

Perspectives on Psychological Science

<http://pps.sagepub.com/>

Learning From Others : The Consequences of Psychological Reasoning for Human Learning
Patrick Shafto, Noah D. Goodman and Michael C. Frank
Perspectives on Psychological Science 2012 7: 341
DOI: 10.1177/1745691612448481

The online version of this article can be found at:
<http://pps.sagepub.com/content/7/4/341>

Published by:



<http://www.sagepublications.com>

On behalf of:



[Association For Psychological Science](http://www.sagepub.com/content/7/4/341)

Additional services and information for *Perspectives on Psychological Science* can be found at:

Email Alerts: <http://pps.sagepub.com/cgi/alerts>

Subscriptions: <http://pps.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Learning From Others: The Consequences of Psychological Reasoning for Human Learning

Patrick Shafto¹, Noah D. Goodman², and Michael C. Frank²

¹University of Louisville and ²Stanford University

Perspectives on Psychological Science
7(4) 341–351

© The Author(s) 2012

Reprints and permission:

sagepub.com/journalsPermissions.nav

DOI: 10.1177/1745691612448481

http://pps.sagepub.com



Abstract

From early childhood, human beings learn not only from collections of facts about the world but also from social contexts through observations of other people, communication, and explicit teaching. In these contexts, the data are the result of human actions—actions that come about because of people’s goals and intentions. To interpret the implications of others’ actions correctly, learners must understand the people generating the data. Most models of learning, however, assume that data are randomly collected facts about the world and cannot explain how social contexts influence learning. We provide a Bayesian analysis of learning from knowledgeable others, which formalizes how learners may use a person’s actions and goals to make inferences about the actor’s knowledge about the world. We illustrate this framework using two examples from causal learning and conclude by discussing the implications for cognition, social reasoning, and cognitive development.

Keywords

Bayesian model, learning, reasoning, social cognition

Children are often compared to scientists, but even a perfect scientist, using experiments alone, would struggle to rediscover all of human knowledge in the span of one lifetime. How then are children able to acquire a good fraction of this knowledge in just a few years? The answer must be that children do not discover everything they learn—they use their ability to reason intuitively about other people to learn what others already know. Our goal in this article is to sketch a formal analysis of learning from knowledgeable others that is based on Bayesian inference and a careful examination of the kinds of goals that give rise to human actions. We begin by addressing the need of learners to consider the particular goals of people in their environment.

Imagine that while living in Paris, you decide to search for the best cup of coffee in the city. As you wander, you find yourself a good distance away from your neighborhood. You observe three pieces of evidence: First, a man wearing a baseball cap and an “I Heart Paris” T-shirt (obviously a tourist) turns into Cafe 1, buys a coffee, and looks down at his cup. Second, Véronique, a woman from your neighborhood, enters Cafe 2 to get a coffee, and looks down at her cup. Third, Madeleine, another woman from your neighborhood, goes into Cafe 3 and buys a cup. Madeleine sees you, and she nods at the coffee.

Which cafe would you think has the best coffee? You can infer very little about the coffee at Cafe 1, because the tourist likely chose the cafe at random. Cafe 2 was visited by a local, but maybe Véronique was just strapped for time and grabbed a

cuppa wherever she could. Without knowing her motivations, you can’t tell whether she had gone out of her way to go to this particular cafe (though you may guess at her motivations, and hence guess her beliefs about the cafe). On the other hand, at Cafe 3, Madeleine telegraphed her intentions to you: She was there for the coffee, and she congratulated you for figuring out a local secret. Although nothing is certain (for example, Madeleine could have terrible taste), Cafe 3 seems likely to have the best coffee, and Cafe 2 is likely to have better coffee than Cafe 1. In this article, we propose a formal framework for understanding why reasoning based on observations of three different types of actions—randomly chosen, goal-directed, and communicative actions—leads to qualitative differences in learning.

Learning From Limited Data

Traditional approaches to understanding how people learn about the world so quickly and robustly have focused on the nature of human representations and a priori biases about the physical world, investigating the biases that allow rapid and accurate learning from a limited amount of data. In concept

Corresponding Author:

Patrick Shafto, University of Louisville, 103 Life Sciences Building, Louisville, KY 40205

E-mail: p.shafto@louisville.edu

learning, the problem has been characterized as “carving nature at its joints,” and the debates have been over the types of representations that support these abilities (Gelman, 1996; Keil, 1989; Mandler, 1992; Medin & Schaffer, 1978; Nosofsky, 1984; Posner & Keele, 1968; Rosch, 1978; Rosch & Mervis, 1975). Similarly, in causal learning, the learning problem has been viewed as one of discovering the laws that govern physical world (Michotte, 1963), and debates have been over the representational and inferential mechanisms that support these abilities (Cheng, 1997; Gopnik, Glymour, Sobel, Schulz, & Danks, 2004; Griffiths & Tenenbaum, 2005; Rescorla & Wagner, 1972). These approaches have successfully described reasoning in contexts in which the data are (assumed to be) observed objectively. For example, in logical inference—if A , then B , such that the observation of A implies that B is true—the conditions of observation of A are assumed to be irrelevant to the truth of B . A is simply a given.

Much of the evidence observed in human learning does not have this character. Evidence is often provided *by someone, for some purpose*. A potter throwing a pot, a friend fiddling with her iPod, a parent demonstrating how to tie a shoe, and a teacher conveying a mathematical concept are all creating evidence with a particular goal in mind. The goals of the pot thrower and the iPod fiddler have to do with the world—they are trying to get matter or artifacts to conform to their desires. The goals of the parent and teacher have to do with the observer; they want the observer to learn a particular fact, skill, or concept. In each of these cases, it is possible for people to learn a tremendous amount from only a small number of data points.

Yet for many traditional formal approaches, learning based on limited data is nearly impossible except in the most circumscribed domains (Gold, 1967; Savage, 1951; Wolpert & Macready, 1997; Zinkevich, 2003). Most famously, Gold (1967) proved that the learning any formal language that is sufficiently broad to express an infinite range of possible sentences (so that the sentences could not possibly be enumerated one by one) is impossible. Strikingly, this proof suggests that human languages are unlearnable! The conflict between human intuition and formal analysis creates a puzzle: How is human learning so quick and successful when formal learning frameworks suggest that it should be slow and fatally conservative?

