys, quarks and quasars. How is this possible? We seem to have abstract, hierarchical, structured representations of the world around us, and those representations allow us to make wide-ranging generalizations and predictions. And yet, we also seem to somehow construct those representations from data that is concrete, messy, and particular.

Going back to Plato and Aristotle, there have been two basic approaches to solving this problem. The nativist option is simply to deny that the abstract representations *are* derived from the data. Instead, they are there innately, from a past life or in the mind of God, for Plato and Descartes, because of evolution for more recent thinkers. The other, empiricist, option is to deny that the abstract representations exist—simply combine enough statistical data and you can do all the same inferential work. This approach goes all the way back to Aristotle and Locke but also underpins many of the most recent approaches to machine learning.

For people who actually study the development of human knowledge, whether as developmental psychologists or philosophers of science, these alternatives have always seemed unsatisfying. When we actually look in detail at the progress of children's thinking, or the progress of science, we do, in fact, see both abstract representations and qualitative changes in those representations in the light of new evidence.

In the past, Jean Piaget, the great founder of cognitive development, argued for "constructivism" as an alternative to nativism and empiricism, and philosophers like Carnap and Kuhn, who were actually both influenced by developmental psychology as well as the history of science, articulated similar ideas. In the 1980s, a number of psychologists including me, Susan Carey, Henry Wellman, and Susan Gelman, articulated the "theory theory"—the idea that children's conceptual development could be understood by analogy to scientific theory formation, explicitly connecting conceptual development and scientific theory change [Carey 1985, Gopnik 1988, Wellman and Gelman 1992]. "Theory theory" researchers could qualitatively describe children's representations as theories and chart the changes in those representations as children learned more. The research program made a great deal of empirical progress, describing the development of intuitive physics [Smith et al. 1985], biology [Carey 1985] and especially intuitive psychology or "theory of mind" [Gopnik and Wellman 1994].

The problem, though, was that there was no computational way of characterizing the constructive process that was responsible for theory formation and change either in childhood or in science. The overarching faith of cognitive science is that the mind is a computational system instantiated in the brain. In some areas of cognitive science, particularly vision science, we really had begun to redeem that faith and solve the problem of induction computationally and even neurally. Building on 100 years of perception and psychophysics, vision scientists could begin to describe how the visual system recovers information about objects and space from the light patterns on the retina, and computer vision systems could start to instantiate those ideas (e.g., Marr [1982]). There has been remarkable progress on this project since; although, of course, there is still much work to be done.

But doing the same thing for theories, whether these were the everyday theories of childhood or the theories of formal science, seemed like an impossibly forbidding task. Indeed, in the early 1990s there was a kind of nihilism about solving such problems, reflected in both the philosophy of science on the one hand, and in statistics, on the other hand. The slogans of the time were "there is a logic of confirmation but no logic of discovery" and "no causation from correlation."

This was where Pearl's work came in. Although theories involve many kinds of representations, certainly causal representations are crucial, both in everyday cognition and in science. And for a long time, going back at least to David Hume, causal knowledge was one of the canonical cases of the problem of induction. As Hume pointed out, it seemed impossible to see how simply observing the constant conjunction of two events could lead you to a causal conclusion, and yet such inferences are ubiquitous in both everyday cognition and in science. The pessimism about scientific induction very much extended to causation. Bertrand Russell famously said, "The law of causality, I believe, like much that passes muster among philosophers, is a relic of a bygone age, surviving, like the monarchy, only because it is erroneously supposed to do no harm."

In the 1990s, two developments coming from very different directions restored the reign of causality and articulated a computational account of causal inference in the form of "causal Bayes nets." One set of developments came from Pearl's initial work on expert reasoning [Pearl 1988]. Initially, Pearl's project was to find a way a computer could generate complex judgments and predictions about conditional dependencies in the way that experts, like doctors, do (if, but only if, the patient has a fever and green phlegm as well as a cough, and tests for viruses are negative, antibiotics will help). It turned out that the best way to do this was to equip the system with causal models, integrating ideas about probability, intervention, and counterfactual inference. In parallel, the philosophers of science Peter Spirtes, Richard Scheines, and Clark Glymour at Carnegie Mellon University formulated very similar mathematical ideas [Spirtes et al. 1993]. Moreover, philosophers like

James Woodward, working in a more traditional philosophical framework, used these ideas to characterize the very nature of causation [Woodward 2003].

