Computational Cognitive Science

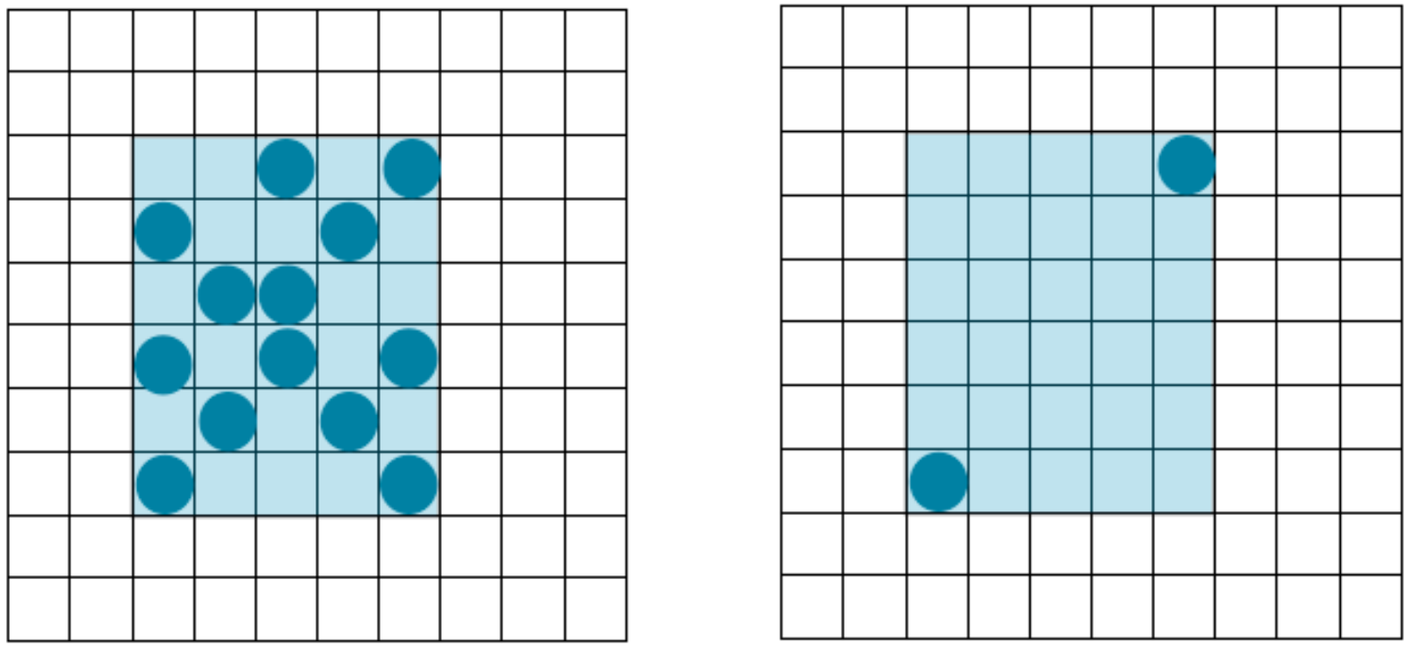

Lecture 21: Pedagogical sampling

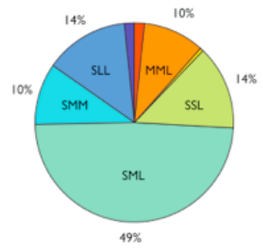

(b) Pedagogical model predictions

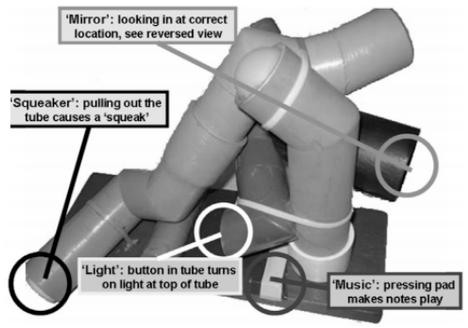

Fig. 1. Novel toy stimuli used in Experiments 1 and 2.

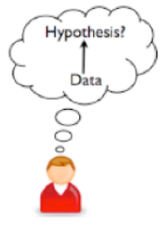

Last time

- Difference between strong and weak sampling

Strong

- Items generated from concept
- Additional items lead to tighter generalisations

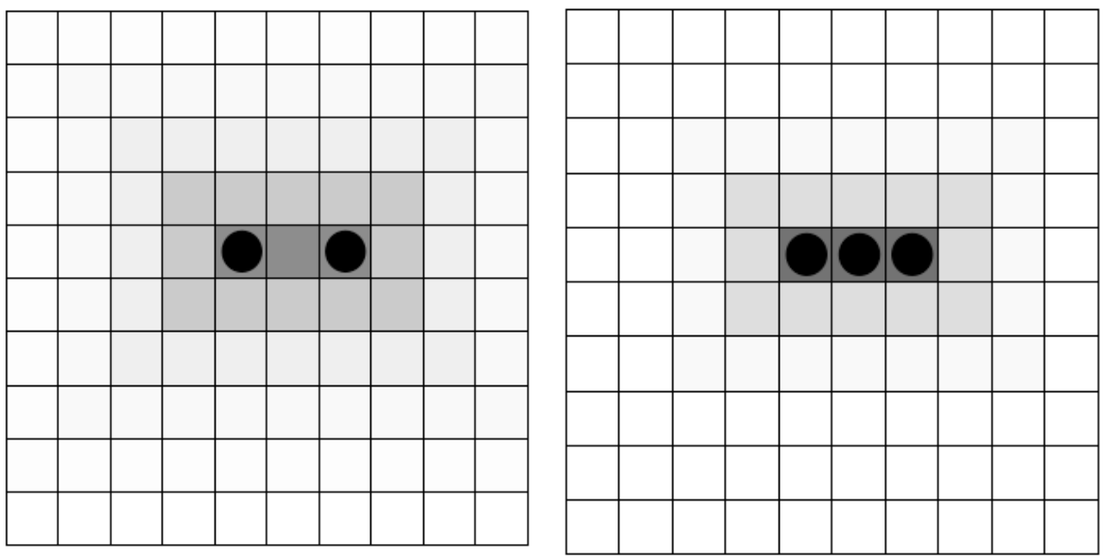

Weak

- Items generated from world and then labelled
- Additional items do not lead to tighter generalisations

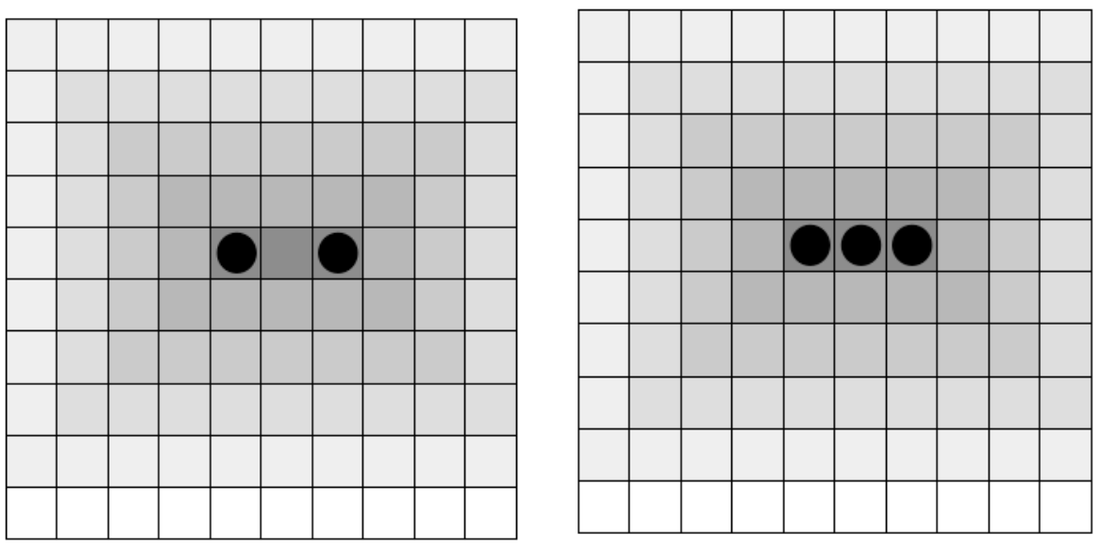

Last time

- Difference between strong and weak sampling
- People pay attention to how data were sampled when figuring out how to generalise words