Social Learning Contexts

We believe that the key to this puzzle lies in the assumptions that traditional learning theory makes about the *conditions of observation*. For instance, Gold’s proof assumes that the data points for learning are selected by an adversary. Imagine trying to learn which cafe has the best coffee when everyone is deliberately trying to mislead you! When this assumption is relaxed even slightly and data are assumed to be sampled randomly, the language-learning problem studied by Gold is found to be considerably less difficult (Horning, 1969). Nonetheless, learning from randomly sampled data—as opposed to

data from an adversary who wants to “fool” the learner (in Gold’s words)—can still be quite difficult. Paris is a big city—imagine trying to learn where the best coffee in Paris is by randomly sampling cafes. It would take a very long time. Yet most models assume that learning depends on this sort of random sampling (Anderson, 1991; Goodman, Tenenbaum, Feldman, & Griffiths, 2008; Kruschke, 1992; Love, Medin, & Gureckis, 2004; Medin & Schaffer, 1978; Nosofsky, 1984; Nosofsky, Palemeri, & McKinley, 1994; Pothos & Chater, 2002).

Research on human learning has painted a very different picture of how data points are selected. A wide variety of approaches have pointed to people’s intentions as an important factor in learning, highlighting the fact that data are chosen rather than random (Bruner, 1966; Vygotsky, 1978) and that observed data are often the consequence of goal-directed actions (Bandura, Ross, & Ross, 1961; Gergely & Csibra, 2003; Meltzoff & Moore, 1977) or of intentional communication or teaching (Coady, 1992; Csibra & Gergely, 2009; Harris, 2002; Tomasello, 1999; Tomasello, Carpenter, Call, Behne, & Moll, 2005). For instance, Csibra and Gergely (2009) suggested that young children’s interpretation of observed data changes fundamentally according to whether the demonstrator engages the child with ostensive cues—saying the child’s name, using child-directed speech, shifting gaze between the child and the object of the demonstration—prior to the demonstration. According to this account, these cues lead children to assume that the demonstrated data are not randomly sampled but purposefully sampled to support broad generalizations.

The Goals of This Article

Psychological intuitions about what makes human learning so effective will remain exactly that—intuitions—until we can formalize and test whether they actually lead to faster, more robust learning. Indeed, one explanation for the disagreement is that there has not been a framework in which these different proposals can be formalized and their implications tested. We describe a computational theory of learning from other people’s actions, with the goal of reconciling the formal literature on learning with psychological accounts about learning from other people. Our approach is inspired by classic work on attribution theory that characterized social attributions as inferences about unseen traits or goals (Jones & Davis, 1965; Kelley, 1967), and it builds on recent work that has formalized “theory of mind” as the inverse of rational decision making (Baker, Saxe, & Tenenbaum, 2009). Bayesian methods describe how learners can work backward from an agent’s actions to his or her goals and beliefs; when a learner has reason to believe that the agent is knowledgeable, the learner can use this information to strengthen his or her inferences about the structure of the world.

We use three different learning contexts to illustrate how learning may differ as a consequence of this kind of intuitive psychological reasoning: (a) learning from physical evidence,

(b) learning from goal-directed action, and (c) learning from communicative action (see Fig. 1). We discuss how to formalize different kinds of goal-directed actions in our framework, as well as how different assumptions about the agent's goals may cause learners to make qualitatively different inferences. We conclude by discussing implications for theories of learning, cognitive development, and the relationship between social and cognitive psychology.

A Framework for Modeling Goal-Directed Actions

We want to formalize how an actor's goal can be used by learners to guide inference: Why does observing Madeleine's nod lead to a strong inference about the quality of the coffee she is drinking, whereas observing a tourist look down while drinking his coffee does not? We can formalize this inference in a way that depends on the actor's knowledgeability, as a relation among the actions that learners observe, a , the effects of those actions, e , the actor's goal, g , and the learner's hypotheses about the world, h .

The objective of the learner is to infer the probability of a hypothesis being true—in this case, given a set of actions, events, and goals (the *posterior probability*). In our example, what should we believe about the coffee? Using Bayes's rule, the posterior probability of the hypothesis, $P(h | a, e, g)$, can be factored into three terms.

$P(h)$ is the prior probability of the hypothesis, which represents the learner's expectations entering the situation. Is good coffee common in Paris, or is it rare?

$P(a | g, h)$ is the likelihood of the action given the goal, assuming the hypothesis is true. This likelihood represents the degree to which the action is consistent with the hypothesis and the actor's goal. For example, the tourist's action—purchasing coffee at this shop—is consistent with the goal of drinking coffee and with either hypothesis: The coffee could be either good or bad.

$P(e | a, h)$ is the likelihood of the effect given the action and hypothesis. This likelihood depends on the underlying structure of the world; assuming the hypothesis is true, how likely is the effect to follow from the action? In our example, the tourist would likely obtain coffee of uncertain quality.

We can then use Bayes's rule to express the relationship between these factors:

$$P(h | a, e, g) = \frac{P(e | a, h) P(a | g, h) P(h)}{\sum_{h'} P(e | a, h') P(a | g, h') P(h')} \quad (1)$$

The degree to which we believe the hypothesis depends on our prior beliefs, the choice of actions in light of the goal and the hypothesis, and the effects of the chosen actions. The denominator in the fraction is the sum over all possible hypotheses—often referred to as the *normalizing constant*—which ensures that the posterior probability of each hypothesis reflects its probability relative to all other possibilities.

In this article, we are concerned with the third term, $P(a | g, h)$, the likelihood of an action given the hypothesis and the actor's goals. This term is fundamentally *psychological*: It encodes beliefs about the actor's knowledge and different assumptions about the goals underlying his or her actions.

We can formalize how, for example, the minimally goal-directed (and likely naive) actions of the tourist convey meaning that is different from the actions of Véronique and Madeleine; in other words, we can formalize how these actors' goals lead learners to different interpretations of superficially similar actions. We are specifying how people's intuitive theories of psychology can be leveraged to facilitate learning (for full mathematical details, see Appendix A).

