Computational Models of Development, Social Influences

Elizabeth Bonawitz & Patrick Shafto

Address for correspondence:

Elizabeth Bonawitz

334 Smith Hall

Department of Psychology

Rutgers University - Newark

Newark, NJ 07102 USA

Phone: 973-353-3944

Email: elizabeth.bonawitz@rutgers.edu

Abstract

In the article we argue that past Bayesian approaches that model children's learning from data are missing an important element – the role of other people in generating that data. We propose that children take the origin of data into account when learning, which can be understood through ideal observer analyses of the social situation. Moreover, when observing evidence, children are not just learning *from* others, but also *about* others. We review recent literature suggesting that children can make inferences about the knowledge and goals of the individual selecting the data and use this knowledge to bolster learning from this evidence.

Conventional wisdom often points to the myriad ways in which children have unique limitations and cognitive capacities from adults – children may be lacking the knowledge, skills, and logical reasoning abilities to perform as competently as their parents (Diamond & Kirkham, 2005; Flavell, 1971; Inhelder & Piaget, 1958; Kuhn, 2000; Kuhn, Amsel, & O’Laughlin, 1988; Klahr & Dunbar, 1989; Piaget, 1965; Zelazo, 2004; Zelazo, Frye, & Rapus, 1996). However, children must be powerful learners. They enter the world with few expectations about objects, events, others and themselves, but within a year they begin express complicated concepts like “all done”, “uh oh”, or “mine”. To explain such remarkable learning, researchers have proposed elements of the learning toolkit including statistical learning mechanisms (Kirkham, Slemmer, & Johnson, 2002; Saffran, 2003; Saffran, Johnson, Aslin, & Newport, 1999) and the ability to integrate these statistics with inductive constraints (Gopnik & Schulz, 2004; Tenenbaum, Kemp, Griffiths, & Goodman, 2011; Perfors, Tenenbaum, Griffiths, & Xu, 2011; Schulz, Bonawitz, & Griffiths, 2007; Bonawitz & Lombrozo, 2013; Griffiths, Chater, Kemp, & Perfors, 2010; Xu & Tenenbaum, 2007).

One way to understand how even very young children can attend to the evidence in their environment and – from this data – rapidly update their beliefs, is through computational models based on ideal observer analyses. Ideal observer analyses answer the question, what is the learning problem that is being solved? The answer is formalized as a computational model that assumes ideal performance. In cognitive development, the learning problems being solved are things like learning language and learning the causal structure of the environment. The computational models formalize learning as Bayesian inference over structured representations. Bayesian inference captures ideal learning in the face of uncertainty and noise, and structured representations captures the underlying regularities of, for example, the grammar of language or the causal structure of the

environment that must be inferred. The result of this marriage is “Theory-based Bayesian” models of learning (Tenenbaum, Griffiths, & Kemp, 2006).

In recent years, myriad studies have demonstrated how ideal observer approaches can be used to describe how children might learn language (Kemp, Perfors, & Tenenbaum, 2007) and develop causal beliefs (Gopnik & Wellman, 2012; Kemp & Tenenbaum, 2008). Although these models provide important insight into how we can capture the remarkably sophisticated learning in childhood, we suggest that they miss an important aspect of learning: learning from others. Evidence does not appear in a social vacuum; children often learn from observing others. When a word is uttered, it comes from the mouth of a person trying to communicate a concept. When a light switch is flicked, it is almost always because a person wants there to be light.

Computational models of learning have demonstrated the importance of considering social information in language acquisition (e.g. Yu & Ballard, 2007), social referencing (e.g. Thomaz, Berlin, & Breazeal, 2005), and imitation (for review see, Breazeal & Scassellati, 2002). These past approaches use social information as an additional source of evidence, independent from other streams of information. Here, we focus on how social inferences shape the *interpretation* of evidence. The critical contribution of *socially generated data* – that people are part of the generative process – is that it can lead to even stronger inferences and more rapid learning. We will stress is that children are not only paying attention to the evidence (e.g. the light switch in the up position and the light being on), but that children are paying attention to the people that are demonstrating those events (e.g. that even if the person tries, but fails, to flip the light switch up, there was a reason that they tried to flip it). We will discuss recent probabilistic models that include this socially generative component. These models relate the *knowledge* and *goals* of a demonstrator to actions they choose, and formalize how these purposeful actions can influence the interpretation of the data being demonstrated. We will present research

suggesting that children can make inferences about the knowledge and goals of the individual selecting the data and use this knowledge to bolster learning from this evidence.