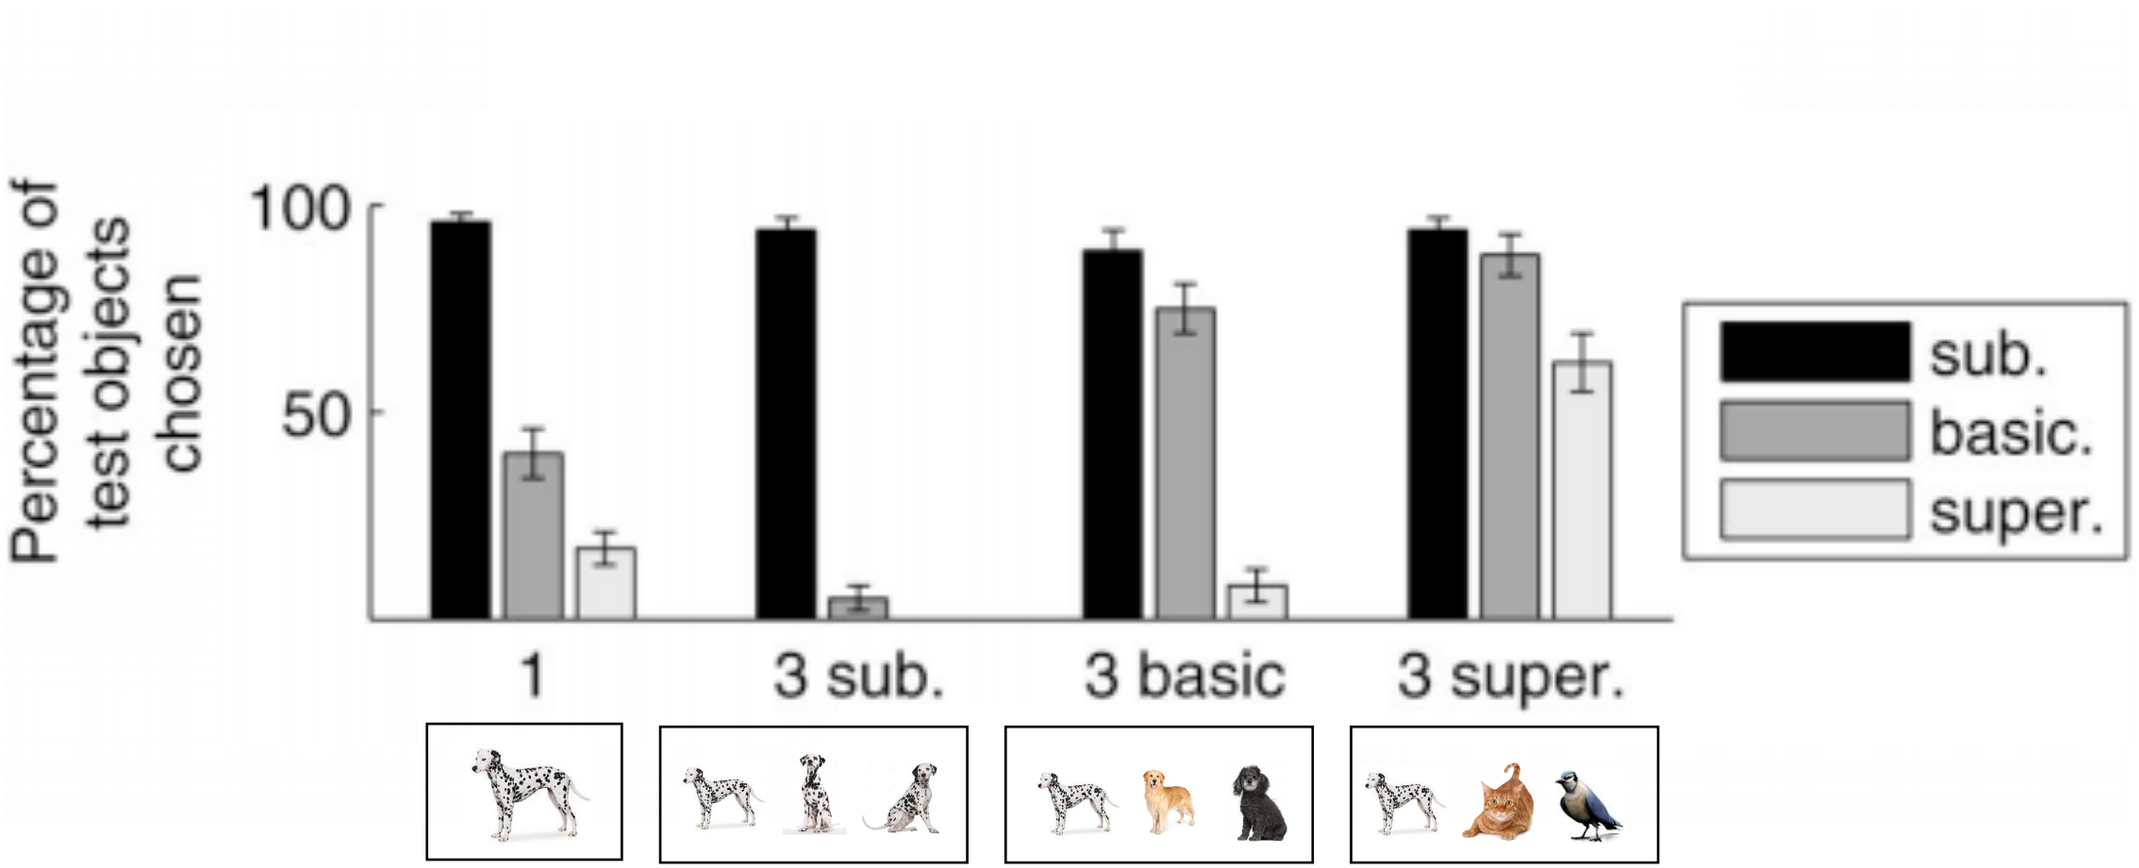

Last time

- Difference between strong and weak sampling
- People pay attention to how data were sampled when figuring out how to generalise words, changing their generalisations if they did not come from the concept

Learner-driven

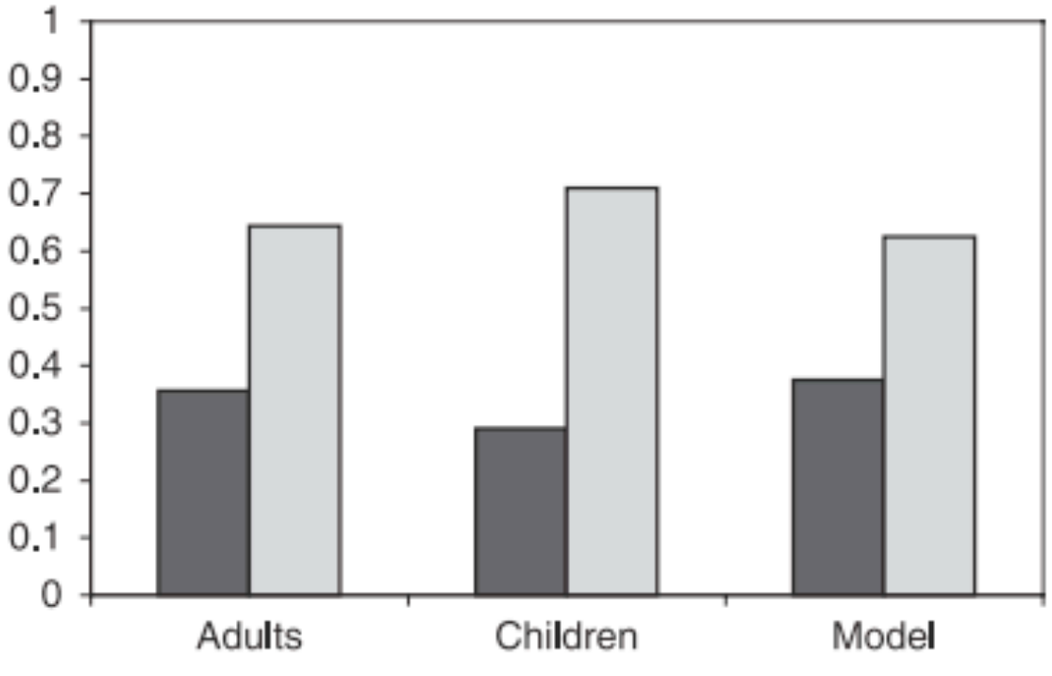

Teacher-driven

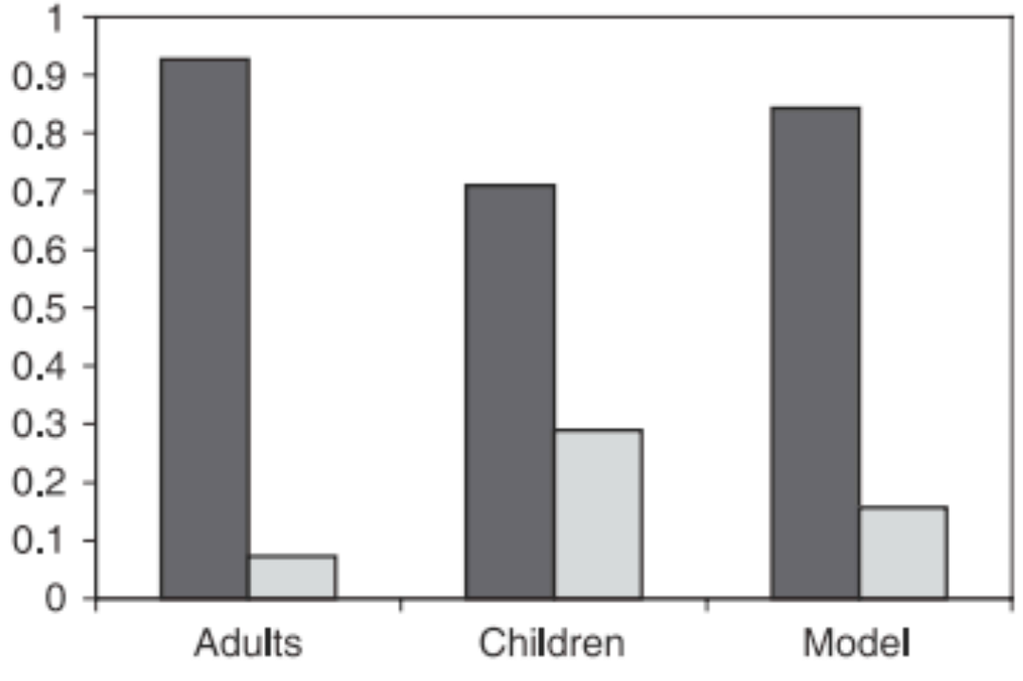

Last time

- Difference between strong and weak sampling
- People pay attention to how data were sampled when figuring out how to generalise words, changing their generalisations if they did not come from the concept
- Even infants do this, and with novel features

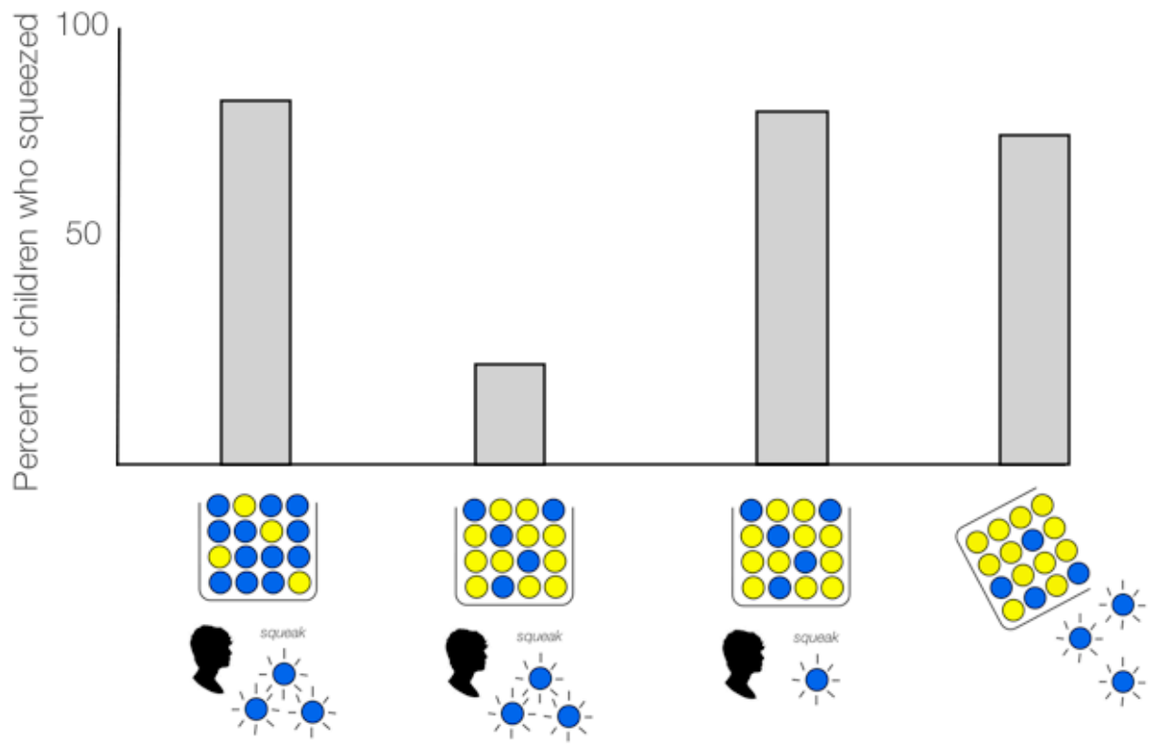

Last time

- Difference between strong and weak sampling
- People pay attention to how data were sampled when figuring out how to generalise words, changing their generalisations if they did not come from the concept
- Even infants do this, and with novel features
- People show strong individual differences in the amount to which they assume strong sampling, but almost always they tighten at least somewhat

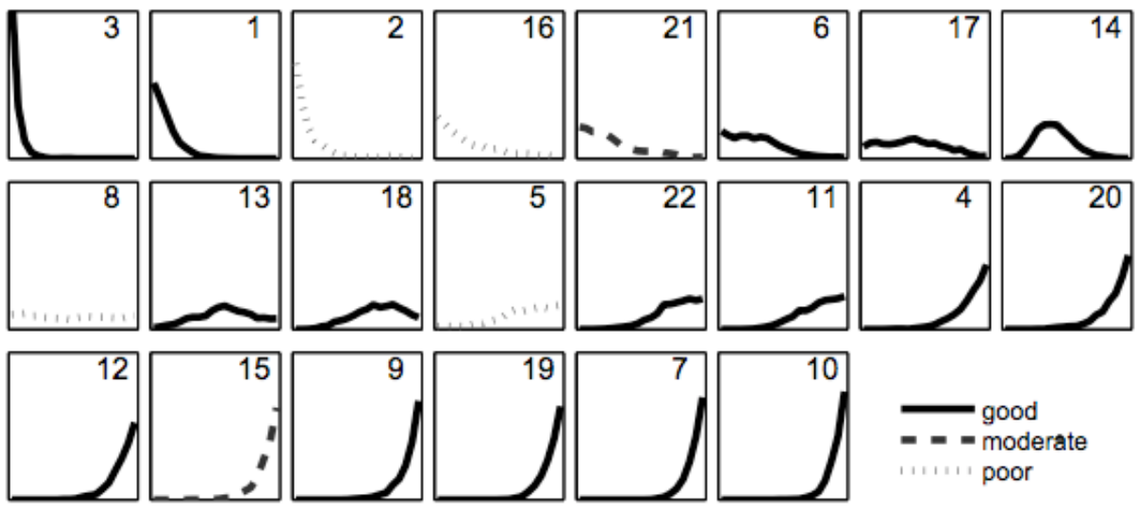

But.... note one thing

- Data here aren't actually being generated from the concept. They are being provided by a helpful teacher. Does this matter?

