Day 2: Online experiments

Amy Perfors

DAY 2: EXPERIMENTS

Tentative plan

1. Experiment logic, motivation, and design
2. R basics for coding: branching, functions, lists
3. Creating a template experiment with jaysire and putting it online
4. Making a more complex experiment

DAY 2: EXPERIMENTS

Tentative plan

1. Experiment logic, motivation, and design
2. R basics for coding: branching, functions, lists
3. Creating a template experiment with jaysire and putting it online
4. Making a more complex experiment

A SAMPLE EXPERIMENT

The scientific problem: how do people generalise from individual category examples?

These are edible…

A SAMPLE EXPERIMENT

The scientific problem: how do people generalise from individual category examples?

Can I eat this…?

A SAMPLE EXPERIMENT

This is well studied, often in a framework called a category induction task

Premise: EAGLES have more than one fovea per eye

Conclusion: HAWKS have more than one fovea per eye

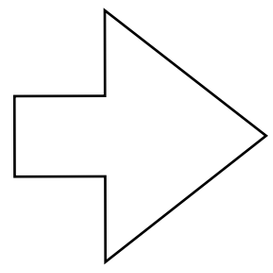

Osherson et al., 1990; Medin et al., 2003

GENERALISING FROM A FEW EXAMPLES

Premise monotonicity: Adding premises to an argument typically strengthens it

EAGLES have more than one fovea per eye

FALCONS have more than one fovea per eye

HAWKS have more than one fovea per eye

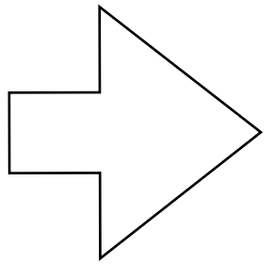

GENERALISING FROM A FEW EXAMPLES

Premise monotonicity: Adding premises to an argument typically strengthens it

more likely that hawks have multiple fovaea

GENERALISING FROM A FEW EXAMPLES

Premise non-monotonicity: Occurs, but more rarely (when adding premises to an argument weakens it)

less likely for buffalo to have the property

GENERALISING FROM A FEW EXAMPLES

Premise non-monotonicity

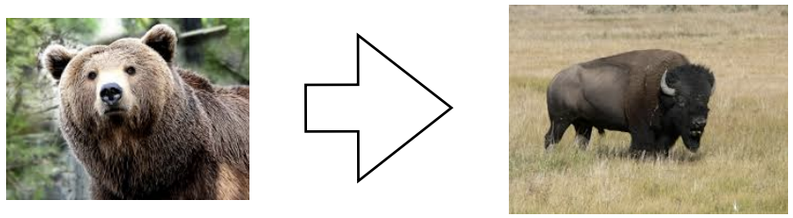

We can account for this by assuming that people are thinking about how the premises were generated (or sampled) and figuring out the implications of that

A MODEL OF CATEGORY-BASED INDUCTION

The world consists of a set of things which may or may not have some property *P*

related to Heit, 1998; Sanjana & Tenenbaum, 2003

A MODEL OF CATEGORY-BASED INDUCTION

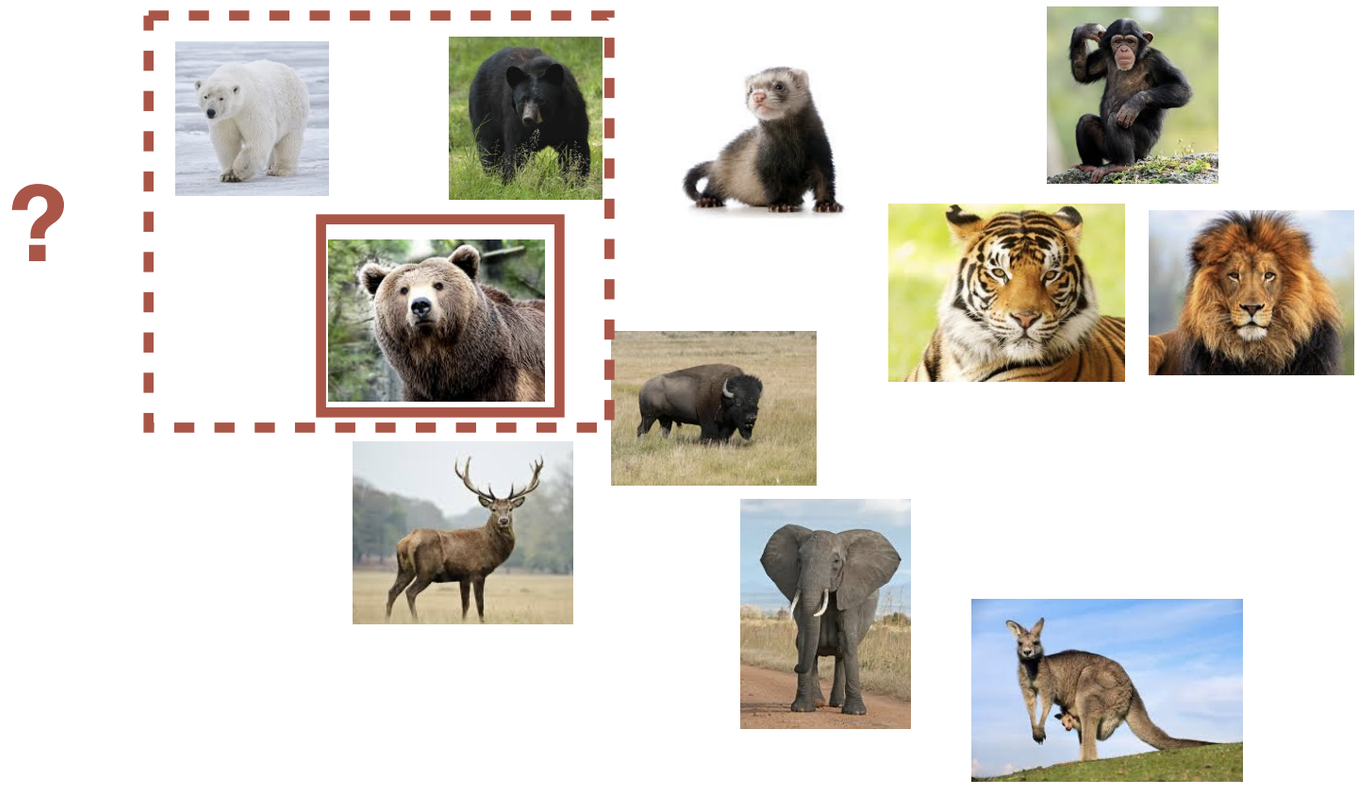

Each hypothesis *h* captures how far a property should be extended

A MODEL OF CATEGORY-BASED INDUCTION

?

Each hypothesis *h* captures how far a property should be extended

A MODEL OF CATEGORY-BASED INDUCTION

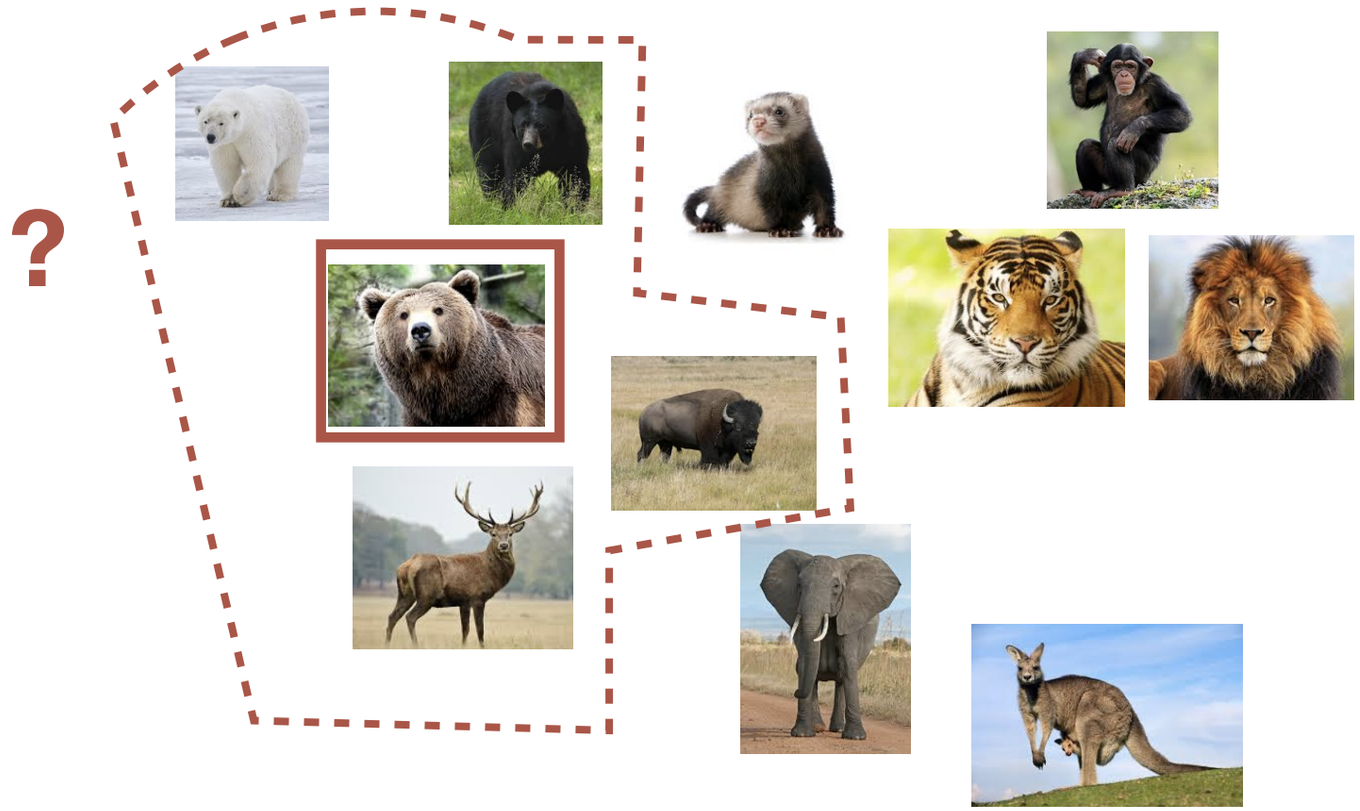

Each hypothesis *h* captures how far a property should be extended

A MODEL OF CATEGORY-BASED INDUCTION

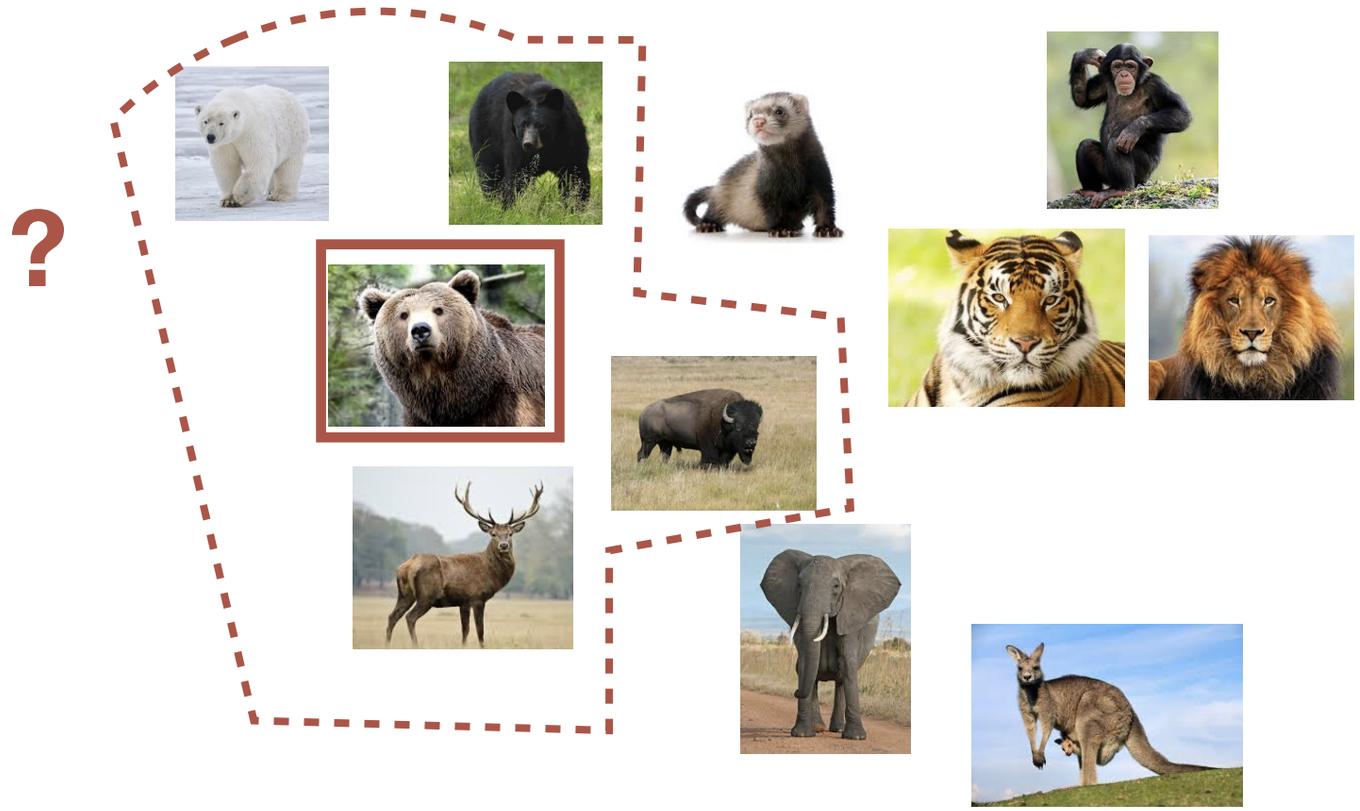

Belief in *h* after having seen data *x* is given by Bayes' Rule

$$P(h \mid x) = \frac{P(x \mid h)P(h)}{\sum_{h'} P(x \mid h')P(h')}.$$

… but how far to generalise depends on the assumptions about how the data were generated

A MODEL OF CATEGORY-BASED INDUCTION

Strong sampling: Picking instances from the concept (having *P*), as one would in order to communicate about it

$$P(x \mid h) = \begin{cases} \frac{1}{|h|} & \text{if } x \in h \\ 0 & \text{otherwise} \end{cases}$$

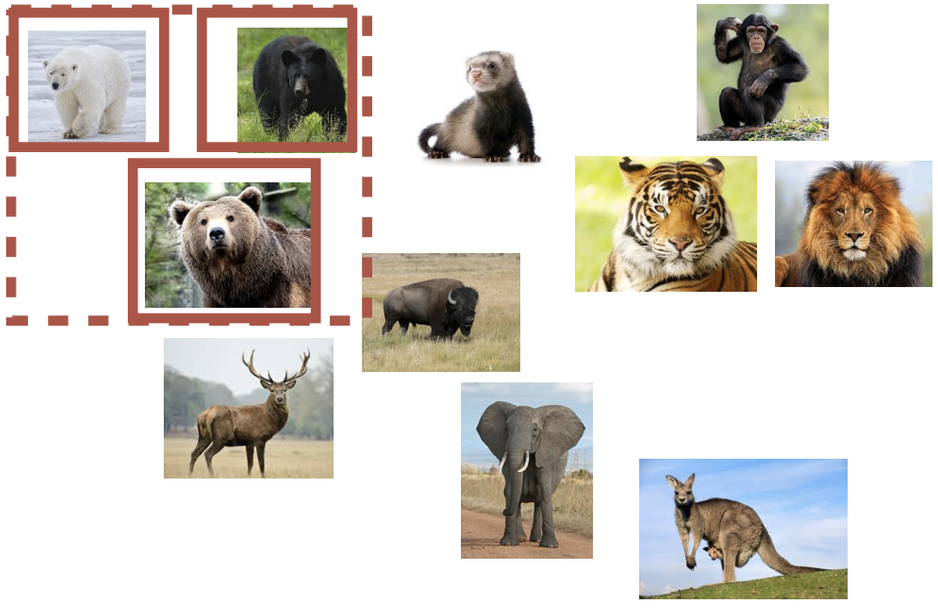

Licenses non-monotonic reasoning: otherwise, poor communication

… but how far to generalise depends on the assumptions about how the data were generated

A MODEL OF CATEGORY-BASED INDUCTION

Weak sampling: Picking instances from the world at random, and then labeling them as having property *P* or not

$$P(x \mid h) \propto \begin{cases} 1 & \text{if } x \in h \\ 0 & \text{otherwise} \end{cases}$$

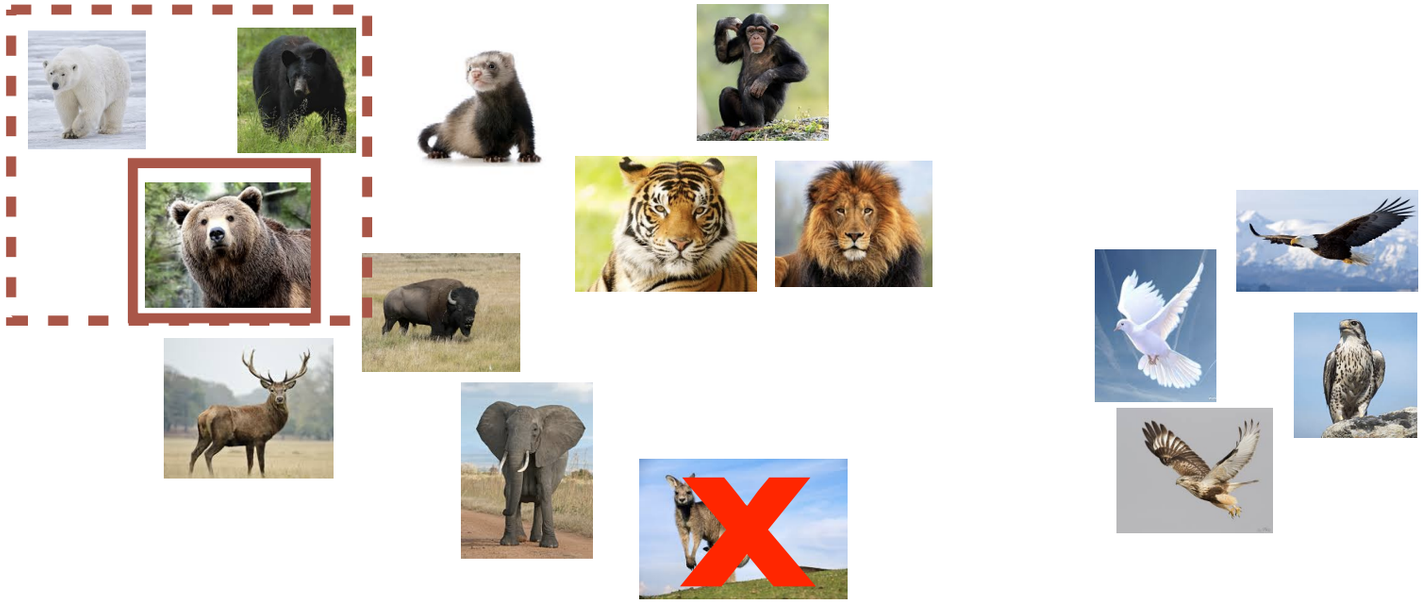

... but how far to generalise depends on the assumptions about how the data were generated

A MODEL OF CATEGORY-BASED INDUCTION

Weak sampling: Picking instances from the world at random, and then labeling them as having property *P* or not

$$P(x \mid h) \propto \begin{cases} 1 & \text{if } x \in h \\ 0 & \text{otherwise} \end{cases}$$

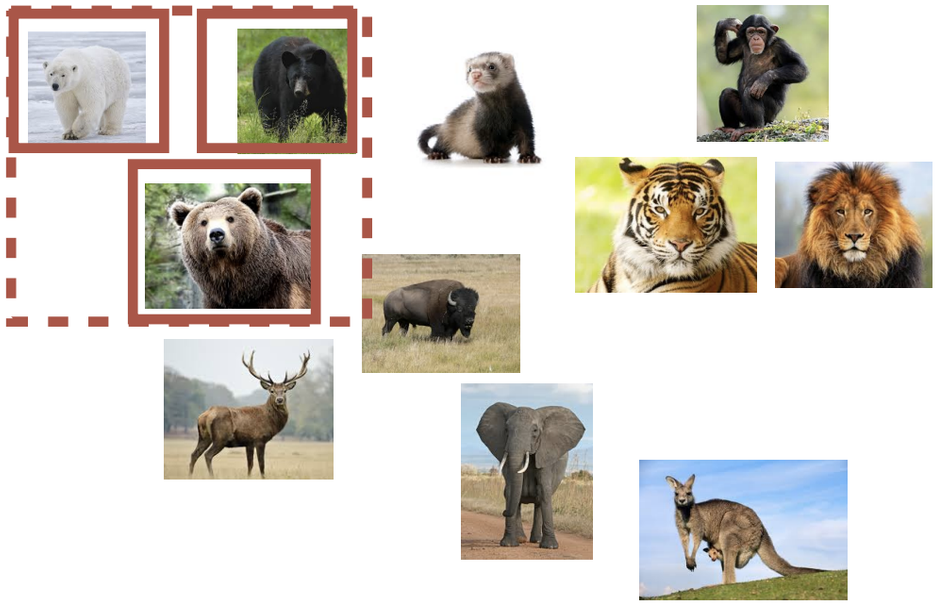

Does not license non-monotonic reasoning: just happened to be that way (i.e., the selection of items is not meaningful)

… but how far to generalise depends on the assumptions about how the data were generated

DIFFERENT SAMPLING ASSUMPTIONS YIELD DIFFERENT PREDICTIONS

TARGET 1

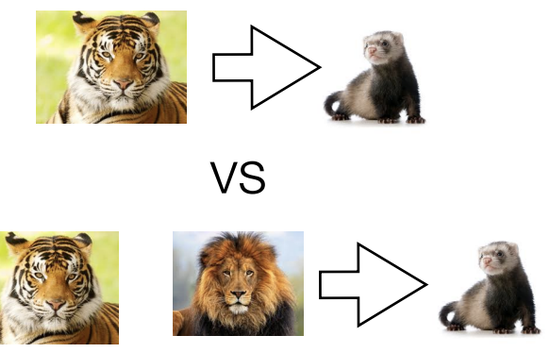

TARGET 2

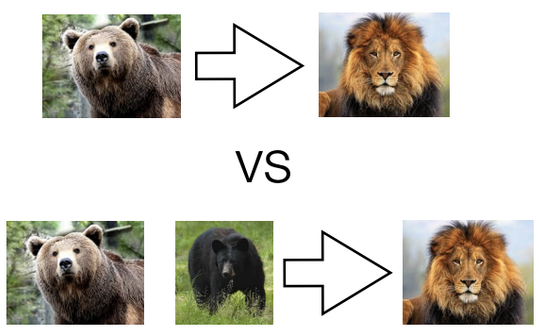

Non-monotonic: Additional argument should make conclusion weaker (if strong sampling, not if weak)

CONTROL

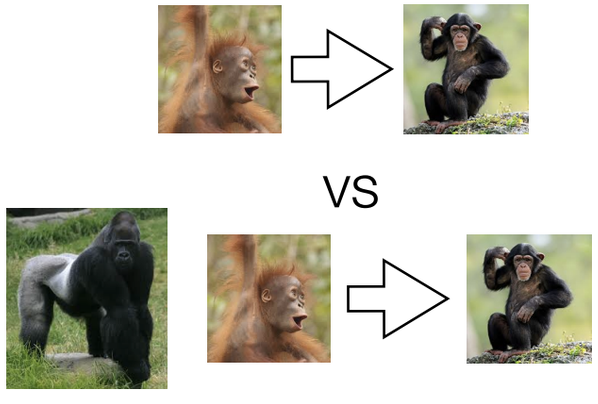

Monotonic: Additional argument should make conclusion stronger (if strong sampling, not if weak)

