Probabilistic Models, Learning Algorithms and Response Variability: Sampling in Cognitive Development

Elizabeth Bonawitz¹, Stephanie Denison², Thomas L. Griffiths³, & Alison Gopnik³

¹ Rutgers University – Newark; ² University of Waterloo; ³ University of California, Berkeley

Address for correspondence:

Elizabeth Bonawitz

334 Smith Hall

Department of Psychology

Rutgers University - Newark

Newark, NJ 07102 USA

Phone: 617-893-0125

Email: elizabeth.bonawitz@rutgers.edu

Abstract:

Although probabilistic models of cognitive development have become increasingly prevalent, one challenge is to account for how children might cope with a potentially vast number of possible hypotheses. We propose that children might address this problem by “sampling” hypotheses from a probability distribution. We discuss empirical results demonstrating signatures of sampling, which offers an explanation for the variability of children’s responses. The sampling hypothesis provides an algorithmic account of how children might address computationally intractable problems and suggests a way to make sense of their “noisy” behavior.

Probabilistic Models in Development

In the course of development, children's beliefs about the world undergo substantial revision. Probabilistic models of cognitive development provide a potential account of some aspects of this remarkable learning process [1]. These models can rigorously characterize the structure of early representations and their revision. On this account, children's beliefs, such as their "intuitive theories", can be formally described as generative models, for example, as causal graphs, grammars, or taxonomies. A generative model predicts some patterns of evidence and not others. For example, a particular graphical model of a causal system will predict that some patterns of contingency between events are more likely than others; a grammar will predict that certain sentences will be more likely to be acceptable than others.

If theories are expressed as probabilistic generative models, then the process of revising those theories can be formally described as Bayesian inference. Different generative models systematically generate some patterns of data rather than others, so a learner can start with the data and infer which model was most likely to have generated those data, guided by prior knowledge. Formally, prior knowledge is expressed in a "prior" probability distribution over hypotheses, and Bayes' rule indicates how to compute a "posterior" distribution that incorporates the data. This approach can thus provide a desirably precise account of how prior knowledge and new evidence may be combined to update a set of beliefs.

The probabilistic modeling approach is not without critique [e.g. 2]. Many of the critiques stem from the fact that the Bayesian view, just by itself, is extremely flexible and can accommodate a very wide range of data patterns (just as connectionist or production system models can, in principle, accommodate any data pattern). In order to be informative, probabilistic model accounts must specify the nature of the generative models, and the likelihood functions in detail. Indeed an advantage of this approach is that it requires the theorist to make this specification in a precise and detailed way, and therefore generates precise quantitative predictions about that particular probabilistic model.

Most probabilistic models operate at what Marr [3] called the *computational* level of analysis. Computational-level models provide clear descriptions of the problems the learner faces and describe ideal solutions for those problems. Probabilistic models at this computational level can characterize how children infer beliefs from evidence in at least some cases, such as causal learning tasks [4-7]. In these studies, researchers assess the state of children's prior beliefs, carefully control the evidence they receive, and then examine which hypotheses they endorse. Children tend to choose the hypotheses that have the greatest posterior probability according to a Bayesian analysis.

Children's responses on these tasks on average look like the posterior distributions predicted by these computational-level models. However, that does not necessarily imply that learners are working through the calculations prescribed by Bayes' rule at the *algorithmic* level. A very large number of hypotheses may be compatible with any pattern of evidence, and it would be impossible to assess each of these hypotheses individually. This problem might be particularly challenging for

young children who might have more restricted memory and information-processing capacities than adults.

So how do learners behave in a way that is apparently consistent with probabilistic models when it is unlikely that they are actually assessing all the possible hypotheses in practice?

Approximating probabilistic models: The Sampling Hypothesis

Applications of probabilistic models in computer science must also tackle the problem of evaluating large spaces of hypotheses. They often do so by randomly but systematically sampling a few hypotheses rather than exhaustively considering all possibilities. These calculations use “Monte Carlo” methods. They obtain the equivalent of samples from the posterior distribution without computing the whole posterior distribution itself. A system that uses this sort of sampling will be variable, as it will entertain different hypotheses apparently at random.

However, importantly, this variability will also be systematic. The system will sample more probable hypotheses with greater frequency than less probable ones, so the distribution of responses will reflect the probability of the hypotheses. And most importantly, such a system will be efficient, as it trades approximation error for computing time. The success of Monte Carlo algorithms in computer science and statistics suggests an exciting hypothesis for cognitive development. The algorithms children use to perform inductive inference might also involve sampling. We call this the Sampling Hypothesis.

Some recent work supports the idea that adults may sometimes approximate Bayesian inference through psychological processes that are equivalent to sampling. Participants in a simple judgment task provided responses that suggested they sampled their judgments from an internal distribution, rather than providing a single best guess [8]. It is often observed that people produce responses proportional to the Bayesian posterior probabilities [e.g., 9]. While producing just a few samples may lead to behavior that appears suboptimal [10], it is a rational strategy for compromising between the cost of errors and the opportunity cost of taking more samples [11].

