Teaching Games:
Statistical Sampling Assumptions for Learning in Pedagogical Situations

Patrick Shafto
University of Louisville
`p.shafto@louisville.edu`

Noah Goodman
Massachusetts Institute of Technology
`ndg@mit.edu`

Abstract

Much of learning and reasoning occurs in pedagogical situations – situations in which teachers choose examples with the goal of having a learner infer the concept the teacher has in mind. In this paper, we present a model of teaching and learning in pedagogical settings which predicts what examples teachers should choose and what learners should infer given a teachers' examples. We present two experiments using a novel experimental paradigm called the rectangle game. The first experiment compares people's inferences to qualitative model predictions. The second experiment tests people in a situation where pedagogical sampling is not appropriate, ruling out alternative explanations, and suggesting that people use context-appropriate sampling assumptions. We conclude by discussing connections to broader work in inductive reasoning and cognitive development, and outline areas of future work.

Much of human learning and reasoning goes on in pedagogical settings. In schools, teachers impart their knowledge to students about mathematics, science, and literature through examples and problems. From early in life, parents teach children words for objects and actions, and cultural and personal preferences through subtle glances and outright admonitions. Pedagogical settings – settings where one agent is choosing information to transmit to another agent for the purpose of communicating a concept – dominate human learning and reasoning.

If learners' assumptions about how teachers sample information reflected this purposeful sampling, then learners might be able to make much stronger inferences in pedagogical situations. Sampling assumptions are assumptions that a learner makes about the source of data, in order to better interpret the evidence for statistical learning. Recent research suggests that even infants are sensitive to the sampling processes that underlie observed data (Xu & Tenenbaum, 2007) and young children make qualitatively different inferences when data are sampled by a teacher (Gergely, Egyed, & Kiraly, 2007).

Computational models of human learning have not focused on the sampling processes that generate observed data, and when they do the assumed processes are relatively simple. For example, Fried and Holyoak (1984) modeled category learning by assuming examples are generated uniformly at random (weak sampling) and Tenenbaum (1999) modeled learning from positive data by assuming examples are generated at random from the true concept (strong sampling).

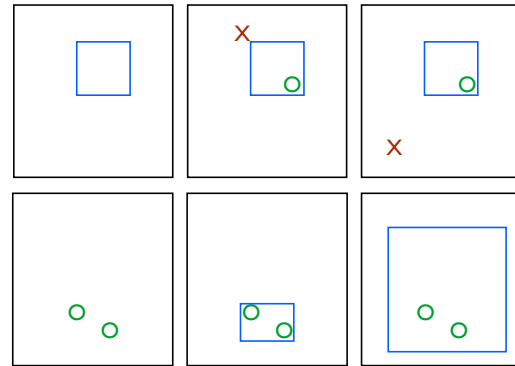Even simple assumptions gain a lot for a learner. For ex-

Figure 1: Possible rectangle game scenarios. The top row shows a possible rectangle concept, and two possible pairs of examples that a teacher might choose to communicate to a learner. The bottom row shows possible examples a learner may observe, and two possible guesses about what rectangle the teacher had in mind. The middle column shows better choices than the right column.

ample, Xu and Tenenbaum (2007) showed that the strong sampling assumption allows learning the meanings of words from positive examples alone. However, neither weak nor strong sampling captures the purposeful sampling that underlies pedagogical situations. In this paper, we introduce a pedagogical sampling model that formalizes which examples teachers should give learners to most help them, and what learners may infer from these helpful examples.

Consider a simple example which we call the rectangle game: a game where the teacher thinks of a rectangle on a board, and tries to communicate that concept to a learner by choosing to label points inside and/or outside the rectangle (cf. Tenenbaum, 1999). In the rectangle game, the learner's job is to try to infer, given the labeled examples chosen by the teacher, what rectangle the teacher is thinking of. Figure 1 presents potential teacher and learner scenarios. In each case, there seem to be choices which are obviously better than others. As a person trying to teach someone the rectangle in blue (top left), the examples in the middle panel seem better than those on the right. Similarly, as a learner, given the examples in the bottom left, the rectangle in the middle panel seems like a better guess than that on the right. Notice that in both cases the examples on the top and the rectangles on the bottom are possible, however, our intuition tells us that the middle panels are better guesses than the right panels. Peda-

gogical sampling results in samples that are representative of the concept (Tenenbaum & Griffiths, 2001), in contrast with weak and strong sampling, which choose examples randomly.