DIFFERENT SAMPLING ASSUMPTIONS YIELD DIFFERENT PREDICTIONS

TARGET 1

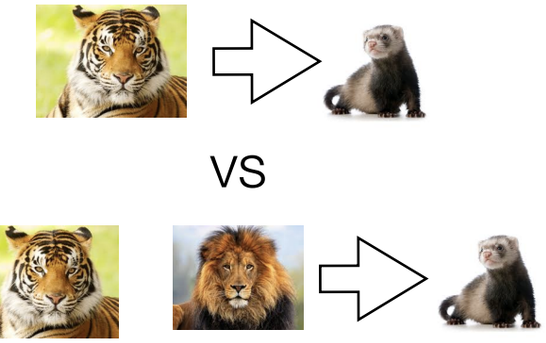

TARGET 2

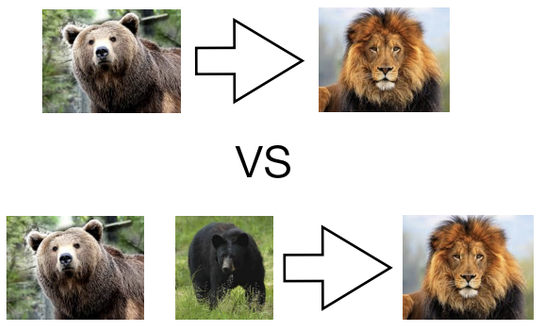

CONTROL

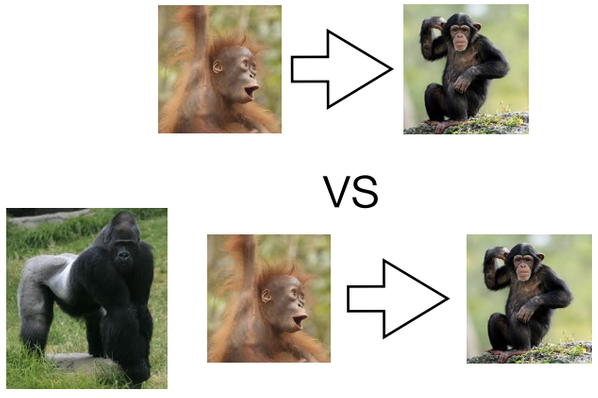

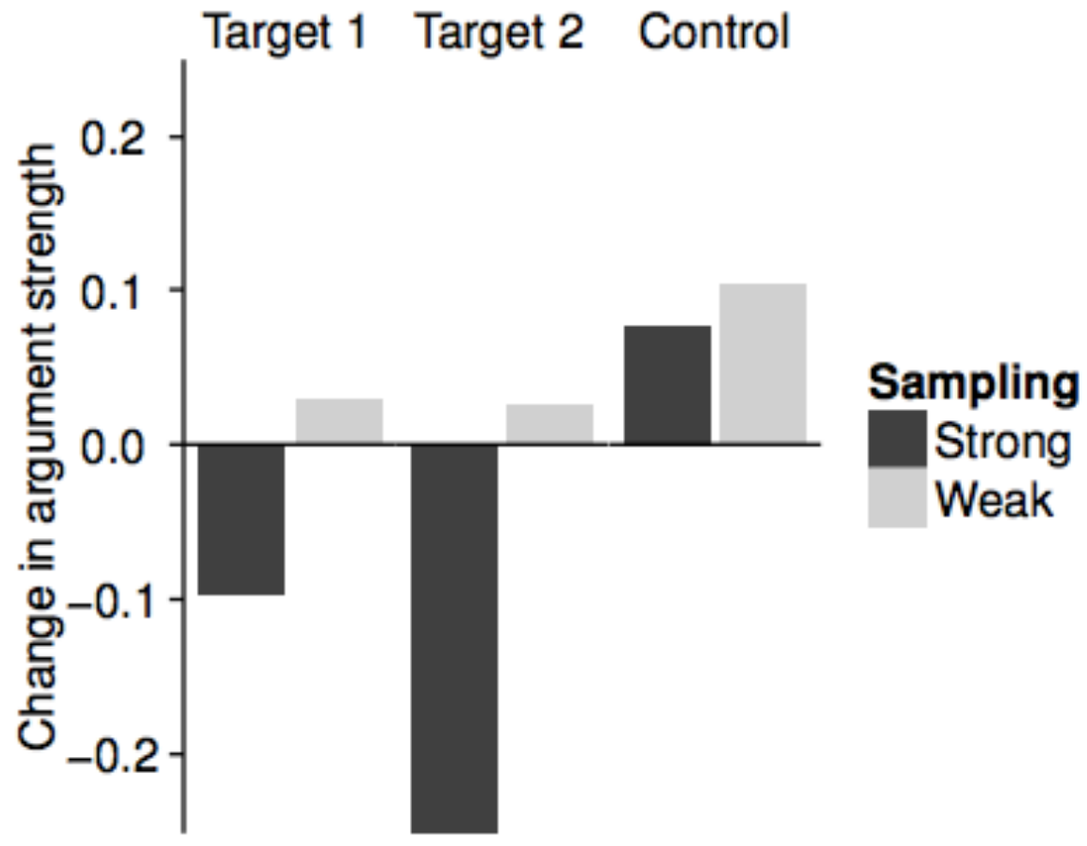

DIFFERENT SAMPLING ASSUMPTIONS YIELD DIFFERENT PREDICTIONS

Do people change their pattern of reasoning based on mainpulating the cover story about how the data were generated (socially, or not)?

COVER STORY MANIPULATION

HELPFUL: People were told that the second fact in each trial was generated by a past player of the game who was trying to be helpful

RANDOM: People "drew" the second fact randomly from a set of cards drawn on the screen, one for each animal

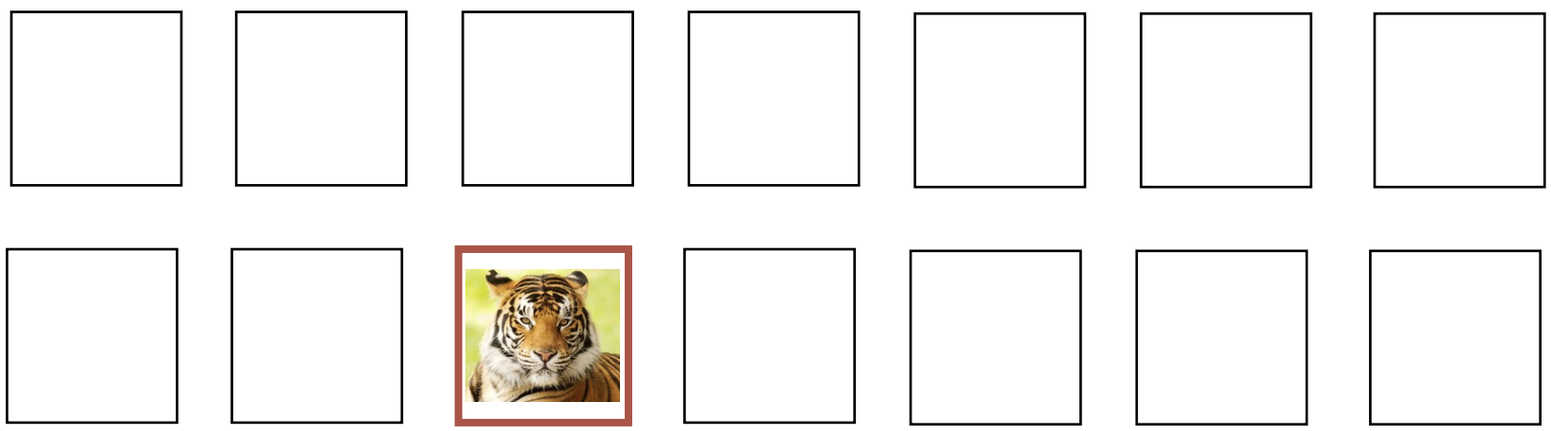

CHANGING THE SOCIAL STORY CHANGES THE PATTERN OF PEOPLE'S REASONING

People

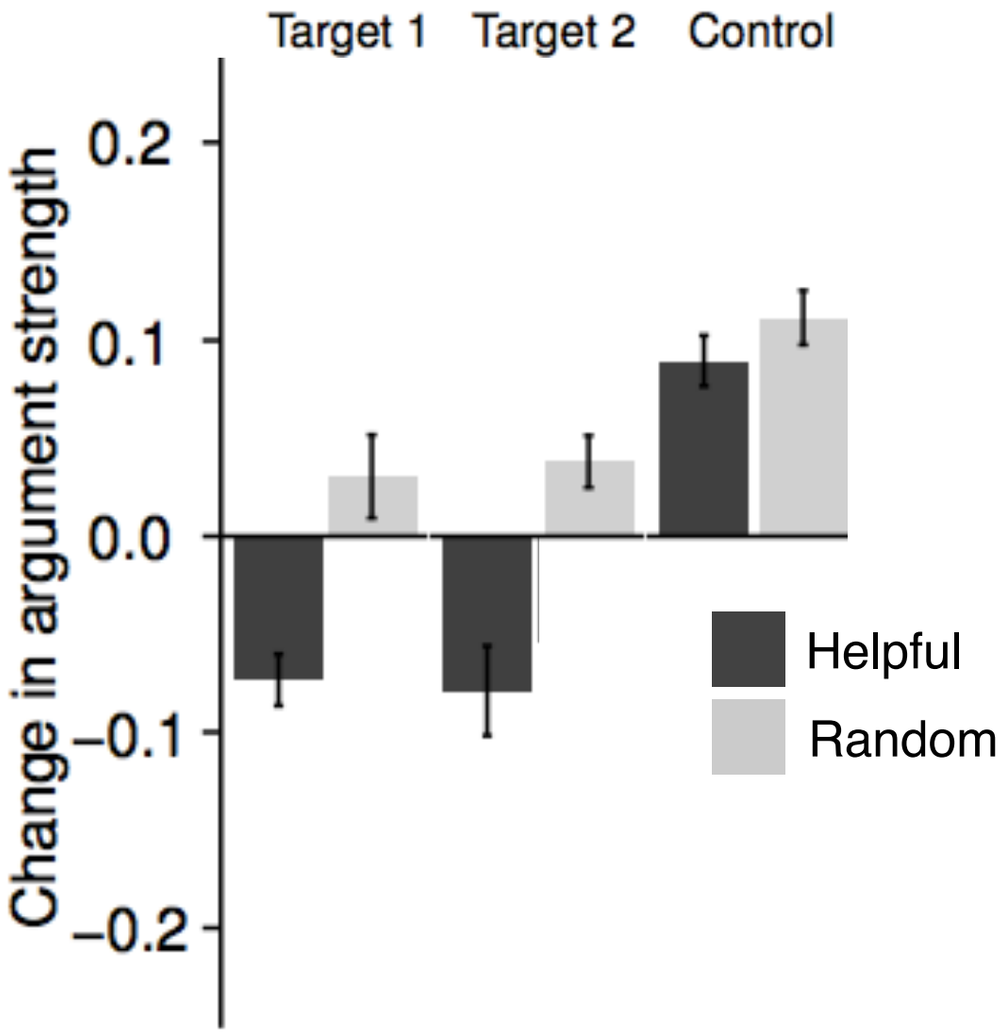

Model

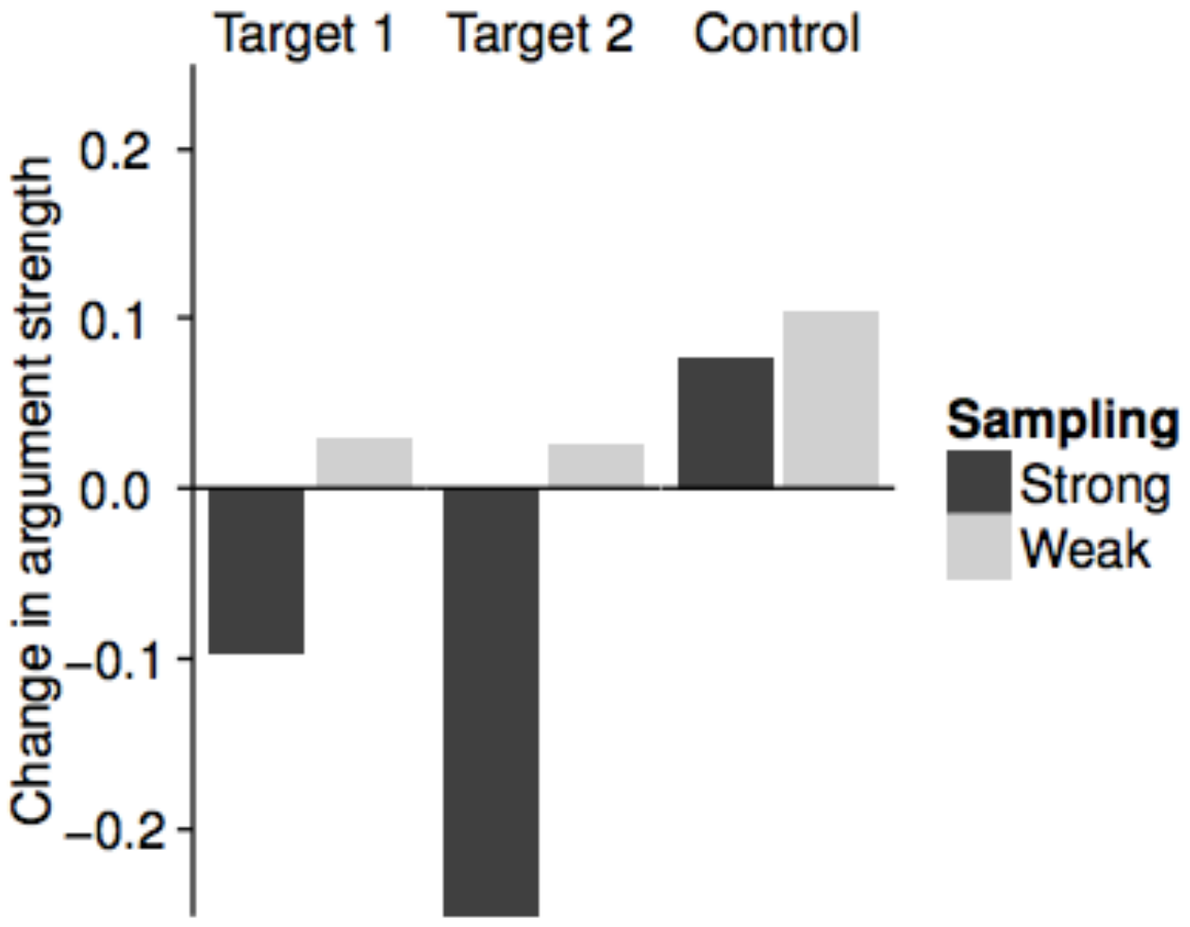

SAMPLING ALSO AFFECTS HOW YOU SHOULD RESPOND TO ADDITIONAL DATAPOINTS

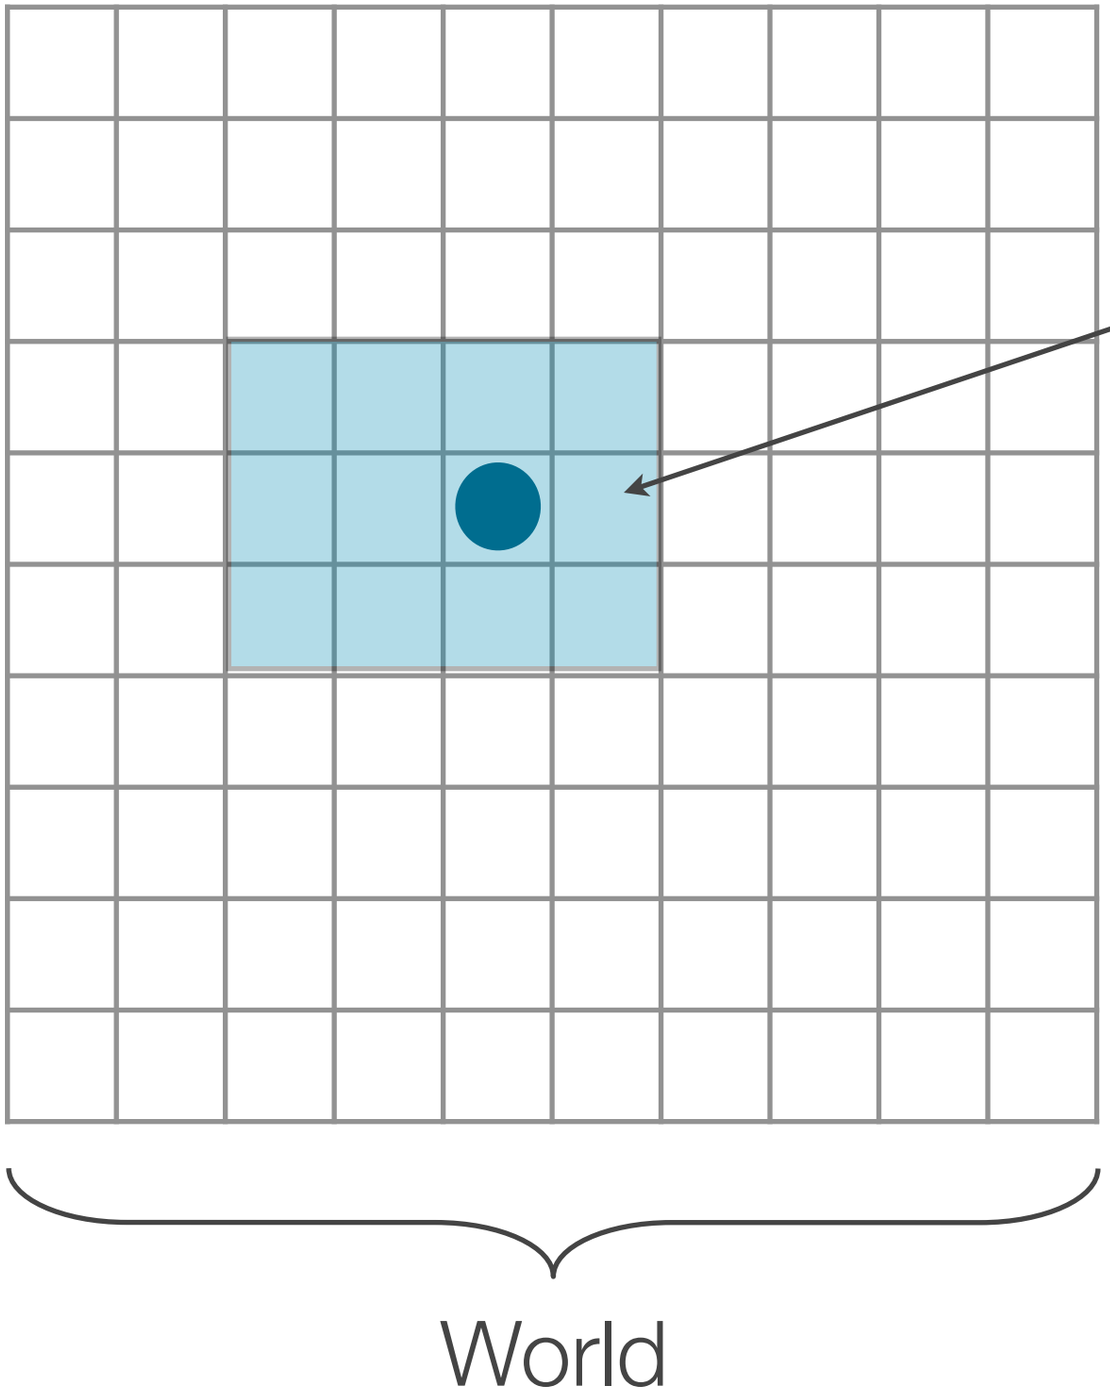

Strong sampling

Hypothesis of size n

$$p(d|h) = 1/n$$
$$= 1/12$$

This is known as the size principle

CONSEQUENCE OF SIZE PRINCIPLE

- It is due to the size principle that additional data points will cause generalisation curves to tighten

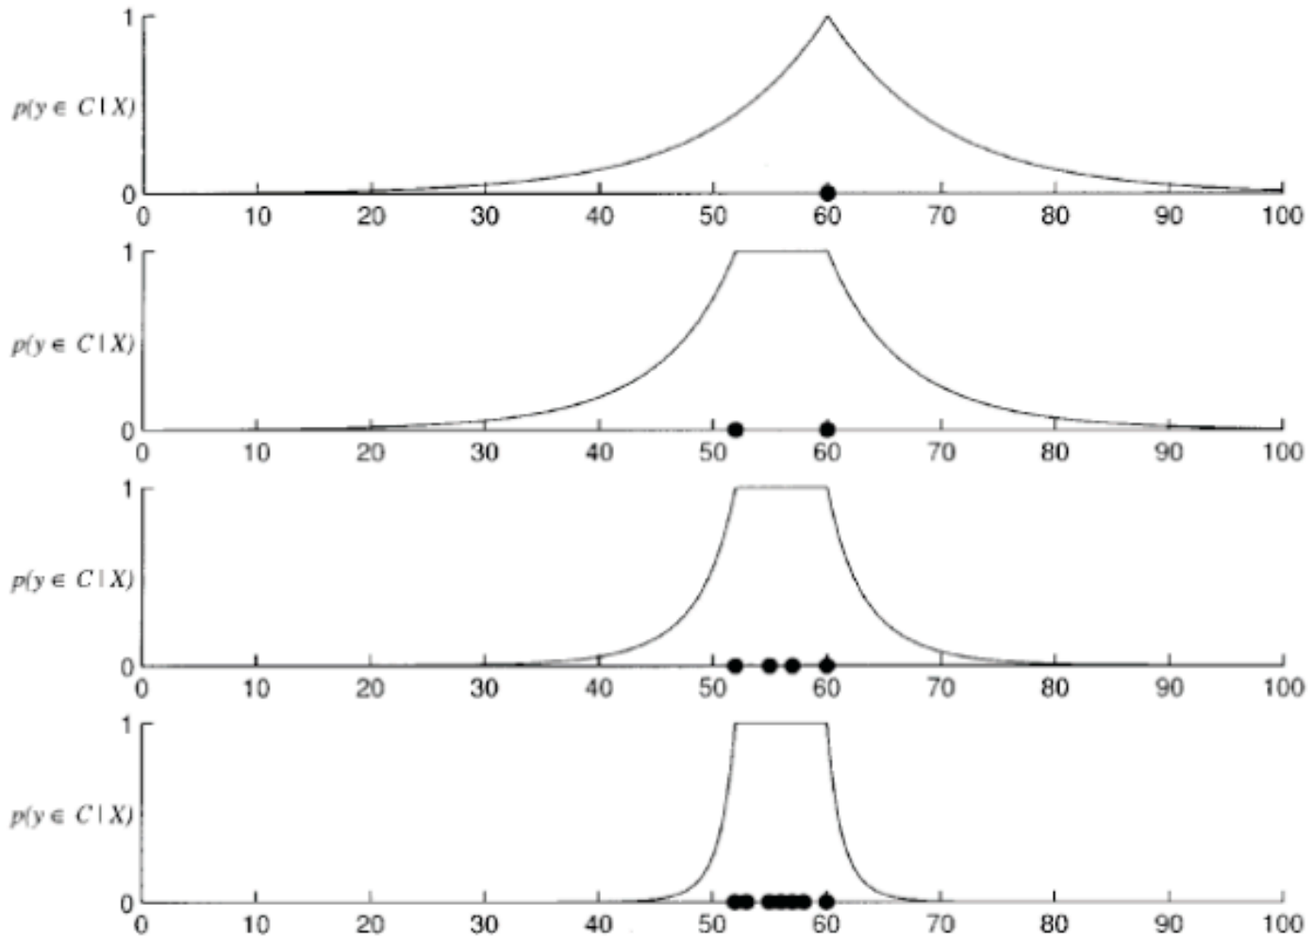

CONSEQUENCE OF SIZE PRINCIPLE

- It is due to the size principle that additional data points will cause generalisation curves to tighten