... well, yes.

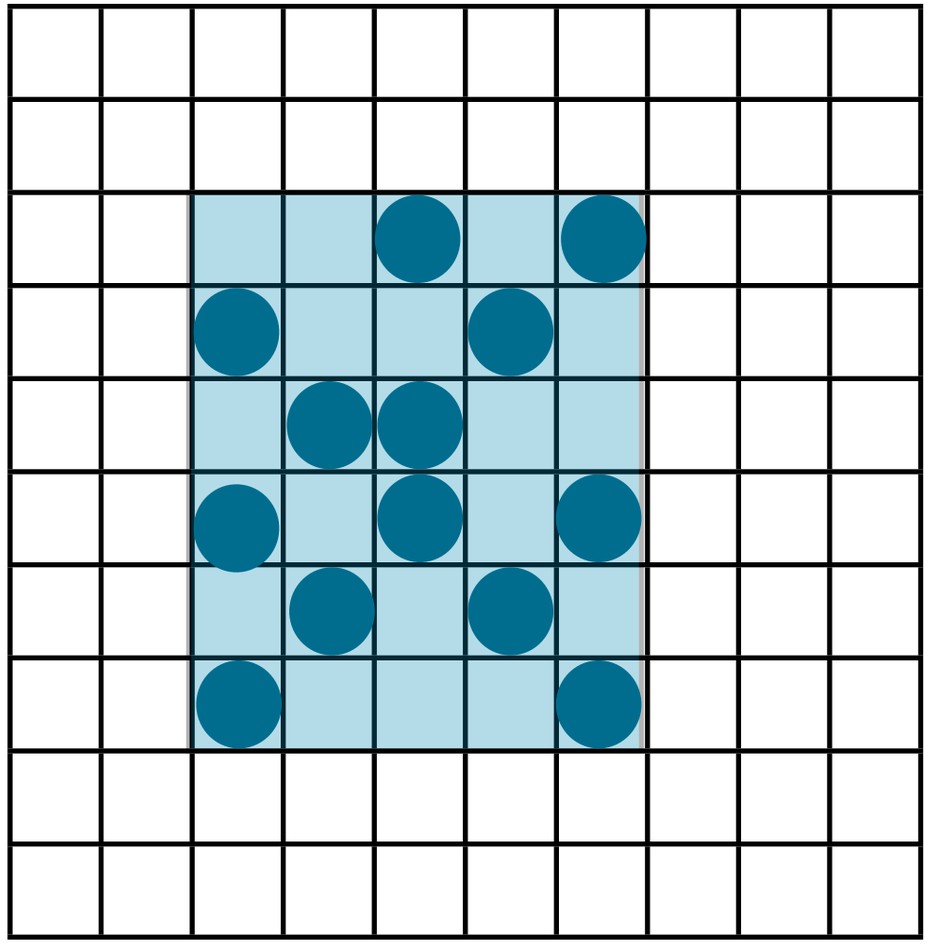

Suppose you want people to guess the bounds of the rectangle.

If data are generated randomly from it, you'll get there eventually

But.... note one thing

- Data here aren't actually being generated from the concept. They are being provided by a helpful teacher. Does this matter?

... well, yes.

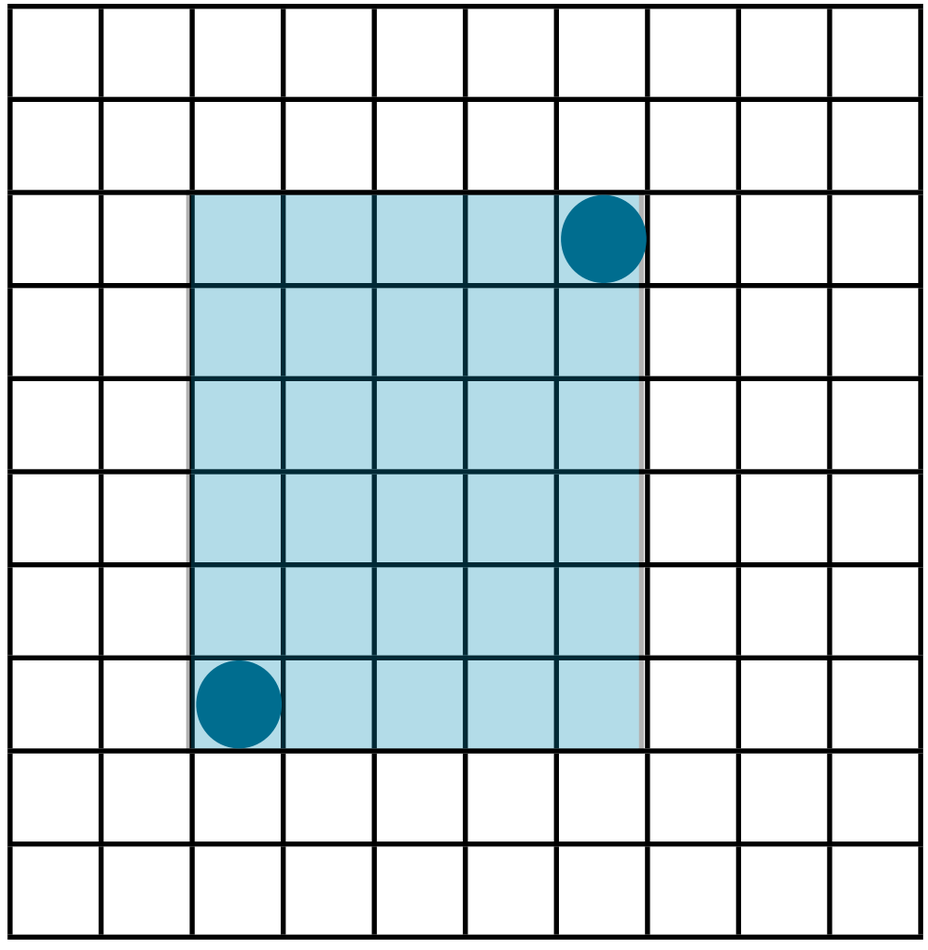

Suppose you want people to guess the bounds of the rectangle.

If a teacher chooses points specifically (and you know that's what they're doing), you will get there much faster

Plan today

- Lecture 1 [short]: Another kind of sampling: pedagogical
  - Model for pedagogical sampling
  - Sensitivity to pedagogical data
  - Double-edged sword of pedagogy
- Lecture 2 [long]: Sensitivity to the informativeness of the data
  - Confirmation bias at the positive test strategy (PTS)
  - When is positive evidence most useful?
  - When do people use a PTS?
  - Are people sensitive to evidence utility in general?

Pedagogical sampling

- Assume data are generated from another person, with beliefs and desires.

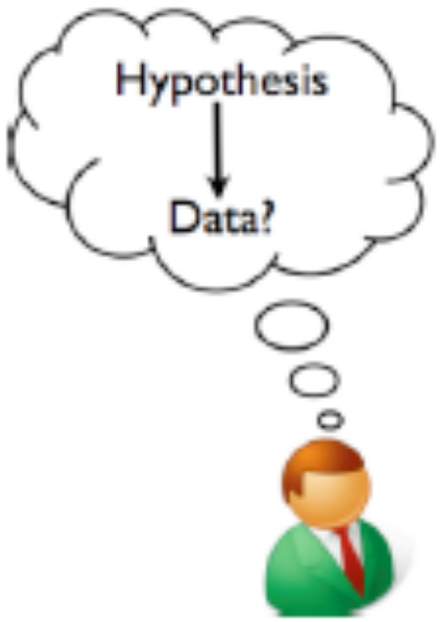

Teacher: knows the hypothesis, generates data

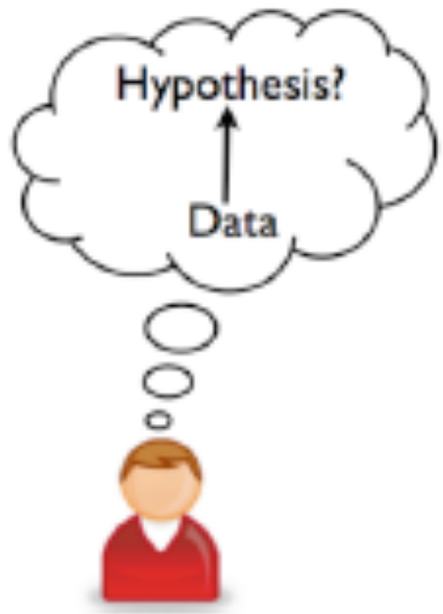

Learner: knows the data, figures out hypothesis

Pedagogical sampling

- Assume data are generated from another person, with beliefs and desires.

Now it is a system of yoked equations

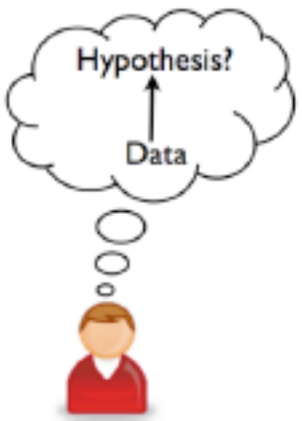

Learner (recipient) assumes that the teacher (provider) is generating data based on what they know and what their goal is (i.e., to help the learner or not). This is given by:

$$P_{\text{rec}}(h|x) = \frac{P_{\text{prov}}(x|h)P(h)}{\sum_{h'} P_{\text{prov}}(x|h')P(h')}.$$

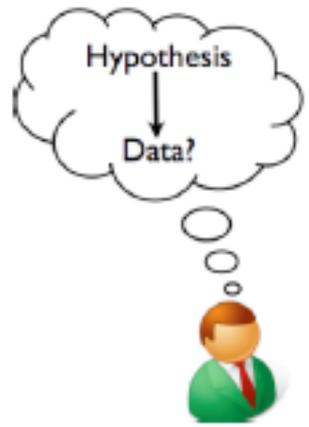

Teacher (provider) provides data so as to manipulate beliefs of the learner. If they are helpful (α>0), they are trying to make the learner infer the correct belief. If not, they want to mislead.

$$P_{\text{prov}}(x|h) = \frac{P_{\text{rec}}(h|x)^{\alpha}}{\sum_{x} P_{\text{rec}}(h|x)^{\alpha}}$$

Pedagogical sampling

- This predicts that some choices of data are better (more informative) than others