In this paper, we formalize the problem of pedagogical reasoning from the perspectives of teacher and learner. For the teacher, the problem is to choose the examples that will most help the learner infer the correct concept. For the learner, the problem is, given that the teacher is choosing helpful examples, infer the correct concept. In two experiments, we compare model predictions to human behavior in a novel experimental paradigm (the rectangle game), showing strong correspondence between model predictions and human data. We conclude by addressing implications of pedagogical sampling and identifying areas of future work.

A computational model of pedagogical sampling

We formalize pedagogical reasoning as an inference problem based on the twin assumptions that learners and teachers act as (approximately) rational agents. The *rational learner assumption* is that a learner will update their beliefs, given new examples, according to Bayesian inference (a description of optimal belief updating),

$$p(h|d)_{learner} \propto p(d|h)_{teacher} p(h). \tag{1}$$

The *rational teacher assumption* is that teachers choose examples that tend to increase the learner's beliefs in the correct hypothesis. We may formalize this via a Luce decision rule (Luce, 1959),

$$p(d|h)_{teacher} \propto (p(h|d)_{learner})^\alpha, \tag{2}$$

where the steepness parameter α governs the greediness of the teacher. (When $\alpha=0$, pedagogical sampling recovers weak sampling, as α becomes large the teacher chooses the best examples.) Because Equations 1 and 2 are linked (with the optimal teaching behavior depending on the learner, and vice versa), rational pedagogical reasoning is a solution to this *system* of equations. [1]

To understand this model it helps to consider one way of solving the system of equations: fixed-point iteration. Imagine that you are the learner, and wish to update your beliefs. To do so you will need an estimate of the likelihood $p(d|h)_{teacher}$ of seeing the examples you are given. You can estimate this likelihood by assuming the teacher is rational—Eq. 2—but to do this you need an estimate of the $p(h|d)_{learner}$ used by the teacher. If you assume the teacher assumes that you are rational, you can use Equation 1 as such an estimate.... This recursive reasoning could carry on forever, but

eventually the estimated values from repeatedly using equations 1 and 2 will no longer change—we then say that the process has iterated to a fixed point, and this fixed point will necessarily be a solution to the system of equations defining rational pedagogical reasoning. Thus we can understand the model as capturing the outcome of a recursive mental reasoning process, based on twin rational agent assumptions. However, it is worth emphasizing that rational pedagogical reasoning describes the *outcome* of this process (or rather the solution to the system of equations), and it is entirely possible that this reasoning may be implemented by a psychological process that doesn't require any explicit recursive reasoning. [2]

There is one additional complication in modeling the teaching games used below. In addition to observing the location of positive and negative examples, learners also observe the strategy the teacher used – how many positive and negative examples they use. The learner can observe the strategy chosen by the teacher, and we simplify by treating this choice of strategy as uninformative. Inference for both teacher and learner are then simply conditioned on the observed strategy (number and kind of examples).

Applying this model to the examples in Fig. 1 results in predictions that are consistent with intuition. In the case of two positive examples, the prediction is that the teacher will generally place examples in opposite corners of the rectangle, and the learner will infer a rectangle such that the examples are near opposite corners. To understand why this is a solution to equations 1 and 2, consider the recursive reasoning described above (idealized to avoid complications of uncertainty): if the learner assumes that the teacher will choose examples at opposite corners, then Eq. 1 implies that the unique inference made by the learner is the tightest rectangle around two examples; if the teacher assumes that this is what the learner is doing, then, according to Eq. 2, the teacher will usually choose examples in opposite corners of the true rectangle. In the case of one positive example and one negative example, the reasoning is similar: the learner assumes that the teacher will choose a negative example close to the boundary of the rectangle, enabling the learner to rule out larger rectangles; in turn the teacher will chose such examples under the assumption that this is how the learner will reason. We test the model predictions in the following experiment.