Intrinsic to our proposal is the idea that children are able to represent and reason about others' beliefs. These are non-trivial abilities, combining Theory of Mind and planning capabilities, and the development of these skills may extend into middle childhood (e.g. Talwar & Lee, 2008; Wellman & Liu, 2004). However, many of the types of social learning we discuss do not necessarily require competency with false beliefs; many other forms of social learning based on reasoning about others' beliefs and goals can be observed much earlier in development (Gergely & Csibra, 2003; Repacholi & Gopnik, 1997; Woodward, 2005), and even infants may understand false-beliefs (Baillargeon, Scott, and He, 2010).

Strong Sampling in word learning

Word learning is arguably the canonical case of social learning. It is also famously wrought with uncertainty (Quine, 1960): How should one infer the precise extension of a word's meaning given a few observations of a label? When a scene is labeled for a learner, the possible meanings are endless, and yet even children are likely to infer that the speaker intends something like "rabbit", and not "rabbits and also tomatoes but not squirrels". Words are also learned with very few positive examples and rarely any negative examples; words can carve up the world in overlapping categories; and multiple words can be used to express similar concepts (Bloom, 2000; Carey, 1978; Gold, 1967; Markman, 1990; Regier, 1996). How might children accomplish such a feat of inductive prowess? As Quine (1960) famously noted, "Language is a social art"; the learner's reasoning depends on making inferences about the speaker's intentions. Might such intuitions leverage inductive inference?

Theory-based Bayesian models can begin to offer a solution to the riddle of induction and a precise way to characterize the contribution of social inferences in word

learning (Xu & Tenenbaum, 2007). Imagine learning about a concept, such as 'dalmation' by observing someone randomly labeling objects from the world, and in each case stating whether the drawn object is an instance of the concept or not. Given these examples, the learner must infer which of many possible hypotheses is most likely to account for the data produced by these haphazard observations.

Of course, this is not how people use language. Intuitively, we know that, in general, labeling any given object as 'not a dalmation' provides nearly no information about what 'dalmation' actually is. Instead, we can consider language learning as a different sort of problem, one in which people almost exclusively label objects with their correct label. The examples of 'dalmation' may still be haphazard, but they are all dalmations. Thus, when a learner sees three dalmations drawn from a box and labeled with a novel word, she can reason about its extension via a sort-of counterfactual logic: "what are the chances that I would have seen exactly these three objects if the concept I'm learning is 'dalmations', 'dogs', or 'animals'?" The narrower the concept fits to the examples, the more likely the set will have been drawn. One is most likely to see three dalmations randomly sampled from a box only containing dalmations, less likely to see such evidence if the box contains all dogs, and least likely if the box contains all animals.

Computationally, researchers have modeled these different potential selection processes. The former, called weak sampling, provides only the information present in the object-label pairing. This would lead to strong limitations on learning because the sampling process—random sampling—is unrelated to the true concept. Indeed, mathematical demonstrations formalize the idea that, if this were true, language would not be learnable (Gold, 1967). In contrast, strong sampling formalizes the idea that if a person intends to communicate 'dalmation', they will choose dalmations. This does not completely solve the problem: if the person wishes to communicate 'animal' they may also choose a dalmation. However, if as a language learner we aggregate over multiple labelings, then the probability

of someone choosing 3 dalmations at random from all animals is very, very small; dalmation must mean 'dalmation' (Xu & Tenenbaum, 2007ab).

Learning from teaching

Another especially compelling case for social learning comes from teaching. According to an ideal observer analysis, learning from teaching ("pedagogical learning") differs from standard learning in one critical way—the data that the learner observes are chosen by a knowledgeable and helpful person, not selected at random. That is, in pedagogical sampling the learner assumes that the teacher is not merely sampling examples from within the concept, but is purposefully sampling the best possible data in order to help the learner rapidly converge on the correct concept given the learner's current beliefs. Thus, the learner not only asks herself "what are the chances I would have observed these examples, given this category" but she asks "what are the chances that *the teacher would have chosen these examples for me if this was the true category.*" A knowledgeable and helpful teacher would be especially unlikely to provide an overly restrictive set of examples, such as three Dalmations, if the concept to be conveyed was really something broader, such as "dogs". Thusly, pedagogical sampling allows the learner make even stronger inferences about the data than strong sampling would warrant (Shafto & Goodman, 2008; Shafto, Goodman, & Frank, 2012).