The central idea was to use graphical models as a way of representing the relations among variables in a causal system, and to systematically relate those relations to the conditional probability of the variables. Within this system it was possible to define the effects of interventions (in what Pearl called the "do-calculus") as well as counterfactuals. The distinction between associations and predictions, on the one hand, and interventions and counterfactuals on the other hand, is the crucial distinction that separates mere correlation from causation. I might notice correlations both between having yellow nicotine-stained fingers and getting lung cancer, and between smoking and getting lung cancer, and make the appropriate prediction that someone who has yellow fingers or who smokes is more likely to have cancer. But an intervention to wash the yellow off a patient's fingers won't have any effect on the probability of cancer, while an intervention to stop smoking will. Similarly, the counterfactual that if the patient had washed his hands he wouldn't have gotten cancer is false while the similar counterfactual about smoking is true.

The entire causal Bayes net system allowed for interwoven inferences about probability, intervention, and counterfactuals in a way that captures many of the central elements of causation both in everyday life and in science. If you knew the causal structure you could make accurate predictions, interventions, and counterfactual inferences, and significantly, the formalism naturally distinguished between these different kinds of inferences.

The Bayes net formalism also had important implications for causal learning and the problem of causal induction. The formalism made systematic connections between the structure of the causal graphs and data about the conditional probability of events and the outcomes of interventions. This meant that in principle the inferences could be reversed—if we knew about the conditional probability of variables and the outcomes of interventions on those variables, we could accurately infer the causal structure. And this, in turn, suggested a computational solution to the classic problem of theory induction. Scientists had always used evidence from statistics (i.e., patterns of conditional probability) and experiments (i.e., systematic interventions on variables) to infer causal structure. But thanks to the new work on causal Bayes nets we could begin to explain mathematically how and why this actually worked.

Initially, nobody thought of these systems as potential models of everyday human cognition, let alone children's cognition. (I have an email exchange with Clark Glymour from 1989, where I suggested children might be doing something similar to Bayes nets, his initial response was that these systems were precisely designed to do things that humans couldn't). By the late 1990s though, this idea had come to seem more appealing, and at least worth testing. Even if children couldn't make inferences about complex systems with hundreds of variables, would they use the same basic principles to uncover causal structure?

At the time essentially all of the work of children's causal reasoning, and adults', for that matter, fell into one of two camps, either researchers who emphasized the role of reasoning about physical mechanisms in causal understanding, or those who saw causal reasoning as merely an extension of simple association. The combination of graphical models, probability, intervention, and counterfactuals was an entirely new way of approaching the subject.

Glymour and I decided to test whether children might do something like causal Bayes net inference with a new method—the blicket detector—a machine that lights up and plays music when you put some things on it and not others. The first question, which we tested with my student David Sobel, one of the participants in the *Causality* reading group and now at Brown, was whether children could make any causal inferences with this method (they could) [Gopnik and Sobel 2000]. By 2000, we realized that we could use simple methods like this to test more complex inferences, of the sort that Pearl described. In particular, could children use conditional probability and intervention to make inferences? (they could) [Gopnik et al. 2001]. After one of the reading group meetings, my student Laura Schulz, now at the Massachusetts Institute of Technology, raced excitedly down the hill from the computer science department to the hardware store, where she constructed a toy with two gears and a switch to test whether children could infer different causal structures (chains vs. common causes, for instance) from the pattern of interventions and answer counterfactual questions (they could) [Schulz et al. 2007]. By 2004, we had shown that preschoolers could determine the direction of causal arrows, infer unobserved variables, and design novel interventions, and that they did so in a way that fit much more naturally with Pearl's and Spirtes, Scheines, and Glymour's ideas than any of the traditional views of causal knowledge [Glymour 2002, Gopnik et al. 2004]. In 2005, the McDonnell Foundation funded a large interdisciplinary grant combining developmental psychologists, philosophers, and computationalists to work more on these ideas [see Gopnik and Schulz 2007, Gopnik 2012, Gopnik and Wellman 2012].