Experiment 1: The rectangle game

In this experiment, people played the rectangle game described in the introduction. People played the roles of teacher, choosing the examples given a rectangle, and learner, guessing a rectangle given examples. In our analyses, we will investigate how the models predict the qualitative and quantitative patterns observed in the human data. For the teaching task, we will ask whether the pattern of responses is ran-

---

[1]Solutions to this system of equations will depend on the learner's prior beliefs, $p(h)$. For the current purposes, where we are dealing with a completely novel game, we assume both the teacher and the learner know that all hypotheses are equally likely. In general, it is reasonable to assume that the teacher makes a reasonable assumption about the learner's prior. Interesting questions arise when the teacher is uncertain or incorrect about the learner's prior; these questions are beyond the scope of this paper.

[2]To compute model predictions we use the fixed-point iteration formulation of the model. That is, we initialize with the learner assuming weak sampling, then iterate Equations 1 and 2 until convergence. This is a straightforward, though computationally expensive, algorithm.

dom, as predicted by weak and strong sampling. We will also characterize people's preferred pattern of responses, and ask whether the pedagogical sampling model correctly predicts the pattern of examples generated by people. For the learning task, we will ask comparable questions: are learner's inferences consistent with the assumption that examples are randomly sampled? If not, do learner's inferences reflect knowledge about the pedagogical nature of the data?

Method

Participants. 18 University of Louisville undergraduates participated in exchange for course credit.

Design. The experiment consisted of two parts: teaching and learning. Because pre-testing showed that participants understood the learning task better when presented with the teaching task first, the two parts were presented in fixed order.

Procedure. Experiments were run using MATLAB. Participants were seated at a computer and told that they were going to learn about a game called the rectangle game. In the game there is a teacher and a learner. It is the teacher's job to help the learner guess what rectangle he is thinking about by choosing helpful examples, points that can be inside or outside the true rectangle.

In the teaching task, participants were shown a rectangle and asked to choose the best one, two, or three examples by clicking on the screen. A green circle automatically appeared if the click was inside the rectangle, and a red X if the click was outside the rectangle. Participants were shown 90 rectangles of various sizes and positions. Order of presentation of sizes and positions of the rectangles varied randomly.

In the learning task, participants were shown between one and three examples and asked to draw the rectangle that they thought the teacher was thinking of, by clicking on the screen and dragging with the mouse. Examples were pre-chosen to include all combinations of positive and negative examples. Examples were generated based on a small set of patterns, and position on the screen and distance between the examples was randomized. The learning task consisted of 92 trials in total. When participants completed the task, they were debriefed and thanked.

Model implementation. To model the rectangle game, we used a hypothesis space of rectangles based on a 6x6 discretized approximation to the full board. Hypotheses included all rectangles from 2x2 up to 5x5, at all possible locations on the grid. 36 possible points, corresponding to all possible blocks on the 6x6 board were considered. Examples were sets of positive, negative, or mixed pairs of these points. To attempt to minimize strong boundary effects resulting from our discrete approximation, we modeled inferences on an extended board upon which the rectangle or examples were centered.

Results & Discussion

This experimental paradigm generates an extremely rich data set. For teaching data, people could choose any point(s) on

(a) Positive examples
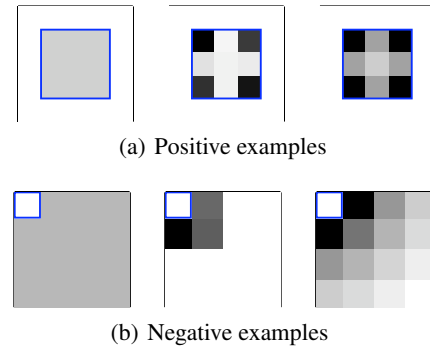

(b) Negative examples

Figure 2: Distributions of examples in the teaching task for (a) positive examples and (b) negative examples. Pictured from left to right in each panel are the predictions of strong and weak sampling, the observed human data, and the predictions of pedagogical sampling. For the models, figures display the probability of an example in each block. For human data, the proportion of positive examples in each location is plotted. People strongly preferred to give positive examples in the corners of the rectangle and negative examples near the boundaries, as predicted by pedagogical sampling.