The bottom line here is that to formalize learning in social contexts, we must specify what the goals are and how likely different actions are to lead to the desired goal. Critically, in our framework, different kinds of goals will lead to different choices of actions; actions are purposeful. A learner who is aware of the actor's goals can use this information, together with the choice of actions, to infer what the actor knows about the world. Indeed, even a guess about the actor's goal is often sufficient for the learner to infer what the actor knows.

People may have a great variety of different goals. Returning to our initial example, a goal may be merely to bring about an effect. I may walk to a far-off neighborhood because I want some good coffee. Alternately, my goal may be to communicate or teach someone something: that *this* is good coffee (Clark, 1996; Grice, 1975; Sperber & Wilson, 1986). In the



Fig. 1. A schematic of three social learning contexts. When learning from physical evidence (a), learners reason about the implications of the data for different hypotheses. When learning from goal-directed action (b), learners reason about the actor's goal with respect to the world. When learning from communicative action (c), learners reason about which belief the actor intends the learner to have.

first case, my goal is to affect a state of the (physical) world. In the second case, my goal is to affect a change in another person's beliefs. These two different kinds of goals have different consequences for learning. In the following section, we introduce two scenarios from related experimental work, which we use to investigate the distinct consequences of rational and communicative goals.

Implications of Goal-Directed Actions for Learning

We illustrate the effects of intuitive psychological reasoning on learning with two simple causal-learning scenarios: one involving inferring the causal structure of a device from observed data, and the other involving inferring the number of causal properties of a novel object from observed data. The first scenario, taken from Goodman, Baker, and Tenenbaum (2009), we call "Bob's Box." This scenario involves a learner observing a box with two buttons on the top. Bob presses the buttons

simultaneously, and a light illuminates (see the left column of Fig. 2). The learner must infer the causal structure—which buttons are necessary to cause the effect: one button, the other button, both buttons, or neither button.

The second scenario, taken from Bonawitz et al. (2011), we call "Tim's Toy." This scenario involves a learner observing a complex toy with an unknown number of causal properties. Tim reaches into the purple tube and pulls a knob that elicits a squeak, and then stops (see Fig. 2, right column). With Tim's Toy, the learner observes a set of cause-effect relationships and must infer the causal structure of the toy.

In the following sections, we discuss the intuitions that underlie the predictions of our computational framework (for quantitative predictions, see Fig. 2; for full mathematical details, see Appendix B). We consider actors who choose actions randomly, actors who choose actions to bring about effects, and actors who choose actions to communicate knowledge about the world to others (e.g., to teach). To highlight the role of social inference in each of these scenarios, we assume




Fig. 2. Two causal learning scenarios. In "Bob's Box," the learner observes a box with two buttons. Bob presses both buttons simultaneously and a light illuminates. In "Tim's Toy," the learner observes a complex toy with an unknown number of non-obvious causal properties. Tim pulls a knob on the purple tube, which elicits a squeak, then does not perform any additional actions. The graphs show learning predictions for the two scenarios in three different contexts. The x-axes show possible hypotheses; the y-axes indicate probabilities. Note that in each scenario, learners are assumed to see the same evidence; only the context varies. Across all contexts, we assume that effects must have a cause and that causal relationships are deterministic. The black bars in each graph indicate the true state of the world. For Bob's Box, both goal-directed-action and communicative-action contexts lead to stronger inferences than does learning from evidence alone. For Tim's Toy, communicative-action contexts lead to stronger inferences than do either goal-directed action or evidence alone. Thus, we see qualitative dissociations in learning across social contexts.

that all hypotheses are equally likely and that causal relationships are deterministic. These assumptions are not meant to represent people's beliefs about these particular causal-learning problems; rather, these assumptions allow for pedagogical and computational simplicity and allow us to highlight the dramatic effects that social inference can have on learning.

Bob's Box: inferring the causes of an effect

Consider Bob's Box (left column of Fig. 2). Possible hypotheses, h , are that Button A causes the light, that Button B causes the light, that Buttons A and B together cause the light, and that neither button causes the light. With these details, we can ask the following question: What would a learner infer from watching Bob press both buttons?

Physical evidence. Based on physical evidence alone, all three hypotheses are equally likely. (See the top left graph in Fig. 2.) This reflects the fact that the observed actions and effects are ambiguous (i.e., the evidence is confounded). Given that Bob pressed both Button A and Button B and elicited the effect, all hypotheses are possible: Button A could be the cause of the effect, as could Button B or Buttons A and B together.

Goal-directed action. Knowing that Bob knows how the toy works and that his goal is to turn on the light allows the learner to disambiguate the possible hypotheses. (See the middle left graph in Fig. 2.) After all, if Bob was knowledgeable about the toy and wanted to turn on the light, he would definitely press both buttons if the hypothesis that Buttons A and B cause the effect were true. Of course, there is also some chance that, if the true hypothesis was that Button A alone (or Button B alone) causes the effect, he would press both buttons—that action would still elicit the effect—but in either of those cases, because pressing both buttons would not be Bob's only option, that action would be improbable (see Appendix A, Equation 3; more actions that lead to the effect results in a larger denominator and lower overall probability). Because the learner knows Bob's goal, he or she infers that Bob's choice of action is most consistent with the hypothesis that Buttons A and B are both necessary. This leads to the inference that the most likely hypothesis is that Buttons A and B together cause the effect, though the hypotheses that Button A alone or Button B alone causes the effect remain plausible. This is an important qualitative shift; unlike physical evidence, goal-directed action supports learning about the true hypothesis from this single observation.

Communicative action. Communicative action leads to the strongest inferences about the true hypothesis. (See the bottom left graph in Fig. 2.) Bob has looked at you, drawn your attention to the toy, then pressed both Button A and Button B. Why should inferences about this action differ from inferences about a goal-directed action? As above, the fact that pressing

Buttons A and B led to the effect is consistent with all three hypotheses. However, Bob is teaching you how the toy works; he is inviting you to think about how the toy works. Bob could press both Button A and Button B if the true cause of the effect was only Button A; however, Bob chooses his action on the basis of how it would affect the learner's beliefs. This means that Bob's choice (see Appendix A, Equation 3) depends on the learner's inferences (see Appendix A, Equation 2). To avoid confusing the learner, Bob would choose to press Button A alone only if the true cause of the effect was Button A alone. For the learner, this means that, given Bob's intention to teach, his actions are consistent only with the hypothesis that Buttons A and B are both necessary to elicit the effect, leading to a strong inference that this is the correct hypothesis (see Appendix A, Equation 2).