These pedagogical inferences also extend beyond word learning to other kinds of causal reasoning. For example, imagine you are a child, and a person walks into the room with a complex looking novel toy. The person says, "This is my toy. I'm going to show you how it works" then pulls a non-obvious yellow knob, eliciting a squeak, before handing the toy to the child. What should a child infer about the number of causal functions of the toy? Contrast this with a situation in which the person says, "Look at this toy I just found" and, as they are handing the toy to the child, they accidentally pull the knob, eliciting a squeak,

before handing the toy over to the child. What should the child infer about the number of causal functions of the toy in this other situation? It seems obvious that, in the second case, the child can infer that while there is evidence for one function, there could be more functions. However, in the first case, by reasoning about why the teacher chose to demonstrate only one function, the child can infer that there is in fact only one function: if there were more functions, a knowledgeable and helpful teacher would have chosen to demonstrate those too. Recent research using this and related paradigms have shown that, before they are arriving at formal schooling, 4- and 5-year-old children can use the knowledge and intent of the teacher to draw stronger inferences than are afforded by the data alone (Bonawitz, Shafto et al., 2011; Buchbaum et al., 2011). Indeed, parallel research in education has shown similar implications of evidence selection on children's learning of the mathematical concept of equality (McNeill, 2008; McNeill and Alibali, 2005, 2004; Behr, 1980; Seo and Ginsburg, 2003).

Recent computational research has formalized a framework for modeling why other people's purposeful data selection affects learning (Shafto, Goodman, Frank, 2012).

Consider the case of learning from teaching, where the data are chosen by a knowledgeable and helpful informant ("pedagogical reasoning"; Shafto & Goodman, 2008; Shafto, Goodman, Griffiths, 2014). The model captures the deceptively simple intuition that a teacher will choose the evidence that tend to increase the learner's belief in the correct hypothesis; it makes precise the implications of that for learners who expect this sort of purposeful selection of data. Formally, we can describe the probability of a teacher selecting data given a hypothesis, $P_T(d|h)$, and the relationship of the teacher's choice to the learner's posterior beliefs after observing the data, $P_L(h|d)$. To capture the intuition that a teacher will choose data that tend to lead the learner to the correct hypothesis, we state: $P_T(d|h) \propto P_L(h|d)$.

Learning from teaching is similar to standard learning based on Bayes rule, where the

sampling process is modified to take into account the fact that the data are chosen by the teacher: $P_L(h|d) \propto P_T(d|h)P_L(h)$. (See Figure 1.)

The model has generated a variety of testable predictions about how learning from a teacher should differ from situations in which the observed data is generated by accidental or other random selection methods. Indeed, the model was the inspiration for the novel toy experiment described above (Bonawitz, Shafto, et al., 2011; see also Buchbaum et al., 2011), and explains why children have difficulty learning the concept of equality (McNeill, 2008; McNeill and Alibali, 2005, 2004; Behr et al., 1980; Seo and Ginsburg, 2003). Even children's causal inferences are sensitive to the goals and knowledge of the person generating the data, and such inferences can be modeled with this framework.

Epistemic trust

If, as indicated by research on direct instruction, learners are going to treat people as knowledgeable and helpful, then it is necessary that one consider the alternative: some people are not knowledgeable nor helpful; they can even be deceptive. Research on epistemic trust investigates whom children choose to provide information. The literature has focused on a host of superficial cues including accuracy (Pasquini et al., 2007), familiarity (Corriveau & Harris, 2009), consensus (Corriveau et al., 2009), and accent (Kinzler et al., 2011) to name a few. There is robust evidence that children by at least the age of 4 (and on some tasks much earlier, e.g. Tummelshammer et al., 2014) children are discerning about whose information they will endorse and whom they will ask.

Building on the learning *from* others framework, probabilistic models of learning *about* others provide a framework for understanding children's remarkable proficiency, and the developmental changes in their abilities, to assess who to trust for information. Shafto, Eaves, et al., (2012, cf. Bovens & Hartmann, 2004; Butterfield et al., 2009) formalized a model that inferred, based on the evidence provided by an informant, whether the

informant was knowledgeable or not and helpful or deceptive. The choice of who to trust is then an inference about which informant's information is most likely to be correct.