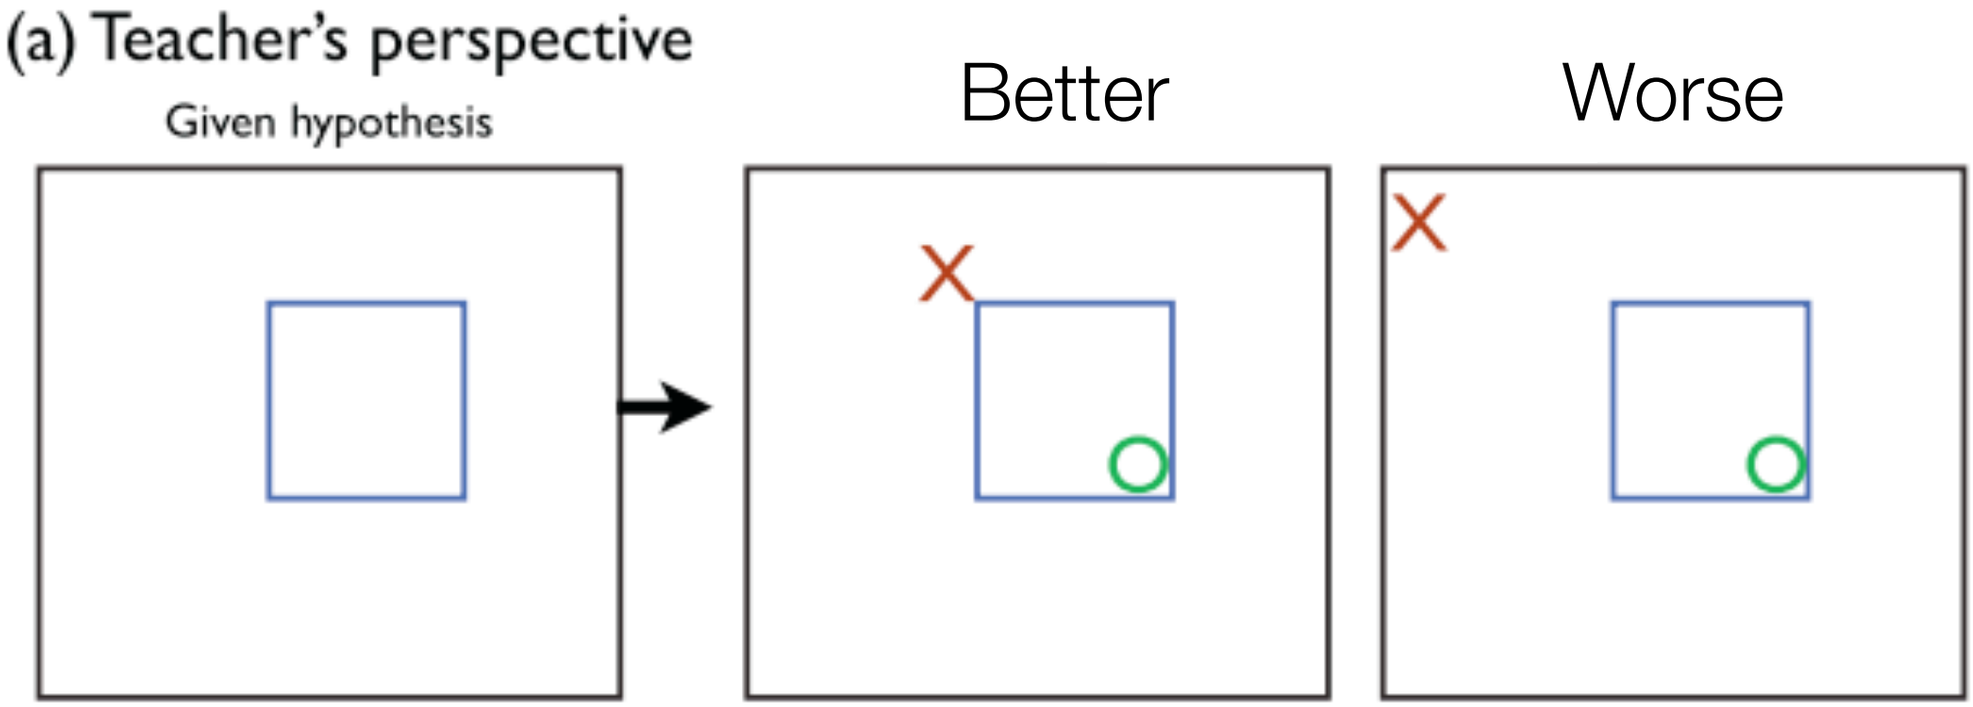

Pedagogical sampling

- This predicts that some choices of data are better (more informative) than others

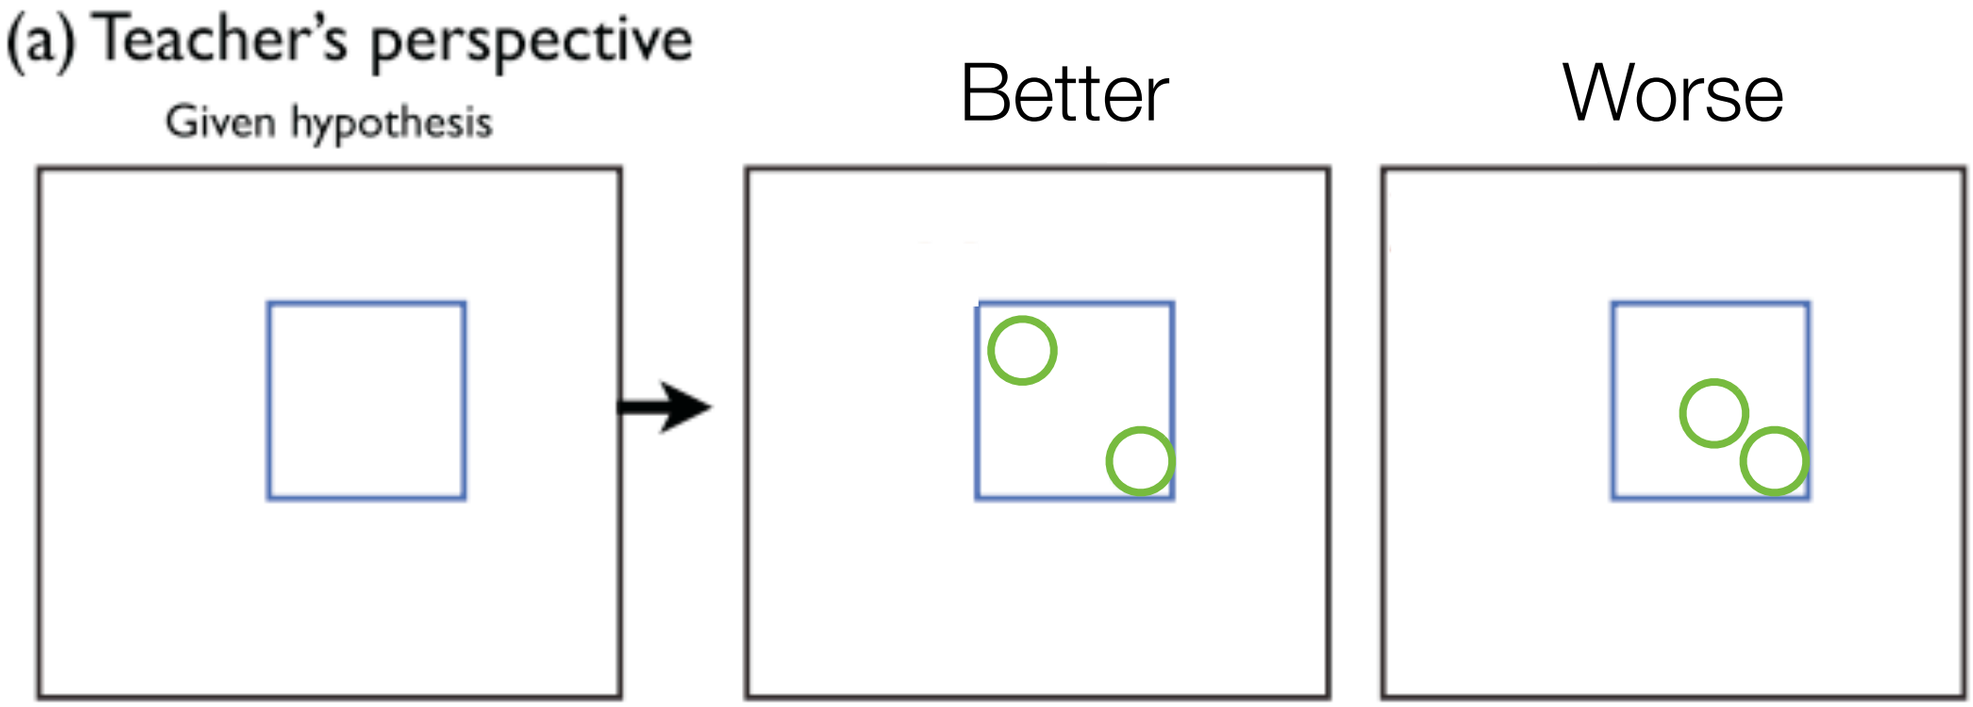

Pedagogical sampling

- This predicts that some choices of data are better (more informative) than others

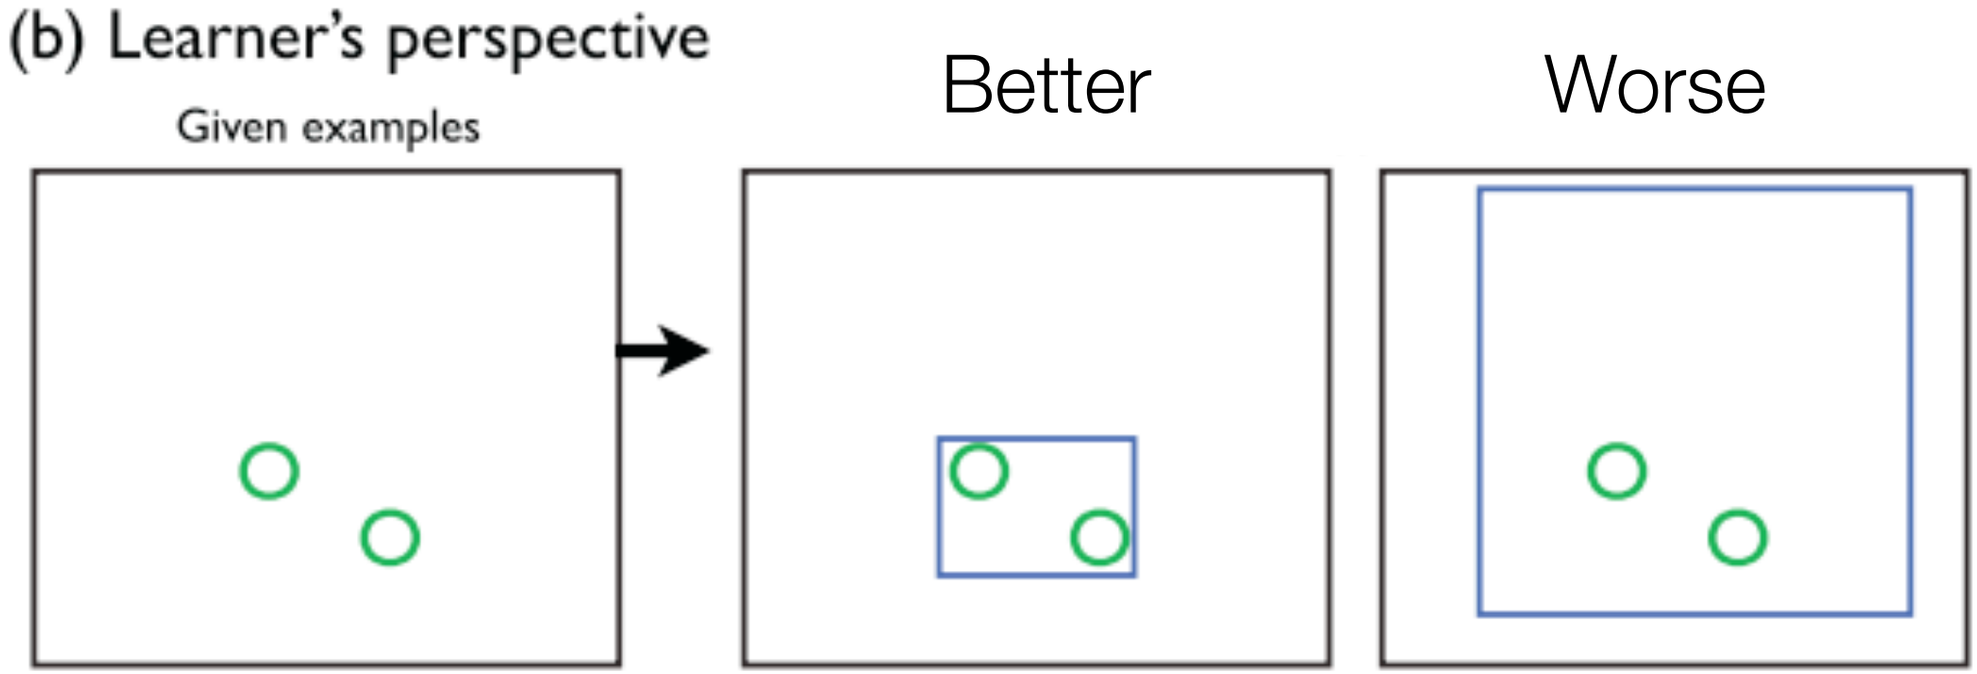

Pedagogical sampling

- The model supports this intuition: some data are more informative for teachers to provide