Tim's toy: learning the number of latent causes

Next, consider Tim's Toy (see the right column of Fig. 2), a complicated collection of tubes and coils and things that appear to be knobs and buttons. How many cause-effect relationships are there? Possible hypotheses are that there are no cause-effect relationships, that there is one cause-effect relationship, and that there are two or more cause-effect relationships. For purposes of demonstration, we consider only zero, one, or two relationships, all of which have equal prior probabilities.

Physical evidence. Physical evidence alone leaves considerable uncertainty about the true hypothesis. (See the top right graph in Fig. 2.) The observation that one action leads to an effect rules out the hypothesis of no action-effect pairs. However, the learner cannot be certain that this is the only action-effect pair, and is thus uncertain about the two remaining hypotheses.

Goal-directed action. Goal-directed actions do not change learners' inferences. (See the middle right graph in Fig. 2.) Why not? Imagine that Tim states, "I love squeaking—I'm going to squeak my toy," and then pulls a knob, eliciting a squeak. Like physical evidence, this goal-directed action rules out the hypothesis of no cause-effect pairs. However, Tim's intent to bring about squeaking does not provide information about whether other (relatively unloved) effects do or do not exist.

Communicative action. Communicative action does lead to confident inferences about the true hypothesis. (See the bottom right graph in Fig. 2.) Tim clearly intends to teach you about the toy and chooses to demonstrate that pulling the knob causes squeaking, but he demonstrates nothing else. As in the previous scenarios, the hypothesis that there are no action-effect relationships is ruled out. If the true hypothesis was a single action-effect pair, then Tim could not perform any additional demonstrations. Alternatively, if the true hypothesis was

two action-effect pairs, Tim's demonstration would be consistent with the hypothesis but inconsistent with his goal—to teach you how the toy works—so this choice would be less probable (see Appendix A, Equation 3; note that the goal $P(g | a, h)$ is for the learner to infer the correct hypothesis; i.e., $P(g | a, h)$ is described by Equation 2). Tim's reasoning about the learner's inferences should lead Tim to discount alternative, but possible, choices of actions that would provide the learner with evidence. Therefore, from the learner's perspective, Tim's choice to demonstrate only a single action-effect pair is sensible only if there really is a single cause-effect relationship—his demonstration should lead to the inference that there are no other pairs to be discovered.

Empirical Evidence

Recent studies have tested the predictions of the computational framework. Specifically, the Bob's Box example was based on a study by Goodman et al. (2009) that investigated the different implications of goal-directed action and physical evidence for learning. Similarly, the Tim's Toy example was based on a study by Bonawitz et al. (2011) that investigated the different implications of communicative action, goal-directed action, and physical evidence for learning. We review those studies and their findings, as well as convergent findings that support the predictions of our computational framework.

Goodman et al. (2009) investigated adults' inferences from goal-directed action and physical evidence. As in our Bob's Box example, participants were presented with scenarios in which a person interacted with a toy that had two buttons and a light. Either the person was knowledgeable about the toy and decided to turn the light on (goal-directed action), or the person was not knowledgeable about the toy. In both cases, the person pressed both buttons and the light turned on. Participants were asked what made the light turn on. Results showed that people in the goal-directed-action condition inferred that both buttons A and B were necessary to elicit the effect, whereas participants in the physical-evidence condition were relatively unsure about the cause of the light.

This distinction between physical evidence and goal-directed action helps to explain findings of pervasive overimitation in the developmental literature (Goodman et al., 2009). A variety of recent studies have shown that, when learning, children appear to be overly faithful to actions produced by demonstrators, even repeating actions that are clearly superfluous to eliciting the desired outcome (Horner & Whiten, 2005; Lyons, Young, & Keil, 2007; Meltzoff, 1995). Although researchers have offered a variety of explanations for such findings, the explanations rely on ad-hoc mechanisms (e.g., automatic causal encoding) to explain behavior. Our framework offers a different account, whereby learners leverage intuitive psychological inferences to support learning in otherwise ambiguous situations. From our perspective, rather than being a strange anomaly in the human cognitive system, overimitation is a sensible approach to learning when learners are

surrounded by knowledgeable others (cf. Krueger & Funder, 2004).

Bonawitz et al. (2011) investigated children's inferences from communicative action, goal-directed action, and physical evidence. As in our Tim's Toy example, children observed scenarios in which a person interacted with a novel, complex-looking toy. In one study, the person either was knowledgeable about the toy and engaged the child via ostensive cues (see Csibra & Gergely, 2009) or was not knowledgeable about the toy. In both conditions, the child observed that pulling a purple knob led to squeaking, either as a result of an intentional demonstration (communicative action) or an accidental one (physical evidence). Children were then given the toy and allowed to play with it. If the communicative action led children to infer that only one cause-effect relationship was present, then they should have engaged in less exploratory play than children exposed only to physical evidence. The results, including the number of actions tried and the number of nondemonstrated causal relations discovered by children, confirmed that communicative action led to decreased exploration. A subsequent study contrasted goal-directed action and communicative action, showing that communicative action led to decreased exploration relative to goal-directed action. Together, these results confirm the prediction that communicative action did in fact lead to a robust inference that only the demonstrated causal relationship exists.

Other recent studies have also investigated the implications of teaching and communicative actions. Shafto and Goodman (2008) showed that adults draw different inferences in concept learning when a knowledgeable teacher, as opposed to a naïve demonstrator, selects the data. Their results suggested that learners infer that teachers are selecting diverse examples, which in turn supports the inference that narrower concepts are more probable than are concepts that include examples not chosen by the teacher. Similarly, recent findings from research on mathematics learning have suggested that the use of nondiverse mathematics problems may be the cause of erroneous inferences about the concept of equality. McNeill (2008) notes that early math problems are typically presented in the $X + Y = Z$ format, as opposed to the equally correct $Z = X + Y$ format (e.g., $2 + 1 = 3$ and $3 + 2 = 5$, but not $3 = 2 + 1$ and $5 = 3 + 2$). The use of nondiverse problems leads children to make sensible, but ultimately incorrect, inferences about what the equals symbol means (McNeill, 2008).