Over the next 10 years, this work continued and expanded. Although the causal Bayes net formalism was particularly elegantly designed and relatively easy to implement, it was to begin with, at least, rather limited in scope. The causal graphs were limited to describing systems of variables at a single level of description. A number of psychologists and cognitive scientists, notably Josh Tenenbaum, Tom Griffiths, and Noah Goodman, who were all involved in the McDonnell collaborative, argued for a much more expansive and general version of the project that

Pearl started, including a wide range of probabilistic generative models with different kinds of logical structure and including hierarchical as well as single-level models [see Griffiths and Tenenbaum 2007, 2009, Griffiths et al. 2010, Goodman et al. 2011, review in Tenenbaum et al. 2011].

This became an important and pervasive movement within cognitive science. It is often described as the "Bayesian" approach but this is something of a misnomer. The Bayesian part of the idea is simply this. If you have a probabilistic generative model, like a causal Bayes net, and can therefore systematically predict the probability of a pattern of evidence given that model, then you can invert this inference in a Bayesian way to infer the probability of the model given the evidence. But all the work is done by the specifics of the generative model, how well it is linked to the data, and how feasible it is to perform the Bayesian inversion and solve the search problems that result. Causal Bayes nets were and remain one of the best examples of how a probabilistic generative model could actually work.

Fei Xu, another developmental psychologist who pioneered the idea of probabilistic generative models [Xu and Tenenbaum 2007], came up with the term "rational constructivism" [Xu and Kushnir 2012], which is perhaps the best way of describing the enterprise. I suspect that the popularity of the Bayesian terminology partly reflects a principle I call The Tyranny of the Euphonious Monosyllable—if Kolmogorov had discovered Bayes' rule it wouldn't have taken off as a descriptor. But it certainly could, and perhaps should, be called Pearl-y Cognitive Science instead.

Further work in my lab and others over the next 15 years showed that very young children could make Pearl-y causal inferences across a wide range of domains, including "theory of mind." Tamar Kushnir, now at Duke, yet another student who had been part of the Pearl reading group, showed that even 18-month-olds could use Pearl-y methods to infer other people's preferences and desires [Kushnir et al. 2010]. One interesting body of work has argued that children use something like an intuitive utility calculus—a representation of the causal relationships between goals and actions—to understand other people [Hamlin et al. 2013, Lucas et al. 2014]. Kushnir and I and others showed that children and even infants were remarkably skilled at tracking and using conditional probabilities [Saffran et al. 1996, Kushnir and Gopnik 2005, Xu and Garcia 2008]. We and others also showed that children were not limited to making inferences about specific causal relationships. Instead, they could also infer quite abstract features of causal structure, such as whether causal structures were disjunctive or conjunctive. In fact, in some circumstances they could do this better than adults [Dewar and Xu 2010, Lucas et al. 2014, Gopnik et al. 2017]. Moreover, we recently showed that low-income children in Peru and in Head Start programs in the USA were just as good at making these inferences as the usual middle-class American samples [Wente et al. 2019].

In short, across what are now hundreds of studies from dozens of labs with thousands of children, it turns out that if you give children a particular pattern of data they can infer which causal structure was most likely to have generated that data, and can design new interventions and counterfactuals on that basis, in precisely the way that Pearl described.

So far, this is a largely triumphal story. But as always in science, advances lead to new problems and much of the most interesting recent work in cognitive science focuses on those problems.

One of the strengths of probabilistic generative models such as Pearl's is precisely that they are probabilistic. Earlier attempts to solve the problem of induction, such as Noam Chomsky's theory of how children infer grammars from linguistic data, were deterministic. Either a grammar was supported by the data or it wasn't. This also meant that induction was radically underdetermined—there was almost never a way of definitely ruling a grammar in or out given the data, and that led to Chomsky's nativist conclusions. The Bayesian probabilistic model approach in contrast, considers a wide range of hypotheses and tries to determine how likely each hypothesis is given the data and your prior knowledge.

But there's a catch. The catch is that for the Bayesian inversion trick to work you need to have some way of searching among the possible hypotheses and testing them against the data. Even for a relatively restricted set of representations like simple causal graphs with a limited number of variables, this problem quickly becomes untenable—there are simply too many possibilities to consider. And as the range of representations we consider becomes more abstract and complex, as with hierarchical Bayes nets, for example, or "language of thought" probabilistic logics, the search problem just becomes hairier.