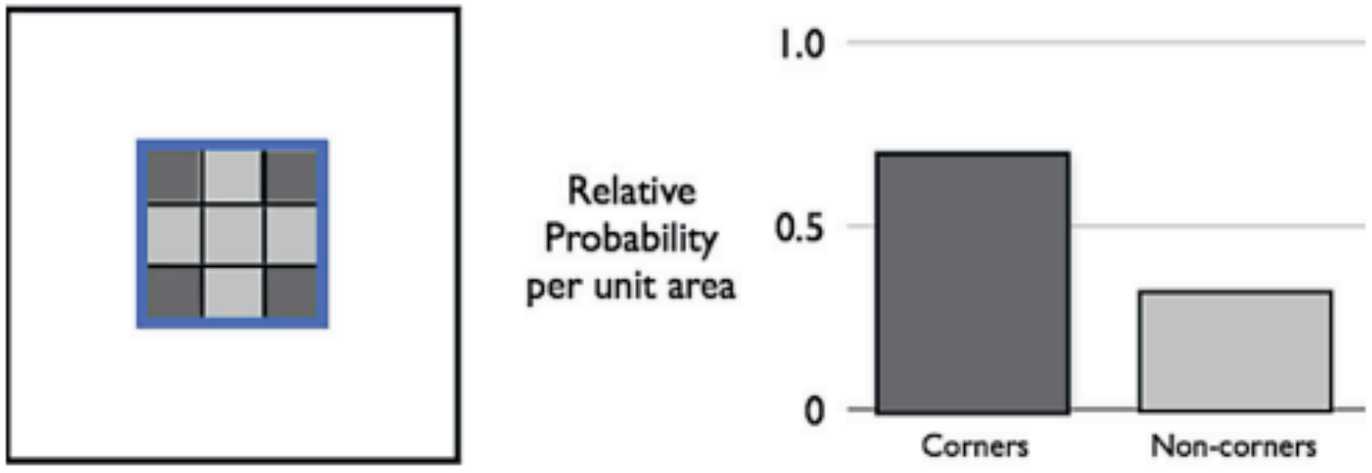

(a) Positive examples.

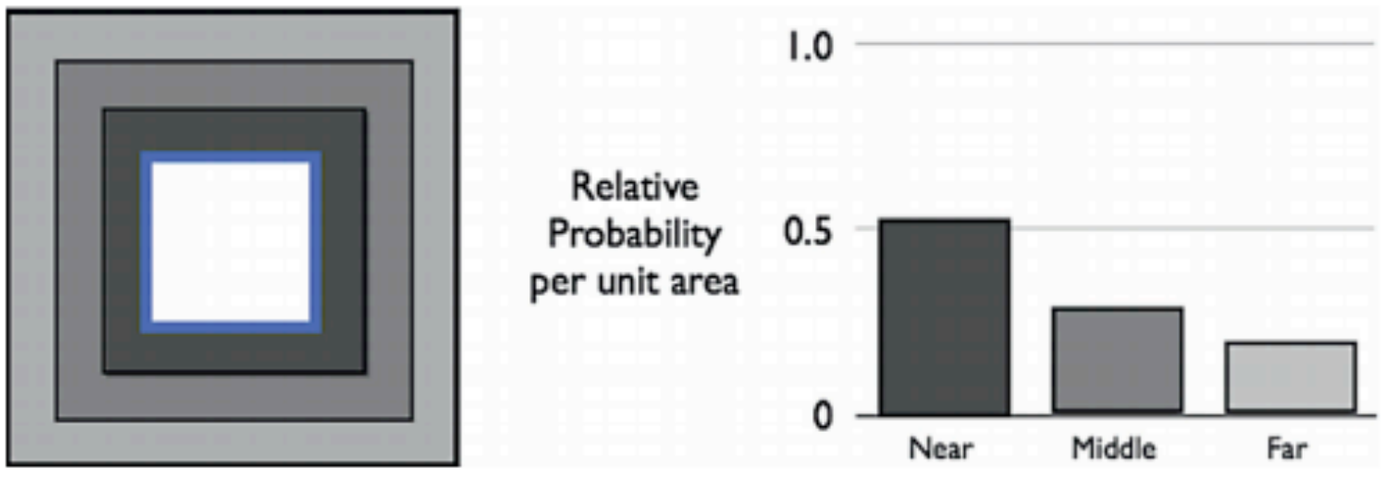

(b) Negative examples.