Language comprehension and language learning

Although the framework described above is only beginning to be applied to language comprehension and language learning, there are a number of examples that support its predictions about learning from communicative actions. Xu and Tenenbaum (2007) described a study in which they presented participants with either examples chosen by an ignorant learner (much like the cafe chosen by our tourist) or examples chosen

by a knowledgeable teacher. Each learner observed as the objects were chosen and labeled. Consistent with the predictions of our framework, results revealed that learners generalized the labels conservatively when a teacher had chosen the examples and generalized the labels more broadly when a naïve learner had chosen the examples.

Ideas about speakers' goals and intentions have also been influential in theories of language acquisition. Many theorists have proposed that an understanding of speakers' communicative intentions is key in understanding (Clark, 1996) and acquiring language (Bloom, 2002; Clark, 2003). Frank, Goodman, and Tenenbaum (2009) used a Bayesian model to capture this intuition in the context of associative word-learning tasks. This model, which attempted to jointly infer speakers' communicative intentions and meanings of words, performed better than simple associative models across a range of factors, including learning from corpus data and fitting human performance.

In addition, in recent work, Frank and Goodman (in press) used a model related to the communicative model described above to capture ideas from Gricean pragmatics. This work demonstrated that the basic principles underlying the communicative model effectively captured the Gricean maxim "be informative." The researchers compared the predictions of this model with adult data on production and comprehension of words in ambiguous situations and found a tight quantitative correspondence. Although this work is relatively new, it nevertheless suggests that the kind of framework we have discussed here can be applied productively to language learning.

Summary: Implications of Goal-Directed Actions for Learning

In each of our scenarios, we held constant the data that the learner observed and derived predictions from the computational framework. We considered the inferences that were afforded in three different contexts: learning from physical evidence, learning from goal-directed action, and learning from communicative action. In each case, we showed how actors' goals affected the inferences that could be drawn from the same evidence. We outlined recent research suggesting that both children and adults use other people's goals to support learning. Furthermore, our computational framework suggested that these inferences were rational consequences of considering actions as being goal-directed. As in our coffee-finding example, intuitive psychological reasoning provided information that dramatically affected learning.

We have focused on actions driven by two kinds of social goals: goal-directed actions and communicative actions. The strength of our framework is in showing that these goals can be formalized. Data that result from other people's actions are ultimately a consequence of their goals. The goal may have to do with a state of the world, as in goal-directed actions. The goal may also have to do with another person.

These are two examples of others' goals that affect inference, but they are not the only possibilities. For instance, a

person may be interested in deceiving or lying instead of teaching or (cooperatively) communicating. These kinds of goals can be straightforwardly formalized within this framework (Shafto, Eaves, Navarro, & Perfors, 2012; Warner, Stoess, & Shafto, 2011). However, we do not yet have definitive lists of the kinds of goals that are relevant. Considerable work remains in identifying, cataloging, and organizing the myriad goals that people may have and their respective implications for learning.

For the purposes of exposition, we have assumed that the actors whom learners observe are knowledgeable, but this assumption is not necessary. Using a model of how knowledgeable and naïve (i.e., not knowledgeable) individuals choose actions, it is straightforward to formalize how learners could infer who is knowledgeable (and, similarly, who is well-intentioned). Indeed, Shafto et al. (2012) proposed a model of this problem and brought it to bear on research on the development of epistemic trust (Corriveau, Fusaro, & Harris, 2009; Corriveau & Harris, 2009; Koenig & Harris, 2005; Pasquini, Corriveau, Koenig, & Harris, 2007). In this literature, children's success in choosing reliable informants has been attributed to their ability to monitor which informants are knowledgeable. By contrasting the predictions of two models that formalize inference about either informants' knowledge or their knowledge and intent, Shafto et al. (2012) argued that 4-year-olds' behavior is best explained by joint inference about informants' knowledge and intent, as opposed to inference about their knowledge alone—the standard interpretation. A similar approach shows that 3-year-olds' behavior is best explained by inference about informants' knowledge alone, suggesting developmental changes in behavior on these tasks.

Recent research has also begun to explore the relationship between epistemic trust and attachment theory. Corriveau et al. (2009) showed, for example, that children with avoidant attachment patterns are equally likely to trust information provided by a stranger as they are to trust information provided by their caregiver, unlike children with secure attachment patterns, who trust their caregiver more than a stranger.

Together, these results suggest that inferences about people's knowledge and intent can be formalized and that our formal approach holds promise for unifying and explaining a broad range of phenomena across development and social learning.

Conclusions

How do people learn and reason about the world so rapidly and effectively? Traditional computational approaches to learning have explained this by focusing on prior biases and knowledge representation. We have argued that intuitive psychological reasoning plays a critical role in the success of human learning. Under our approach, data are not inert observations with a fixed, context-free meaning; rather, they carry information about people's beliefs and intentions. We have presented a formal analysis that explains the effects of intuitive psychological reasoning on

learning as a consequence of rational inference about people's goals and beliefs. This analysis explains why the inferences afforded by two social learning contexts—learning from goal-directed action and learning from communicative action—differ from each other and from inferences afforded by learning in nonsocial contexts.

We have focused on very simple causal-learning scenarios for the purposes of exposition, but our framework naturally applies to noncausal learning domains and to situations of considerably greater complexity. Our recent work has provided applications of these ideas to concept learning (Goodman et al., 2009; Gweon, Tenenbaum, & Schulz, 2010; Lucas, Griffiths, Xu, & Fawcett, 2009; Shafto & Goodman, 2008; Warner et al., 2011), referential ambiguity (Frank & Goodman, in press; Frank, Goodman, & Tenenbaum, 2009), and causal learning in children (Bonawitz et al., 2011, Buchbaum, Griffiths, Gopnik, & Shafto, 2011).