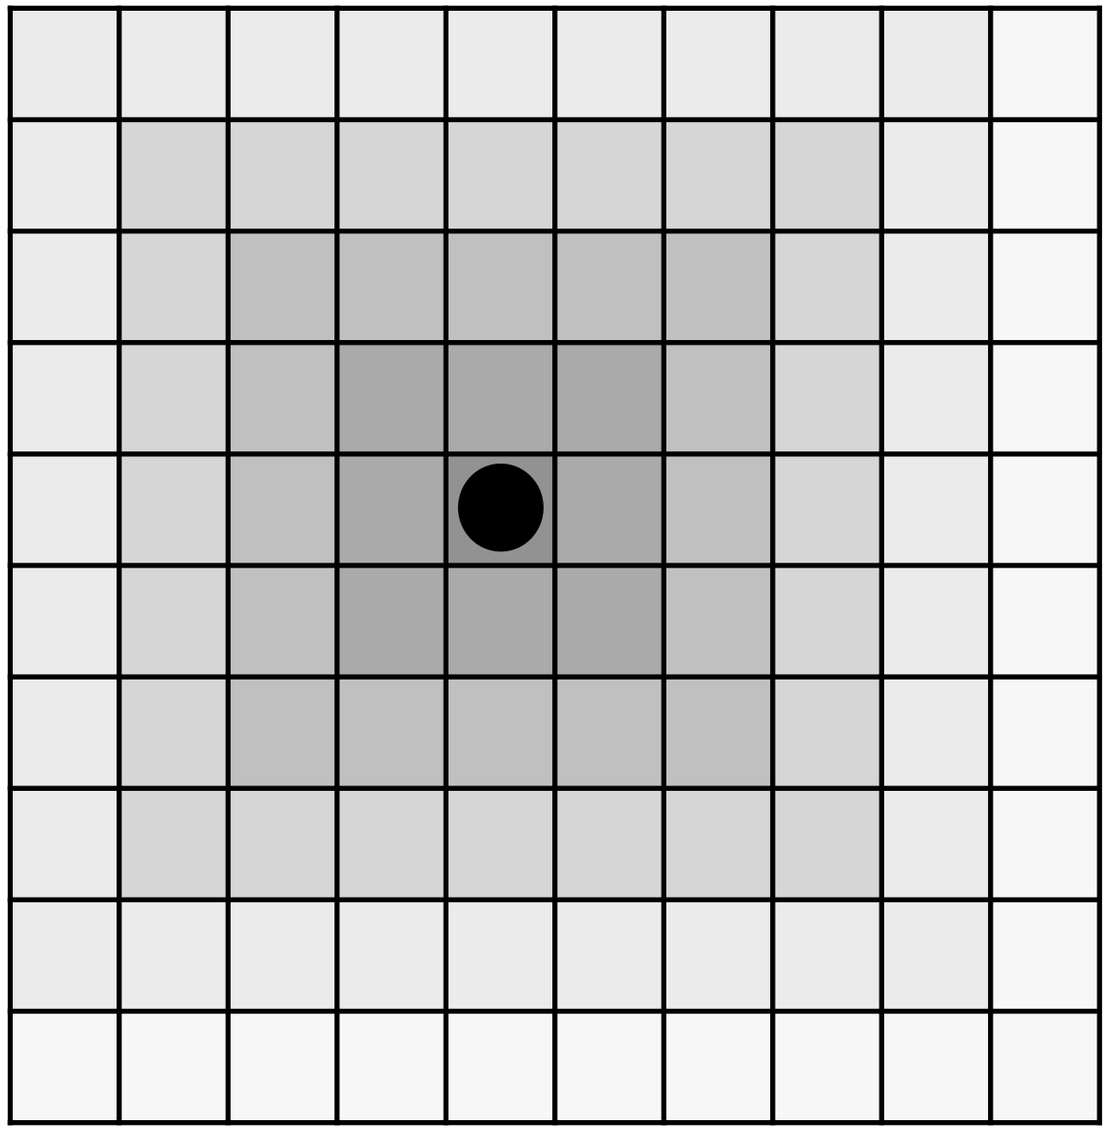

CONSEQUENCE OF SIZE PRINCIPLE

- It is due to the size principle that additional data points will cause generalisation curves to tighten

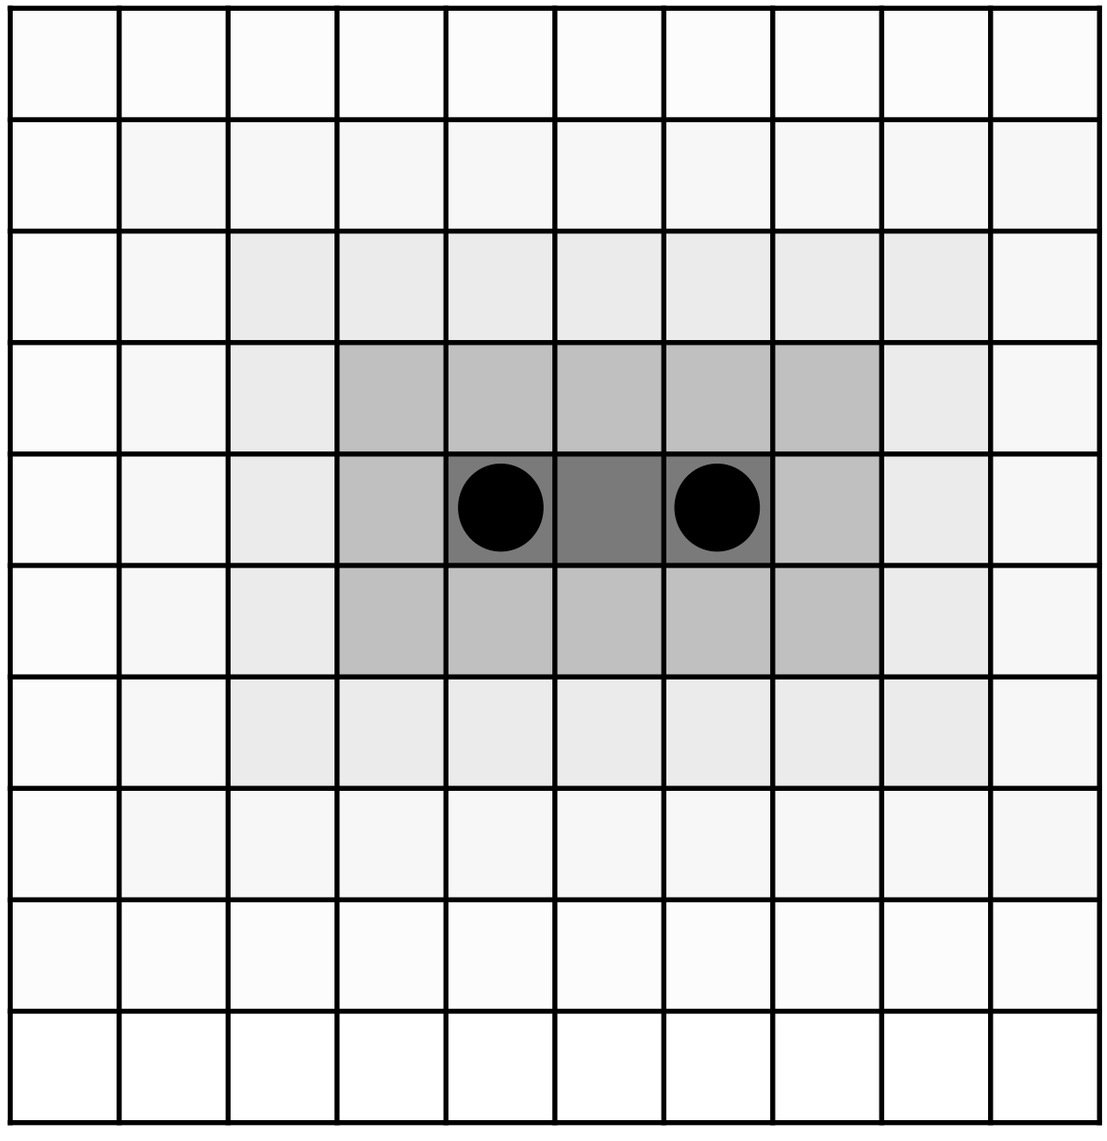

CONSEQUENCE OF SIZE PRINCIPLE

- It is due to the size principle that additional data points will cause generalisation curves to tighten

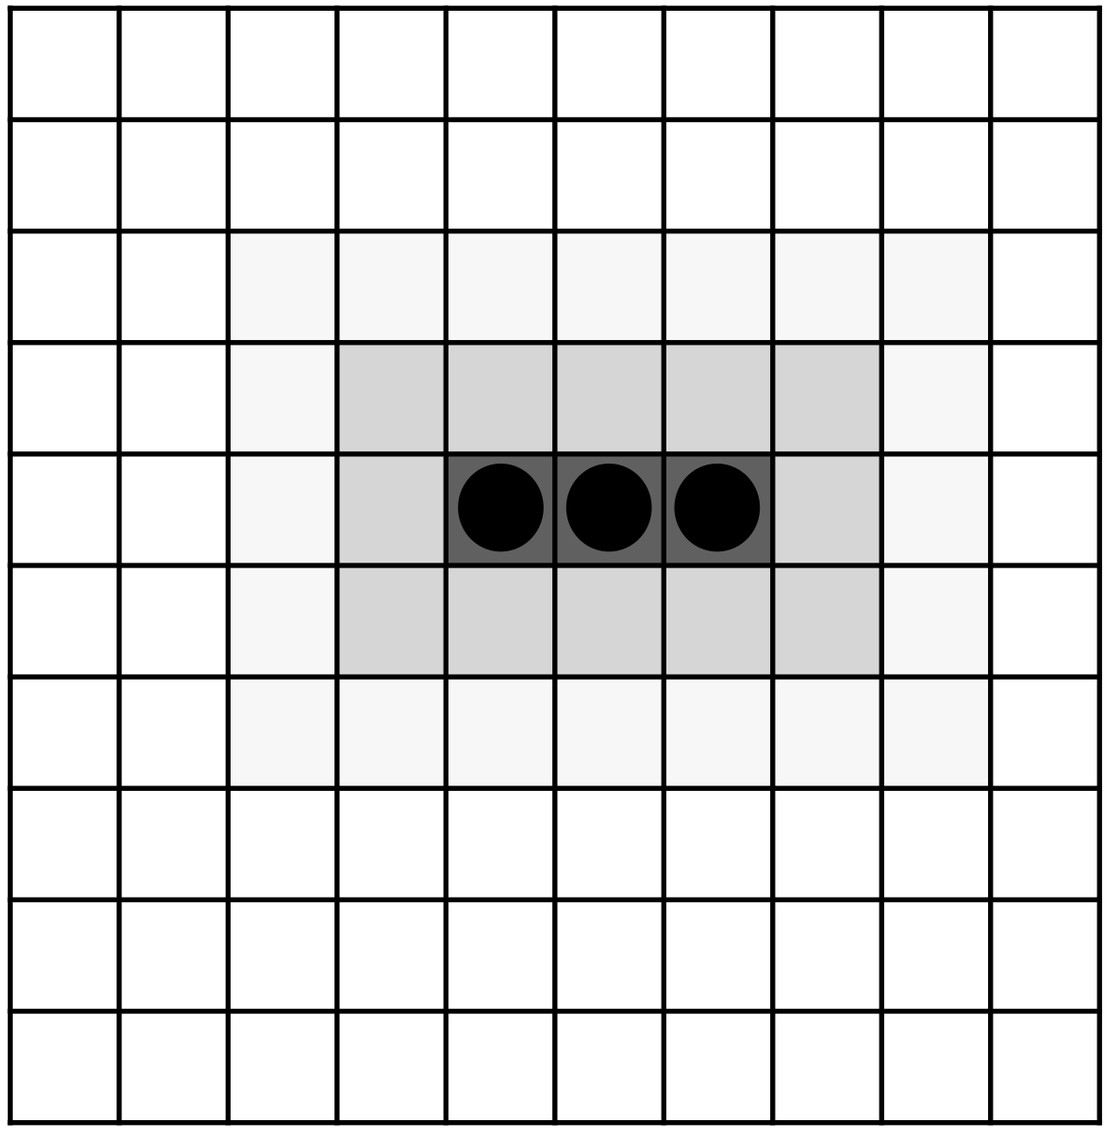

CONSEQUENCE OF SIZE PRINCIPLE

- It is due to the size principle that additional data points will cause generalisation curves to tighten

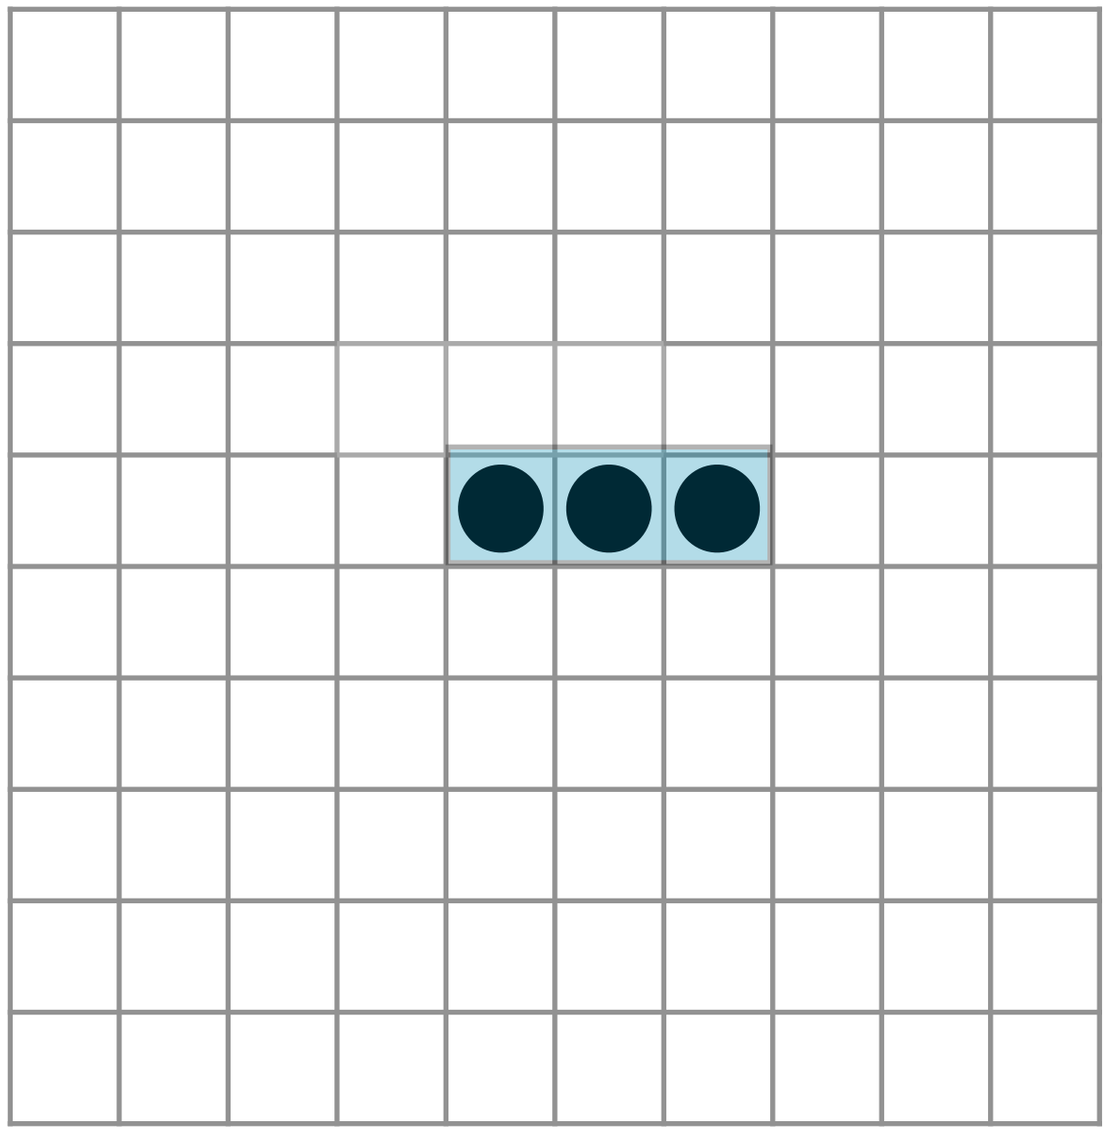

This is because it's quite a suspicious coincidence for these data points to have been generated if the true hypothesis is *not* h

WEAK SAMPLING IS DIFFERENT!

- Weak sampling suggests that data were generated from the world in general, and then only labelled as belonging to the hypothesis (or not)

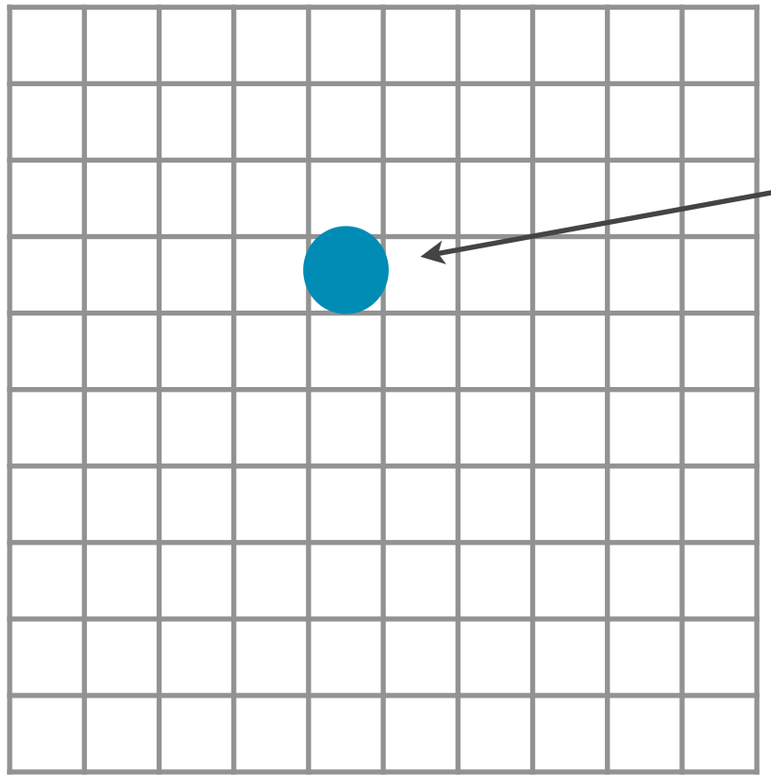

Data sampled from the world at random

Then labelled as in the hypothesis or not

$p(d=\bullet|h) = 1$ *if in the hypothesis*
0 if not

WEAK SAMPLING IS DIFFERENT!

- Weak sampling suggests that data were generated from the world in general, and then only labelled as belonging to the hypothesis (or not)

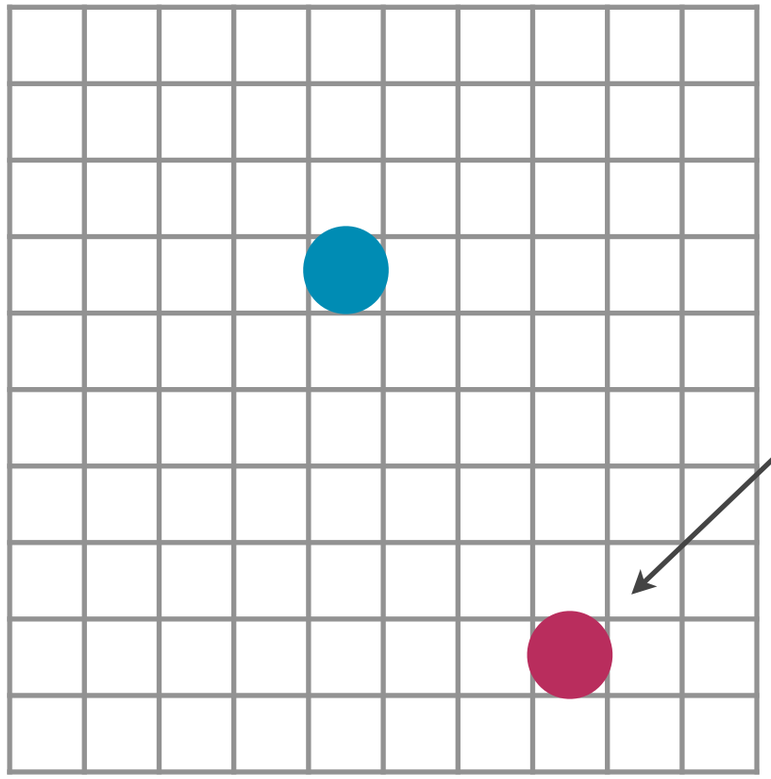

Data sampled from the world at random

Then labelled as in the hypothesis or not

$p(d=\bullet|h) = 1$ *if in the hypothesis*
0 if not

WEAK SAMPLING IS DIFFERENT!

- If data are weakly sampled, the generalisation curves should not tighten -- there is no suspicious coincidence since the data were generated by the *world,* and not from the hypothesis

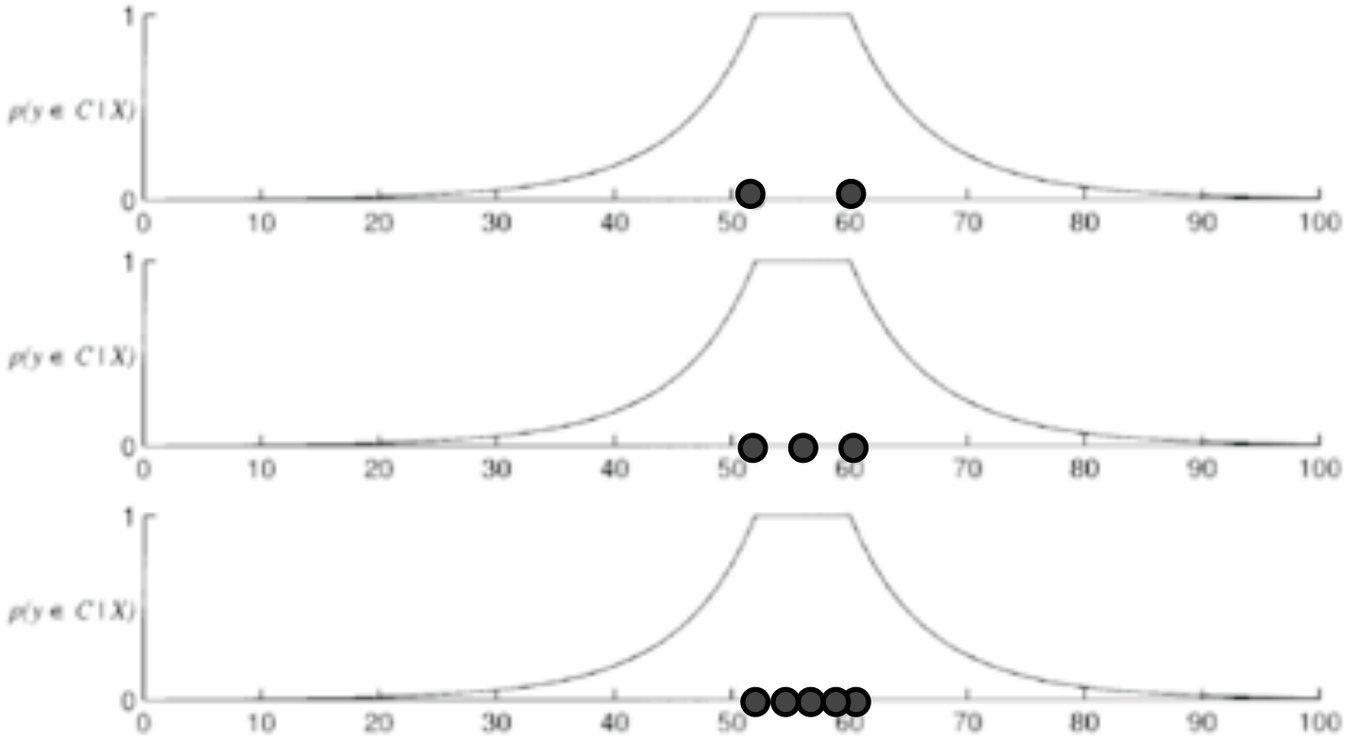

ARE PEOPLE SENSITIVE TO SAMPLING ASSUMPTIONS WHEN REASONING ABOUT ADDITIONAL DATA?