Much of the exciting recent work in cognitive science, following up on Pearl's work, tries to find solutions to the search problem. Two approaches are especially interesting and exciting. First, in the computational literature the search problem is often solved by some form of sampling, randomly but systematically testing some hypotheses rather than others (e.g., Roberts and Casella [1999]). At least in "asymptopia," as one statistician calls it, these sampling methods can approximate full Bayesian inference. My collaborator Tom Griffiths and I and a number of others have shown that both adults and children show the signatures of this kind of sampling [Vul and Pashler 2008, Denison et al. 2012, Ullman et al. 2012, Bonawitz et al. 2014]. How these sampling measures could be extended and how randomness and systematicity are combined are fascinating directions for the future.

Active learning is an even more interesting and underexplored way of solving the search problem. The relationship between causal structure and intervention means that interventions can be deliberately designed to reveal causal structure, as in scientific experiments. In the early work on causal Bayes nets the assumption was that systems passively absorbed patterns of data and matched them against the potential graphical structures. When we began our first blicket detector experiments, I remember remarking that one of the big advantages of working with computers over kids was that computers weren't constantly trying to grab the blocks and try them on the machine! That observation has turned into a very productive research program, particularly as pursued by Laura Schulz and her student and my post-doc Elizabeth Bonawitz, now at Harvard. Schulz and Bonawitz have shown that children's spontaneous play often involves active interventions that are designed to resolve causal ambiguities and recover causal models [Schulz and Bonawitz 2007, Schulz et al. 2008, Schulz 2012]. The philosopher of science Frederick Eberhardt, now at the California Institute of Technology, another product of the McDonnell collaborative, has pursued a similar project in the context of science—systematically using the formalism to describe how experiments can reveal causal structure [Eberhardt and Scheines 2007].

A final frontier is the integration of causal inference and the more empiricist and statistical forms of learning, such as "deep learning" and "deep reinforcement" learning that have led to the very recent renaissance of artificial intelligence (AI), and were the subject of the 2018 Turing prize. Although these techniques have turned out to be surprisingly effective, they are beginning to come up against significant limitations. In particular, they allow only limited kinds of generalizations, and they require very large data sets and supervised forms of learning.

Increasingly, AI researchers are turning back to combine the neural network techniques with Pearl's work on causal models and the empirical work in cognitive development to try to design systems that have the power and flexibility of children's learning. For example, causality and cognitive development both play a central role in the recent DARPA machine common sense program, which we are part of at Berkeley.

Perhaps it is symbolic that the Berkeley Artificial Intelligence Research unit, of which I am now a member, just moved into the same building as the Developmental Psychology group. Both geographically and intellectually, the distance between the two fields is beginning to disappear. We very much have Judea Pearl to thank for that.

References

E. Bonawitz, S. Denison, T. Griffiths, and A. Gopnik. 2014. Probabilistic models, learning algorithms, response variability: Sampling in cognitive development. *Trends Cogn. Sci.* 18, 497–500. DOI: https://dx.doi.org/10.1016/j.tics.2014.06.006.

S. Carey. 1985. *Conceptual Change in Childhood*. MIT Press, Cambridge, MA.

S. Denison, E. Bonawitz, A. Gopnik, and T. L. Griffiths. 2012. Rational variability in children's causal inferences: The Sampling Hypothesis. *Cognition* 126, 285–300. DOI: https://doi.org/10.1016/j.cognition.2012.10.010.

K. M. Dewar and F. Xu. 2010. Induction, overhypothesis, and the origin of abstract knowledge. *Psychol. Sci.* 21, 12, 1871–1877. DOI: https://dx.doi.org/10.1177/0956797610388810.

F. Eberhardt and R. Scheines. 2007. Interventions and causal inference. *Philos. Sci.* 74, 5, 981–995. DOI: https://doi.org/10.1086/525638.

C. N. Glymour. 2002. *The Mind's Arrows: Bayes Nets and Graphical Causal Models in Psychology*. MIT Press, Cambridge, MA. DOI: https://doi.org/10.1093/mind/112.446.340.

N. D. Goodman, T. D. Ullman, and J. B. Tenenbaum. 2011. Learning a theory of causality. *Psychol. Rev.* 118, 1, 110–119. DOI: https://doi.org/10.1037/a0021336.

A. Gopnik. 1988. Conceptual and semantic development as theory change: The case of object permanence. *Mind Lang.* 3, 3, 197–216. DOI: https://doi.org/10.1111/j.1468-0017.1988.tb00143.x.