Pedagogical sampling

- Teachers have the same intuition, providing sensible data

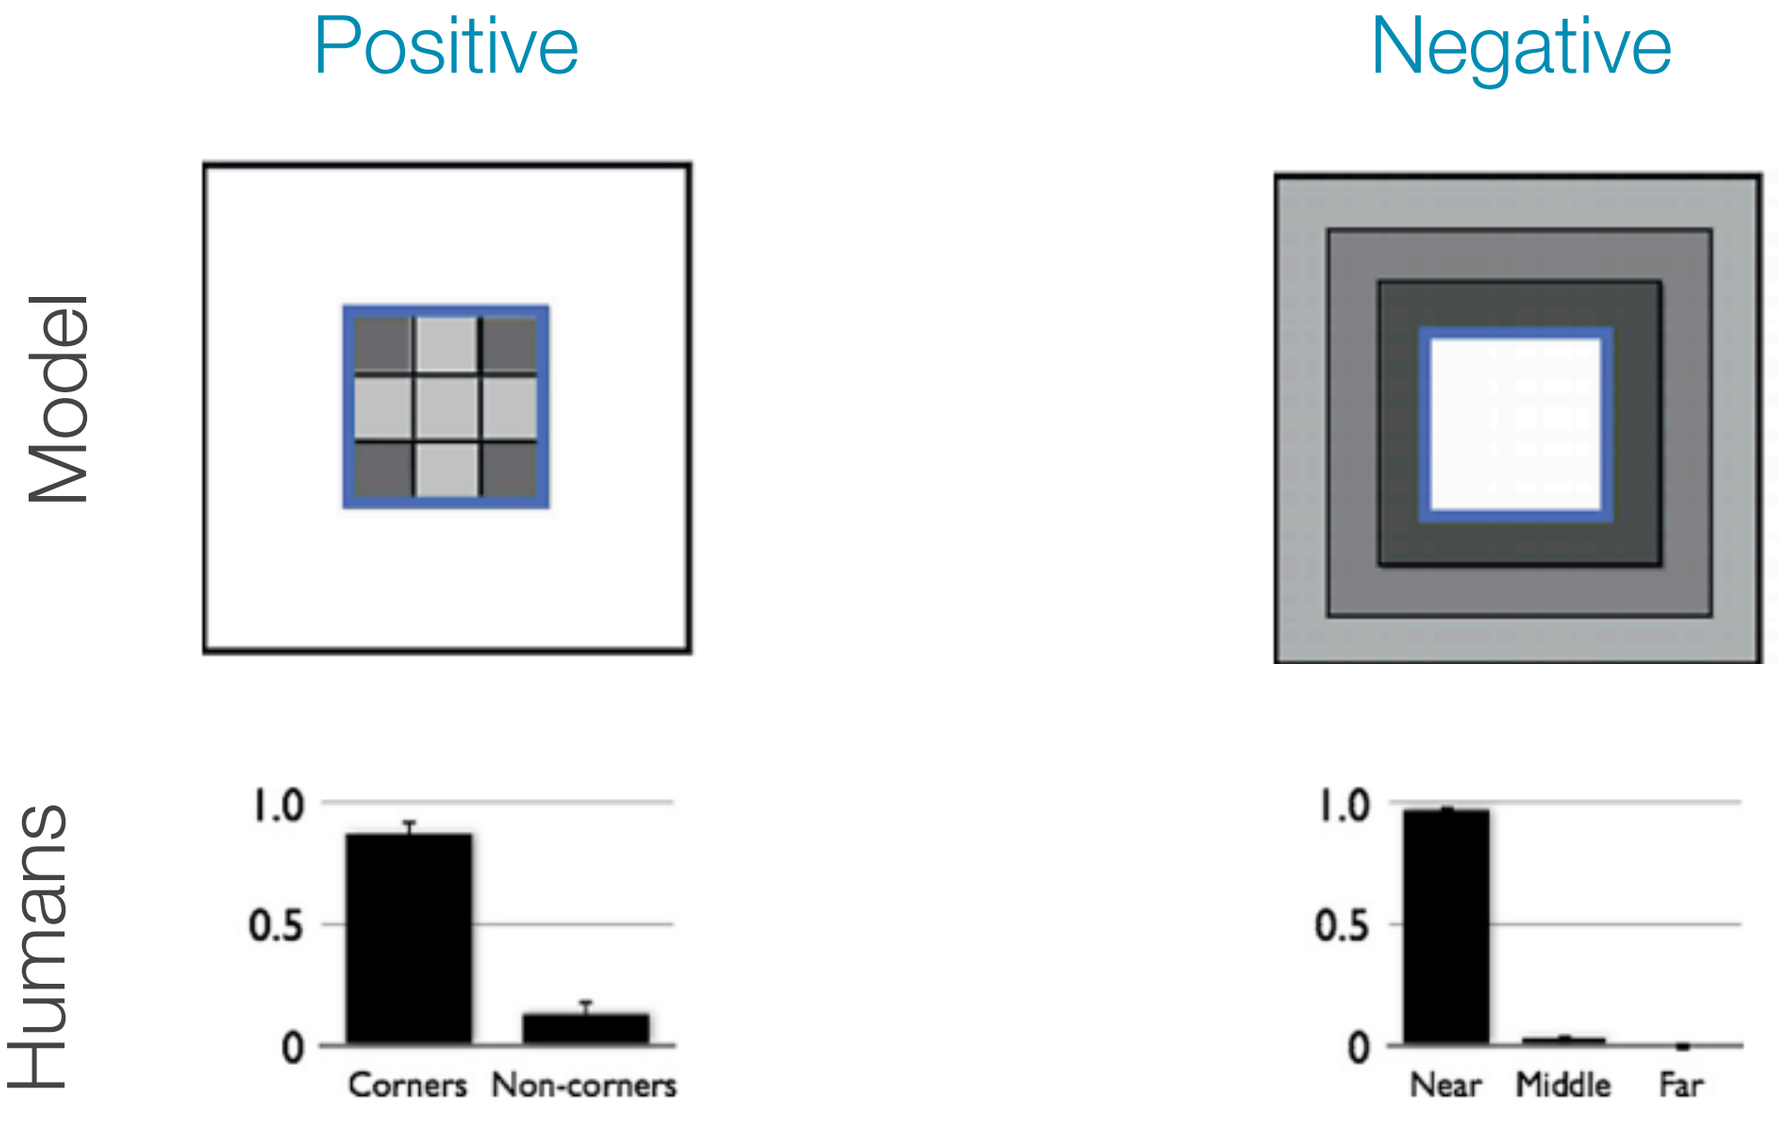

Pedagogical sampling

- Learners make sensible inferences based on teacher's data

y axis: proportion making corresponding inference

	2 Positive	2 Negative
From teacher	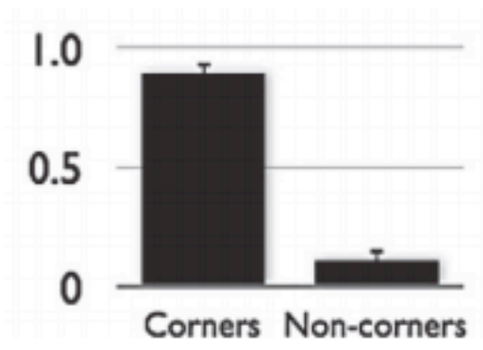	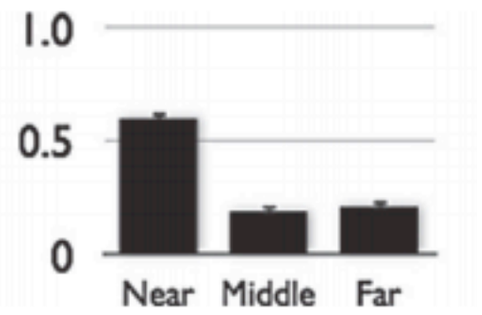
Self-generated	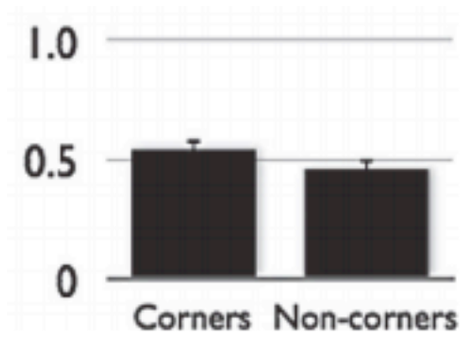	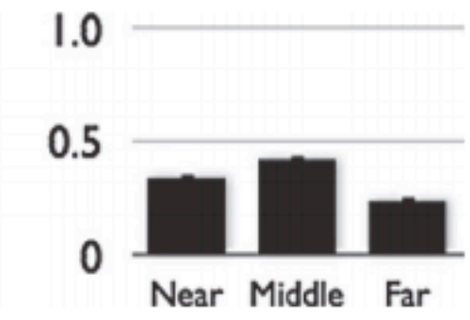

Pedagogical sampling

- What if the situation is less contrived / more realistic?

People are given 27 examples of a
concept based on line length

Pedagogical sampling

- What if the situation is less contrived / more realistic?

People are given 27 examples of a concept based on line length

Pedagogical sampling

- What if the situation is less contrived / more realistic?

Teachers are asked to provide three examples to teach the concept

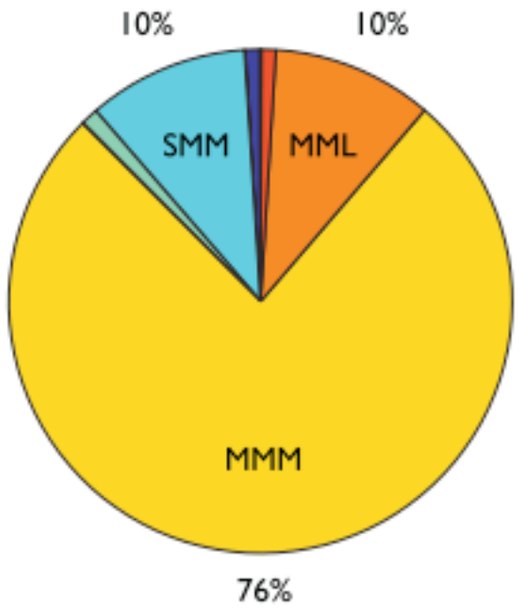

(a) Strong sampling predictions

MMM SMM MML

Pedagogical sampling

- What if the situation is less contrived / more realistic?

Teachers are asked to provide three examples to teach the concept

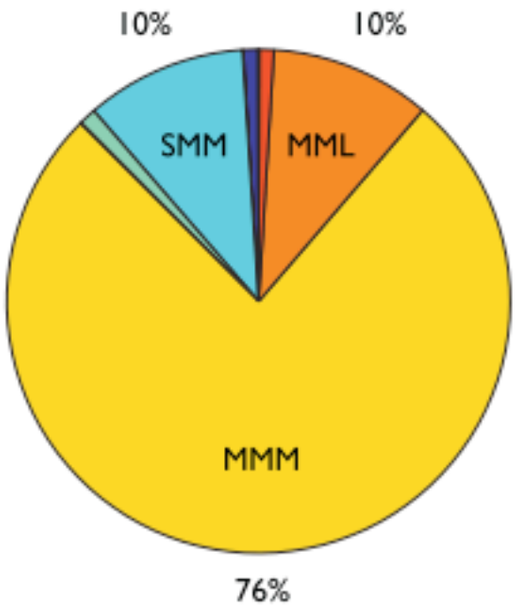

(a) Strong sampling predictions

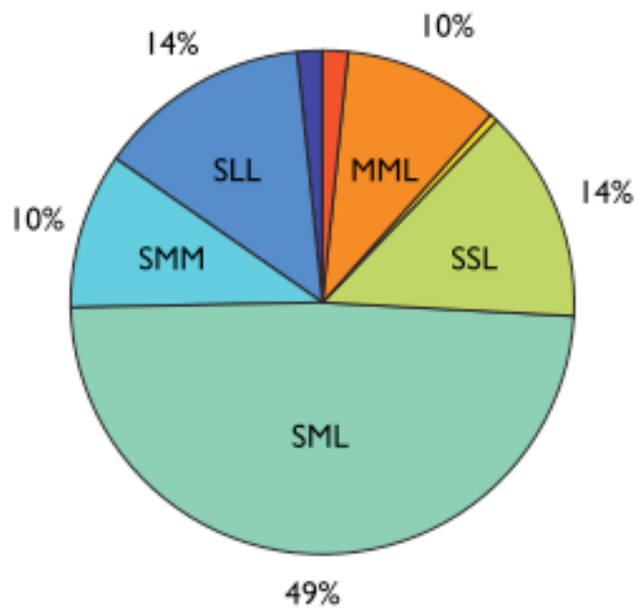

(b) Pedagogical model predictions

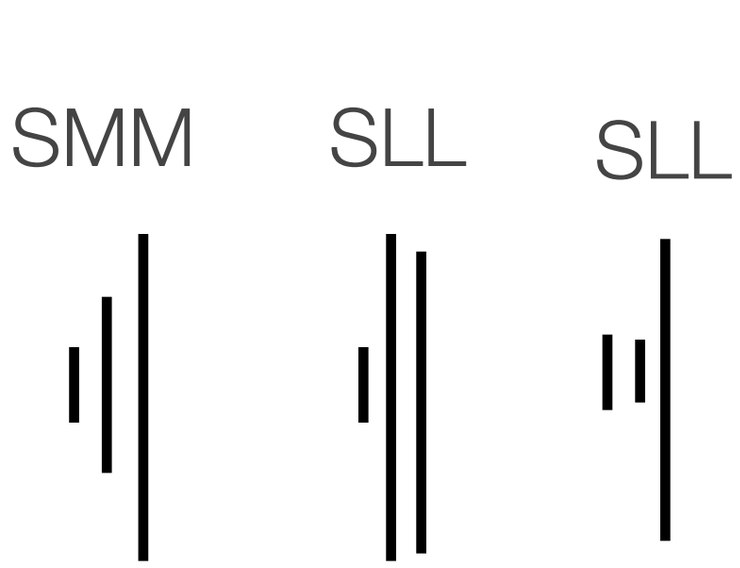

Pedagogical sampling

- What if the situation is less contrived / more realistic?

Teachers are asked to provide three examples to teach the concept

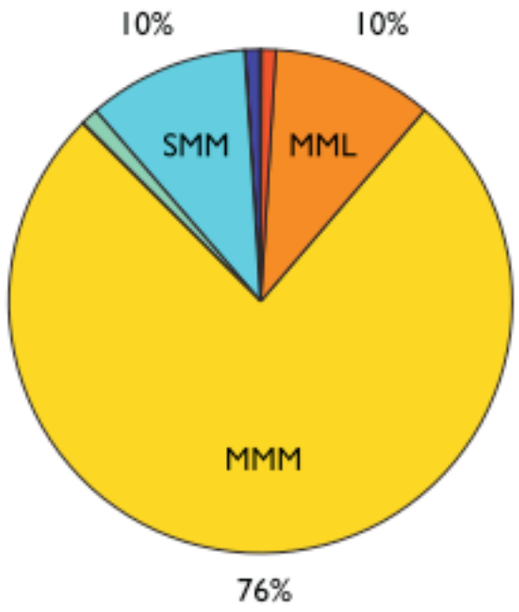

(a) Strong sampling predictions

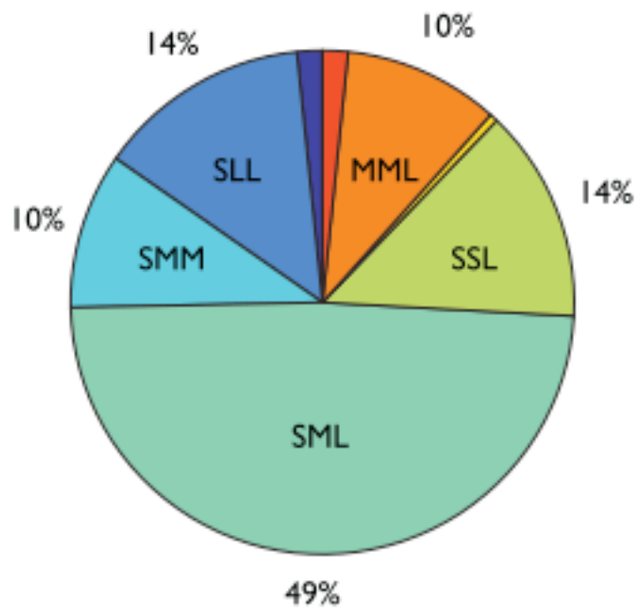

(b) Pedagogical model predictions

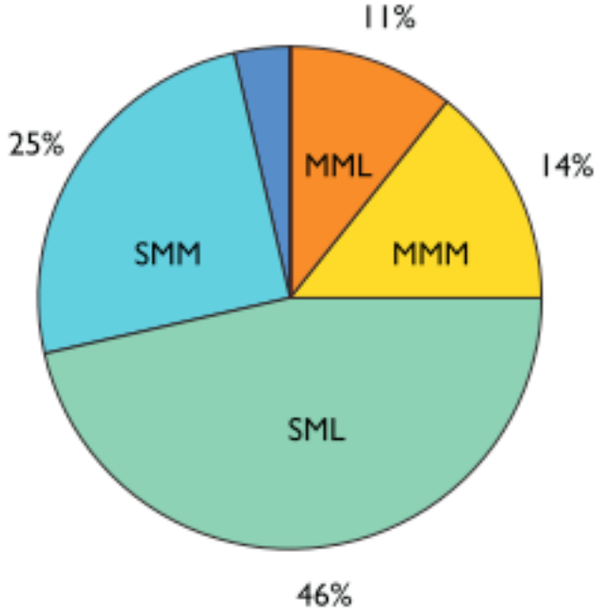

(c) Experiment results

Pedagogical sampling

- Are *children* sensitive to the presence of a helpful teacher?
  - We've already seen evidence that they are - the earlier experiments with children all involved teachers providing the data
  - This resulted in different inferences
- Another possibility: one of the roles and effects of teaching is to tell you what the structure of the hypothesis space is

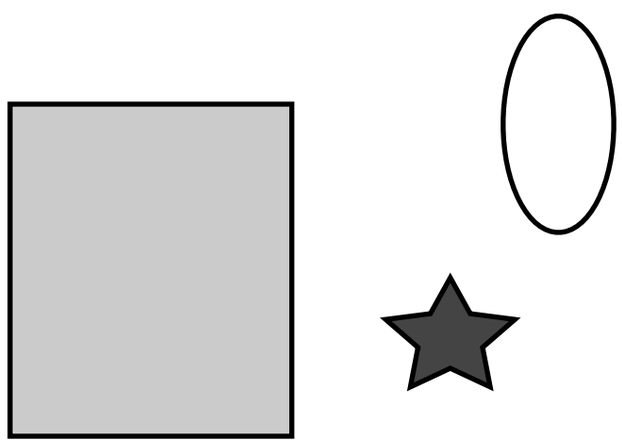

example:
toys out of a
toy box

Pedagogical sampling

- Are *children* sensitive to the presence of a helpful teacher?
  - We've already seen evidence that they are - the earlier experiments with children all involved teachers providing the data
  - This resulted in different inferences
- Another possibility: one of the roles and effects of teaching is to tell you what the structure of the hypothesis space is

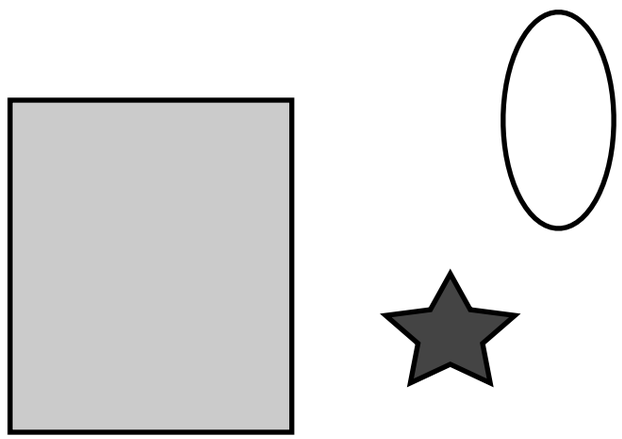

which of other toys would be most surprising?