Theories of cognitive development have almost uniformly claimed that other people are important for learning; however, theorists have differed in the role they assign to other people in learning and in how they assume that role to address basic problems of learning. Proposals have variously suggested that merely allowing learners to explore is enough to support learning (Bruner, 1966; Vygotsky, 1978), that some kind of modeling or imitative learning is key (Bandura et al., 1961; Meltzoff & Moore, 1977), or that reasoning about communicative intent is critical (Csibra & Gergely, 2009; Tomasello et al., 2005).

Our framework shows that these different proposals lead to very different predictions about learning, and it suggests a synthesis of these different views that opens the door for continued exploration of how intuitive psychological reasoning affects learning. The most powerful goals are not ubiquitous in learning—not every action is intended to be communicative. So, even though leveraging the strongest social contexts will lead to more powerful learning, identifying when different assumptions are appropriate is critical as well.

Our approach suggests that intuitive psychological reasoning may provide a strong lever by which learners can capitalize on others' knowledge to learn about the world (see also Coady, 1992; Csibra & Gergely, 2009; Harris, 2002; Tomasello et al., 2005). The key departure from previous formal approaches is in how our approach views other people. If other people are viewed to act in random or even malicious ways, then learning on the basis of their actions will likely be very difficult. However, if people are viewed as approximately rational, goal-directed agents or as knowledgeable and helpful teachers, then the problem of learning becomes much more tractable. These assumptions make it possible for learners to learn what others already know, rather than rediscovering all knowledge from the ground up.

Appendix A: Model Details

In contrast to the psychological formulation outlined here, typical Bayesian models of learning (and standard models of

learning more generally) assume that the evidence is independent of actions by any agents, $P(e | a, h) = P(e | h)$, or that the actions (also known as *interventions*) are generated uniformly at random, $P(a | g, h) \propto 1$. In the first case, we recover the usual formulation of Bayesian learning:

$$P(h|e, g) = \frac{P(e|h) P(h)}{\sum_{h'} P(e|h') P(h')} \quad (2)$$

In the second case, we recover the standard formulation of Bayesian causal learning, where effects depend on actions, $P(e | a, h)$, but the actions themselves are random. Under either formalization of learning, there is no role for other people in learning; actions are taken as a given.

To develop our understanding of the implications of psychological reasoning—the $P(a | g, h)$ term—we build on the idea that actions are chosen by a knowledgeable person to produce intended goals (see, e.g., Dennett, 1987; Gergely & Csibra, 2003). This idea is formalized via Luce's choice axiom (Luce, 1959),

$$P(a|g, h) = \frac{P(g|a, h)}{\sum_{a'} P(g|a', h)} \quad (3)$$

Intuitively, the left-hand side asks which actions should be chosen, given the actor's goals and beliefs. The right-hand side provides the answer: Actions should be preferred to the degree that they are likely to lead to desired goal, given the actor's beliefs and the chosen action.

Appendix B: Mathematical Details for Results in Figure 2

Bob's Box: inferring the causes of an effect

Possible hypotheses, h , are that Button A causes the light, that Button B causes the light, that Buttons A and B together cause the light, or that neither button causes the light; we assume that the prior probabilities, $P(h)$, are all equal to .25. For simplicity, we will consider actions, a , and effects, e , in pairs, focusing on a subset of the total possibilities. Specifically, we consider the following observations: an observation of no action, leading to no effect; an observation of the actor pressing Button A, leading to the light; an observation of the actor pressing Button B, leading to the light; and an observation of the actor pressing Buttons A and B, leading to the light. We assume that causes deterministically bring about their effects; that is, $P(e | a, h)$ is 1 if a is a cause of e according to hypothesis h , and zero otherwise.

Physical evidence. Assuming the actor's choice of behavior was random, the probability of choosing the action "press Buttons A and B", $P(a | g, h)$, is .25, because that action is one of

four possible actions. Because the effect is deterministic, $P(e | a, h) = 1$, and prior beliefs in the hypotheses are the same, $P(h) = .25$, the numerator of Equation 1 is $(.25) \times (.25) = .0625$. But both of the other viable hypotheses (that pressing Button A only leads to the effect and that pressing Button B only leads to the effect) have the same probability! Thus, by Equation 1, the probability of each hypothesis is .33.

Goal-directed action. If Bob knows how the machine works, we assume that he has chosen his action rationally to bring about the effect via Equation 3. If Buttons A and B together cause the light, the probability of pressing both buttons is 1, because that is the only action that would lead to the goal. Thus, the learner can infer that the probability of Buttons A and B jointly causing the light is proportional to $P(e | a, h)P(a | h, g)P(h) = 1 \times 1 \times (.25) = .25$. In contrast, the probability that Button A alone is the cause would be proportional to $1 \times (.5) \times (.25) = .125$ (the probability of Button B alone causing the effect is the same). Thus, when normalized to consider the possible alternative hypotheses, the probability that Buttons A and B jointly cause the light is $.25 / (.25 + .125 + .125 + 0) = .5$.

Communicative action. In this case, Bob chooses his action to maximize the learner's belief in the correct hypothesis. Recall that if the action was chosen randomly, $P(a | g, h) = .25$, the total probability that Buttons A and B together were the cause would be .33, as would the probability of Button A alone or Button B alone. At first glance, the effects of Buttons A and B appear ambiguous; however, the learner must reason about why the teacher chose these actions as opposed to other possibilities. What if the true cause was Button A alone? In that case, the teacher could have pressed only Button A, or both Button A and Button B, to elicit the effect. However, pressing Button A alone would lead the learner unambiguously to the inference that Button A alone caused the effect; there would be no other explanation. So, $P(h | a, e, g)$ for the hypothesis that Button A alone causes the effect would be 1 if the teacher pressed Button A, and $P(h | a, e, g)$ for the hypothesis that Buttons A and B together cause the effect would be .33 (the cause could be Button A alone, Button B alone, or Buttons A and B together). Up to this point, the reasoning is based completely on the assumption that actions were chosen randomly.

