Unit 3: Learning from other people

1. Learning from teaching

10/29/2020

Comparing models

1. Learning from teaching is different from learning from observation

2. A rational model of teaching and learning from teaching

3. Teaching has tradeoffs.
   Some are predictable from this rational model

The number game (Tenenbaum, 2000)

An unknown computer program that generates from 1 to 100.
You get some random examples from this program.

2 4 8

What program do you think the computer is running?

The number game (Tenenbaum, 2000)

Posterior: $P\left(h \mid X\right) \propto P\left(X \mid h\right) P\left(h\right)$

Likelihood: $P\left(x \mid h\right) = \begin{cases} \frac{1}{|h|}, & x \in h \\ 0 & \text{otherwise} \end{cases}$

Prior: $P\left(h\right) = \begin{cases} \frac{\lambda}{N}, & N \text{ mathematical hypotheses} \\ \frac{(1-\lambda)}{M}, & M \text{ interval hypotheses} \end{cases}$

Powers of 2 is a good guess because:

1. it has a high prior
2. it has a high likelihood

The number teaching game

Someone has an unknown rule that generates numbers from 1 to 100. They want teach you by giving you examples

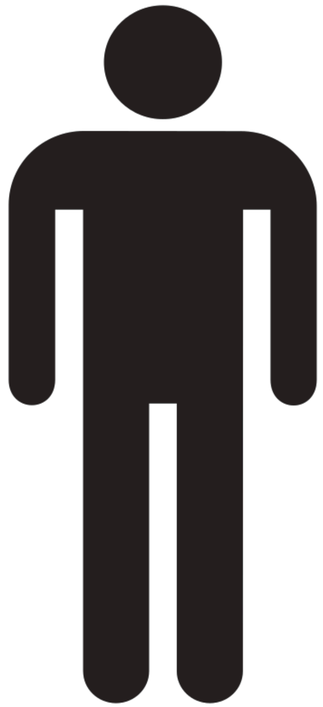

x1 x2 x3

The number teaching game

An unknown computer program generates numbers from 1 to 100.
The computer wants to teach you the program by giving you examples

2 4 8

What program do you think the computer is running?

The number game teaching game

Posterior: $P\left(h \mid X\right) \propto P\left(X \mid h\right) P\left(h\right)$

Likelihood: $P\left(x \mid h\right) = \begin{cases} \frac{1}{|h|}, & x \in h \\ 0 & \text{otherwise} \end{cases}$

Prior: $P\left(h\right) = \begin{cases} \frac{\lambda}{N}, & N \text{ mathematical hypotheses} \\ \frac{(1-\lambda)}{M}, & M \text{ interval hypotheses} \end{cases}$

Which of these parts of the model is wrong?

What is the goal of teaching?

Intuition: give examples that would lead the learner to infer the right hypothesis for the data

If you were teaching **evens,** you could choose less ambiguous examples (e.g. **2 38 94**).

Why did you pick numbers in this tight range?

A rational model of teaching (Shafto, Goodman, & Griffiths, 2014)

$$P_{teacher}\left(d \mid h\right) \propto P_{learner}\left(h \mid d\right)^{\alpha}$$

Teachers should generate data that will make learners likely to infer the true hypothesis.

$\alpha \in \left[0, \infty\right)$ is a rationality parameter.

What happens as $\alpha \to \infty$?

What happens when $\alpha = 0$?

A rational model of learning from teaching

$$P_{teacher}\left(d \mid h\right) \propto P_{learner}\left(h \mid d\right)^{\alpha}$$

$$P_{learner}\left(h \mid d\right) \propto P_{teacher}\left(d \mid h\right) P\left(h\right)$$

This is a recursive reasoning process!