the board. For the learning data, people could choose any number of rectangles that vary in size and location such that they are consistent with the observed examples. Analyzing the data required collapsing across examples to achieve adequate summaries of the patterns of behavior elicited. Also, due to space constraints, we present only the subset of the teaching and learning data that include two examples.

We will consider performance on the teaching task first, then the learning task. For the teaching task, we will consider whether people's data conform to the predictions of weak and strong sampling or pedagogical sampling by separately analyzing the distributions of the positive and negative examples. We will then investigate the specific pairs of examples chosen by people, and compare these to the pedagogical sampling model.

Pedagogical sampling predicts that for positive examples, examples in the corners are more informative than examples in the middle, or along a single side. Both weak sampling and strong sampling predict that positive examples are distributed uniformly at random. To test these predictions, we analyzed the examples that people generated, dividing each rectangle into a 3x3 grid. Grids were normalized based on the size of the rectangle, so the grid was finer for smaller rectangles than for larger ones. This allowed us to ignore the size of the rectangle and focus on the relative position of the examples. Frequencies of examples in each area of the grid were tallied and the proportion of examples in each location are shown in Figure 2a, middle panel. The left and right panels show the locations predicted by weak sampling and strong sampling (left) and pedagogical sampling (right). Note that strong and weak sampling predict no differences across locations – data should be distributed uniformly at random. The examples people generated were highly non-random, $\chi^2(8) = 645.05$, $p < 0.0001$. The pedagogical sampling predictions, right panel, are based on an average over all 3x3 hypotheses, and show that the model predicts a strong
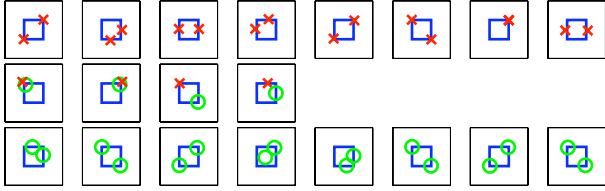
Figure 3: Pairs of examples chosen in the teaching task for a randomly chosen rectangle, divided by type of pair. The top row contains pairs of negative examples, which are characterized by their close proximity to the boundaries, generally on opposite sides or corners. The middle row shows mixed examples marking a single or opposite corners. The bottom row shows pairs of positive examples, which are predominantly examples in opposite corners.

Figure 4: Results from the learning task. Plots show the positions of the teacher's examples, relative to the rectangles drawn by learners for positive (left) and negative (right) examples. The results show that learners clearly understand that teachers are sampling data pedagogically – positive examples indicate corners of the correct rectangle, and negative examples indicate the boundaries.

preference for positive examples in the corners of the rectangle. A qualitatively similar pattern is observed in the human data. Indeed, this qualitative correspondence is reflected in a strong correlation between model predictions and human data, $r = 0.98$.

For negative examples, pedagogical sampling predicts strong effects of distance – the most helpful negative examples are those that are near the boundaries. Weak sampling predicts that examples should be distributed uniformly at random, while strong sampling makes no prediction. We analyzed people's choices by classifying examples based on the relative distance from the boundary of the rectangle to the outside of the board. For each of the four sides, we divided the area from the edge of the rectangle to the edge of the board into three bins. For one quadrant, the result is 15 bins, 3 to each side of the rectangle of the rectangle, and 9 for the intersection of these bins extending diagonally from the rectangle. The results are shown in Figure 2b. The examples people chose were not randomly distributed, $\chi^2(14) = 1268.82, p < 0.0001$. Again, the apparent qualitative agreement is supported by a strong correlation between model predictions and human data, 0.86.