In communicative contexts, teachers and learners are yoked—each reasons about the other—suggesting that the learner's inferences are the *input* to the teacher's choices. If the teacher's goal is for the learner to infer the correct hypothesis, and Button A alone is the cause of the effect, then the probability of the teacher pressing Button A is $1 / (1 + .33) = .75$. Similarly, the probability of the teacher pressing Buttons A and B is $(.33) / (1 + .33) = .25$. On the other hand, if Buttons A and B jointly cause the effect, then the probability of choosing to press Buttons A and B is 1 because that is the only action that will lead to the effect. In a communicative context, learners who observe that pressing Buttons A and B leads to the effect will infer that pressing Button A is less likely to be the

cause, $P(e | a, h)P(a | h, g)P(h) = 1 \times .25 \times .25 = .0625$ (the same as probability that pressing Button B is the cause), and pressing Buttons A and B is more likely to be the cause, $P(e | a, h)P(a | h, g)P(h) = 1 \times 1 \times .25 = .25$. Thus, the probability of Hypothesis A and B after one step is $(.25) / (.25 + .0625 + .0625 + 0) \approx .67$.

Of course, because this is a recursive inference, one could continue reasoning about how the teacher and learner would update their actions. The result of this process is that the only reason a teacher would choose to press both Button A and Button B is if both buttons were both necessary to elicit the effect (i.e., $P(h | a, e, g) = 1$; see Fig. 2).

Tim's Toy: learning the number of latent causes

Possible hypotheses include that there are no cause-effect relationships, that there is one cause-effect relationship, and that there are two (or more) cause-effect relationships. For purposes of demonstration, we consider only zero, one, or two relationships, all of which have equal prior probabilities, $P(h) = 1/3$. The action-effect pairs we consider include no interventions leading to no effect, one intervention leading to one effect, or two interventions leading to two effects.

Physical evidence. Assuming that actions are chosen randomly, the probability of a single action is $P(a | g, h) = 1/3$. Thus, given one action that elicits an effect, the probability of one cause-effect pair is proportional to $P(e | a, h)P(a | g, h)P(h) = 1 \times .33 \times .33 \approx .11$. This is the same as the probability of two cause-effect pairs; thus $P(h | a, e, g)$ for each is .5.

Goal-directed action. Imagine that Tim's goal is to elicit an effect; in this case, the probability of him choosing one action would have been $P(a | g, h) = .5$, because either one or two actions would have lead to the goal. The probability of inferring that there is only one cause-effect relationship would be proportional to $P(e | a, h)P(a | g, h)P(h) = 1 \times .5 \times .33 = .165$. Of course, the probability of two causal relationships would be the same, because there is no need to elicit the other, given the goal; the $P(h | a, e, g) = .5$ for one hypothesis and two hypotheses.

Communicative action. If the learner observed zero actions and the choice was randomly sampled, then all hypotheses are equally probable: $P(h | a, e, g) = .33$ for all hypotheses. If the learner observed one action leading to an effect, then the possibility of zero cause-effect pairs is ruled out; however, the learner remains uncertain whether there are one or two pairs, $P(h | a, e, g) = .5$ for both hypotheses. If the learner observed two actions leading to two effects, the only possibility is that there are two cause-effect pairs; $P(h | a, e, g) = 1$ for this hypothesis.

Because this is a communicative context, the learner's inferences are inputs to the teacher's inference. If there is truly one cause-effect pair, the teacher chooses between doing nothing, leading to no effect, or performing one action with an effect.

The probability of choosing zero actions is $.33 / (.33 + .5) = .4$, and the probability of choosing one action is $.5 / (.33 + .5) = .6$. Similarly, if there truly are two pairs, then the teacher chooses between zero, one, or two actions. The probability of choosing zero actions is $.33 / (.33 + .5 + 1) \approx .18$, the probability of choosing one action is $.5 / (.33 + .5 + 1) \approx .27$, and the probability of choosing two actions is $1 / (.33 + .5 + 1) \approx .55$.

Because teachers and learners are yoked, the learner also reasons about the teacher. Given the observation of one action leading to an effect, the learner must infer whether there are one or two cause-effect pairs. The probability of one pair would be proportional to $P(e | a, h)P(a | g, h)P(h) = 1 \times .6 \times .33 \approx .198$, and the probability of the hypothesis that there are two pairs would be proportional to $P(e | a, h)P(a | g, h)P(h) = 1 \times .27 \times .33 \approx .089$. The hypothesis that there is only one cause-effect pair is more probable, $P(h | a, e, g) = .198 / (.198 + .089) \approx .69$. Again, because this is a recursive inference, one could continue reasoning about how the teacher and learner would update their actions. The result of this process in this case is that a teacher would choose one action only if there was just one cause-effect pair, i.e. $P(h | a, e, g) = 1$.

Acknowledgments

We thank Jamil Zaki and Greg Walton for helpful comments on an earlier draft of this manuscript.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