The rectangle game: Teacher's perspective

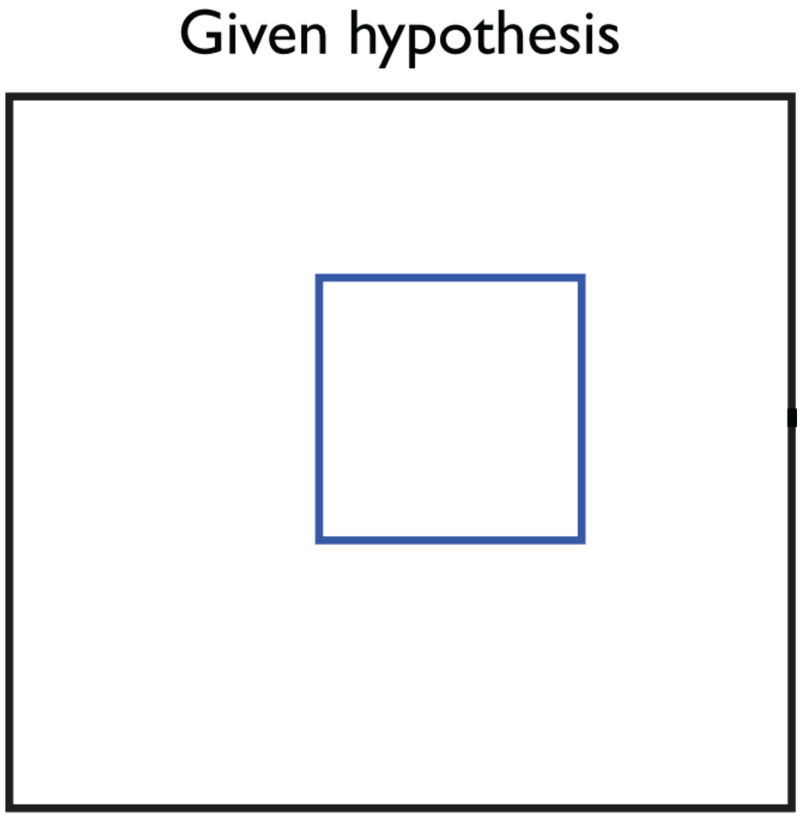

Choose two points (positive or negative) to teach this hypothesis

People should choose examples near the edges

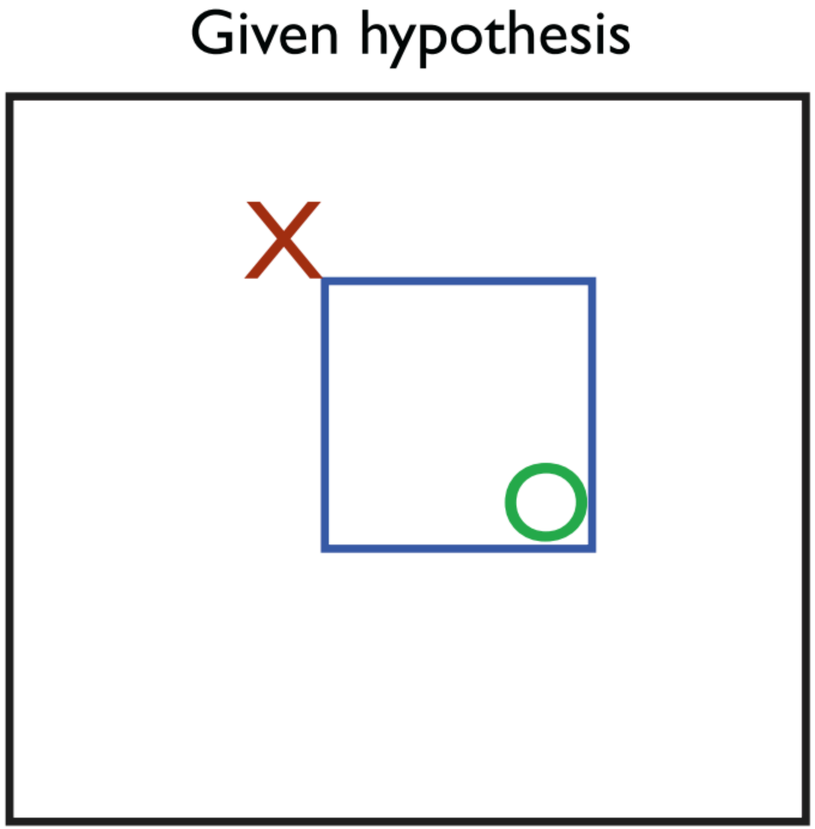

The rectangle game: Learner’s perspective

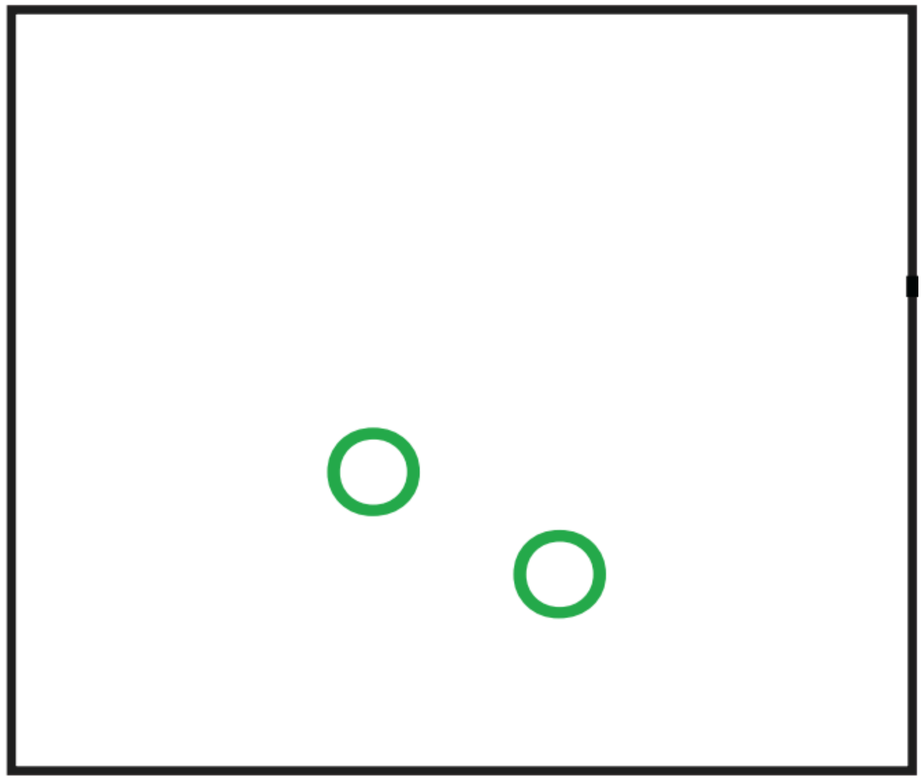

Draw the rectangle that you think these points are an example of