Of primary importance from the perspective of our model is which pairs of examples people choose. Figure 3 shows a random subset of the data from a randomly chosen rectangle. Examples have been divided into pairs of negative examples, mixed example pairs, and pairs of positive examples. For the negative pairs, people tend toward either choosing examples on opposite corners or on opposite sides, and these are the two most likely strategies according to the model. For mixed pairs, people tend to either mark a single corner with a positive example on the inside and a negative example on the outside, or opposite corners. These, along with pairs that mark a side on the inside and outside are most likely according to the model. For positive pairs, people use the examples to mark opposite corners. Pedagogical sampling very strongly predicts a preference for this strategy over other possible pairs of positive examples.

The results of the teaching task suggest that people choose particular examples to communicate different concepts. Of equal importance is whether learners take advantage of this knowledge to make stronger inferences. To answer this ques-
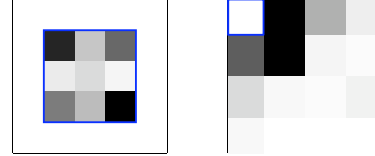
tion, we turn to the data from the learning task.

Pedagogical sampling predicts that learners' should know the strategies that the teacher will use to communicate different concepts. Therefore, if learners use pedagogical knowledge to guide inferences, we expect learners should draw rectangles that recover the patterns observed in the teaching data (see Figure 2). However, weak and strong sampling predict that positive examples should be randomly distributed with respect to the rectangles inferred by learners. Figure 4 shows the observed distributions for positive and negative examples, relative to the rectangles inferred by learners. To test whether the examples where randomly distributed, we computed the frequencies of different example locations relative to the inferred rectangle, as described for the teaching task. The results indicate that the inferred rectangles where not randomly positioned with respect the the positive examples, $\chi^2(8) = 104.37, p < 0.001$, or negative examples, $\chi^2(14) = 465.10, p < 0.001$.

To test whether people's inferences were in accordance with the predictions of pedagogical sampling, we correlated the frequencies of examples in the grid with the probabilities predicted by the pedagogical sampling model. For both the positive and negative examples, we found a strong positive correlation between the model predictions and human data, $r = 0.90$ and $r = 0.87$, respectively.

Finally, we present representative examples of people's inferences in the learning task and *maximum a posteriori* model predictions in Figure 5. The results show that inferred rectangles are consistent with the generating process that people showed in the teaching task. In particular, given positive examples, learners infer rectangles for which the examples are in opposite corners. Similarly, when given a mixed pair, people tend to infer that the rectangle extends up to the negative example and when given two negative examples, people infer rectangles whose corners are just inside the examples. The model predicts these effects for the positive and mixed cases, while in the negative case the model does not capture peoples inferences with the strength that people show.

Experiment 2: Random rectangle game

One possible alternative explanation for the results of the learning task (entertained by psychologists and other cynics) is that people simply prefer scenarios where examples lie at the corners of the rectangles for perceptual or other reasons
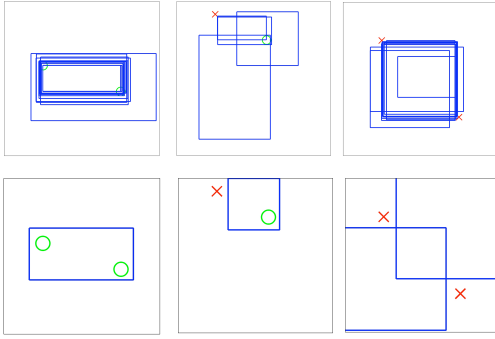
Figure 5: A sample of people's inferences in the learning task (top row) and the MAP solution according to the model (bottom row), divided by type of pair. The model clearly predicts the results for positive examples and mixed examples. For negative examples, people's guesses suggest that they assume negative examples mark corners, while the pair of best results from the model does not capture this prediction, though averaging over hypotheses closely corresponds to peoples resposes.

not related to how the data are sampled. This experiment introduces a variant of the game where the learner chooses the examples to ask a teacher about. The teacher then labels the examples as inside or outside the rectangle. The learner then infers the rectangle the teacher is thinking of. Importantly, in this version of the game, pedagogical sampling is no longer appropriate. Because the learner does not know where the rectangle is when they are guessing, the examples are effectively random with respect to the true rectangle. Thus, we do not expect to find the pedagogical sampling effects observed in Experiment 1.