A. Gopnik. 2012. Scientific thinking in very young children: Theoretical advances, empirical discoveries and policy implications. *Science* 337, 6102, 1623–1627. DOI: https://doi.org/10.1126/science.1223416.

A. Gopnik and H. M. Wellman. 1994. The theory theory. In L. Hirschfeld and S. Gelman (Eds.), *Domain Specificity in Cognition and Culture*. Cambridge University Press, New York, 257–293. DOI: https://doi.org/10.1017/CBO9780511752902.011.

A. Gopnik and D. Sobel. 2000. Detecting blickets: How young children use information about novel causal powers in categorization and induction. *Child Dev.* 71, 5, 1205–1222. DOI: https://doi.org/10.1111/1467-8624.00224.

A. Gopnik and L. Schulz. (Eds.). 2007. *Causal Learning: Psychology, Philosophy, and Computation*. Oxford University Press, Inc., New York. DOI: https://doi.org/10.1093/acprof:oso/9780195176803.001.0001.

A. Gopnik and H. M. Wellman. 2012. Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychol. Bull.* 138, 6, 1085–1108. DOI: https://doi.org/10.1037/a0028044.

A. Gopnik, D. M. Sobel, L. E. Schulz, and C. Glymour. 2001. 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, 5, 620–629. DOI: https://doi.org/10.1037/0012-1649.37.5.620.

A. Gopnik, C. Glymour, D. M. Sobel, L. E. Schulz, T. Kushnir, and D. Danks. 2004. A theory of causal learning in children: Causal maps and Bayes nets. *Psychol. Rev.* 111, 1, 3–32. DOI: https://doi.org/10.1037/0033-295X.111.1.3.

A. Gopnik, S. O'Grady, C. G. Lucas, T. L. Griffiths, A. Wente, S. Bridgers, R. Aboody, H. Fung, and R. E. Dahl. 2017. Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proc. Natl. Acad. Sci. U S A* 114, 30, 7892–7899. DOI: https://doi.org/10.1073/pnas.1700811114.

T. Griffiths and J. B. Tenenbaum. 2007. Two proposals for causal grammars. In A. Gopnik & L. Schulz (Eds.), *Causal Learning: Psychology, Philosophy, and Computation*. Oxford University Press, Inc., New York, 323–345. DOI: https://doi.org/10.1093/acprof:oso/9780195176803.003.0021.

T. L. Griffiths and J. B. Tenenbaum. 2009. Theory-based causal induction. *Psychol. Rev.* 116, 4, 661–716. DOI: https://doi.org/10.1037/a0017201.

T. L. Griffiths, N. Chater, C. Kemp, A. Perfors, and J. B. Tenenbaum. 2010. Probabilistic models of cognition: Exploring representations and inductive biases. *Trends Cogn. Sci.* 14, 8, 357–364. DOI: https://doi.org/10.1016/j.tics.2010.05.004.

K. Hamlin, T. Ullman, J. Tenenbaum, N. Goodman, and C. Baker. 2013. The mentalistic basis of core social cognition: Experiments in preverbal infants and a computational model. *Dev. Sci.* 16, 2, 209–226. DOI: https://doi.org/10.1111/desc.12017.

T. S. Kuhn. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press, Chicago. DOI: https://doi.org/10.1119/1.1969660.

T. Kushnir and A. Gopnik. 2005. Young children infer causal strength from probabilities and interventions. *Psychol. Sci.* 16, 9, 678–683. DOI: https://doi.org/10.1111/j.1467-9280.2005.01595.x.

T. Kushnir, F. Xu, and H. M. Wellman. 2010. Young children use statistical sampling to infer the preferences of other people. *Psychol. Sci.* 21, 8, 1134–1140. DOI: https://doi.org/10.1177/0956797610376652.

C. G. Lucas, T. L. Griffiths, F. Xu, C. Fawcett, A. Gopnik, T. Kushnir, L. Markson, and J. Hu. 2014. The child as econometrician: A rational model of preference understanding in children. *PLoS One* 9, e92160. DOI: https://doi.org/10.1371/journal.pone.0092160.