WORD LEARNING

- Many domains have a hierarchical or tree-based conceptual structure

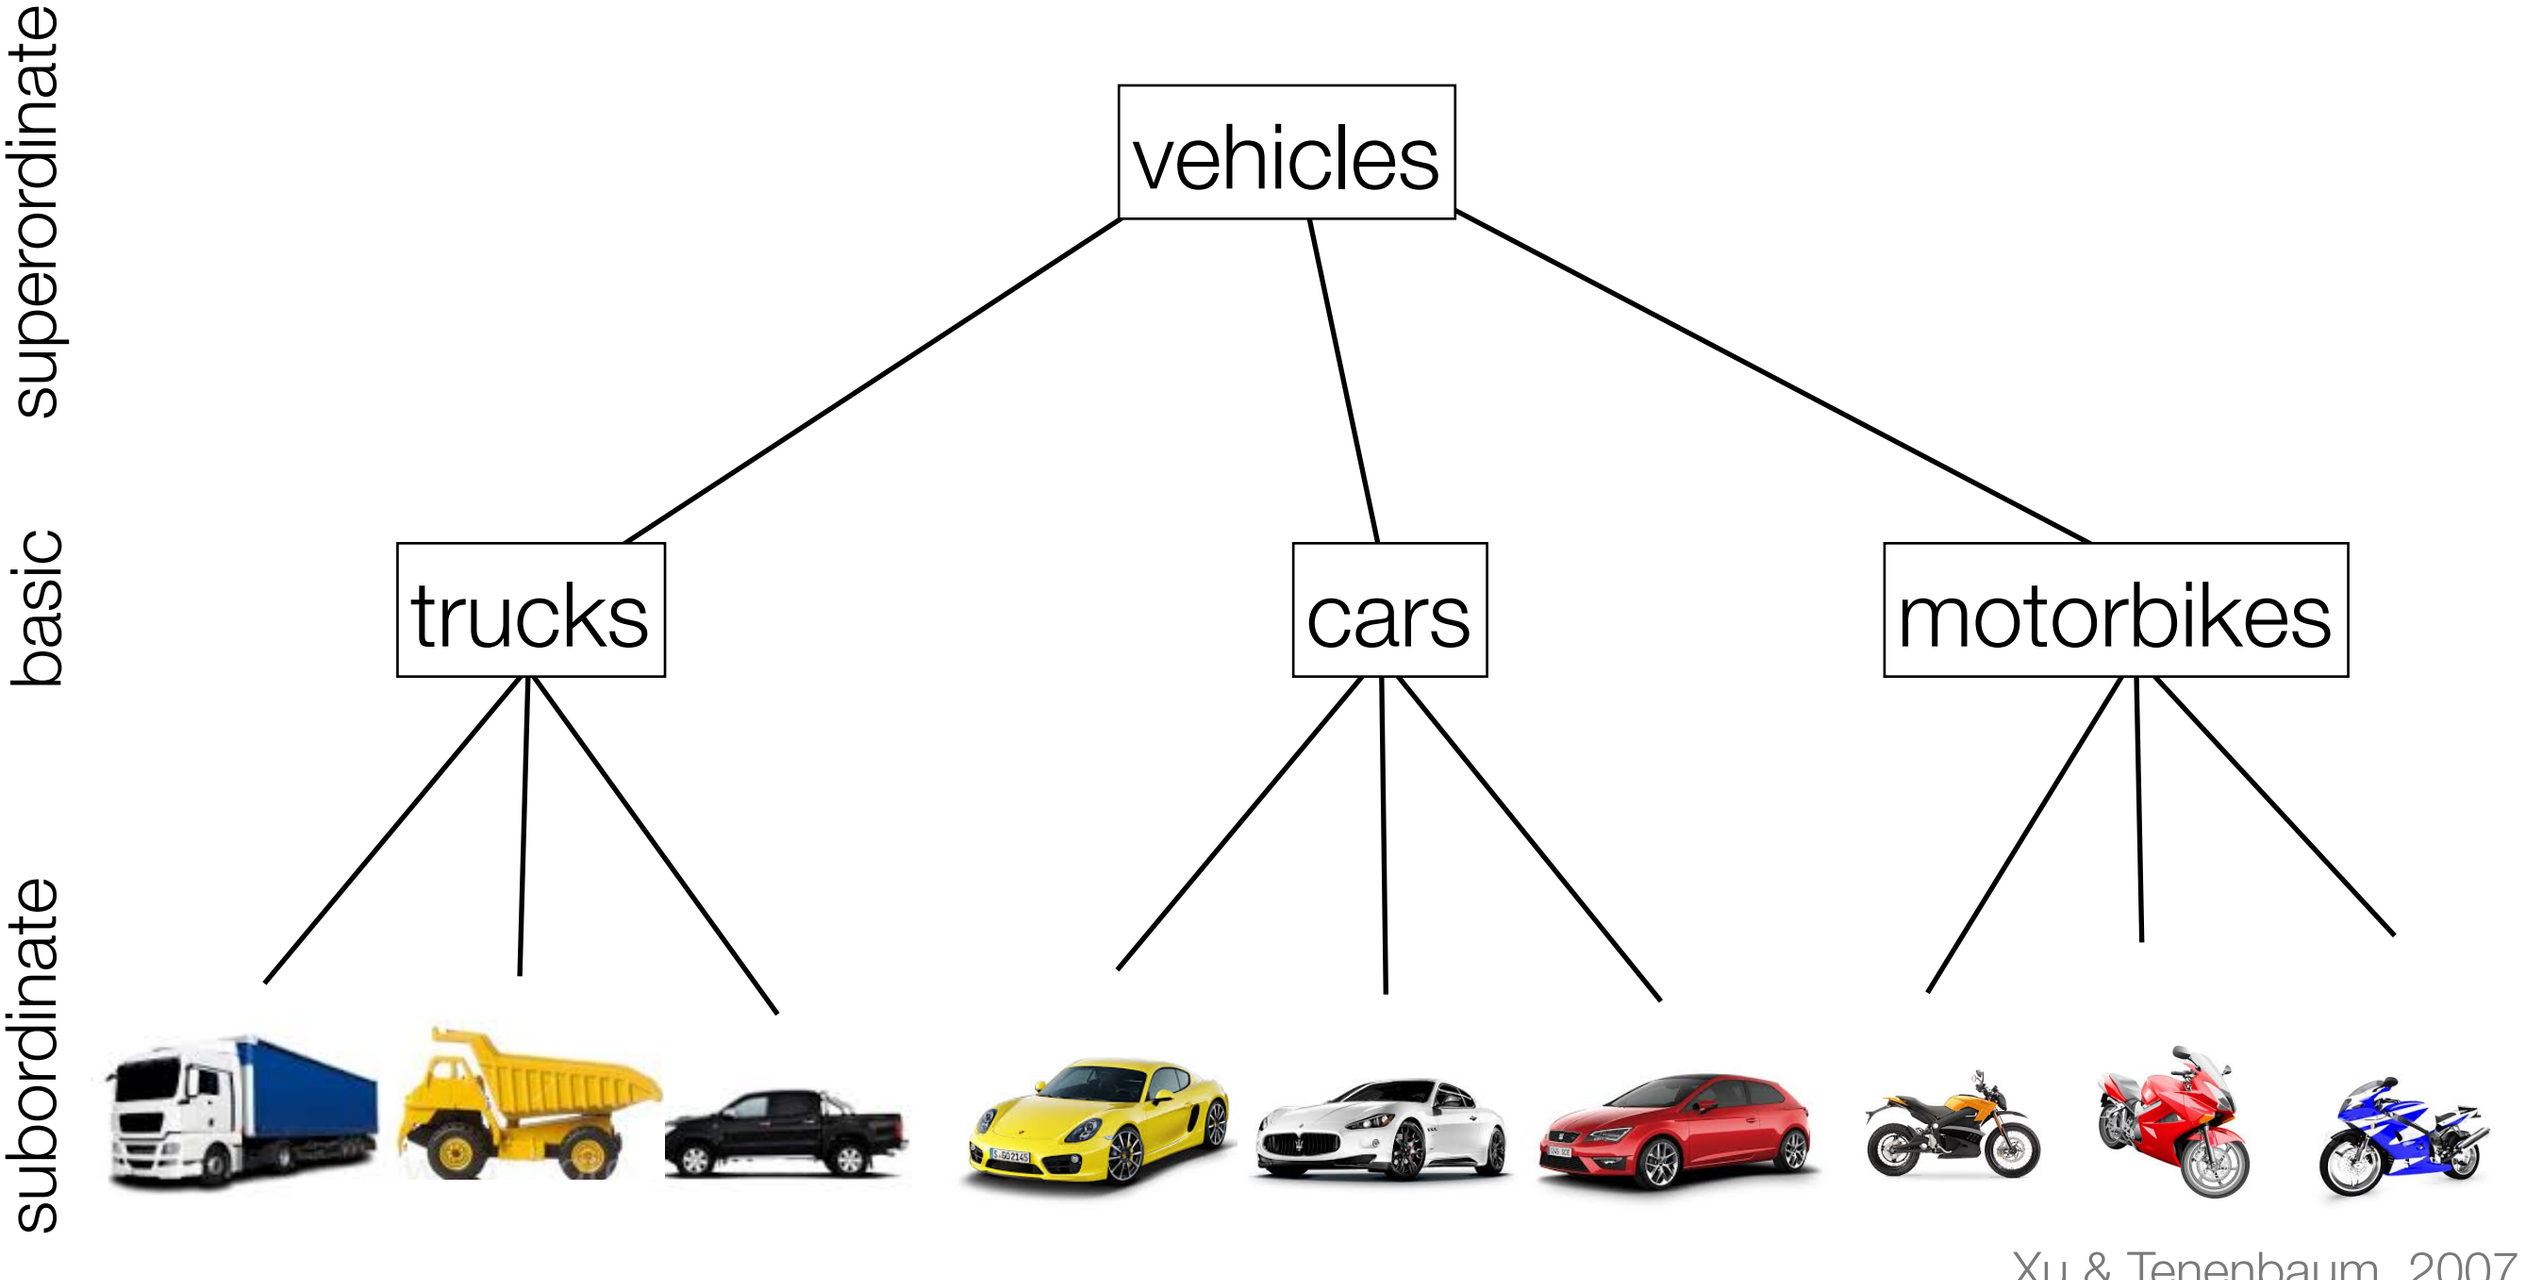

Xu & Tenenbaum, 2007

WORD LEARNING

- Many domains have a hierarchical or tree-based conceptual structure

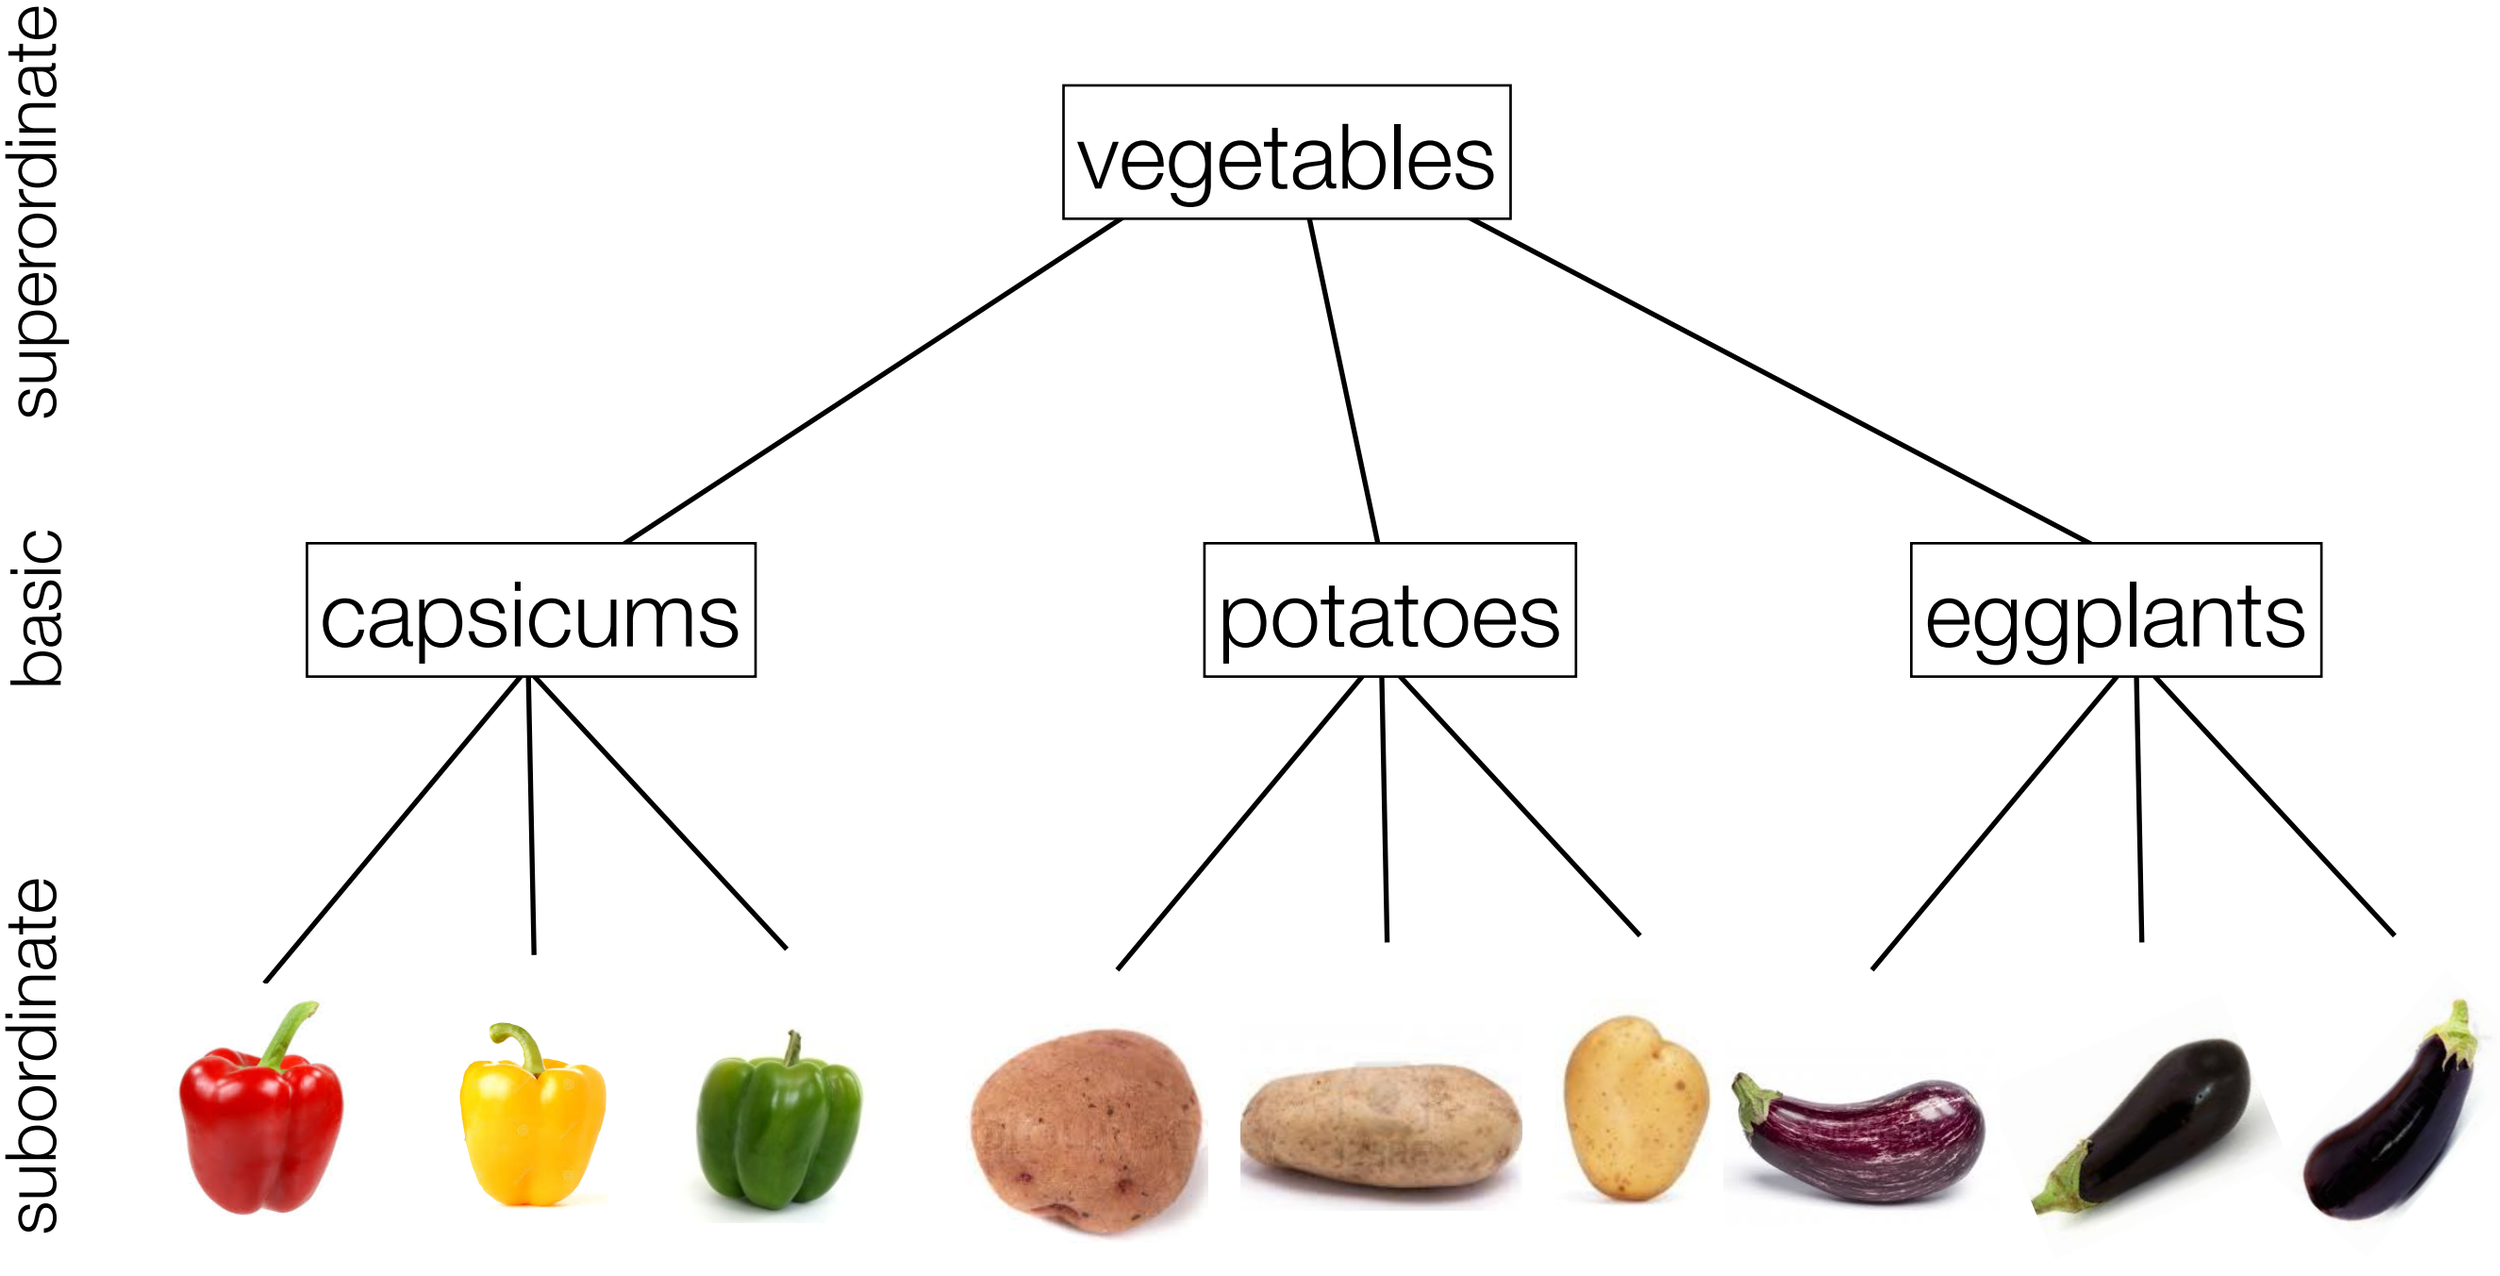

WORD LEARNING

- Many domains have a hierarchical or tree-based conceptual structure

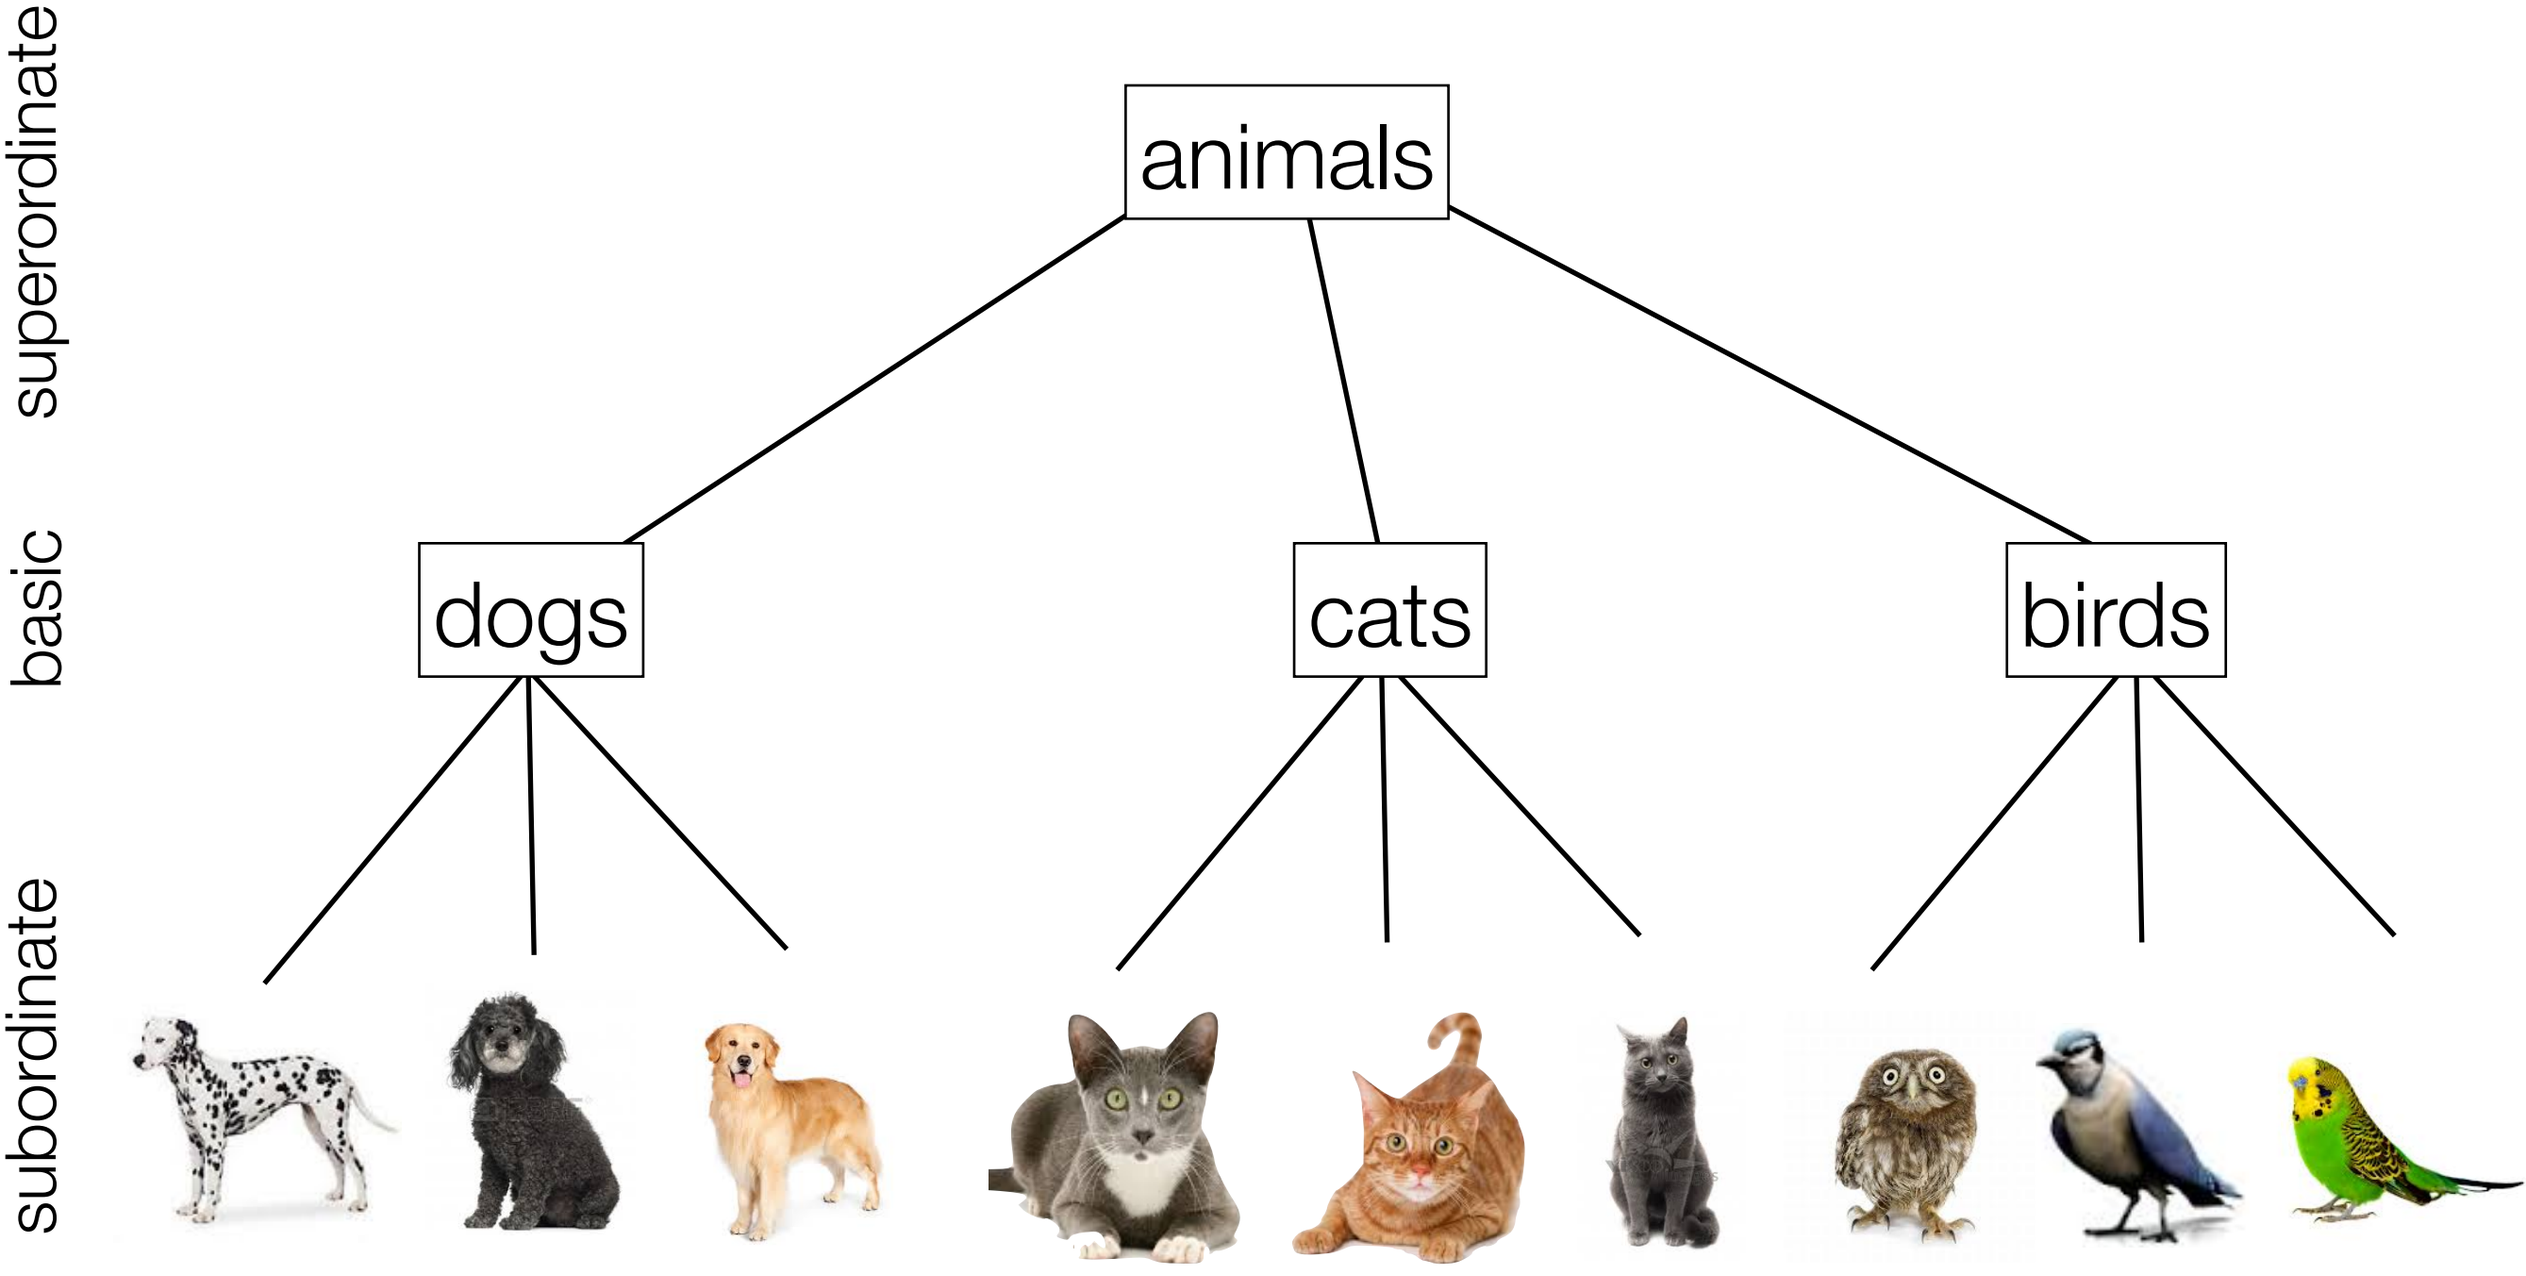

WORD LEARNING

- There is lots of independent evidence that the basic level is privileged: it is what people default to when using names, it has the highest inductive power, etc

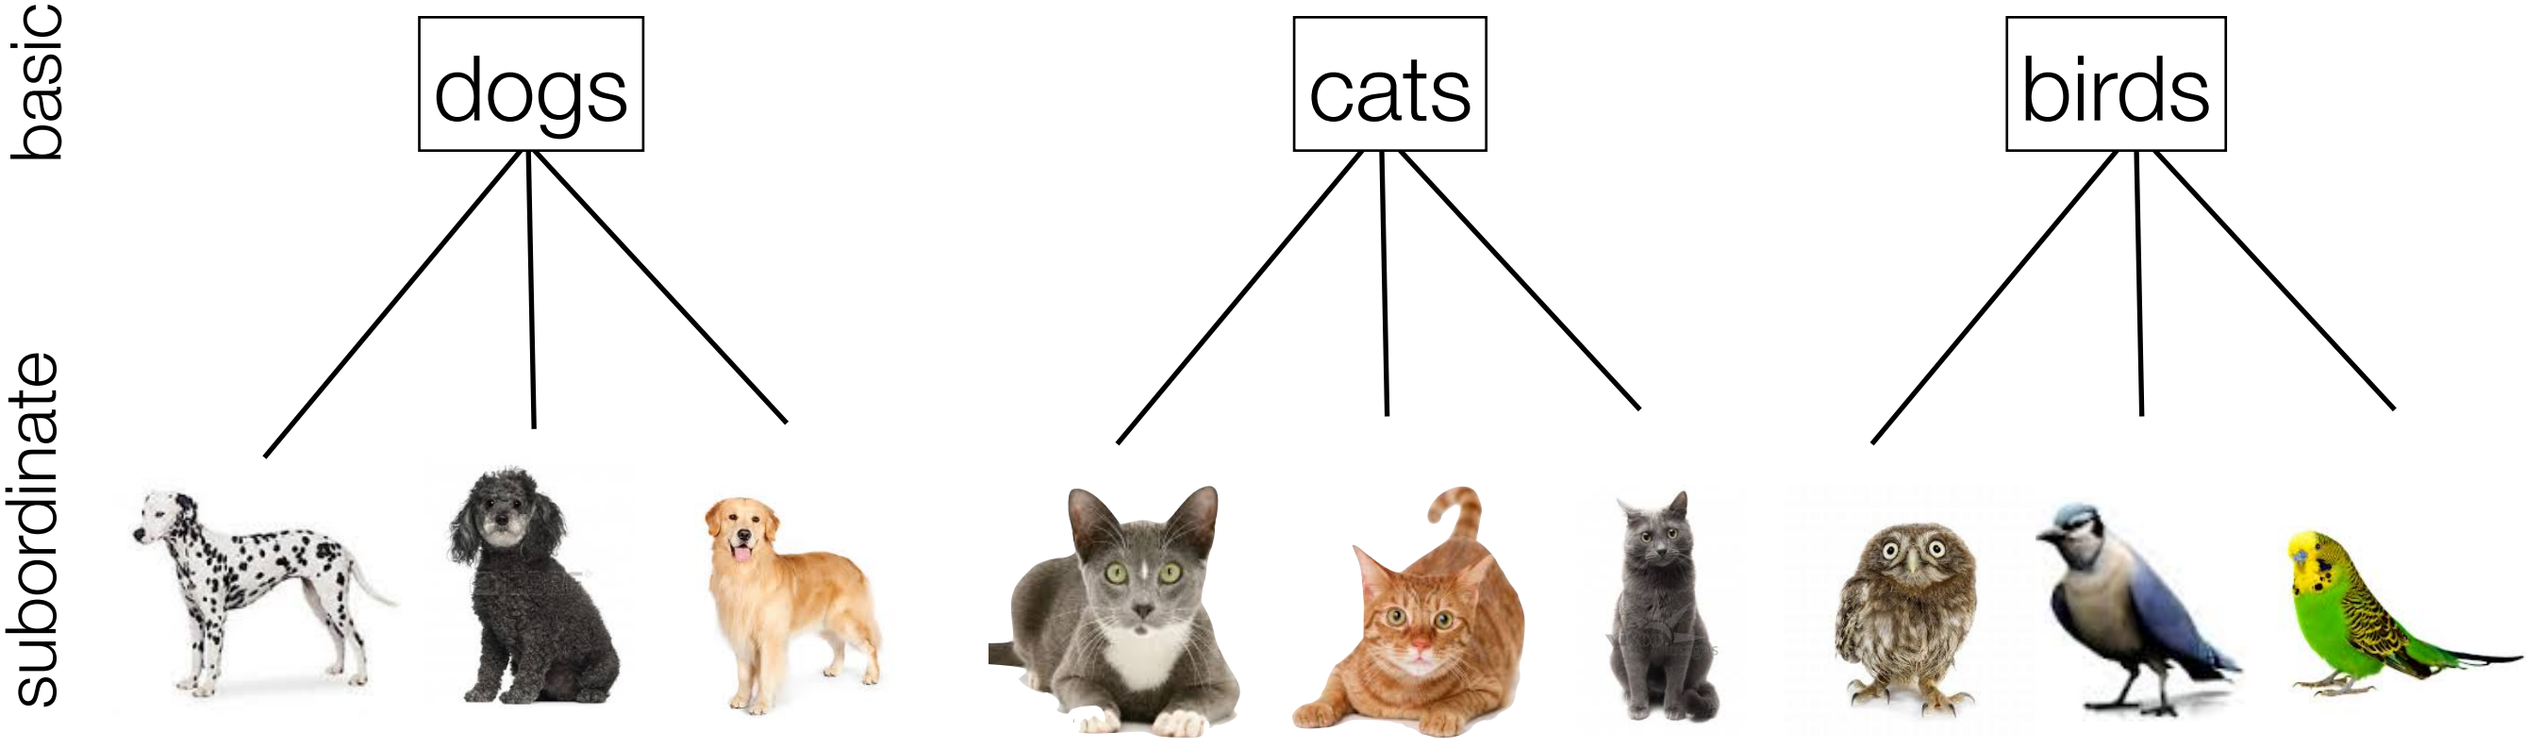