Method

Participants   29 University of Louisville undergraduates participated in exchange for course credit.

Procedure   The design consisted of two parts, as in Experiment 1. In the first part participants were told that they were going to see all of the possible rectangles (the same rectangles that people saw in Experiment 1). They clicked (once) anywhere on the screen to advance to the next rectangle. In the learning task, participants played a version of the rectangle game in which they chose the examples that were labeled, then inferred the true rectangle concept. Participants were told a number of examples and chose by clicking on the screen. Each click was marked with a black dot. After they made the pre-specified number of clicks, the black dots turned into either green circles or red X's depending on whether they were inside or outside the unobserved rectangle. The unobserved rectangles used to label points were the same as those in the teaching section, in random order. Participants were then asked to guess where the true rectangle was by clicking and dragging with the mouse. This manipulation insured that participants were aware that locations of examples they observed were independent of the location of the true concept. Participants were debriefed and thanked upon completion.
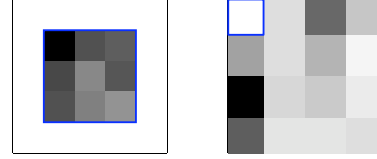

Figure 6: Density plots showing the relative position of the observed examples with respect to the inferred rectangles in Experiment 2 for positive (left) and negative (right) examples. Note that while there is variation across the grids, in neither case is there a discernable pattern of responses (cf. Figure 4). The results show that when the data are sampled by the learner, inferences about the true rectangle do not show the qualitative effects predicted by pedagogical sampling.

Results & Discussion

The focus of this experiment is to investigate learner's inferences when data are not pedagogically sampled. In particular, we are interested in whether the pattern of inferences observed in Experiment 1 are specific to pedagogical situations. Therefore, we will investigate the same qualitative and quantitative effects for the learning task as in the previous experiment.

First, we investigate the relationship between the positive and negative pairs and the inferred rectangles. For positive pairs, the question of interest is whether, when people see a pair of positive examples, they draw a rectangle that places those examples in the corners. That is, do people simply have a preference for drawing rectangles that have positive examples in the corners, or is this strategy observed in Experiment 1 due to the pedagogical nature of the examples? People's rectangles were divided into 3x3 grids, and we tallied the total number of examples in each bin. Figure 6, left side, shows the distribution of positive examples. While they are not completely random, $\chi^2(8) = 16.03, p = 0.04$, they do not show any discernable pattern. Particularly, there is no evidence that people chose rectangles such that the examples were placed in the corners (compare with Figure 4). The lack of correspondence between the predictions of the pedagogical model and observed data is reflected in their low correlation, $r = 0.27$.

For negative pairs, we are interested in whether the inferred rectangles tended to have negative examples near the edges. We categorized examples into bins based on the relative distance from the edge of the rectangle to the side of the board, and tallied counts for each bin. Figure 6, right side, shows the distribution of examples was not completely random, $\chi^2(15) = 323.43, p < 0.0001$, but it does not show the marked distance effects of the previous experiment (see Figure 4). The locations of negative examples do not correspond with the predictions of the pedagogical model, $r = 0.23$. The results show that people did not choose rectangles to be near the negative examples – that the inferences drawn from non-pedagogically sampled data are qualitatively different than those based on pedagogically sampled data.

Together these results show that the predictions of the pedagogical model hold for data sampled by a teacher, but not for a related task where the learner chooses which point they get to observe. The qualitative effects predicted by the model and

observed in Experiment 1 are not the result of simple perceptual preferences. Rather, people are sensitive to when pedagogical sampling applies. People generate examples differentially based on how helpful they are to a learner. Similarly, when people receive data chosen by a teacher, people are able to capitalize on this information to make stronger inferences.