Pedagogical sampling

- Are *children* sensitive to the presence of a helpful teacher?
  - We've already seen evidence that they are - the earlier experiments with children all involved teachers providing the data
  - This resulted in different inferences
- Another possibility: one of the roles and effects of teaching is to tell you what the structure of the hypothesis space is

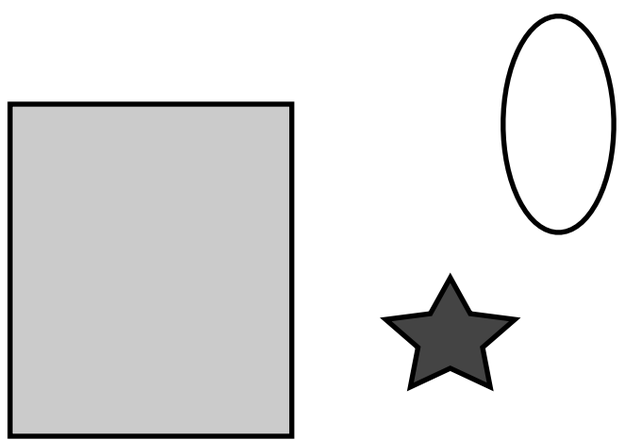

a red toy would be very surprising, because a teacher should choose examples that illustrate the *range* within the hypothesis space

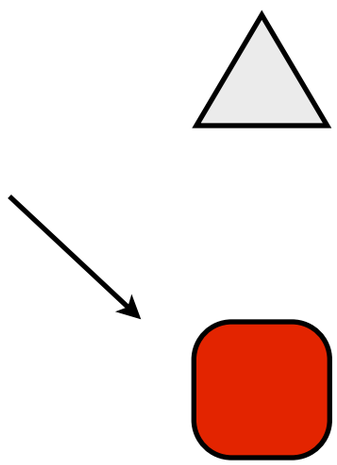

Hypothesis: pedagogy inhibits exploration

- If a teacher illustrates the range of possibilities (i.e., constrains the hypothesis space) then there's no point in exploring more
- By contrast, if you get the same data but not from a teacher, exploration might still be profitable

Test: novel toy with a wide range of affordances

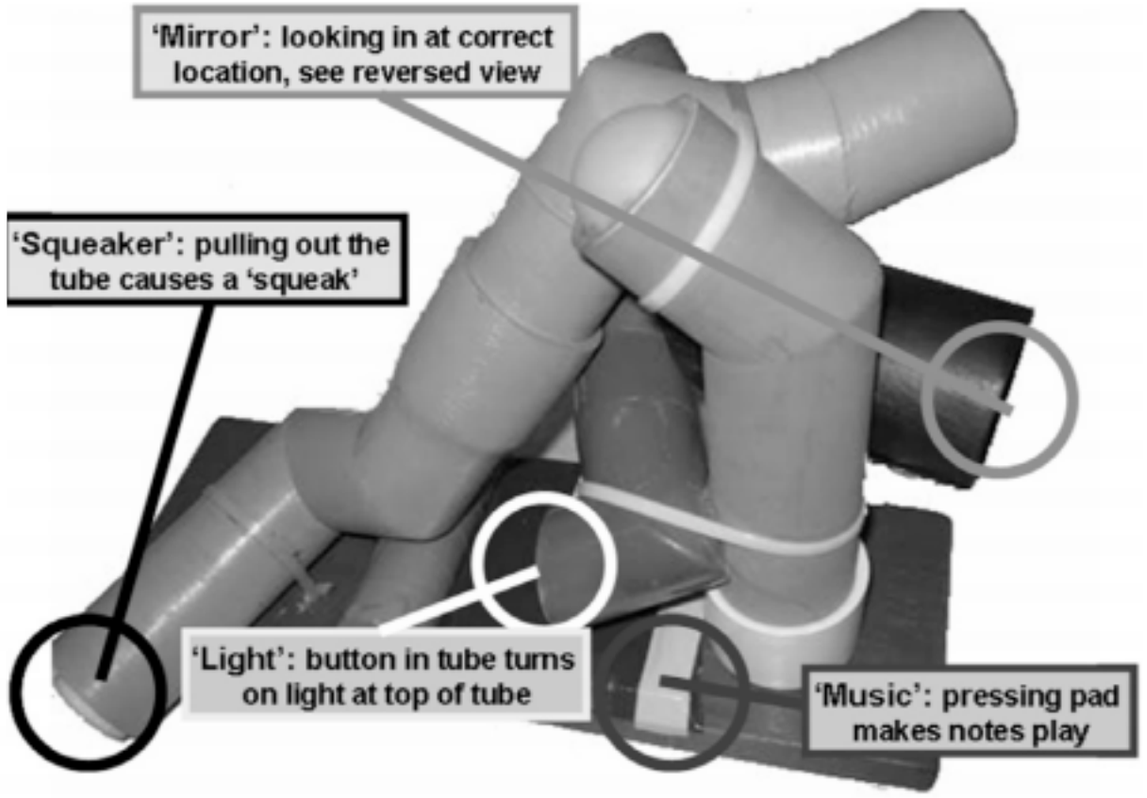

Fig. 1. Novel toy stimuli used in Experiments 1 and 2.

When do children explore a novel toy more?

Baseline

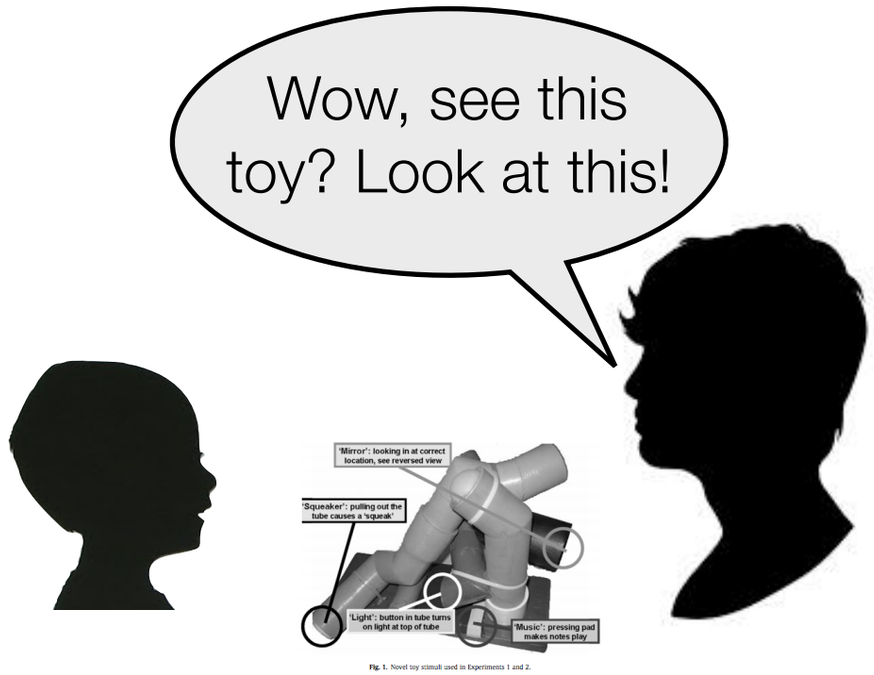

Fig. 1. Novel toy stimuli used in Experiments 1 and 2.

(adult looks at it for 2 sec total)

Pedagogical

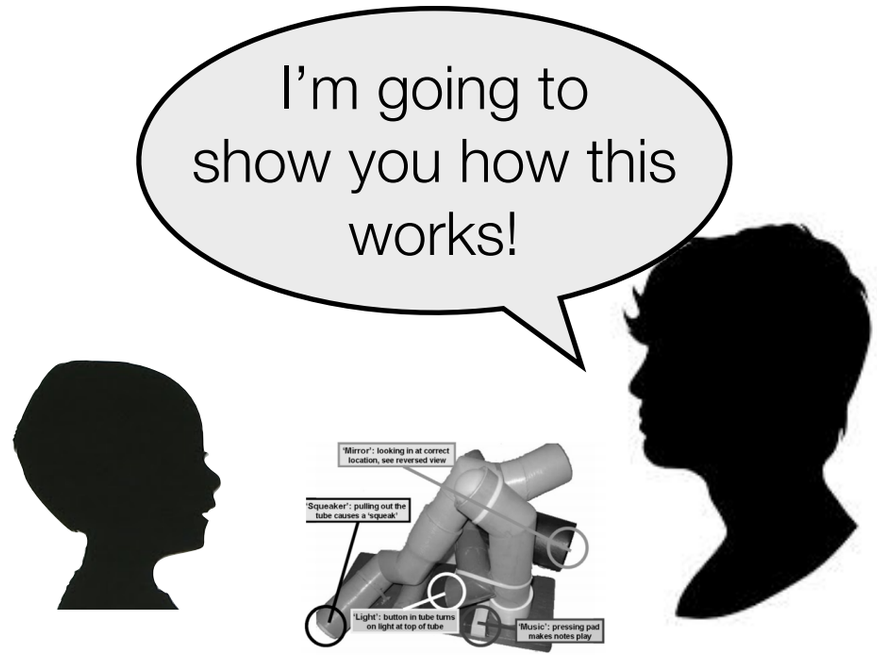

Fig. 1. Novel toy stimuli used in Experiments 1 and 2.

(adult plays with 2 tubes for 2 sec total)

Naive

Fig. 1. Novel toy stimuli used in Experiments 1 and 2.