WORD LEARNING

- We would therefore expect that if people were told that *one* item was a wug, people would guess that all other items at the basic level are wugs too

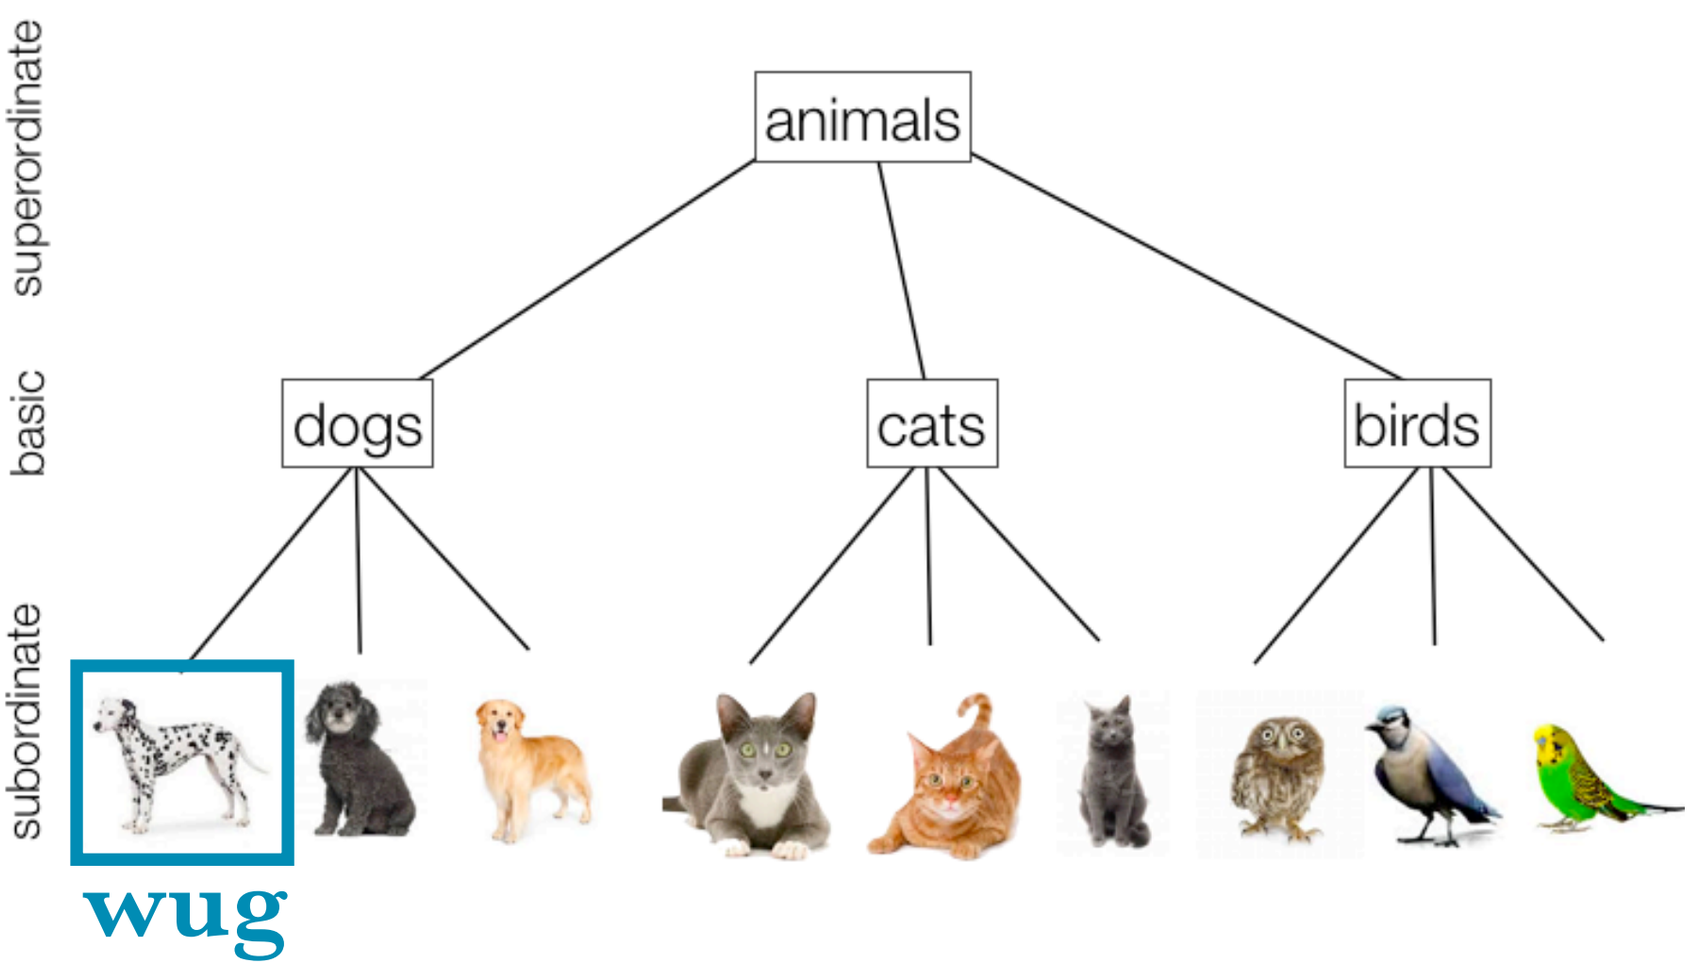

WORD LEARNING

- We would therefore expect that if people were told that *one* item was a wug, people would guess that all other items at the basic level are wugs too

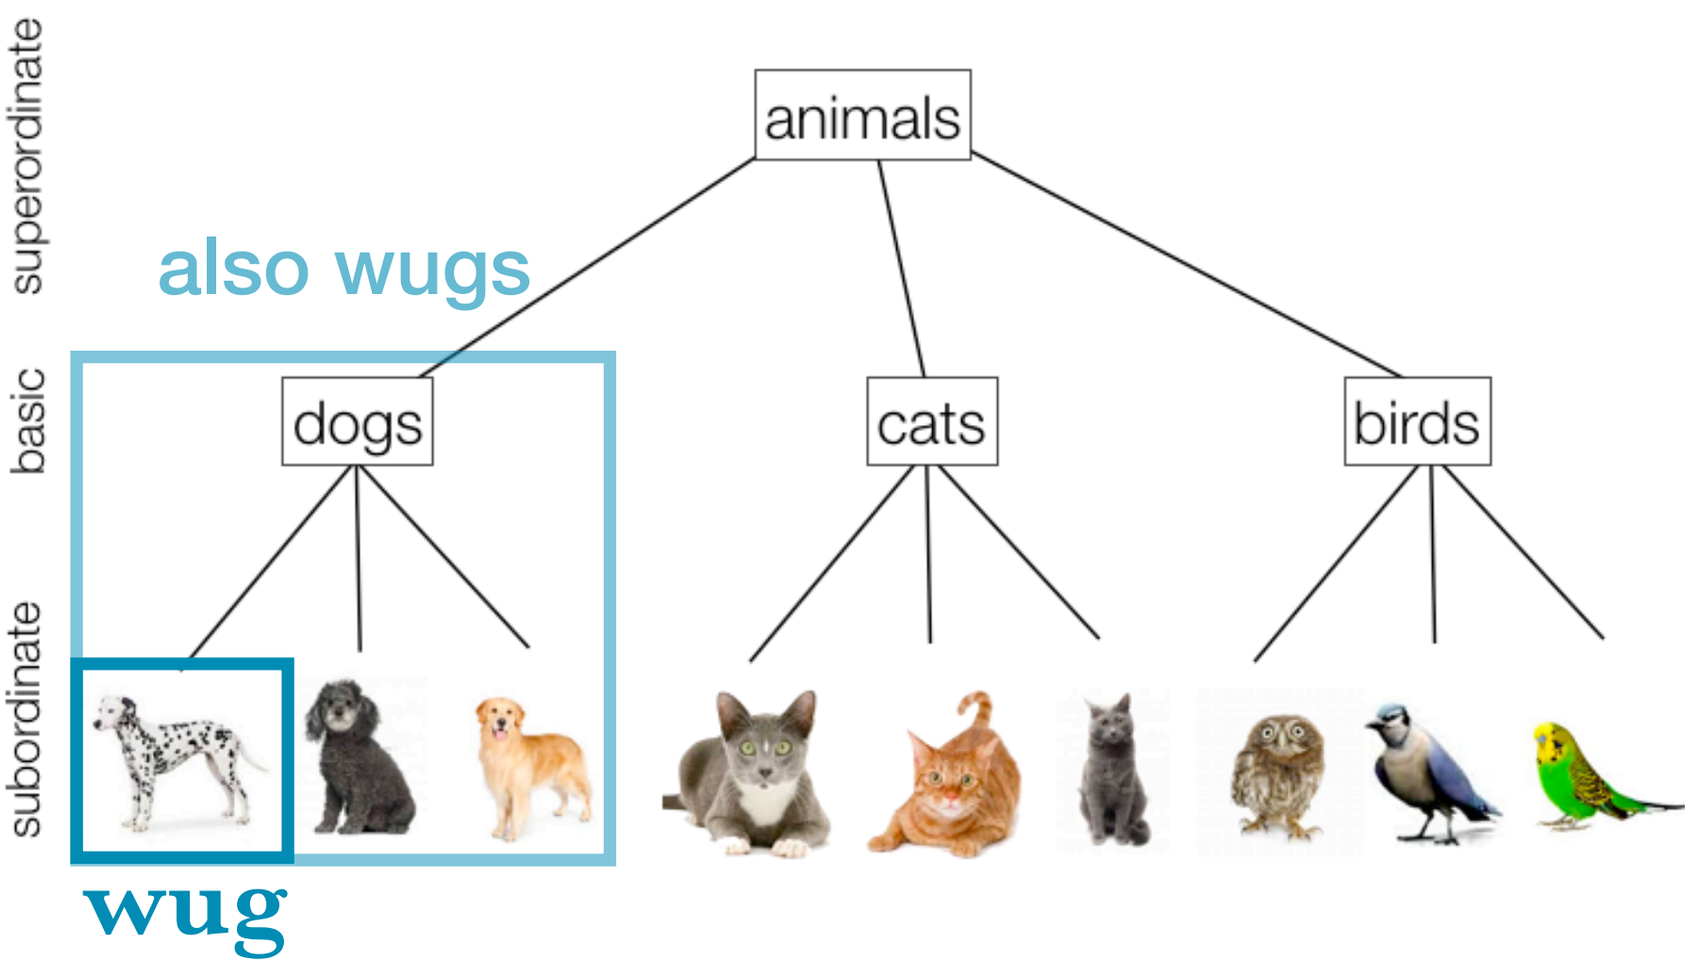

WORD LEARNING

- But what if we are given *three* examples of wugs?
- Then it depends on which three examples, and whether people are reasoning based on the size principle...