Modeling three representative studies from the literature, they showed that 4-year-old's ability to track informant reliability (Pasquini et al., 2007), to balance familiarity and recent accuracy (Corriveau & Harris, 2009), and reason about consensus (Corriveau et al., 2009), was explained by reasoning about the informant's knowledge and intent. The model also revealed developmental changes in reasoning. In contrast with 4-year-olds' behavior, 3-year-olds' behavior was best explained by reasoning only about informants' knowledge, assuming that informants are always helpful (see also Eaves & Shafto, 2012; Landrum, Eaves, & Shafto, 2015; Mascaro & Sperber, 2009).

A convergent line of research has shown that trust is about more than simply tracking informants' accuracy. Consistent with the model, Gweon et al. (2014) presented 6- and 7-year-old with informants who either demonstrated all of the functions of a toy or only one of the functions of a toy. Prior to the demonstration, children had been allowed to explore the toy, and they therefore knew how many functions the toy had. After the demonstration children rated the demonstrator on their effectiveness. Results showed that children rated informants who *omitted* information lower than those who demonstrated all of the functions, and lower than other informants who demonstrated the same number of functions on an identical toy with only one function.

The probabilistic model of epistemic trust thus formalizes social learning as a problem of *simultaneous* learning about the world *and* about informants. Literature on epistemic trust typically focuses on cases where you know the answer (Corriveau & Harris, 2009; Corriveau et al., 2009; Pasquini et al., 2007) and literature on children's pedagogical reasoning tends to focus on the case where you know about the informant (Bonawitz, Shafto, et al., 2011; Buchbaum et al., 2011). Probabilistic modeling of joint inferences about the world and informants provide a framework for modeling learning from and about

informants as a dynamic, interactive Theory-based process (Landrum, Eaves, & Shafto, 2015).

Summary

Developmental science has suggested that children—far from limited, broken versions of adults—are in fact powerful learners capable of remarkable feats of inference. Ideal observer approaches, instantiated in Theory-based Bayesian models, have argued that children succeed by probabilistic inference over structured knowledge representations. We have argued that these approaches are missing an important element—other people.

Research from word learning and pedagogical reasoning shows that children take the origin of data into account when learning. Given the exact same evidence, they draw different conclusions. And, these inferences can be understood through ideal observer analyses of the *social* situation. People choose evidence purposefully – in word learning, by choosing to label instances of the concept, and in pedagogical reasoning, by tailoring the amount of information provided to the complexity of the situation – and children’s behavior is consistent with models that reason about the implications of people’s choices for learning.

Moreover, when observing evidence, children are not just learning *from* others, but also *about* others. Research on epistemic trust has shown that young children select informants based on the goodness of the information that they provide. And, the same models that explain children’s learning from pedagogically selected data have been extended to explain children’s learning about the informants selecting the data.

A consensus is emerging in which children’s understanding of other people plays a key role in explaining children’s powerful learning abilities. This consensus draws strength from the combination of empirical results and computational models derived from ideal observer analyses of social learning. There are many important open questions, such as

whether and how social ideal observer analysis should inform our understanding of children's actions and choices and practices in learning sciences more generally. These will require new research and more elaborate models. In the interim, we can conclude that to understand power of children's learning, it is important to investigate it in a social context.

