Overview

http://compcogscisydney.org/psyc2071/

Danielle Navarro

Cognitive Psychology

- Part 1: Dani Navarro
  - L1: Introduction
  - L2: Attention
  - L3: Similarity
  - L4: Reasoning
  - L5: A case study

- Part 2: Marcus Taft
  - Next!

“Could we have a brief recap?”

Lecture 1: Introduction

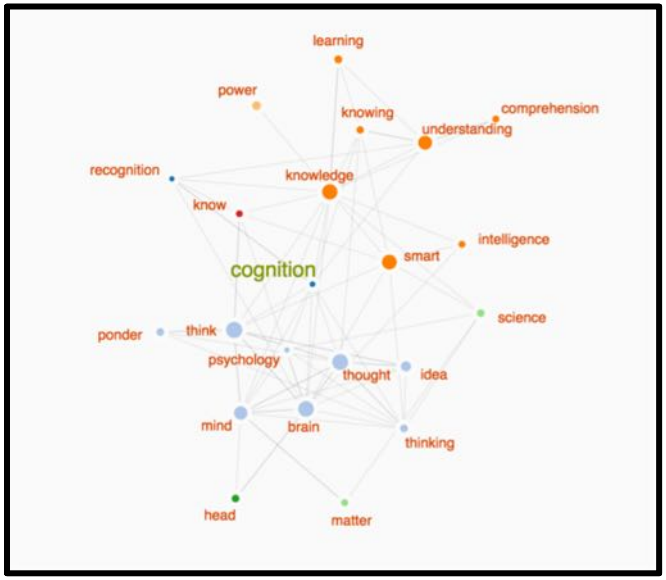

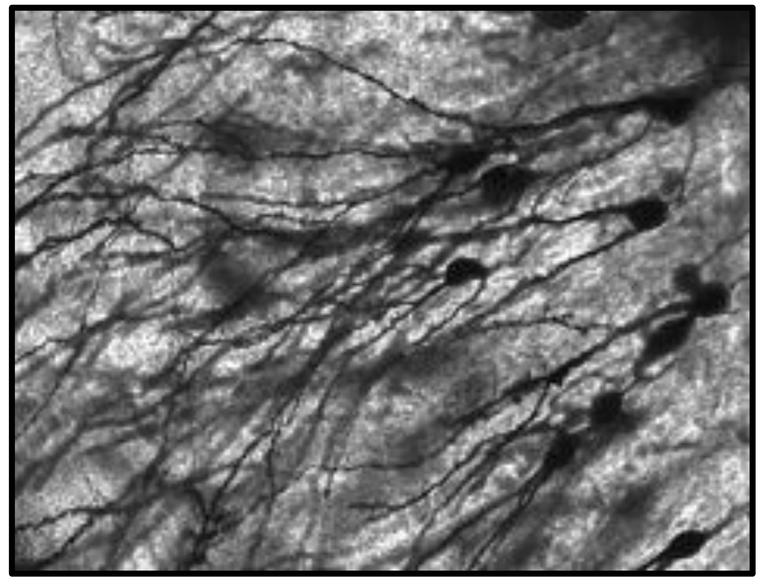

- Goal:
  - To discuss the history of cognitive psychology and introduce key ideas in how it is studied

- Ideas:
  - Behaviourism as a response to introspection
  - Key ideas in behaviourism (methodological & radical)
  - Cognitive revolution as a response to behaviourism
  - Methods for measuring cognition
  - The "computational metaphor" views cognition as "information processing" (does not say mind = laptop!)
  - Marr's levels of analysis

Lecture 2: Attention

- <u>Goal</u>:
  - To discuss different "kinds" of attention, with a focus on results in auditory & visual attention, and visual search

- <u>Ideas</u>:
  - Definitions for different kinds of attention (how many targets, what kind of target, how is attention controlled)
  - Audition: cocktail party problem & early vs late selection
  - Visual analogs (e.g., negative priming vs semantic interference)
  - Visual search: serial search vs parallel search, pop-out effects, feature integration theory & illusory conjunction

Lecture 3: Similarity

- <u>Goal</u>:
  - What is similarity and what is it good for?