IF PEOPLE ARE ASSUMING STRONG SAMPLING…

- Then they should make the tightest possible generalisation

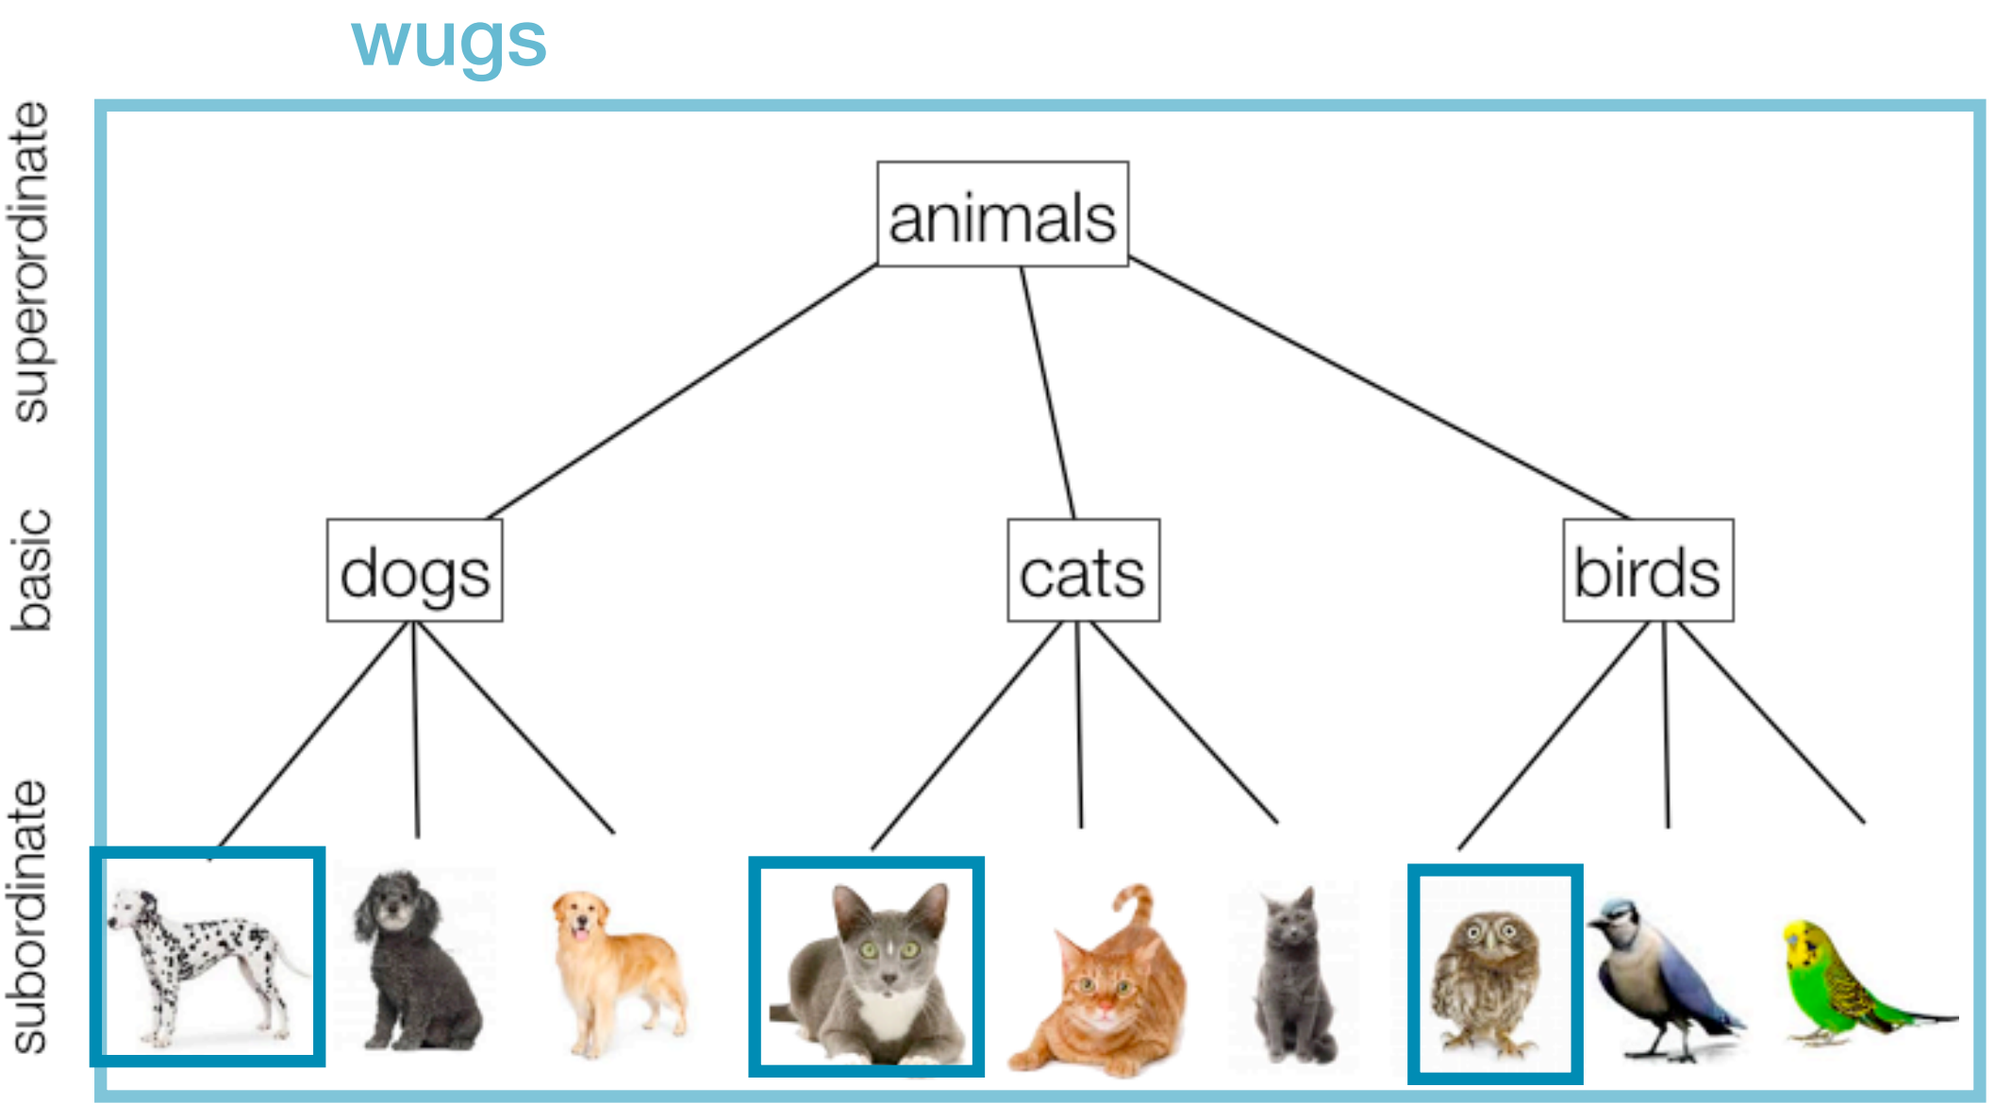

IF PEOPLE ARE ASSUMING STRONG SAMPLING…

- Then they should make the tightest possible generalisation

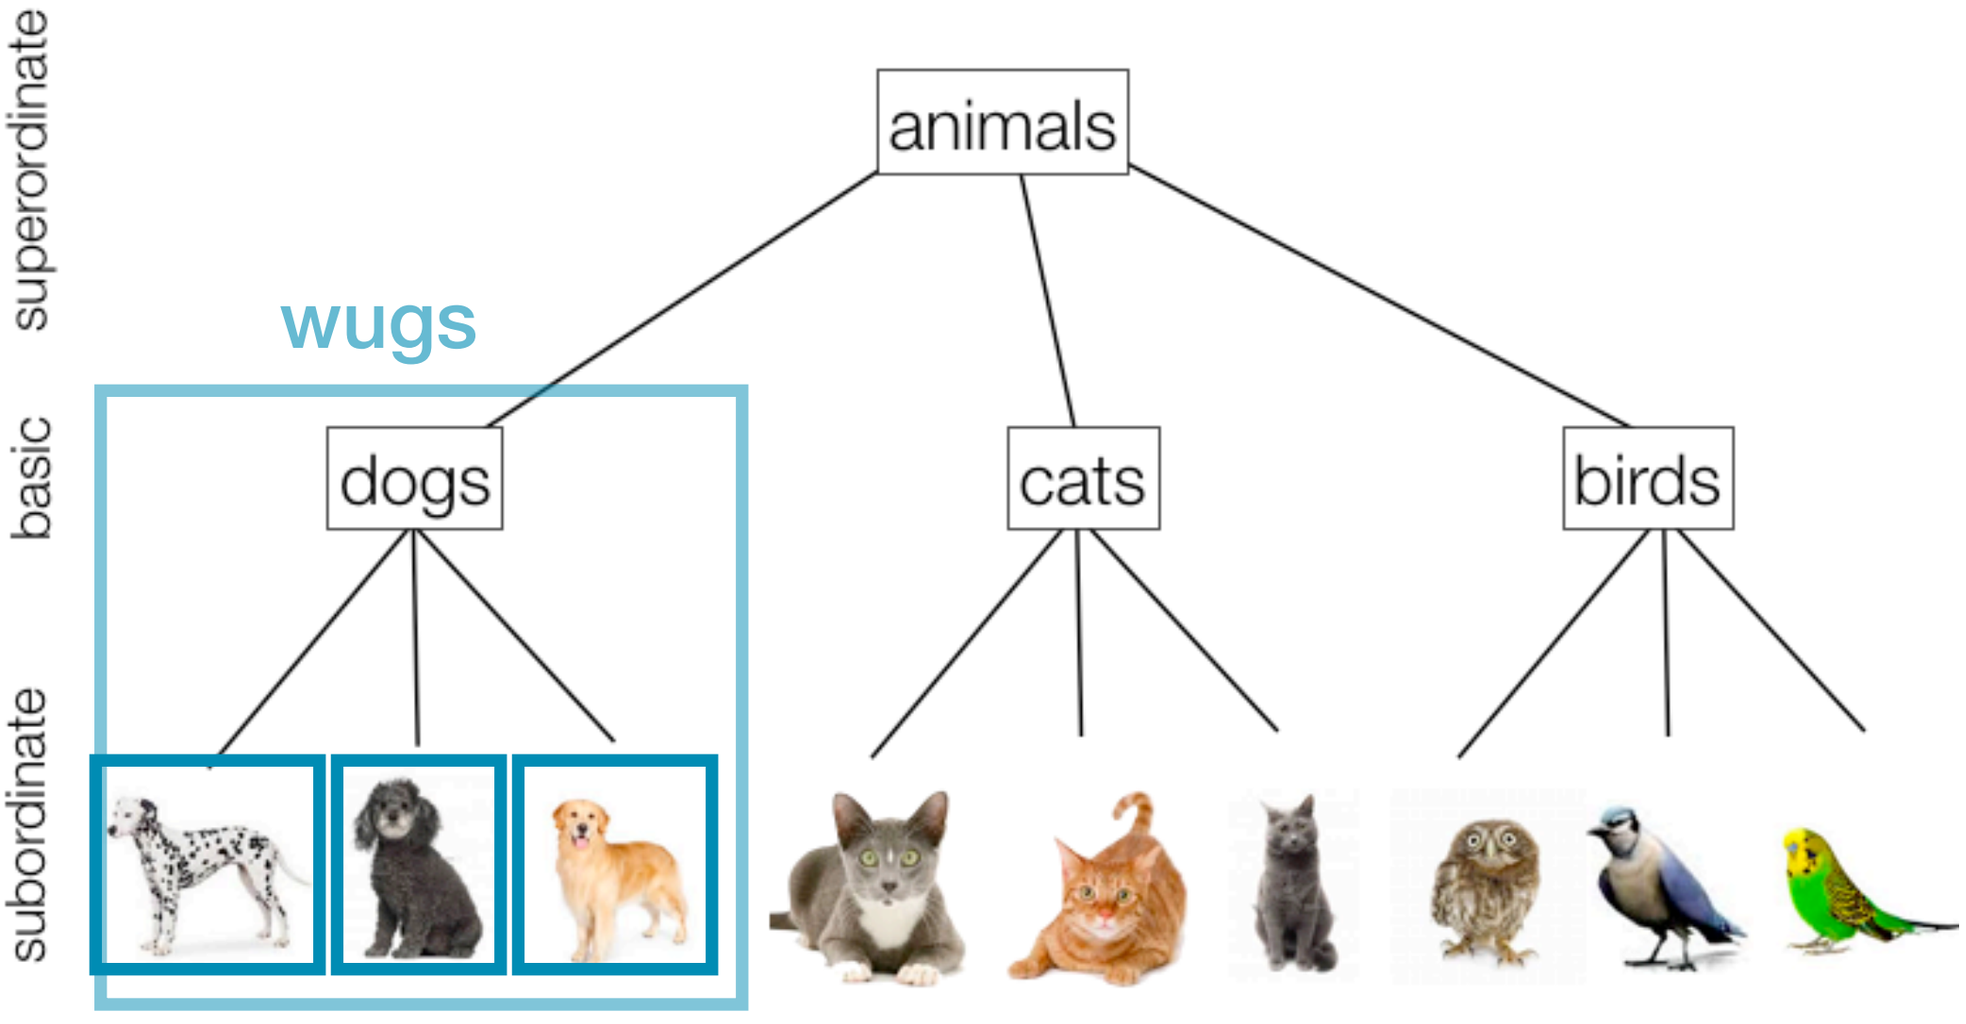

IF PEOPLE ARE ASSUMING STRONG SAMPLING…

- Then they should make the tightest possible generalisation

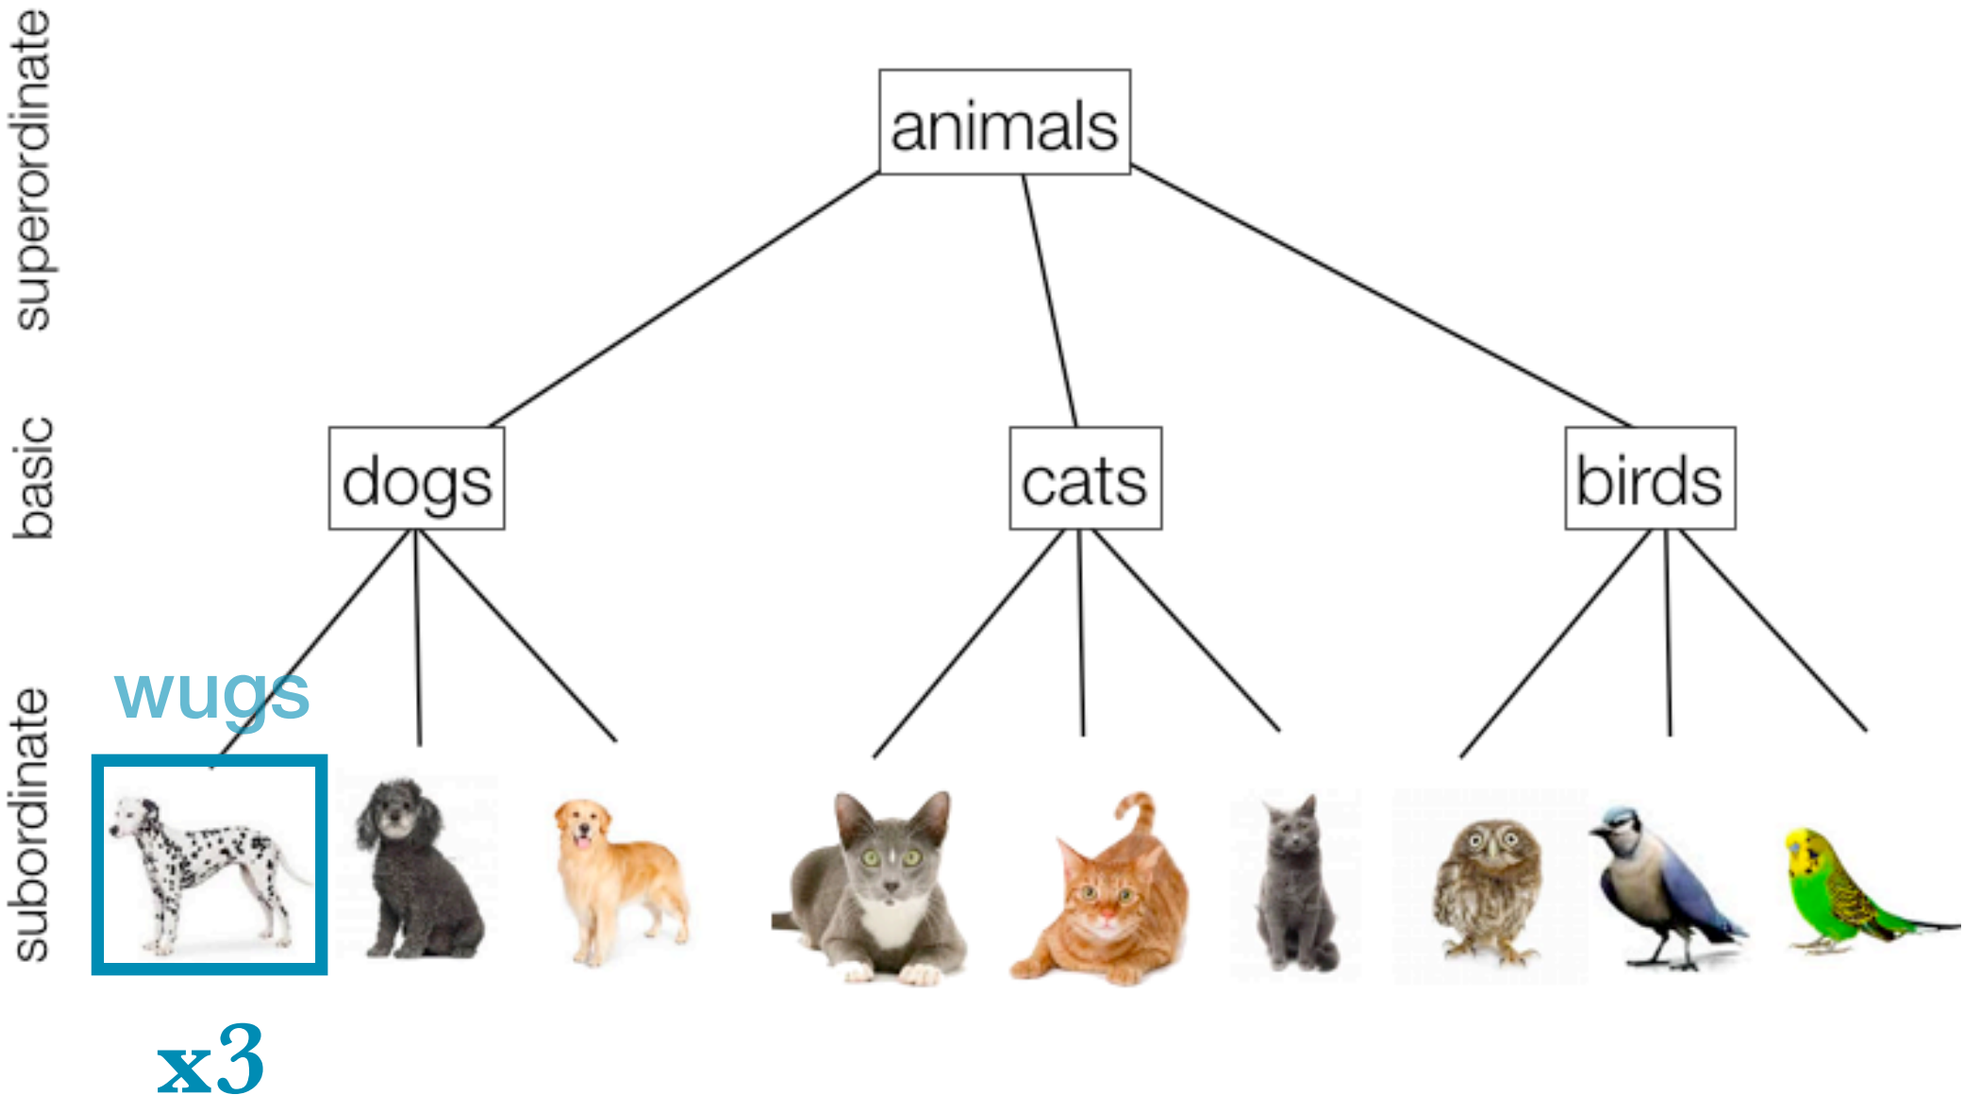

IF PEOPLE ARE ASSUMING WEAK SAMPLING…

- Then they should not tighten their generalisation when given three of the same item - there is no “suspicious coincidence” to explain