References

- Baillargeon, R., Scott, R. M., & He, Z. (2010). False-belief understanding in infants. *Trends in cognitive sciences*, 14(3), 110-118.
- Behr, M. (1980). How children view the equals sign. *Mathematics Teaching*, 92, 13-15.
- Bloom, L. (2000). The intentionality model of word learning: How to learn a word, any word. *Becoming a word learner: A debate on lexical acquisition*, 19-50.
- Bonawitz, E. B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental psychology*, 48(4), 1156.
- *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.
- *This paper provides a computational analysis showing how teaching effects exploration and discovery, and empirical support for the prediction that children tend to limit their exploratory play to a target function highlighted during direct instruction.
- Bovens, L., & Hartmann, S. (2004). Bayesian epistemology. *OUP Catalogue*.
- Breazeal, C., & Scassellati, B. (2002). Robots that imitate humans. *Trends in cognitive sciences*, 6(11), 481-487.
- 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.
- Butterfield, J., Jenkins, O. C., Sobel, D. M., & Schwertfeger, J. (2009). Modeling aspects of theory of mind with Markov random fields. *International Journal of Social Robotics*, 1(1), 41-51.
- Carey, S. (1978). The child as a word learner. In M. Halle, J. Bresnan, & G. Miller (Eds.), *Linguistic theory and psychological reality*. Cambridge, MA: MIT Press.
- Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going with the flow preschoolers prefer nondissenters as informants. *Psychological Science*, 20(3), 372-377.
- Corriveau, K., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and recent accuracy. *Developmental science*, 12(3), 426-437.
- Diamond, A., & Kirkham, N. (2005). Not quite as grown-up as we like to think parallels between cognition in childhood and adulthood. *Psychological Science*, 16(4), 291-297.
- Eaves Jr, B. S., & Shafto, P. (2012). Unifying pedagogical reasoning and epistemic trust. *Rational constructivism in cognitive development. Advances in child development and behavior*, 43, 295-320.

- Flavell, J. H. (1971). Stage-related properties of cognitive development. *Cognitive Psychology*, 2(4), 421-453.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naive theory of rational action. *Trends in cognitive sciences*, 7(7), 287-292.
- Gold, E. M. (1967). Language identification in the limit. *Information and Control*, 16, 447-474.
- *Gopnik, A., & Schulz, L. (2004). Mechanisms of theory formation in young children. *Trends in cognitive sciences*, 8(8), 371-377.
- *A review of the causal learning literature (by two leaders in the field) presenting a summary of evidence showing that young children learn from conditional probabilities and causal interventions in a way consistent with Bayes net learning algorithms.
- Gopnik, A., & Wellman, H. M. (2012). Reconstructing constructivism: Causal models, Bayesian learning mechanisms, and the theory theory. *Psychological bulletin*, 138(6), 1085.
- Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in cognitive sciences*, 14(8), 357-364.
- Gweon, H., Pelton, H., Konopka, J. A., & Schulz, L. E. (2014). Sins of omission: children selectively explore when teachers are under-informative. *Cognition*, 132(3), 335-341.
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning over-hypotheses with hierarchical Bayesian models. *Developmental science*, 10(3), 307-321.
- Kemp, C., & Tenenbaum, J. B. (2008). The discovery of structural form. *Proceedings of the National Academy of Sciences*, 105(31), 10687-10692.
- Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers. *Developmental science*, 14(1), 106-111.
- Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, 83(2), B35-B42.
- Klahr, D., & Dunbar, K. (1989). Developmental differences in scientific discovery processes. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert A. Simon* (pp. 109 -143). Hillsdale, NJ: Erlbaum.
- Kuhn, D. (2000). Metacognitive development. *Current directions in psychological science*, 9(5), 178-181.
- Kuhn, D., Amsel, E., & O'Laughlin, M. (1988). *The development of scientific thinking skills*. Orlando, FL: Academic Press

- Landrum, A. R., Eaves, B. S., & Shafto, P. (2015). Learning to trust and trusting to learn: a theoretical framework. *Trends in cognitive sciences*, 19(3), 109-111.
- Markman, E. M. (1990). Constraints children place on word meanings. *Cognitive Science*, 14(1), 57-77.
- Mascaro, O., & Sperber, D. (2009). The moral, epistemic, and mindreading components of children's vigilance towards deception. *Cognition*, 112(3), 367-380.
- McNeil, N. M. (2008). Limitations to teaching children $2 + 2 = 4$: Typical arithmetic problems can hinder learning of mathematical equivalence. *Child Development*, 79(5), 1524-1537.
- McNeil, N. M., & Alibali, M. W. (2004). You'll see what you mean: Students encode equations based on their knowledge of arithmetic. *Cognitive science*, 28(3), 451-466.
- McNeil, N. M., & Alibali, M. W. (2005). Why won't you change your mind? Knowledge of operational patterns hinders learning and performance on equations. *Child development*, 76(4), 883-899.
- Perfors, A., Tenenbaum, J. B., Griffiths, T. L., & Xu, F. (2011). A tutorial introduction to Bayesian models of cognitive development. *Cognition*, 120(3), 302-321.
- Piaget, J. (1965). The stages of the intellectual development of the child. *Educational psychology in context: Readings for future teachers*, 98-106.
- Pasquini, E. S., Corriveau, K. H., Koenig, M., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental psychology*, 43(5), 1216.
- Quine, W. V. (1969). O.(1960) Word and object. *Cambridge, Mass.*
- Regier, T. (1996). *The human semantic potential: Spatial language and constrained connectionism*. MIT Press.
- Repacholi, B. M., & Gopnik, A. (1997). Early reasoning about desires: evidence from 14-and 18-month-olds. *Developmental psychology*, 33(1), 12.
- Saffran, J. R. (2003). Statistical language learning mechanisms and constraints. *Current directions in psychological science*, 12(4), 110-114.
- Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27-52.
- Schulz, L. E., Bonawitz, E. B., & Griffiths, T. L. (2007). Can being scared cause tummy aches? Naïve theories, ambiguous evidence, and preschoolers' causal inferences. *Developmental psychology*, 43(5), 1124.
- Seo, K. H., & Ginsburg, H. P. (2003). "You've got to carefully read the math sentence...": Classroom context and children's interpretations of the equals sign.