General Discussion

Much of human learning occurs in pedagogical situations – social situations where one person chooses information for the purpose of helping another learn. We have presented a formal model of pedagogical reasoning, addressing which examples teachers should choose to communicate ideas, and what inferences learners should make based on this purposefully sampled data. We also presented two experiments testing the predictions of this model in a simple experimental setting. The first experiment showed that people's choice of data as a teacher, and their inferences as a learner correspond to the qualitative and quantitative predictions of the pedagogical model. The second experiment provided an important control condition, demonstrating that when data are chosen by the learner and hence are not representative of the concept to be learned, the predicted pedagogical effects are not observed. Together these results demonstrate a first step toward capturing the important effects of pedagogical situations on learning and reasoning.

Neither of the well-studied sampling assumptions (weak or strong sampling) captures the qualitative phenomena we have observed because neither captures the intuitions about sampling in pedagogical situations. Both assume that data are sampled randomly, which does not account for the preference for certain examples. As models of learning, both are unable to capture the speed of learning in Experiment 1 – with just two examples, we have a strong intuition for which rectangle the teacher has in mind – in contrast to the many examples necessary for either of these models (Tenenbaum, 1999).

The pedagogical sampling assumption is not appropriate in all circumstances. For many real-world learning situations, strong sampling is the right assumption; for example, when learning about the laws that govern physical objects, we never see situations that cannot happen and the situations we do observe are random – certainly it is not reasonable to assume that nature is helpfully choosing examples to teach us something! Similarly, weak sampling applies to a variety of learning situations we might encounter in casinos. Our evidence, with that of Tenenbaum (1999) and Xu and Tenenbaum (2007) suggests that people are able to distinguish situations in which different sampling assumptions are warranted.

Though modeling pedagogical reasoning in richer domains is a significant challenge, it highlights a great strength of our model. We have formalized pedagogical reasoning in the abstract language of probability and Bayesian reasoning, without reference to the specific details of the particular setting we considered here. As a result we are able to derive predictions in principle for any domain for which we can identify an appropriate set of hypotheses. Interesting domains to pursue include word learning, where speakers choose words to communicate ideas, and causal learning, where a helpful teacher may significantly reduce the number of interventions required to learn latent causal structure. Different domains will have different hypotheses and priors and as a result the model will generate qualitatively different predictions.

One of the things that make people special is that we can teach others what we know (Csibra, 2007). However, even when we are teaching others, we never communicate the complete idea in precise detail. As a result, people must resolve a difficult inference problem in order to capitalize on the information teachers provide. Gergely et al. (2007) have shown that children do capitalize on these situations, and they argue that pedagogical reasoning is among the powerful tools children have for learning about the world. Much work remains before the breadth of the implications of pedagogy for human learning are understood, but our work, providing a computational basis for understanding these kinds of inferences, represents a step in this direction.

References

Csibra, G. (2007). Teachers in the wild. *Trends in Cognitive Sciences*, *11*, 95–96.

Fried, L. S., & Holyoak, K. J. (1984). Induction of category distributions: A framework for classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 234–257.

Gergely, G., Egyed, K., & Kiraly, I. (2007). On pedagogy. *Developmental Science*, *10*, 139–146.

Luce, R. D. (1959). *Individual choice behavior*. New York: John Wiley.

Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. In M. Kear, S. A. Soller, T. K. Leen, & K. R. Müller (Eds.), *Advances in neural processing systems 11* (pp. 59–65). MIT Press.

Tenenbaum, J. B., & Griffiths, T. L. (2001). The rational basis of representativeness. In *Proceedings of the 23nd annual conference of the Cognitive Science Society* (pp. 1036–1041). Hillsdale, NJ: Erlbaum.

Xu, F., & Tenenbaum, J. (2007). Sensitivity to sampling in Bayesian word learning. *Developmental Science*, *10*, 288–297.