- <u>Ideas</u>:
  - The "snowflake problem" means we need similarity
  - How to measure similarity
  - Links between similarity, generalisation & categorisation
  - Geometric theory of similarity & Shepard's law
  - Featural theory of similarity & explanation of asymmetry
  - Structural alignment & the MIPs vs MOPs effect
  - Transformational similarity & explanation of asymmetry (+ experimental evidence)

Lecture 4: Reasoning

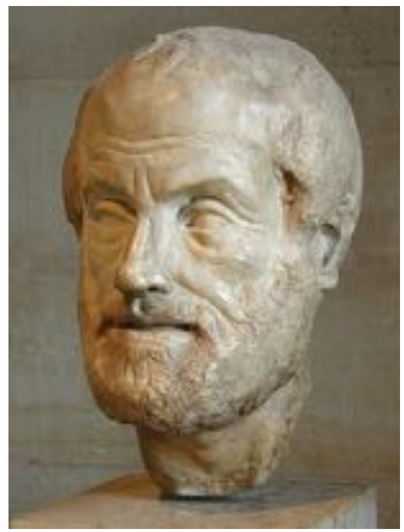

- <u>Goal</u>:
  - How do people reason and evaluate arguments?

- <u>Ideas</u>:
  - Difference between induction and deduction
  - Valid vs invalid arguments: MP, MT, DA, DC
  - Wason selection task (plus the “deontic” version of it)
  - Inductive phenomena: premise-conclusion similarity, premise diversity, premise monotonicity
  - Fallacies: argument from ignorance depends on epistemic closure; circular arguments appeal to explanatory systems and depend on the strength of the alternative