The rectangle game: Model details

$H =$ All possible rectangles from 2 x 2 to 5 x 5 on a 6 x 6 board.

$$P(h) = \frac{1}{196}$$ All hypotheses equally likely prior

$$P(d|h) = \begin{cases} 1, & d \in h \\ 0 & \text{otherwise} \end{cases}$$

Results for teaching

Model predictions

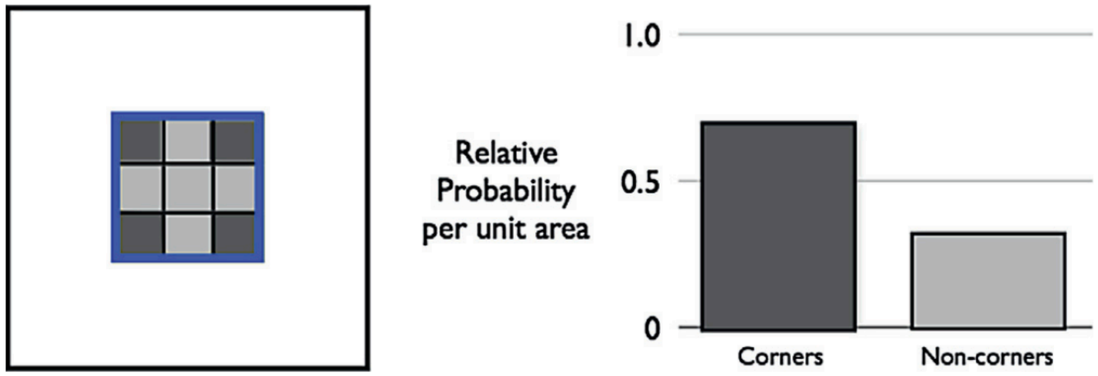

(a) Positive examples.

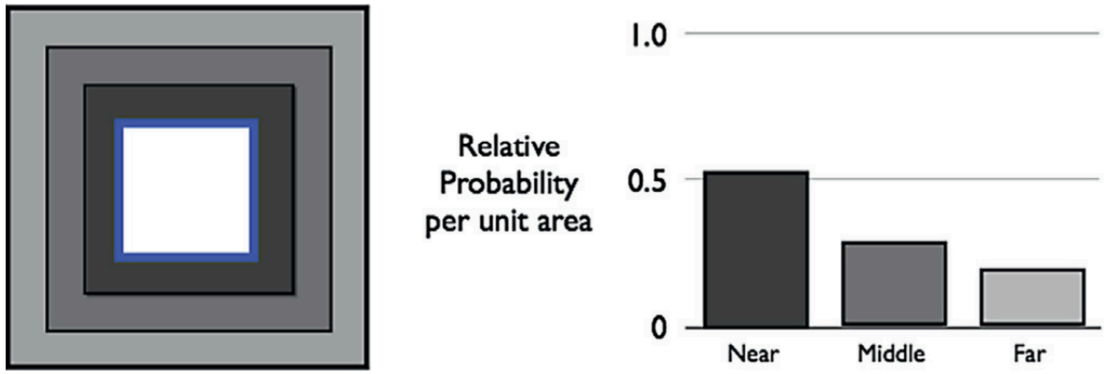

(b) Negative examples.