*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.

* Introduces a computational model of epistemic trust, formalized as a causal Bayes Net relating beliefs, intent, and actions. The model generalizes previous accounts of learning from others to provide an account of learning about others. They model three representative studies from the epistemic trust literature. Results show that behavioral differences between 3 and 4-year-olds are explained by an increased ability to reason about informants' helpfulness.

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.

*Shafto, P., Goodman, N. D., & Frank, M. C. (2012). Learning from others the consequences of psychological reasoning for human learning. *Perspectives on Psychological Science*, 7(4), 341-351.

*Introduces the "learning from others" computational framework. The paper proposes that learners reason about other people's choices. The implications are illustrated by contrasting three levels of a taxonomy of social learning: accidental actions by not knowledgeable actors, purposeful actions by knowledgeable actors, and communicative actions by knowledgeable actors. Simulations show that the exact same evidence leads to very different inferences across these contexts.

Shafto, P., Goodman, N. D., & Griffiths, T. L. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive psychology*, 71, 55-89.

Talwar, V., & Lee, K. (2008). Social and cognitive correlates of children's lying behavior. *Child development*, 79(4), 866-881.

*Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in cognitive sciences*, 10(7), 309-318.

*An important overview by a number of the leaders in the theory-based Bayesian modeling approach, arguing that both strong constraints from structured knowledge and powerful statistical learning can be combined to understand human inductive inference.

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.

Thomaz, A. L., Berlin, M., & Breazeal, C. (2005, August). An embodied computational model of social referencing. In *Robot and Human Interactive Communication, 2005. ROMAN 2005. IEEE International Workshop on* (pp. 591-598). IEEE.

Tummeltshammer, K. S., Wu, R., Sobel, D. M., & Kirkham, N. Z. (2014). Infants track the reliability of potential informants. *Psychological science*, 25(9), 1730-1738.

Wellman, H. M., & Liu, D. (2004). Scaling of theory-of-mind tasks. *Child development*, 75(2), 523-541.

Woodward, A. L. (2005). The infant origins of intentional understanding. *Advances in Child Development and Behaviour*, 33, 229-262.

*Xu, F., & Tenenbaum, J. B. (2007a). Word learning as Bayesian inference. *Psychological review*, 114(2), 245.

*The authors present a Bayesian framework for word learning, exploring both children and adults' predictions about word meanings at multiple levels of a taxonomy. It presents a particularly nice overview of the shortcomings of past theories of word learning, as well as a compelling set of experiments to test quantitative and qualitative predictions of their model.

Xu, F., & Tenenbaum, J. B. (2007b). Sensitivity to sampling in Bayesian word learning. *Developmental science*, 10(3), 288-297.

Yu, C., & Ballard, D. H. (2007). A unified model of early word learning: Integrating statistical and social cues. *Neurocomputing*, 70(13), 2149-2165.

Zelazo, P. D. (2004). The development of conscious control in childhood. *Trends in cognitive sciences*, 8(1), 12-17.

Zelazo, P. D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive development*, 11(1), 37-63.

Figure 1. Depiction of the learner's inferences (on the left) that depend on their beliefs about how the teacher selects data (on the left). The teacher must reason about the beliefs of the learner in order to generate the most effective data.

Acknowledgements

This research was supported by NSF grant DRL-1149116 to PS.