(adult plays with 2 tubes, as if by accident, for 2 sec total)

Children explore more when not taught explicitly

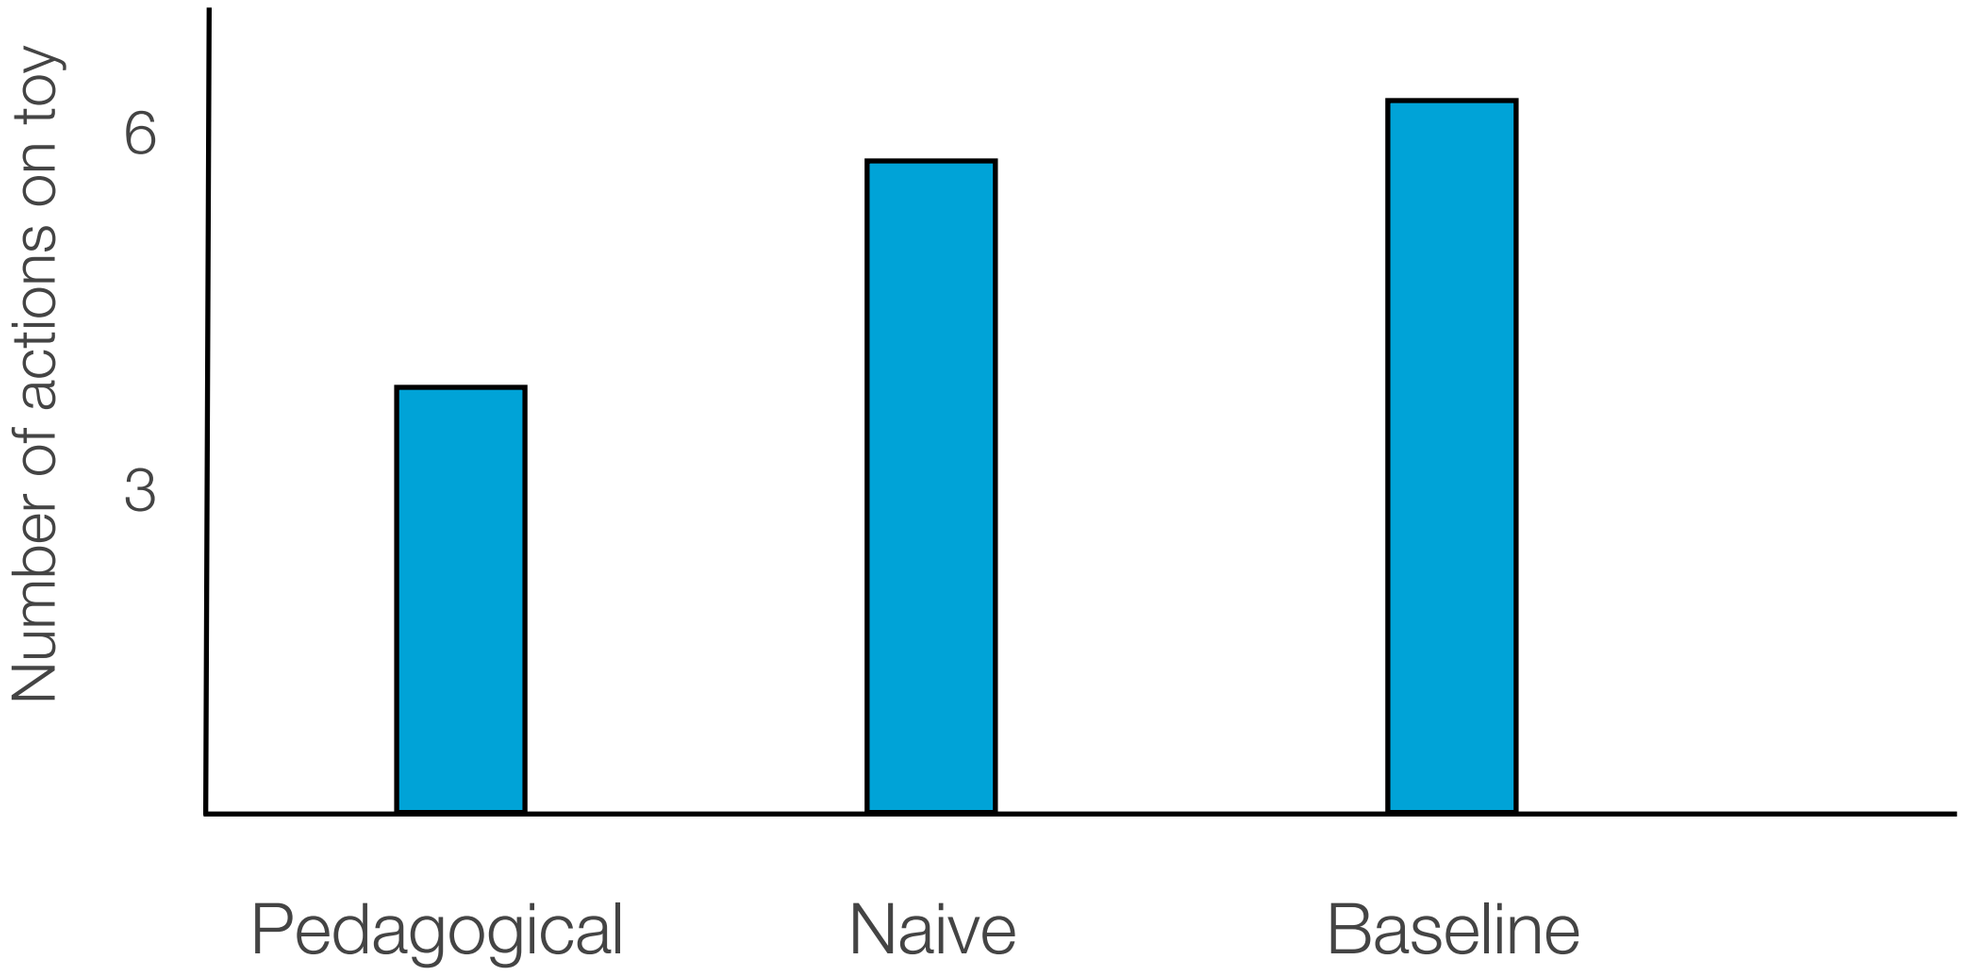

There is also a significant difference in the total time they play with the toy

Summary

- There is a difference between strong sampling and pedagogical sampling: if a teacher provides data helpfully, one can learn even faster and more effectively

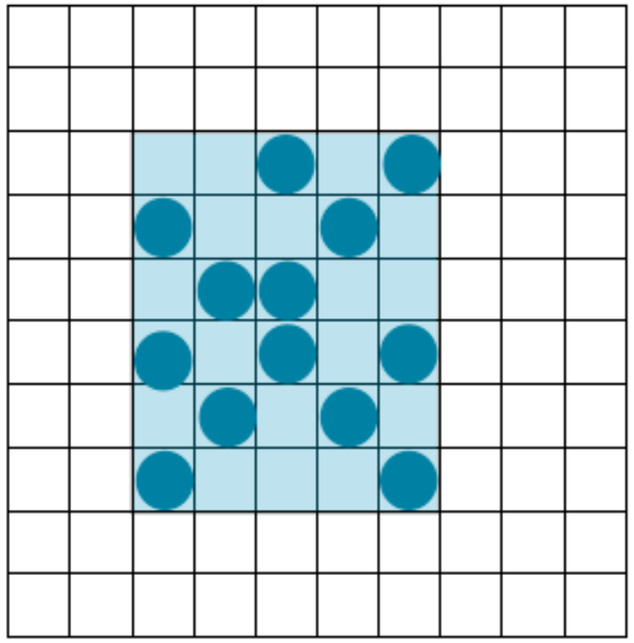

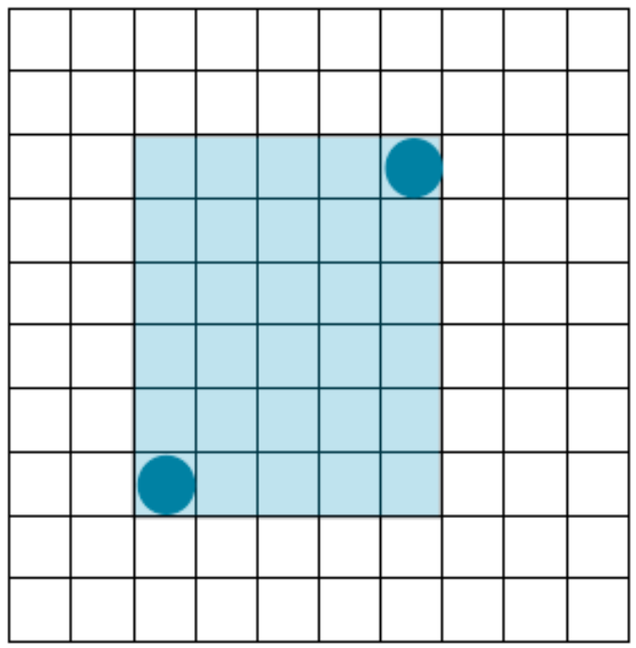

Summary

- There is a difference between strong sampling and pedagogical sampling: if a teacher provides data helpfully, one can learn even faster and more effectively
- At least in simple situations, people are able to provide helpful data and learn appropriately from that data, in line with the pedagogical sampling model

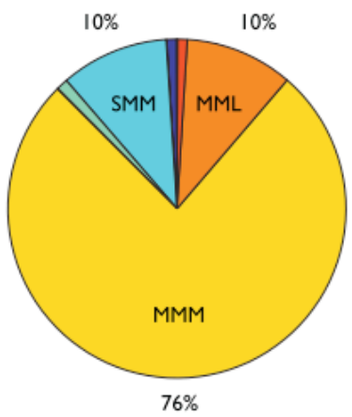

(a) Strong sampling predictions

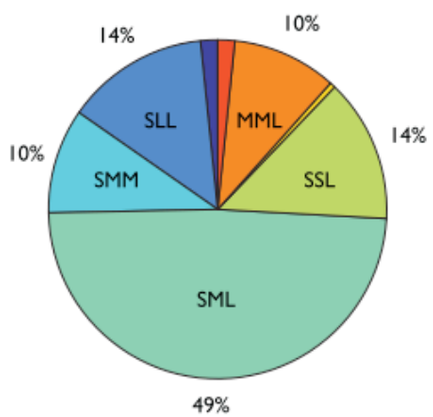

(b) Pedagogical model predictions

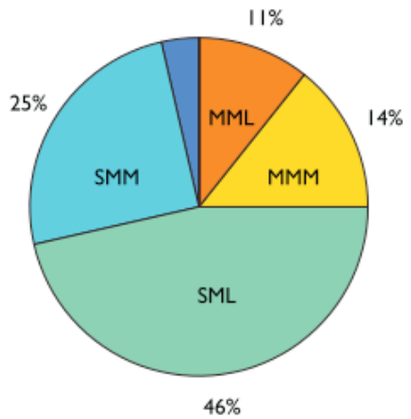

(c) Experiment results

Summary

- There is a difference between strong sampling and pedagogical sampling: if a teacher provides data helpfully, one can learn even faster and more effectively
- At least in simple situations, people are able to provide helpful data and learn appropriately from that data, in line with the pedagogical sampling model
- However, pedagogical learning can be a double-edged sword: it also encourages people to explore less on their own

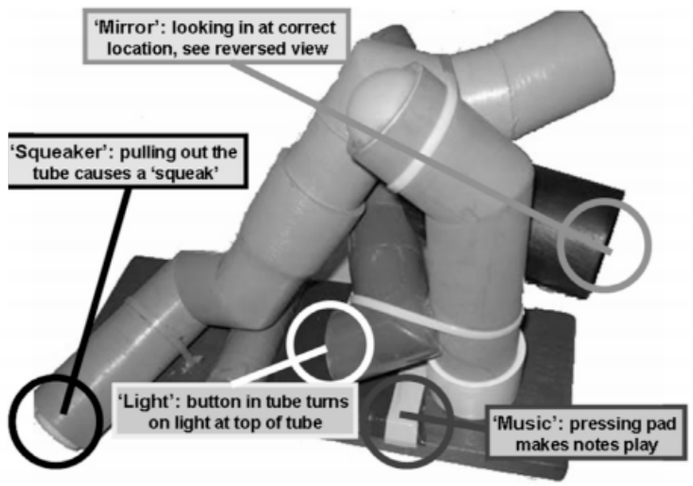

Fig. 1. Novel toy stimuli used in Experiments 1 and 2.

Additional references (not required)

- Bonawitz, E., Shafto, P., Gweon, H., Goodman, N., Spelke, E., & Schulz, L. (2009). The double-edged sword of pedagogy: Instruction limts spontaneous exploration and discovery. *Cognition 120:* 322-330.
- Shafto, P., Goodman, N., & Griffiths, T. (2014). A rational account of pedagogical reasoning: Teaching by, and learning from, examples. *Cognitive Psychology 71*: 55-89