The Sampling Hypothesis is especially interesting from a developmental perspective, since it might explain at least some of the variability in children’s behavior. Developmental studies have pointed to the extensive variability in children’s responses, hypotheses and solutions to problems [e.g., 12]. Often, this variability is assumed to be the result of external “noise” such as attention or memory failures, and this may often be true. But it is also possible that at least sometimes the variability in children’s responding is systematic, as the sampling hypothesis would predict.

For example, in causal learning tasks children tend to pick the hypothesis that is most probable. But not all children pick the most likely response, and an individual child may change responses from trial to trial [e.g., 5-7]. The proportion of times that children select a hypothesis increases as the hypothesis receives more support, but children still sometimes produce alternative hypotheses. That might mean that children are sampling their responses from a posterior distribution.

Alternatively it might be that children aim to produce a best guess all the time, and that the variability in their responses is simply a reflection of stochastically produced errors and “noise” caused by other factors. Children might be “Noisy Maximizers”, producing an error-laden attempt at the most likely answer.

Yet another alternative is that children’s behavior in these tasks does reflect the probability of hypotheses but does so through a simpler process than hypothesis sampling. Children, like adults and even non-human animals, frequently produce a pattern called probability matching in reinforcement learning [e.g. 13]. This “Naïve Frequency Matching” alternative suggests that learners may simply match the frequency of responses to those of rewards.

So the question is, is the variability in children’s answers the result of sampling, is it an error-laden attempt at maximization, or does it involve naïve frequency-matching?

Empirical support for the Sampling Hypothesis

A recent set of developmental studies present the first test of the Sampling Hypothesis, distinguishing sampling from both Noisy Maximizing and Naïve Frequency Matching alternatives. Denison, Bonawitz, Gopnik, and Griffiths [14] explored the degree to which children match posterior probabilities in a causal inference task. Children saw a bin with a varying number of blue and red chips. In the first experiment, children had several chances to guess whether a chip that randomly fell out of the bin was red or blue. The probability that each chip had fallen out of the bin was directly related to the proportion of red and blue chips. When the bin was 80% full of red chips, there was an 80% chance that the randomly selected chip was red. Children’s behavior showed signatures of sampling: They guessed “red” or “blue” in proportion to the probability that a red or blue chip had fallen in the bin.

In a second and third experiment, the proportion of colored chips was systematically varied. Children provided responses that matched the posterior probability of hypotheses; when the probability of a hypothesis decreased, children’s selection of that hypothesis also decreased. Noisy Maximizing would instead predict that children would favor the most likely hypothesis at constant rates across varying probabilities (i.e., whether the probability of a red chip is 95% or 75%, children should guess “red” at near ceiling levels).

In a final experiment the probability did not directly reflect the frequencies of the chips, providing a way to distinguish Sampling from Naïve Frequency Matching. Children saw two blue chips in one bin, and 14 red and 6 blue in the other. Then the bins were obscured and one unknown bin was randomly selected, so that the probability of the blue chip was 65%, while the frequency of blue chips was only 36%. Children did not appear to naively match frequencies as they would in simple probability learning through reinforcement. Consistent with the Sampling Hypothesis, children’s guesses matched the posterior distribution of hypotheses, rather than the simple frequencies of the red and blue chips.

These experiments showed that children were behaving in a way that was consistent with sampling, as well as consistent with probabilistic models. But they did not determine which type of sampling algorithm children might use. A first

challenge in exploring sampling algorithms is to demonstrate that there are psychologically plausible strategies that can produce behavior that is generally consistent with the predictions of probabilistic models. Bonawitz, Denison, Griffiths, and Gopnik [15] mathematically demonstrated that a surprisingly simple version of a single-sample algorithm (based on a win-stay, lose-shift strategy) will produce behavior that is consistent with the exact analytical solution when aggregated across multiple participants. In this algorithm the learner initially chooses a guess at random and then tends to stay with that guess unless it is contradicted by the evidence. As the evidence against the initial guess grows stronger, the learner will be increasingly likely to resample from the distribution and try another guess.

Identifying the sampling algorithm

A second challenge is empirically identifying which sampling algorithm learners might be using on a particular task. An algorithm that samples a new guess from an updated posterior after each observation of data (independent sampling), will behave differently from a “win-stay-lose-shift” algorithm, like the one proposed by Bonawitz et al. [15]. Although both are sampling approaches and would produce behavior consistent with probabilistic models in the long run, they will have different consequences for short-term behavior. The win-stay-lose-shift strategy will lead to a distinctive pattern of dependencies. A learner’s initial guess will shape her immediately subsequent guesses, even if the initial guess was chosen at random.

Bonawitz et al. [15] found that preschool-aged children and adults produce a characteristic pattern of dependencies in their responses to a causal learning task that were consistent with their particular WSLS algorithm. They were able to demonstrate this by using a mini-microgenetic method. They presented children with initial evidence that was compatible with several different hypotheses and asked them to guess which hypothesis was correct. Then on each trial they added evidence that tended to confirm or disconfirm that guess and asked the children to guess again. Even though individual learners might seem to be randomly veering from one hypothesis to the next, on aggregate their responses approximated the exact analytical Bayesian solution. The WSLS algorithm predicted this approximate Bayesian response on aggregate and best captured the trial-by-trial data in the individual responses.