Experimental data

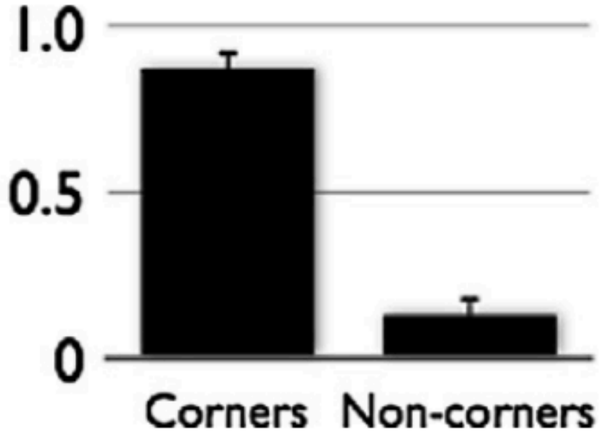

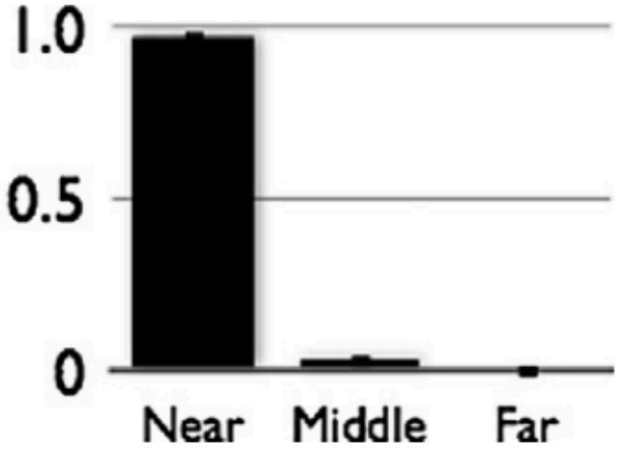

Results for learning

Positive examples

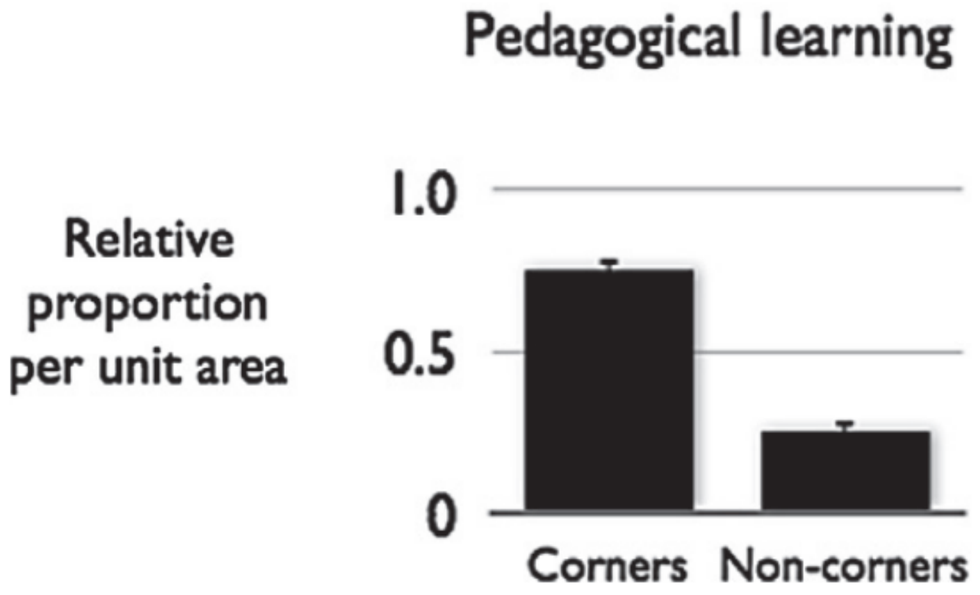

Negative examples

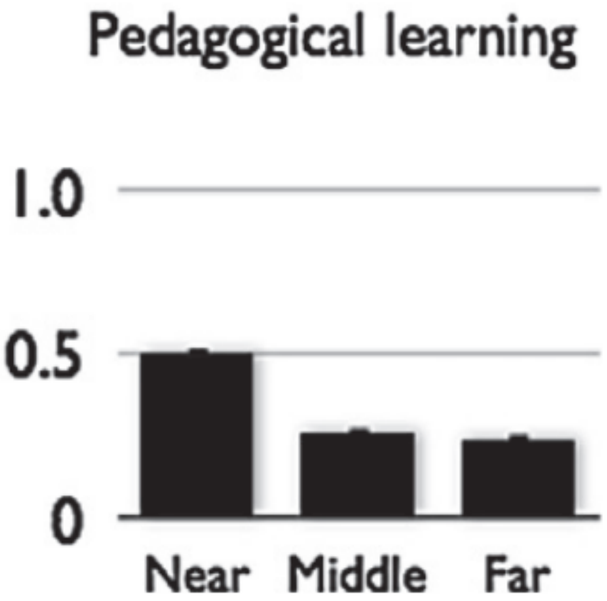

The rectangle game: Non-pedagogical learning

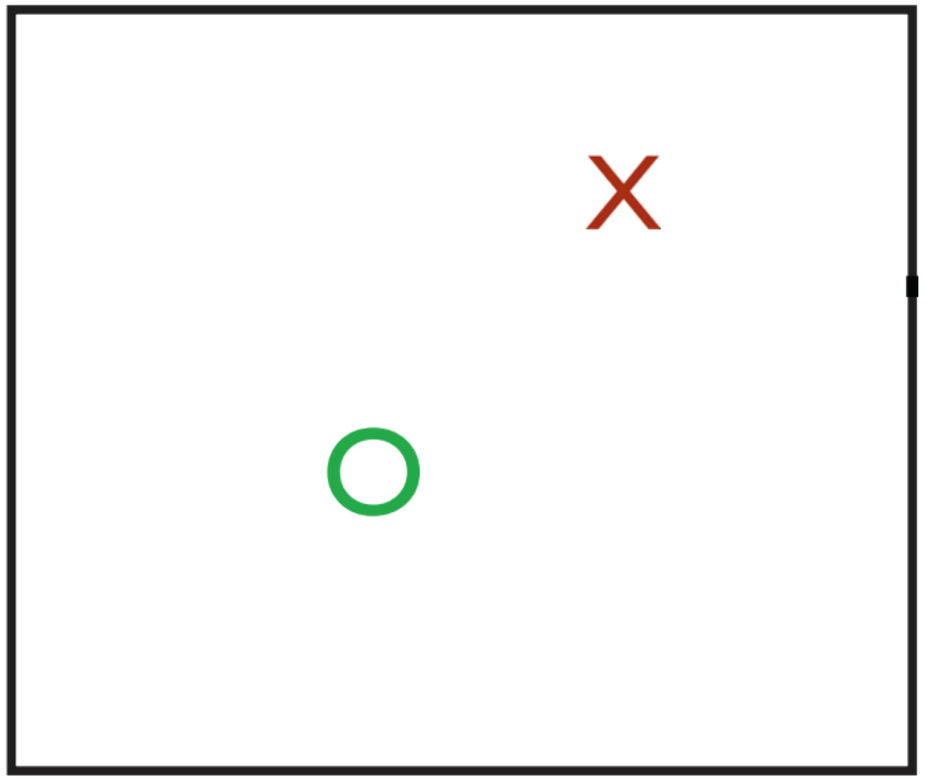

You pick two locations and get evidence for them. What do you think is the rectangle?