References

- Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*, 409–429.
- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*, *113*, 329–349.
- Bandura, A., Ross, D., & Ross, S. A. (1961). Transmission of aggression through imitation of aggressive models. *Journal of Abnormal and Social Psychology*, *63*, 575–582.
- Bloom, P. (2002). Mindreading, communication, and the learning of the names for things. *Mind & Language*, *17*, 37–54.
- Bonawitz, E. B., Shafto, P., Gweon, H., Goodman, N. D., Spelke, E., & Schulz, L. (2011). The double-edged sword of pedagogy: Teaching limits children's spontaneous exploration and discovery. *Cognition*, *120*, 322–330.
- Bruner, J. (1966). *Toward a theory of instruction*. Cambridge, MA: Belknap Press.
- Buchbaum, D., Griffiths, T., Gopnik, A., & Shafto, P. (2011). Children's imitation of causal action sequences is influenced by statistical and pedagogical evidence. *Cognition*, *120*, 331–340.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367–405.
- Clark, H. (1996). *Using language*. Cambridge, England: Cambridge University Press.
- Clark, H. (2003). Pointing and placing. In S. Kita (Ed.), *Pointing: Where language, culture, and cognition meet* (pp. 243–268). Hillsdale, NJ: Erlbaum.
- Coady, C. A. J. (1992). *Testimony: A philosophical study*. Oxford, England: Clarendon Press.
- Corriveau, K. H., Fusaro, M., & Harris, P. L. (2009). Going with the flow: Preschoolers prefer nondissenters as informants. *Psychological Science*, *20*, 372–377.
- Corriveau, K. H., & Harris, P. L. (2009). Choosing your informant: Weighing familiarity and past accuracy. *Developmental Science*, *12*, 426–437.
- Corriveau, K. H., Harris, P. L., Meins, E., Fernyhough, C., Arnott, B., Elliott, L., . . . de Rosnay, M. (2009). Young children's trust in their mother's claims: Longitudinal links with attachment security in infancy. *Child Development*, *80*, 750–761.
- Csibra, G., & Gergely, G. (2009). Natural pedagogy. *Trends in Cognitive Sciences*, *14*, 148–153.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: Bradford Books.
- Frank, M. C., & Goodman, N. D. (in press). Predicting pragmatic reasoning in language games. *Science*.
- Frank, M. C., Goodman, N. D., & Tenenbaum, J. B. (2009). Using speakers' referential intentions to model early cross-situational word learning. *Psychological Science*, *20*, 578–585.
- Gelman, S. A. (1996). *The essential child: Origins of essentialism in everyday thought*. New York, NY: Oxford University Press.
- Gergely, G., & Csibra, G. (2003). Teleological reasoning in infancy: The naïve theory of rational action. *Trends in Cognitive Sciences*, *7*, 287–292.
- Gold, E. (1967). Language identification in the limit. *Information and Control*, *10*, 447–474.
- Goodman, N. D., Baker, C. L., & Tenenbaum, J. B. (2009). Cause and intent: Social reasoning in causal learning. In N. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 2759–2764). Austin, TX: Cognitive Science Society.
- 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.
- Gopnik, A., Glymour, C., Sobel, D., Schulz, L., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 1–31.
- Grice, P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Syntax and semantics, Vol. 3: Speech acts* (pp. 41–58). New York, NY: Academic Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 354–384.
- 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, USA*, *107*, 9066–9071.
- Harris, P. L. (2002). What do children learn from testimony. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The cognitive basis*

- of science (pp. 316–334). Cambridge, England: Cambridge University Press.
- Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (pan troglodytes) and children (homo sapiens). *Animal Cognition*, 8, 164–181.
- Horning, J. J. (1969). *A study of grammatical inference* (Unpublished doctoral dissertation). Stanford University, Stanford, CA.
- Jones, E. E., & Davis, K. E. (1965). From acts to dispositions: The attribution process in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 219–266). New York, NY: Academic Press.
- Keil, F. C. (1989). *Concepts, kinds, and cognitive development*. Cambridge, MA: MIT Press.
- Kelley, H. H. (1967). Attribution theory in social psychology. In D. Levine (Ed.), *Nebraska symposium on motivation* (pp. 129–238). Lincoln: University of Nebraska Press.
- Koenig, M., & Harris, P. (2005). Preschoolers mistrust ignorant and inaccurate speakers. *Child Development*, 76, 1261–1277.
- Krueger, J. I., & Funder, D. C. (2004). Towards a balanced social psychology: Causes, consequences, and cures for the problem-seeking approach to social behavior and cognition. *Behavioral & Brain Sciences*, 27, 313–327.
- Kruschke, J. K. (1992). Alcové: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). Sustain: A network model of category learning. *Psychological Review*, 111, 309–332.
- Lucas, C., Griffiths, T. L., Xu, F., & Fawcett, C. (2009). A rational model of preference learning and choice prediction by children. In D. Koller, Y. Schuurmans, L. Bottou, & A. Culotta (Eds.), *Advances in neural information processing systems* (Vol. 21, pp. 985–992). Cambridge, MA: MIT Press.
- Luce, R. D. (1959). *Individual choice behavior*. New York, NY: John Wiley.
- Lyons, D. E., Young, A. G., & Keil, F. C. (2007). The hidden structure of overimitation. *Proceedings of the National Academy of Sciences, USA*, 104, 19751–19756.
- Mandler, J. M. (1992). How to build a baby: II. Conceptual primitives. *Psychological Review*, 99, 587–604.
- McNeill, N. M. (2008). Limitations to teaching children $2 + 2 = 4$: Typical arithmetic problems can hinder learning of mathematical equivalence. *Child Development*, 79, 1524–1537.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, 100, 254–278.
- Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31, 838–850.
- Meltzoff, A. N., & Moore, M. K. (1977). Imitation of facial and manual gestures by human neonates. *Science*, 198, 75–78.
- Michotte, A. E. (1963). *The perception of causality*. New York, NY: Basic Books.
- Nosofsky, R. M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10, 104–114.
- Nosofsky, R. M., Palemeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.
- Pasquini, E. S., Corriveau, K. H., Koenig, M. A., & Harris, P. L. (2007). Preschoolers monitor the relative accuracy of informants. *Developmental Psychology*, 43, 1216–1226.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 241–248.
- Pothos, E. M., & Chater, N. (2002). A simplicity principle in unsupervised human categorization. *Cognitive Science*, 26, 303–343.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. Black & W. Prokasy (Eds.), *Classical conditioning II*. Norwalk, CT: Appleton-Century-Crofts.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. Lloyd (Eds.), *Cognition and categorization* (pp. 28–48). Hillsdale, NJ: Erlbaum.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 46, 55–67.
- Shafto, P., Eaves, B., Navarro, D. J., & Perfors, A. (2012). Epistemic trust: Modeling children's reasoning about others' knowledge and intent. *Developmental Science*, 15, 436–447.
- Shafto, P., & Goodman, N. D. (2008). Teaching games: Statistical sampling assumptions for pedagogical situations. In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1632–1638). Austin, TX: Cognitive Science Society.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.
- Tomasello, M., Carpenter, M., Call, J., Behne, T., & Moll, H. (2005). Understanding and sharing intentions: The origins of cultural cognition. *Behavioral & Brain Sciences*, 28, 675–691.
- Vygotsky, L. (1978). *Mind in society*. Cambridge, MA: Harvard University Press.
- Warner, R., Stoess, T., & Shafto, P. (2011). Reasoning in teaching and misleading situations. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 1430–1435). Austin, TX: Cognitive Science Society.
- Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1, 67.
- Xu, F., & Tenenbaum, J. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science*, 10, 288–297.
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In T. Fawcett & N. Mishra (Eds.), *Proceedings of the 20th International Conference on Machine Learning* (pp. 928–936). Menlo Park, CA: AAAI Press.