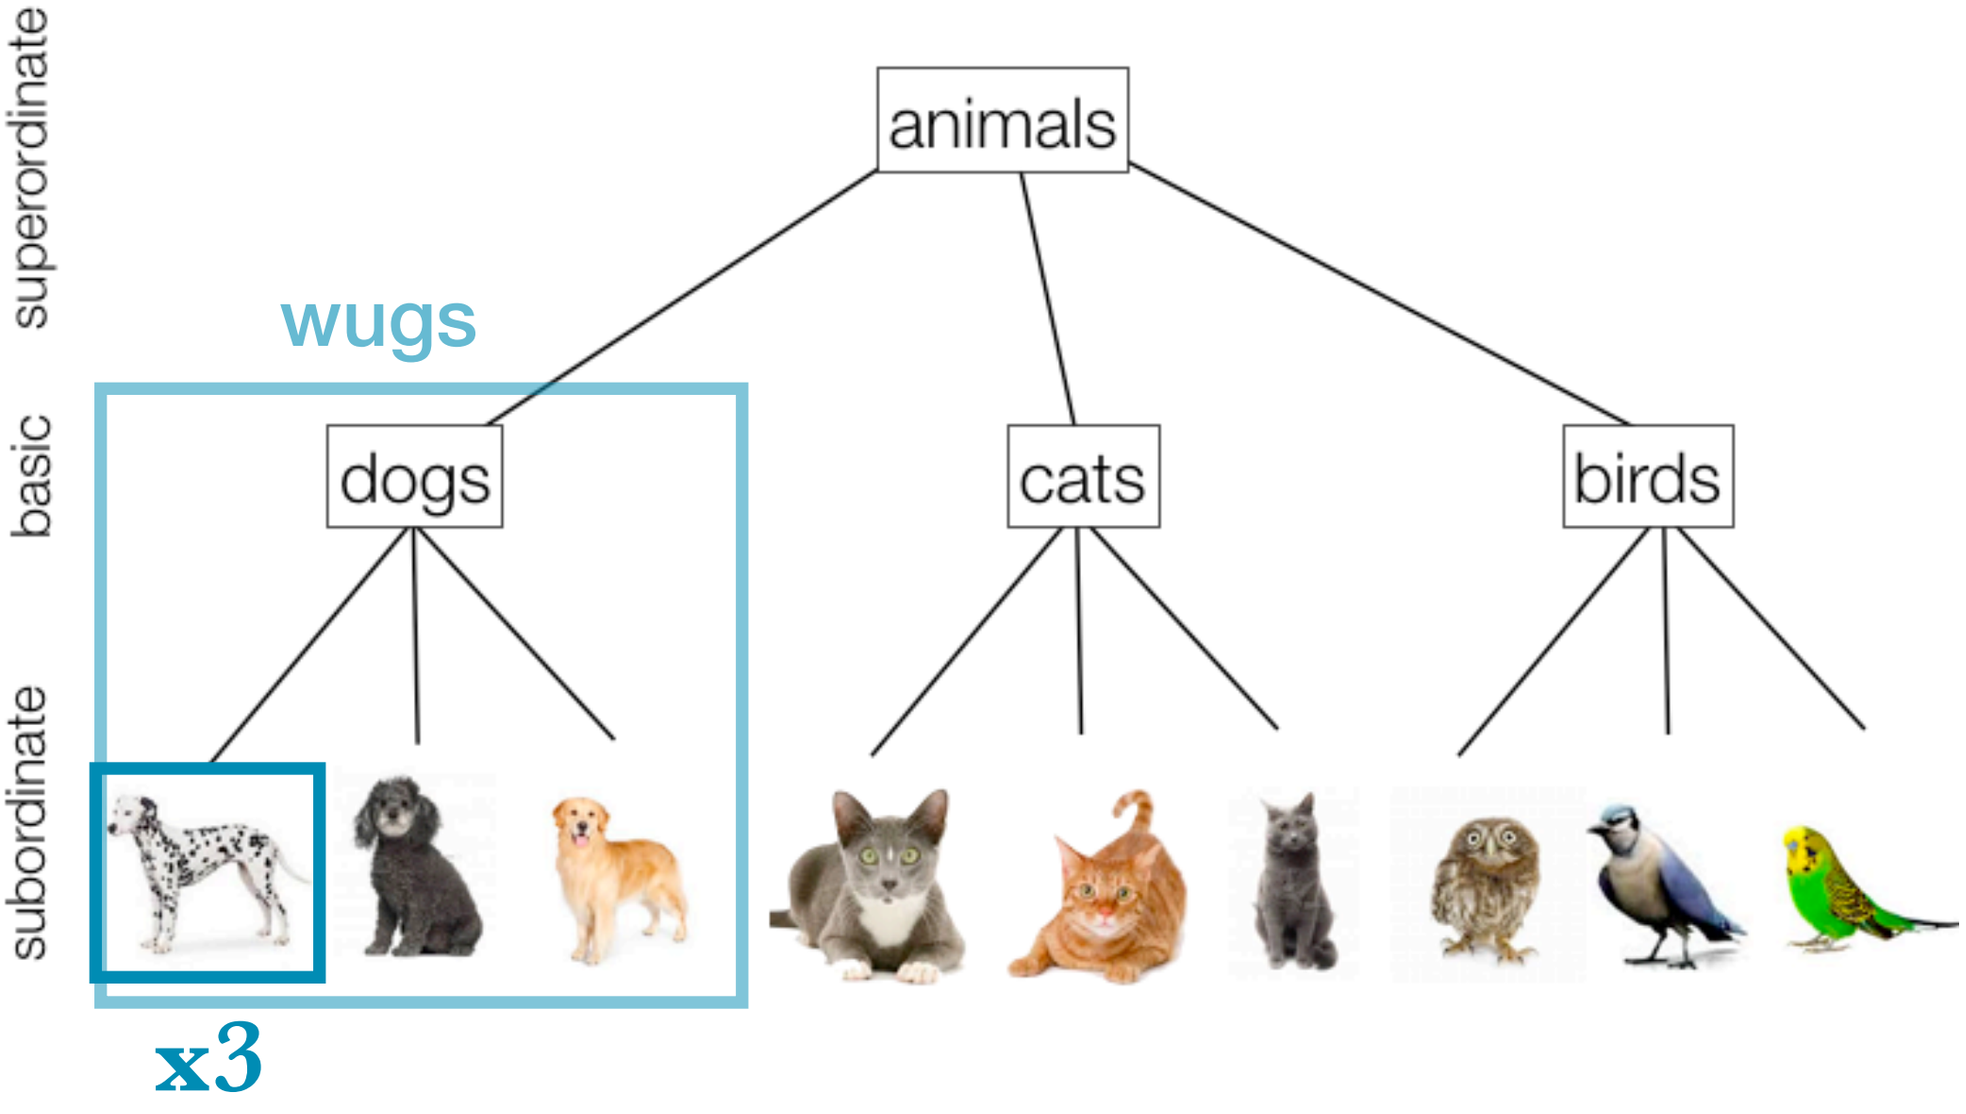

EXPERIMENTAL TEST

- Adults generalise as predicted by the size principle

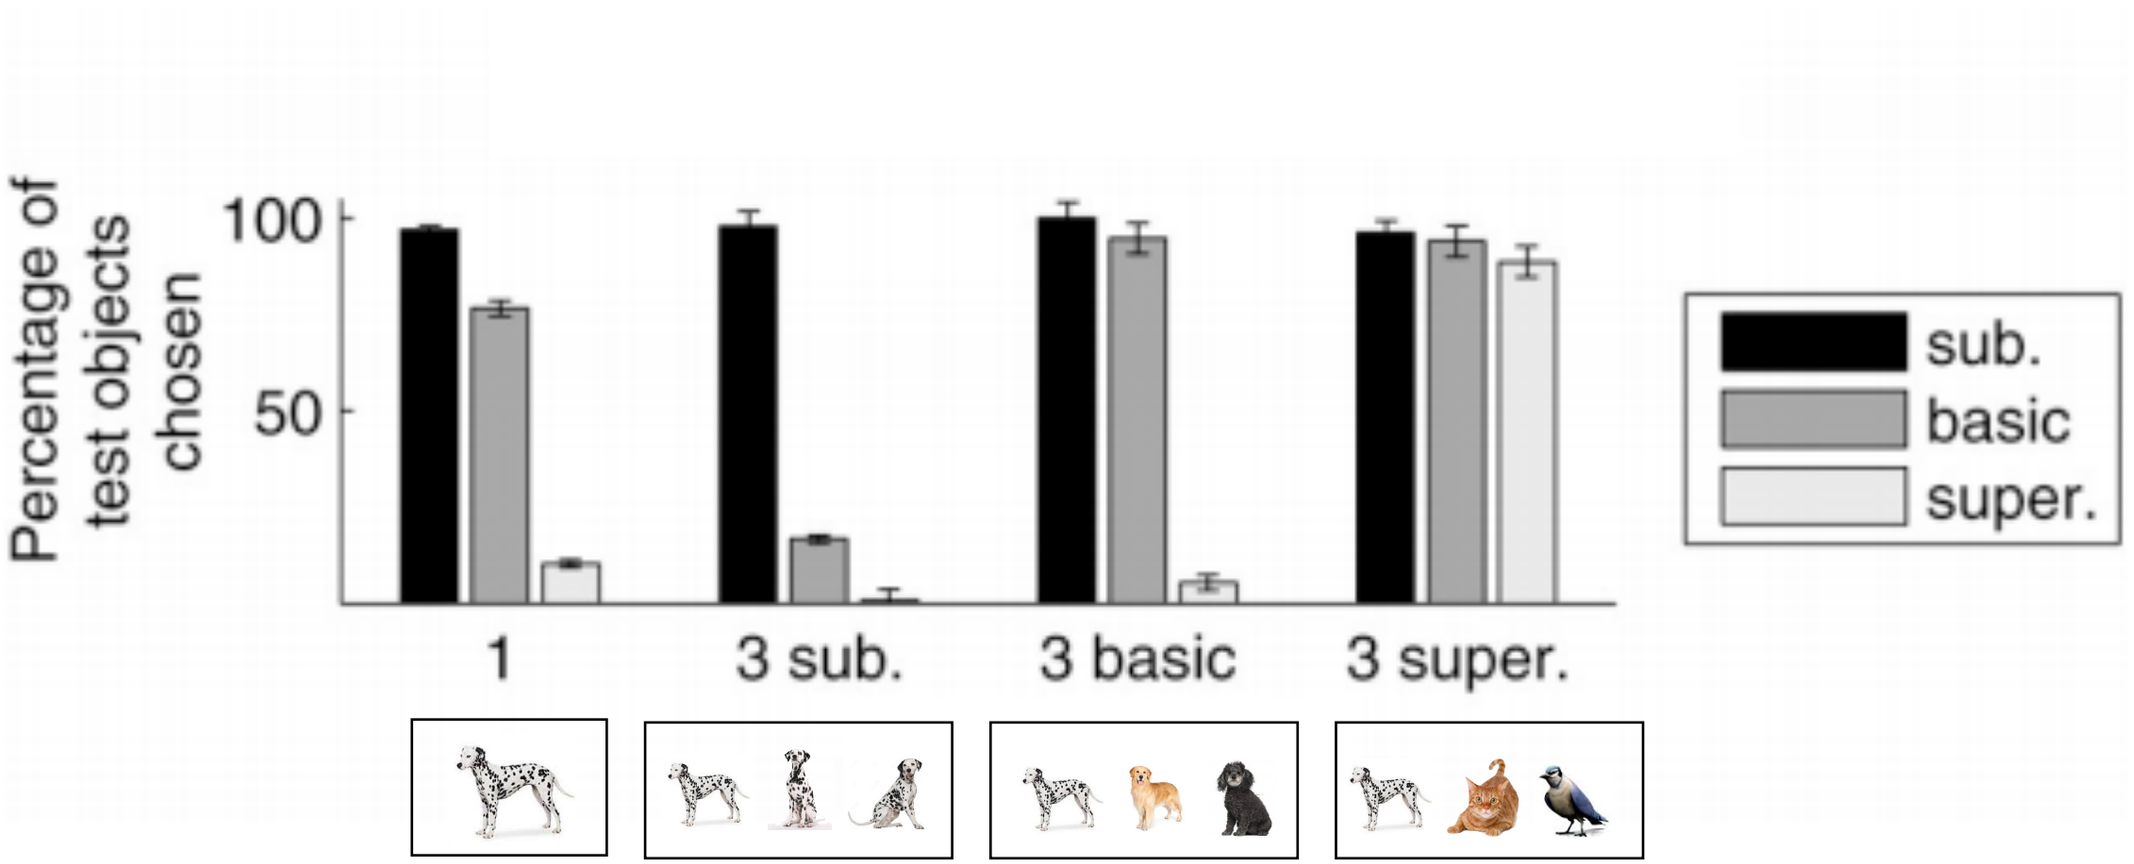

EXPERIMENTAL TEST

- Four-year old children do the same thing!

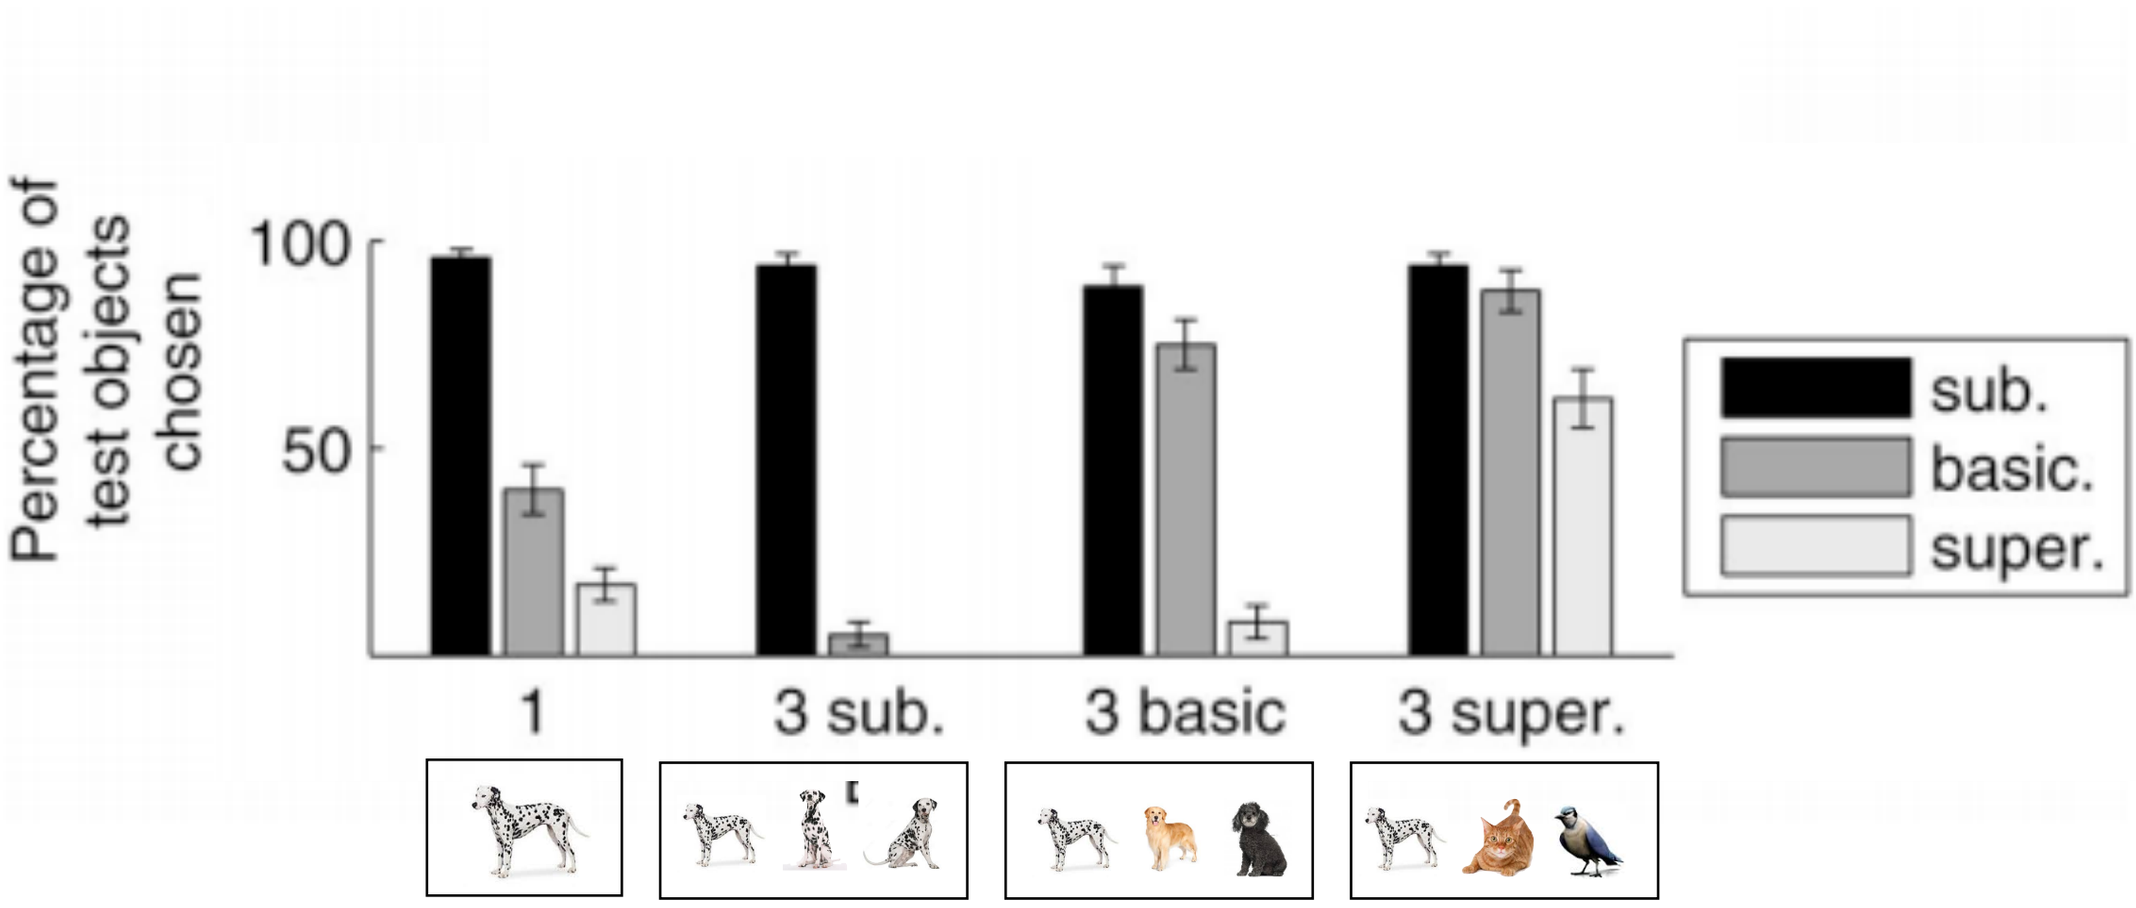

EXPERIMENTAL TEST

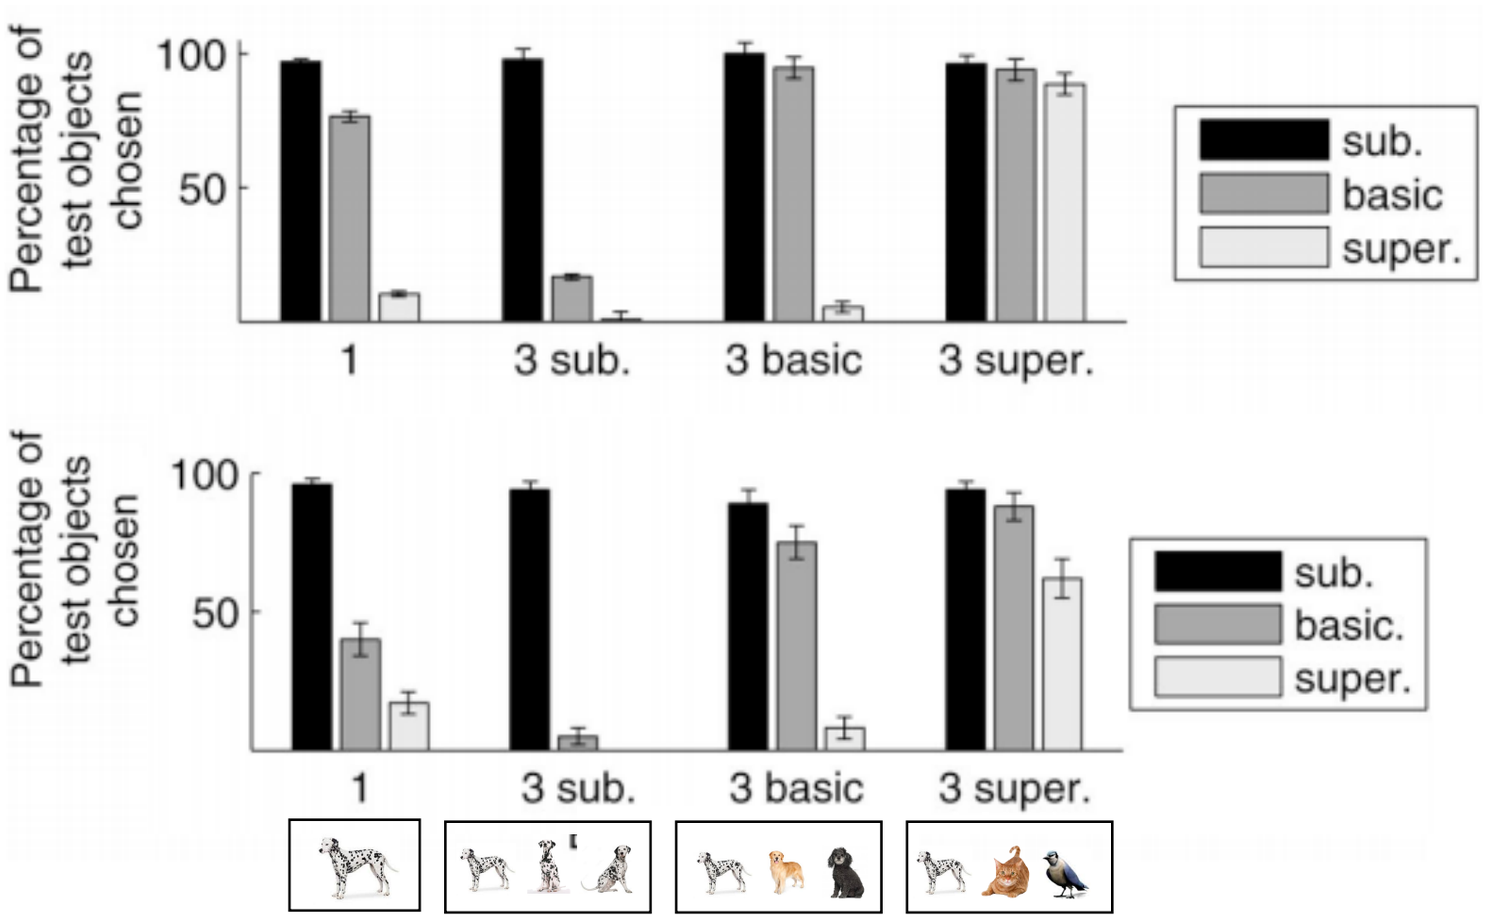

- But so far this just shows that people follow the qualitative pattern predicted by the size principle. It does not imply that they are sensitive to sampling assumptions -- perhaps they would tighten generalisations no matter what

CHANGING SAMPLING ASSUMPTIONS

- This time we vary how data are sampled (also make the objects novel)