Results for non-pedagogical learning

Positive examples

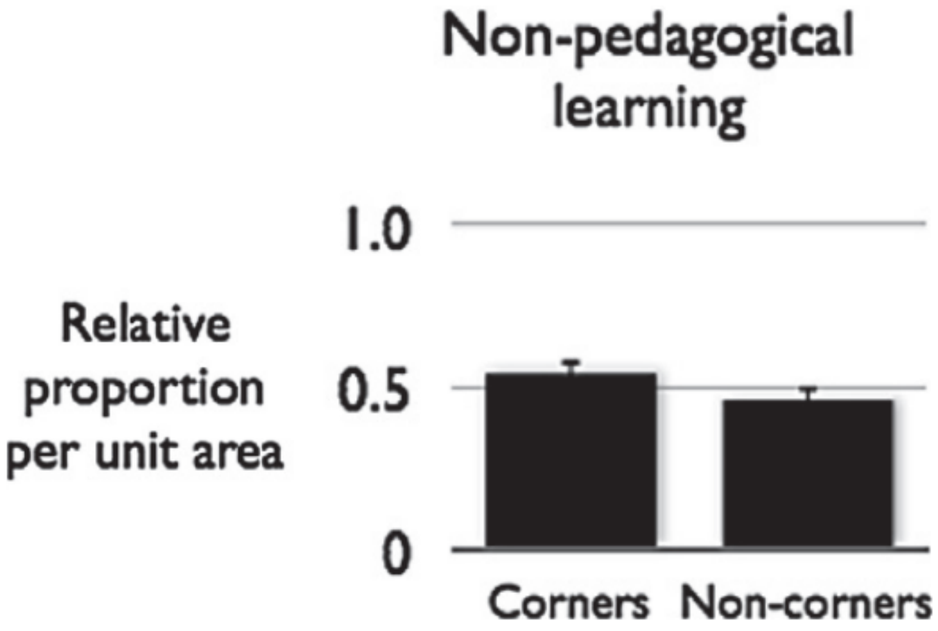

Negative examples

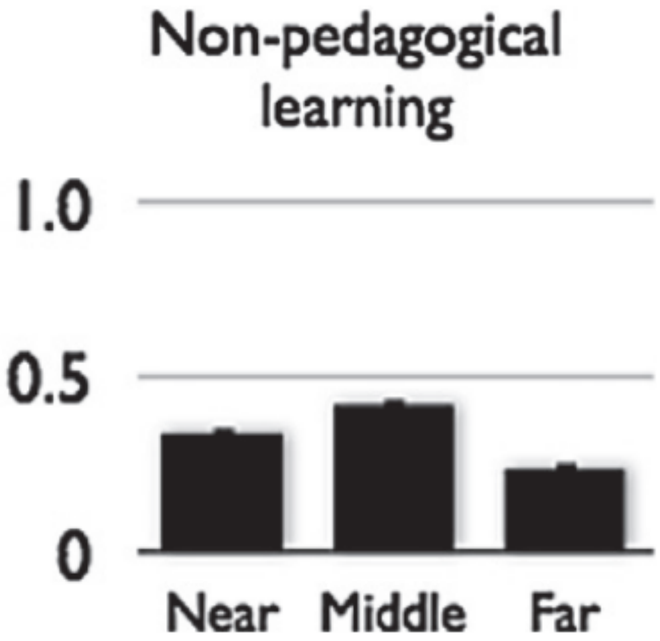

Strong vs. weak sampling in learning word meanings

dax

dalmatian

dog

animal

dax

dax

dax

dalmatian

dog

animal

Xu & Tenenbaum (2007)

Weak sampling in word meanings (Xu & Tenenbaum, 2007)

Here are three sibs. Can you give Mr. Frog all the other sibs?

To give a sib, click on it below. When you have given all the sibs, click the Next button.

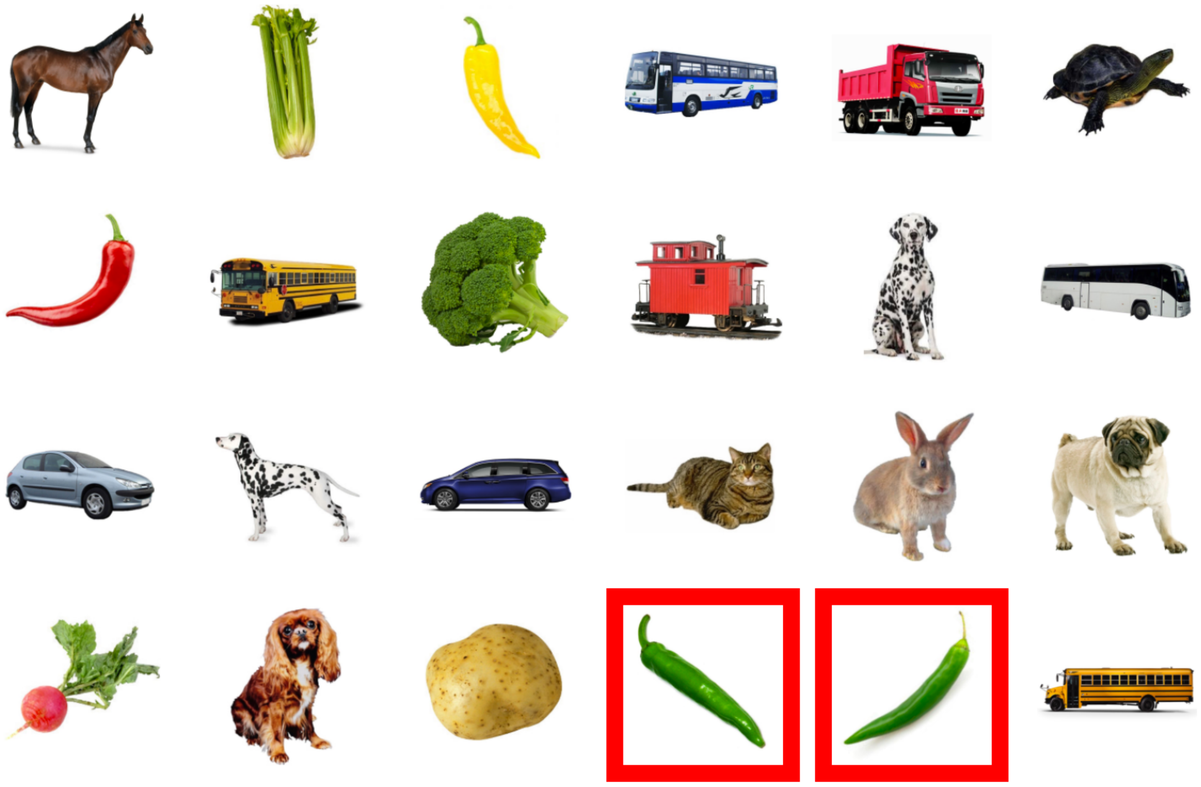

Three different kinds of sampling

Weak sampling: the samples you get are generated from the prior (or from you), and then the machine tells you which ones are consistent with the hypothesis

Strong sampling: the samples you get generated from hypothesis

Pedagogical sampling: the samples you get generated to maximize your likelihood of inferring the hypothesis

Teaching as a special goal-directed action (Shafto, Goodman, & Frank, 2012)

Suppose you want to find the best coffee in Paris

Cafe 1

Tourist buys a cup of coffee, then looks down at it

Cafe 2

Local buys a cup of coffee, then looks down at it

Cafe 3

Local buys a cup of coffee, sees you, then nods at coffee

Teaching as a special goal-directed action

You observed someone take an **action** (a)
and an **effect** (e) occurred.