Conclusions and future work

These studies are just a starting point for asking what algorithms best capture early learning. Sampling algorithms like these may provide a balance between “explore” and “exploit” strategies in learning. They allow the learner to consider potentially unlikely hypotheses on occasion – hypotheses that may prove to be correct later. In aggregate and over time they converge on the hypothesis that is most likely.

So far, we have explored these algorithms in causal learning tasks. We don’t yet know how general such strategies may be. Different types of learning, such as syntactic inference, might employ different approaches. Furthermore, the particular algorithms children employ may depend on task demands, development, or even individual preference.

We have suggested that children may revise their causal beliefs by randomly sampling from a probability distribution. Sampling is an efficient way to search through a space of possibilities while still acting in a way that is consistent with probabilistic inference, and so it can be an algorithmic instantiation of Bayesian inference. The Sampling Hypothesis also suggests that the variability of children's responses may sometimes reflect the use of this type of algorithm rather than being just noise.

References

- [1] Gopnik, A., & Wellman, H.M. (*in press*). Reconstructing constructivism: Causal models, Bayesian learning mechanisms and the theory theory. *Psychological Bulletin*.
- [2] Marcus, G. F. & Davis, E. (2013). How robust are probabilistic models of higher-level cognition? *Psychological Science*, 24(12), 2351-60.
- [3] Marr, D. (1982). *Vision*. Freeman Publishers.
- [4] Bonawitz, E.B., van Schijndel, T., Friel, D., & Schulz, L. (2012) Balancing theories and evidence in children's exploration, explanations, and learning. *Cognitive Psychology*, 64, 215-234.
- [5] Kushnir, T. & Gopnik, A., (2005). Children infer causal strength from probabilities and interventions. *Psychological Science*, 16, 678-683.
- [6] Schulz, L.E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared make your tummy ache? Naive theories, ambiguous evidence and preschoolers' causal inferences. *Developmental Psychology*, 43, 1124-1139.
- [7] Sobel, D. M., Tenenbaum, J. M., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, 28(3), 303-333.
- [8] Vul, E., & Pashler, H. (2008). Measuring the Crowd Within: Probabilistic Representations Within Individuals. *Psychological Science*, 19 (7), 645-647.
- [9] Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science*, 32, 108-154.

- [10] Mozer, M. C., Pashler, H., & Homaei, H. (2008). Optimal predictions in everyday cognition: The wisdom of individuals or crowds? *Cognitive Science*, 32(7), 1133-1147.
- [11] Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive science*.
- [12] Siegler, R. S. (1998). *Emerging minds: The process of change in children's thinking*. Oxford University Press.
- [13] Jones, M. H., & Liverant, S. (1960). Effects of age differences on choice behavior. *Child Development*, 31(4), 673-680.
- [14] Denison, S., Bonawitz, E., Gopnik, A., & Griffiths, T. L. (2013). Rational variability in children's causal inferences: The Sampling Hypothesis. *Cognition*, 126, 285-300.
- [15] Bonawitz, E., Denison, S., Chen, A., Gopnik, A., & Griffiths, T. L. (*in press*). Win-Stay, Lose-Shift: A simple sequential algorithm for approximating bayesian inference. *Cognitive Psychology*.

Glossary of key terms and ideas:

Exact analytical solution (Exact learning): In mathematics, the mechanical steps used to carry out computation lead to a precise single numerical result. This is contrasted with approximate solutions (**Approximate inference**), which provide guesses about likely numerical results. Approximate solutions are employed when computing the analytical solution is intractable or would simply take longer than desirable. Approximate solutions often trade time for accuracy – the longer an algorithm runs, the closer to correct the answer will be.

Monte Carlo methods: Algorithms that depend on repeated stochastic (random) sampling to produce a numerical estimate of the result. The name derives from the fact that casinos are likewise based on sampling from particular probability distributions with every roll of the dice or spin of a wheel.

Posterior probability: The conditional probability that a hypothesis is true, after the evidence is taken into account. The **posterior distribution** is the probability distribution over hypotheses defined by these probabilities.

Probabilistic approaches to cognitive development: Assuming that, at a computational level, processes of inference and learning in cognitive development can be characterized in terms of rationally updating a probability distribution over hypotheses in accordance with Bayes' rule. Strict binary rule-based models of learning are a contrasting example; learners might follow a heuristic that allows them to identify a deterministic outcome (e.g. a yes/no decision on whether an object falls into a particular category given some threshold).

Probability matching in reinforcement learning: In contrast to always producing an action that will most likely bring about the reward, probability matching in reinforcement learning is when the learner matches the proportion of his or her responses to the relative rates of reward. This entails sometimes producing a response that is not the most likely to be rewarded, and hence not maximizing possible rewards.

Acknowledgements.

This research was supported in part by the James S. McDonnell Causal Learning collaborative and a grant IIS-0845410 from the National Science Foundation.

Figure. Example methods, proportion of chips per condition, and results from Denison et al, 2013.

Correlation between Children's % Responding and Posterior Probability