C. Lucas, S. Bridgers, T. Griffiths, and A. Gopnik. 2014. When children are better (or at least more open-minded) learners than adults: Developmental differences in learning the forms of causal relationships. *Cognition*. 131, 2, 284–299. DOI: https://doi.org/10.1016/j.cognition.2013.12.010.

D. Marr. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W.H. Freeman, San Francisco. DOI: https://doi.org/10.7551/mitpress/9780262514620.001.0001.

J. Pearl. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, San Mateo, CA.

J. Pearl. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York. DOI: https://doi.org/10.1017/S0266466603004109.

C. P. Robert and G. Casella. 1999. *Monte Carlo Statistical Methods*. Springer-Verlag, New York.

J. R. Saffran, R. N. Aslin, and R. N. Newport. 1996. Statistical learning by 8-month-old infants. *Science* 274, 5294, 1926–1928. DOI: https://doi.org/10.1126/science.274.5294.1926.

L. E. Schulz. 2012. Origins of inquiry: Inductive inference and exploration in young children. *Trends Cogn. Sci.* 16, 7, 382–389. DOI: https://doi.org/10.1016/j.tics.2012.06.004.

L. E. Schulz and E. B. Bonawitz. 2007. Serious fun: Preschoolers engage in more exploratory play when evidence is confounded. *Dev. Psychol.* 43, 4, 1045–1050. DOI: https://doi.org/10.1037/0012-1649.43.4.1045.

L. E. Schulz, A. Gopnik, and C. Glymour. 2007. Preschool children learn about causal structure from conditional interventions. *Dev. Sci.* 10, 3, 322–332. DOI: https://doi.org/10.1111/j.1467-7687.2007.00587.x.

L. E. Schulz, H. R. Standing, and E. B. Bonawitz. 2008. Word, thought, and deed: The role of object categories in children's inductive inferences and exploratory play. *Dev. Psychol.* 44, 5, 1266–1276. DOI: https://doi.org/10.1037/0012-1649.44.5.1266.

C. Smith, S. Carey, and M. Wiser. 1985. On differentiation: A case study of the development of the concepts of size, weight, and density. *Cognition* 21, 3, 177–237. DOI: https://doi.org/10.1016/0010-0277(85)90025-3.

P. Spirtes, C. Glymour, and R. Scheines. 1993. *Causation, Prediction and Search, Springer Lecture Notes in Statistics* (2nd. ed.). MIT Press, Cambridge, MA, 2000.

J. B. Tenenbaum, C. Kemp, T. L. Griffiths, and N. D. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science* 331, 6022, 1279–1285. DOI: https://doi.org/10.1126/science.1192788.

T. D. Ullman, N. D. Goodman, and J. B. Tennenbaum. 2012. Theory acquisition as stochastic search. *Cogn. Dev.* 27, 4, 455–480. DOI: https://doi.org/10.1016/j.cogdev.2012.07.005.

E. Vul and H. Pashler. 2008. Measuring the crowd within: Probabilistic representations within individuals. *Psychol. Sci.* 19, 7, 645–647. DOI: https://doi.org/10.1111/j.1467-9280.2008.02136.x.

H. M. Wellman and S. A. Gelman. 1992. Cognitive development: Foundational theories of core domains. *Ann. Rev. Psychol.* 43, 337–375. DOI: https://doi.org/10.1146/annurev.ps.43.020192.002005.

A. Wente, K. Kimura, C. Walker, N. Banerjee, M. Fernández Flecha, B. MacDonald, C. Lucas, and A. Gopnik. 2019. Causal learning across culture and SES. *Child Dev.* 90, 859–875. DOI: https://doi.org/10.1111/cdev.12943.

J. Woodward. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, New York. DOI: https://doi.org/10.1093/0195155270.001.0001.

F. Xu and V. Garcia. 2008. Intuitive statistics by 8-month-old infants. *Proc. Natl. Acad. Sci. U S A* 105, 13, 5012–5015. DOI: https://doi.org/10.1073/pnas.0704450105.

F. Xu and J. B. Tenenbaum. 2007. Word learning as Bayesian inference. *Psychol. Rev.* 114, 2, 245–272. DOI: https://doi.org/10.1037/0033-295X.114.2.245.

F. Xu and T. Kushnir. 2012. (Vol. Ed.), *Rational Constructivism in Cognitive Development.* Academic Press, Elsevier Inc., Waltham, MA, 161–189.