What hypothesis (h) should you have about the relationship between the action a and the effect e?

Intuition: It depends on your beliefs about their goal (g)

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Learning about the worlds from observing actions

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

$$P\left(a \mid g, h\right) = \frac{P\left(g \mid a, h\right)}{\sum_{a'} P\left(g \mid a', h\right)}$$

Different kinds of observations

Unintentional Effect,
not Knowledgeable Actor,
Unknown/no goal

Intentional Effect,
Knowledgeable Actor,
Non-social Goal

Intentional Effect,
Knowledgeable Actor,
Social Goal

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Example 1: Bob's box

How do you get the light to turn on?

Physical Evidence

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Example 1: Bob's box

How do you get the light to turn on?

Goal-directed Action

None A B A&B

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Example 1: Bob's box

How do you get the light to turn on?

Communicative Action

None A B A&B

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Example 2: Tim's toy

How do you get the light to turn on?

Physical Evidence

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Example 2: Tim's toy

How do you get the light to turn on?

Goal-directed
Action

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Example 2: Tim's toy

How do you get the light to turn on?

Communicative Action

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

Tim’s toy in real life (Bonawitz, Shafto, Gweon, Goodman, Spelke, & Schulz, 2011)

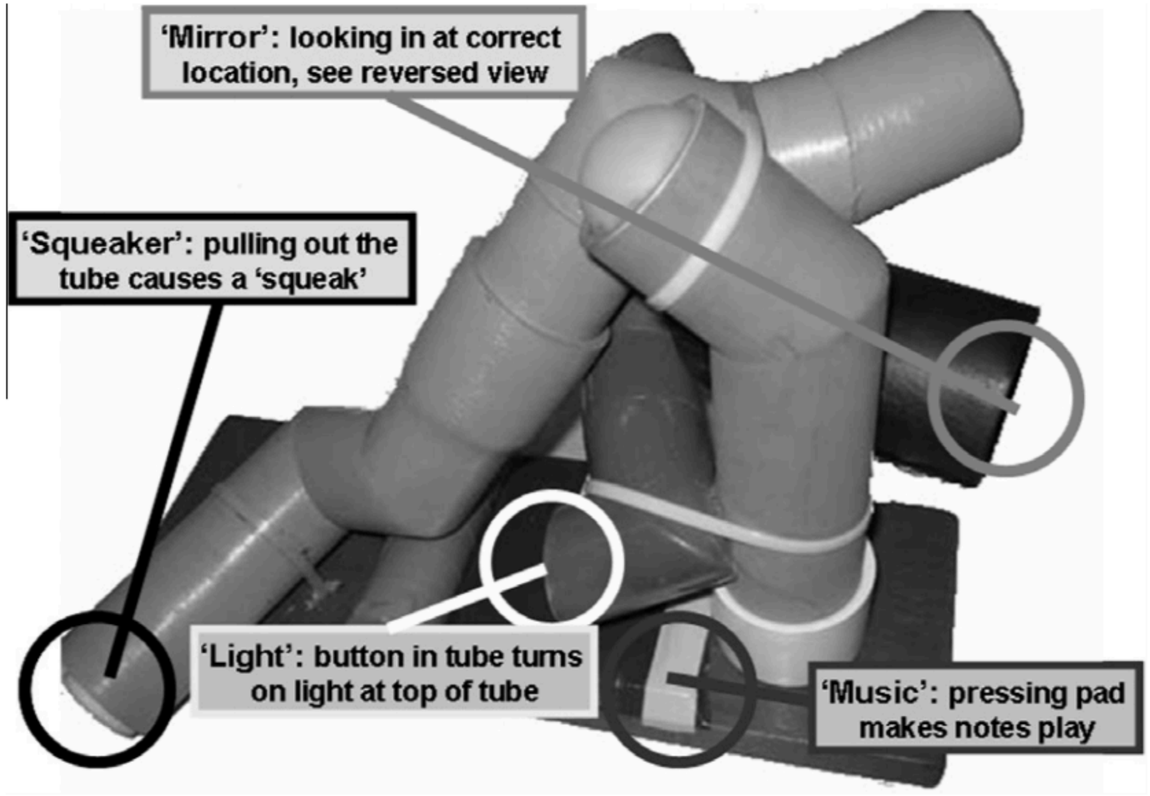

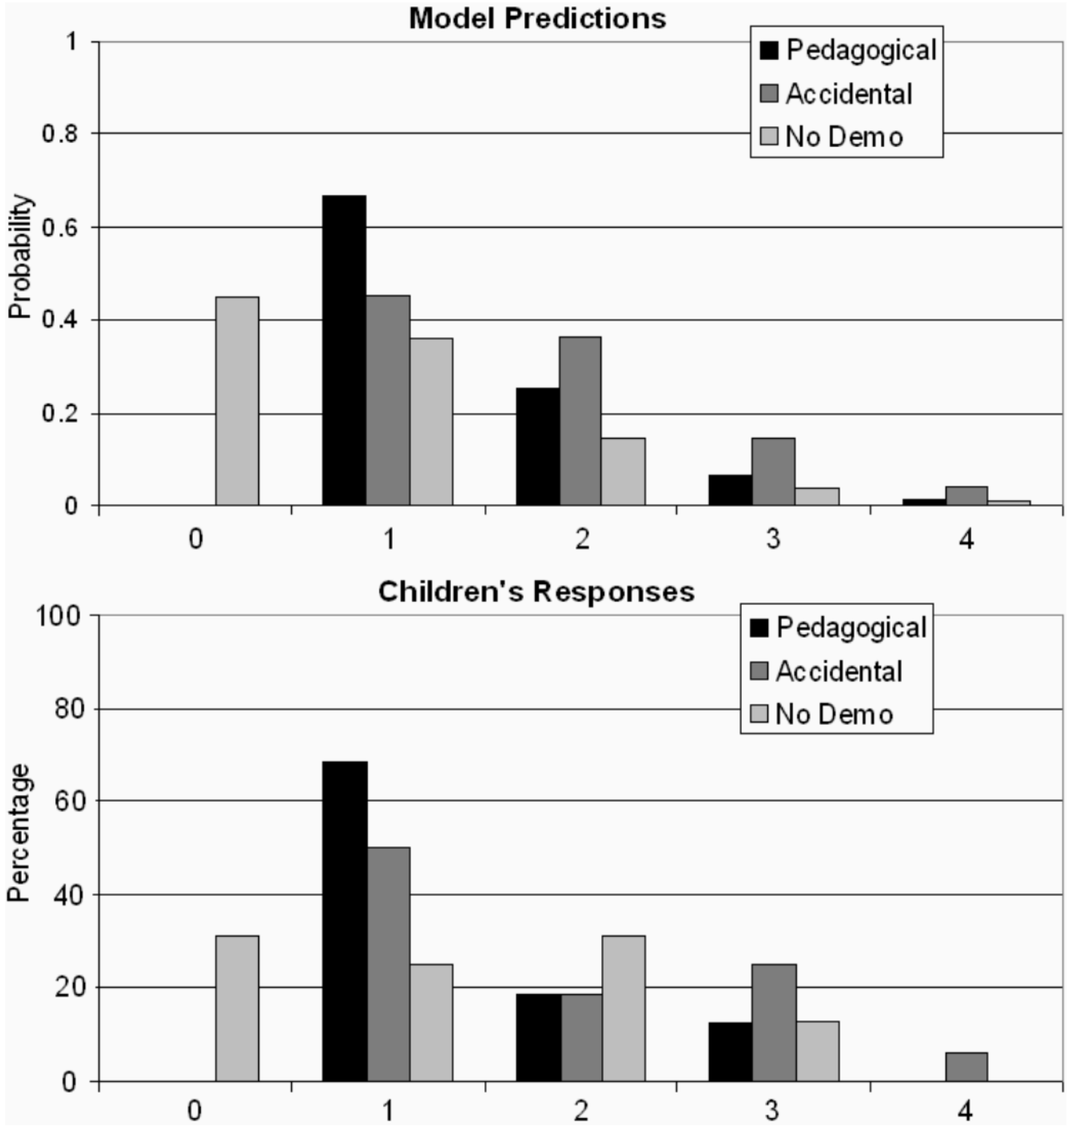

Does this feel right to you? What can the model explain? What can't it?

$$P\left(h \mid a, e, g\right) \propto P\left(e \mid a, h\right) P\left(a \mid g, h\right) P\left(h\right)$$

$$P\left(a \mid g, h\right) = \frac{P\left(g \mid a, h\right)}{\sum_{a'} P\left(g \mid a', h\right)}$$

Comparing models

1. Learning from teaching is different from learning from observation

2. A rational model of teaching and learning from teaching

3. Teaching has tradeoffs.
   Some are predictable from this rational model