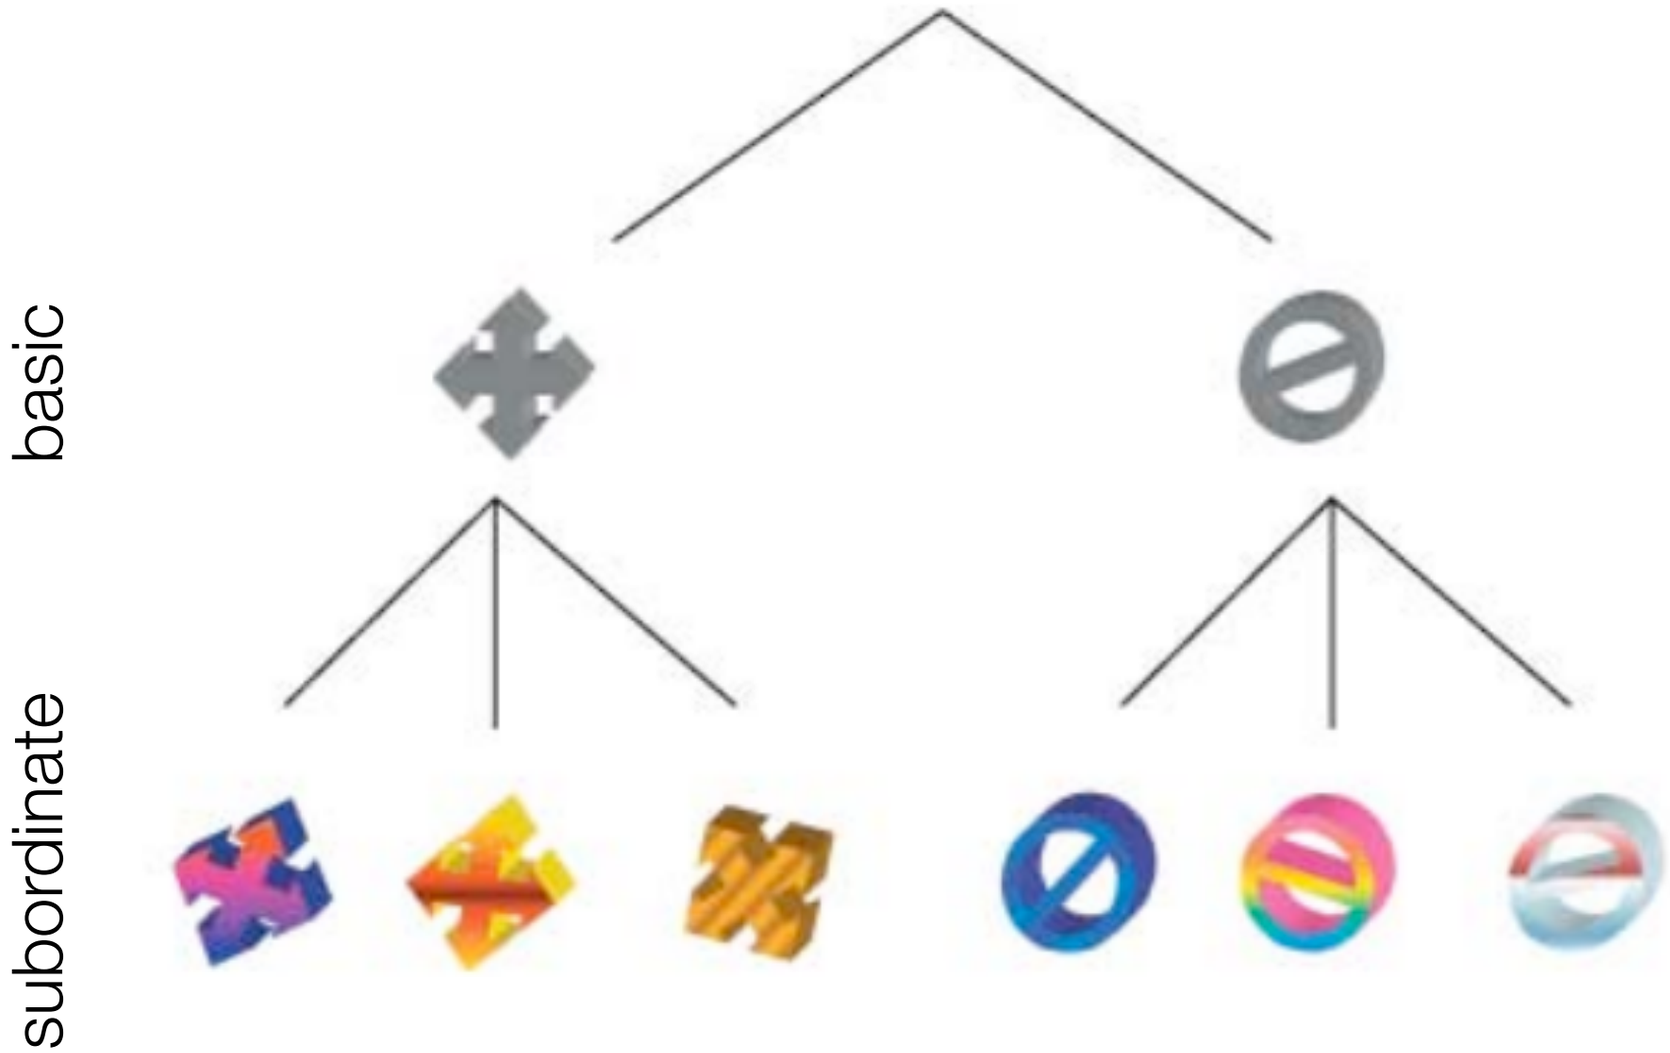

Xu & Tenenbaum, 2007b

CHANGING SAMPLING ASSUMPTIONS

- This time we vary how data are sampled (also make the objects novel)

Teacher-driven

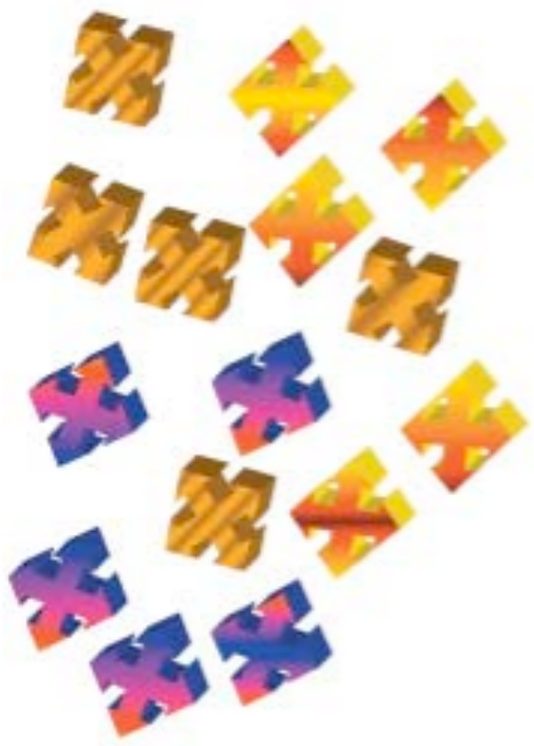

CHANGING SAMPLING ASSUMPTIONS

- This time we vary how data are sampled (also make the objects novel)

Learner-driven

All participants chose two items from the same subordinate category

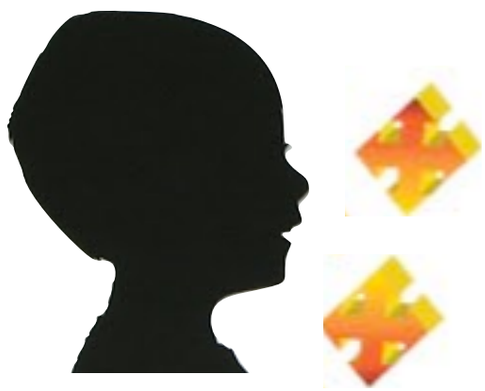

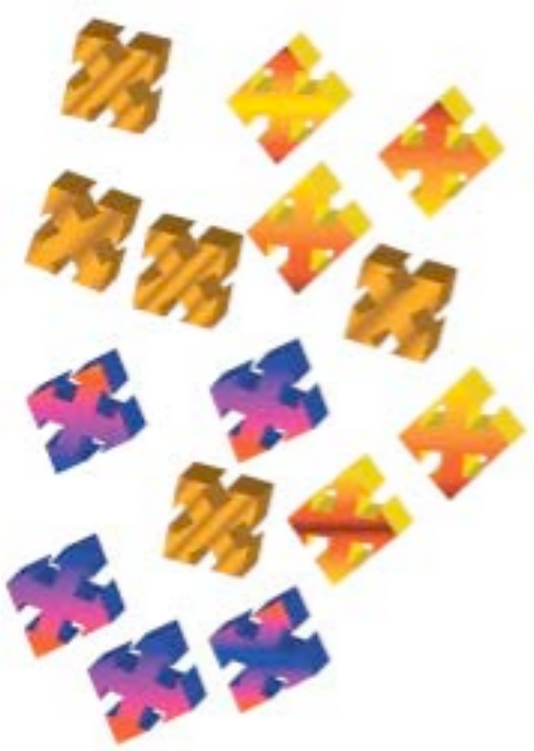

CHANGING SAMPLING ASSUMPTIONS

- This time we vary how data are sampled (also make the objects novel)

Learner-driven

So in this condition people always saw items from the subordinate category, but the 3 items were not chosen by the teacher

Teacher-driven

People saw 3 subordinate items, always chosen by the teacher

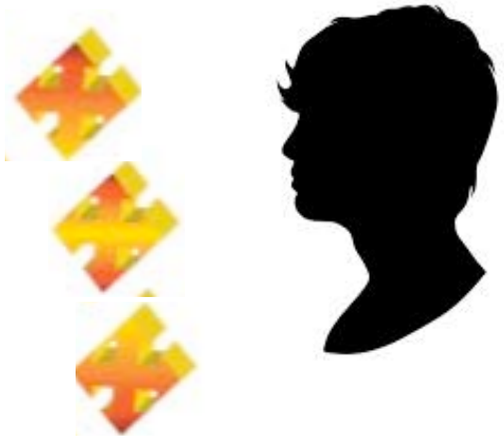

CHANGING SAMPLING ASSUMPTIONS

- People generalise tightly only when the teacher sampled the data

Learner-driven

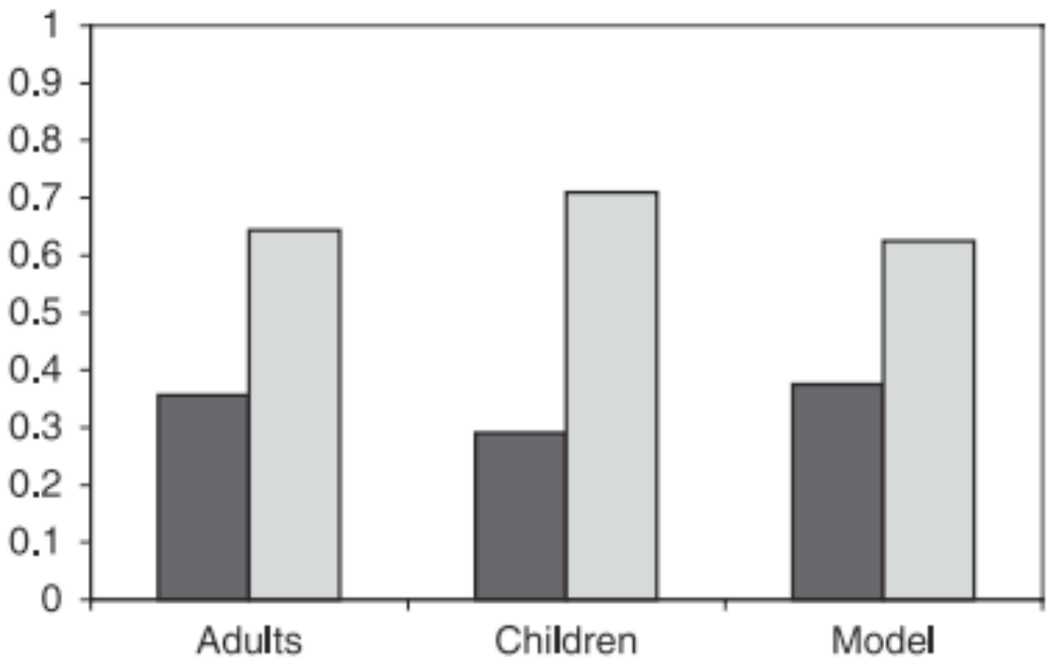

Teacher-driven

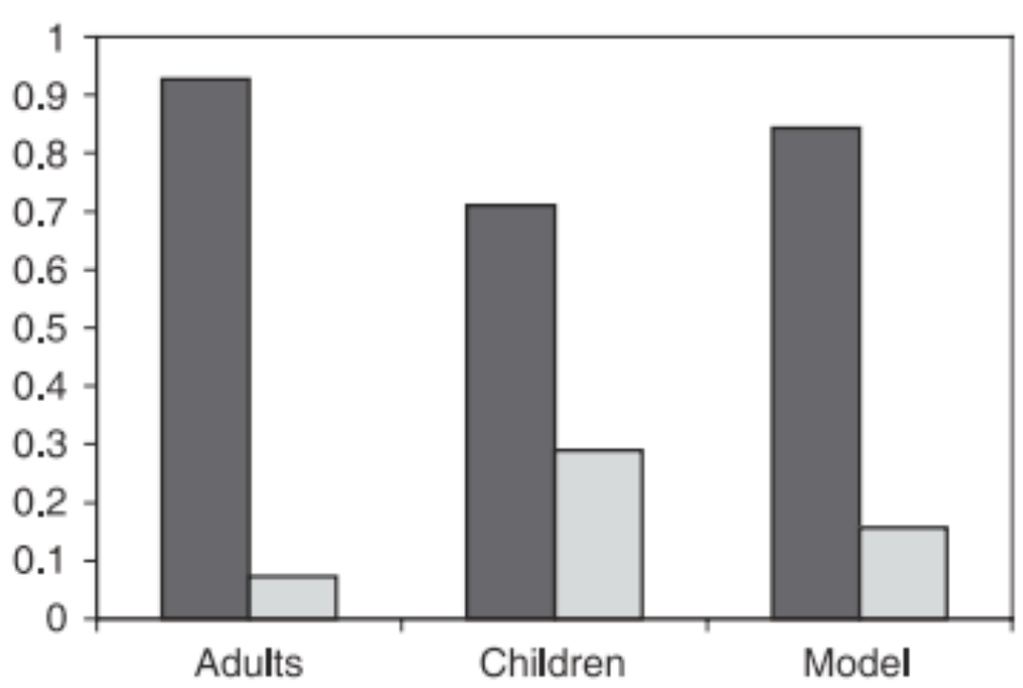

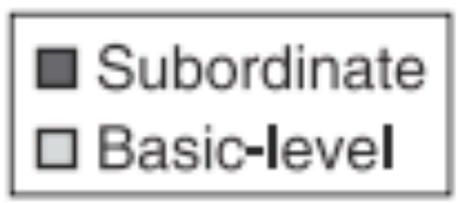

SAMPLING ASSUMPTIONS

So far all of this evidence has shown that people (including children) will tighten their generalisations more if they think the examples were generated from the concept/hypothesis directly.

But we've considered only two different ways data might be generated: strong (helpful) or weak.

In real life, data can be **censored** in many ways that should affect generalisation

CENSORED DATA

Suppose I have a box of clothing accessories, but you don't know what's in it. I like to play a game where I pick examples and you need to predict what colour they will be.

Category sampling:
I choose only hats

Small N

Large N

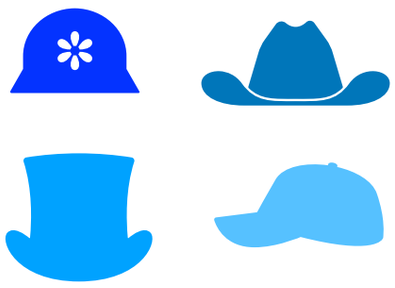

What is the probability that a non-hat is blue?

No size principle: similar with both large and small N

Property sampling:
I choose only blue things

Small N

Large N

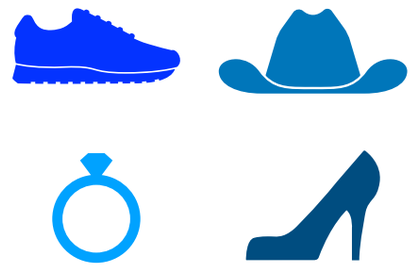

What is the probability that a non-hat is blue?

Intuitively less with large N

CENSORED DATA

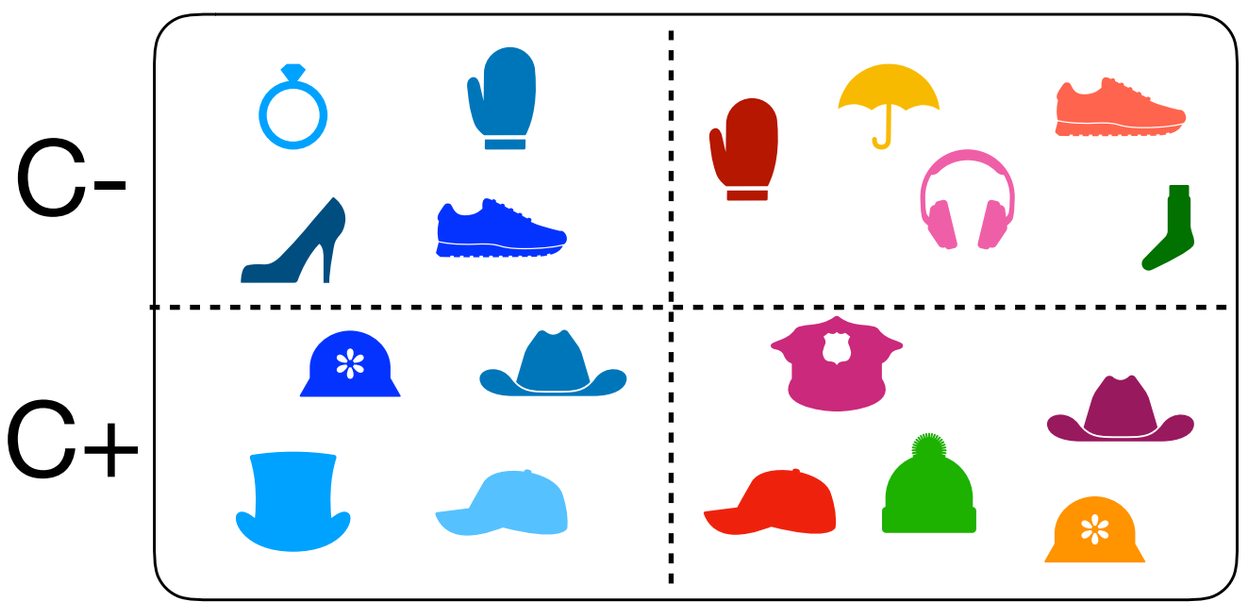

What is the probability that a non-hat is blue?

What is the probability of C-P+?

Prediction of category sampling with increasing N

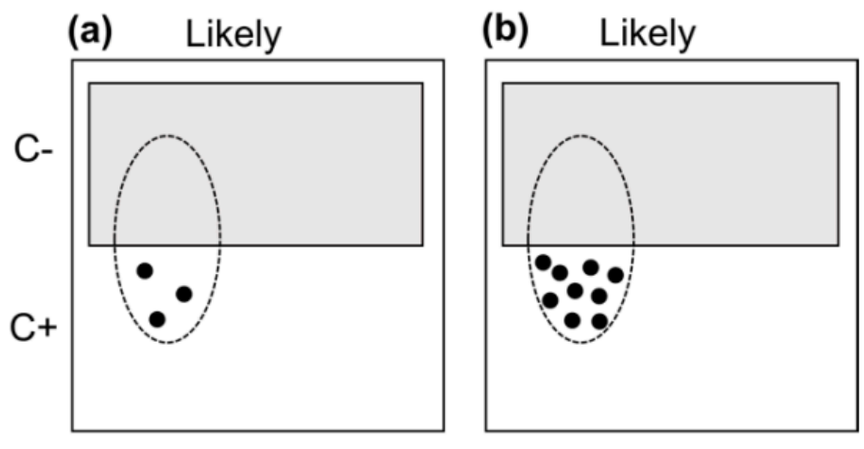

Prediction of property sampling with increasing N

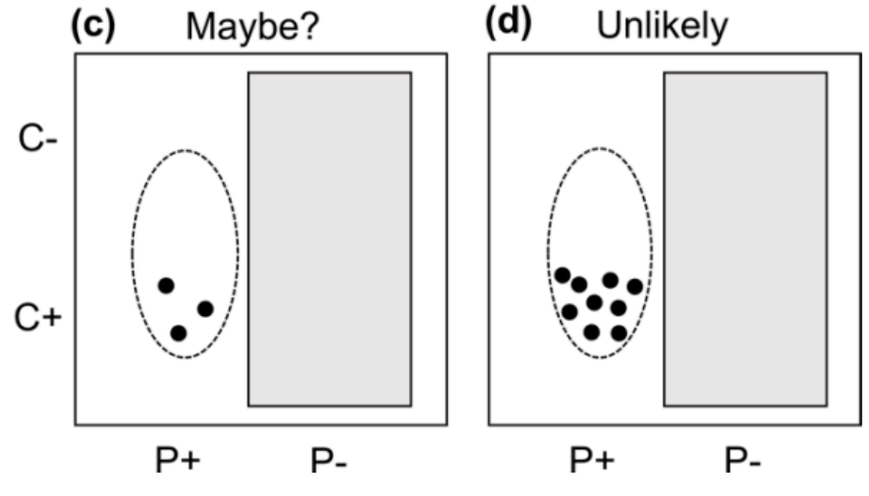

OUR TASK: DESIGN AN EXPERIMENT TO TEST THIS HYPOTHESIS

Prediction of category sampling with increasing N

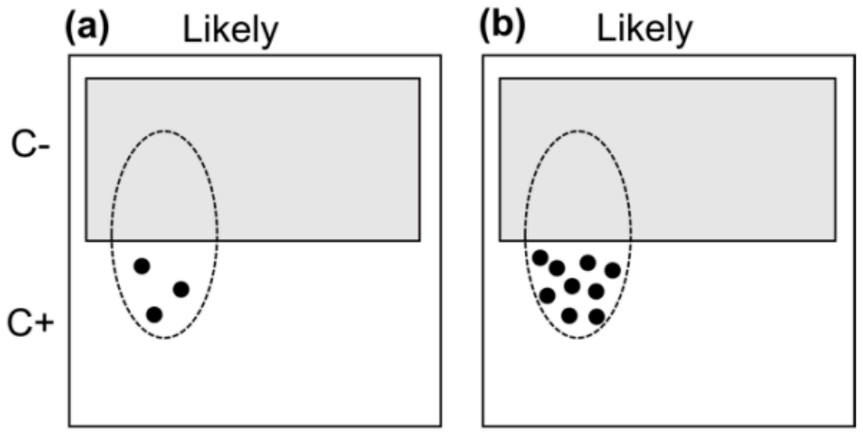

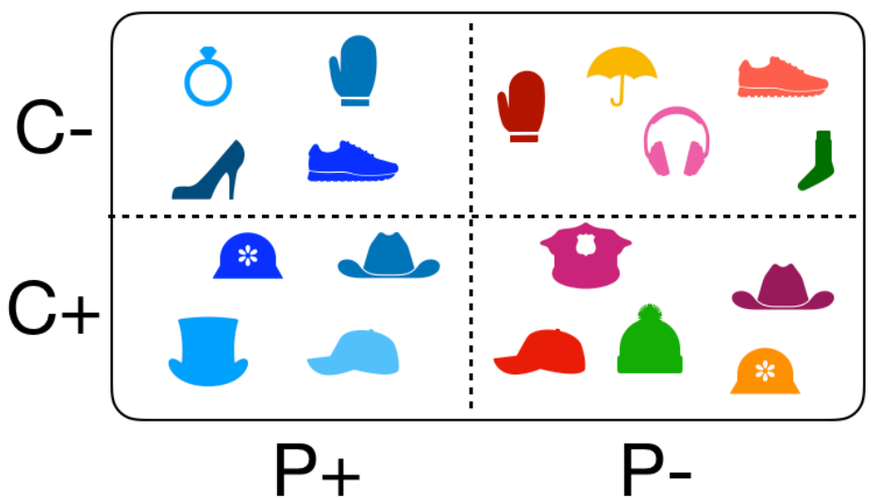

Prediction of property sampling with increasing N

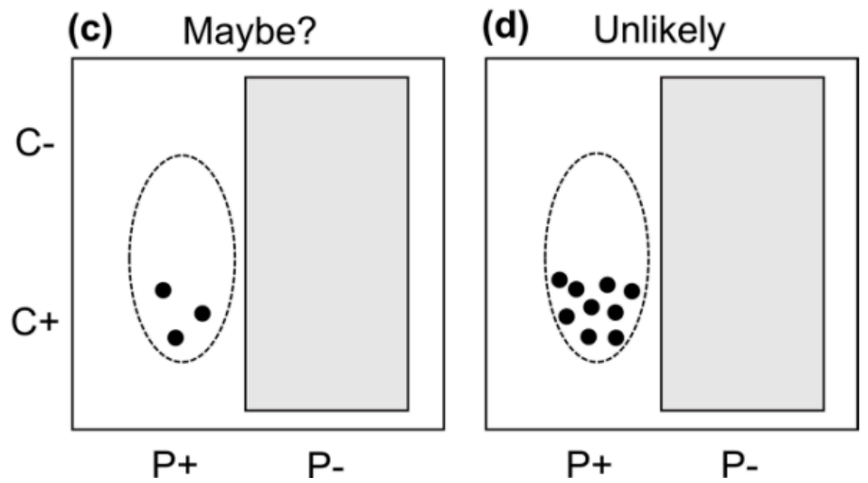

What is the probability of C-P+?

- Conditions / manipulation?
- Task?
- Instructions?