Lecture 5: The case study

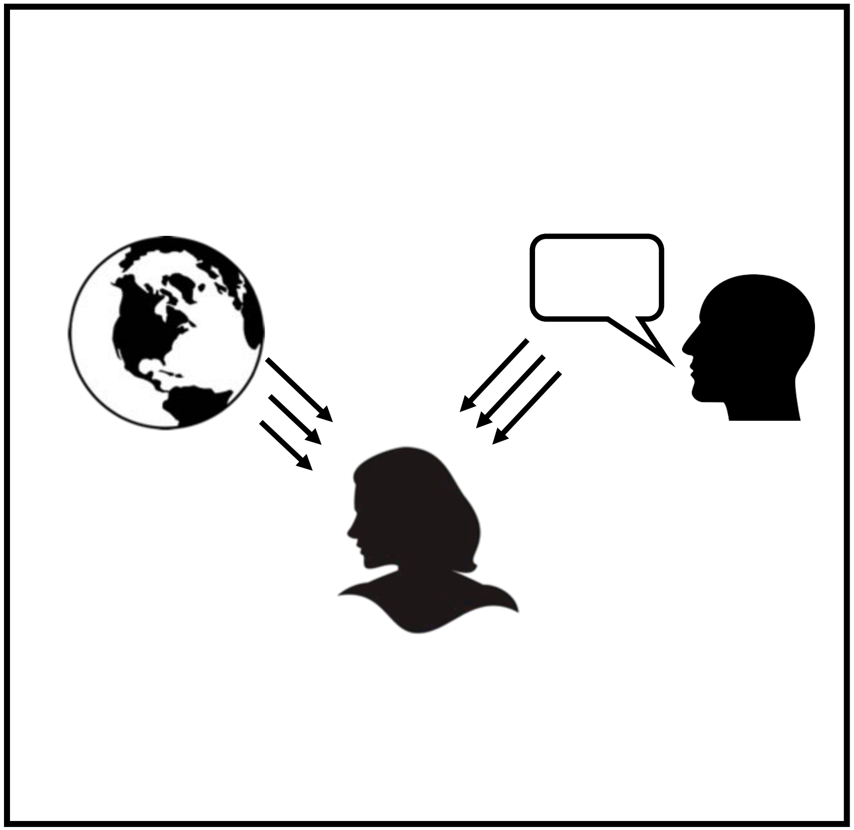

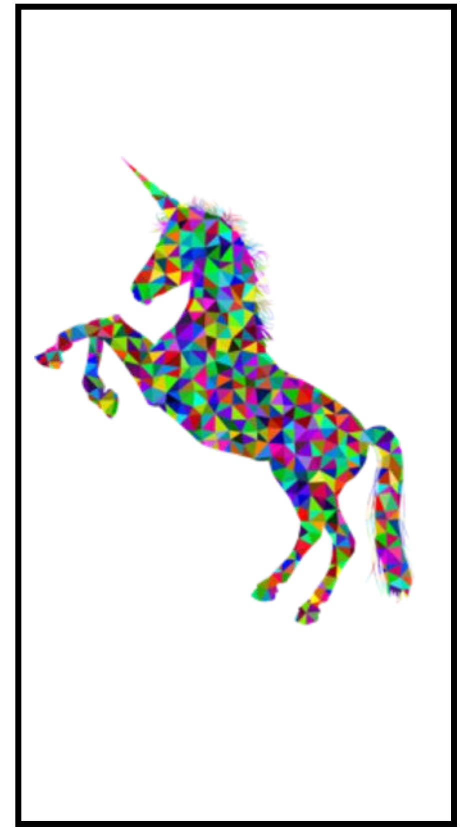

- <u>Goal</u>:
  - Link the previous lectures: show how reasoning uses similarity, attention & social cognition

- <u>Ideas</u>:
  - Similarity calls attention to a target category, which drives the premise non-monotonicity effect
  - Explanation: People use similarity to make persuasive arguments, so this makes sense
  - Prediction: helpful person -> non-monotonicity; unhelpful world -> monotonicity
  - Experiment: Manipulate people's beliefs about the origin of the data and show this changes their reasoning

“Could we have some examples of questions to help us study?”

(besides the quizzes, obviously!)

Lecture 1

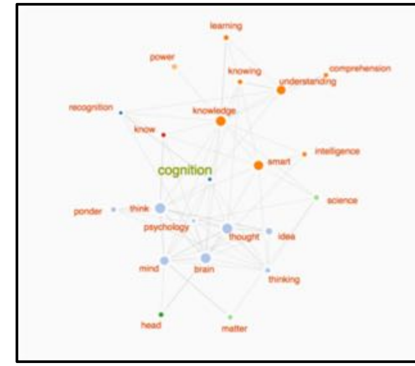

- What is the difference between…
  - … perception and cognition?
  - … theoretical and methodological behaviourism?
  - … behaviourism and cognitivism?
  - … computational, algorithmic & implementation levels?
- What methods are used to measure cognition?
- What is the computational metaphor?
- Why do we use the computational metaphor?
- Is the computational metaphor consistent with behaviourism?

Lecture 2

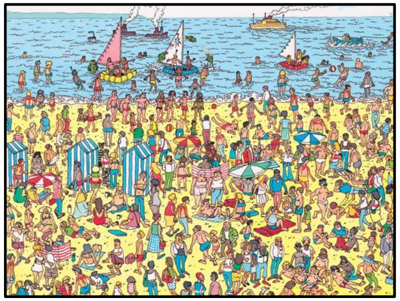

- Can you explain the different "kinds of attention"?
- What do we learn from "shadowing tasks"?
- What is the difference between early and late selection theories?
- Why does reaction time increase with "set size" for serial search but stay flat for "parallel search"?
- What kind of visual searches can we do in parallel?
- How does feature integration theory explain these illusory conjunctions?

Lecture 3

- Why does cognition rely on similarity?
- Explain the difference between …
  - … geometric and featural theories
  - … structural alignment and transformational theories
- Describe different ways to measure similarity?
- What does "the universal law of generalisation" say?
- What's the difference between MIPs and MOPs?
- Why is the similarity from A to B not always the same as the similarity from B to A?
  - Do different theories explain this differently?

Lecture 4

- How are induction and deduction different?
- What is the meaning of "modus ponens", etc?
- Is deductive reasoning always equally easy/hard?
  - … does argument structure matter (e.g., MP, MT)
  - … does it matter if we use an "indicative" or "deontic" conditional? Why?
- Describe the "premise monotonicity" effect
- When is an argument from ignorance acceptable?
- When is a circular argument more acceptable to people?

Lecture 5

- How do similarity & attention relate to reasoning?
- Why do we think social cognition plays a role?
- What does this mean for (non)monotonicity?
- What were the experimental manipulations?
- What was the dependent variable?
- What were the results of the study?
- What can be concluded from it?
- What are the limitations of the study?

The relationship between reaction time and stimulus transformations

Transform(A →B) = 2 steps

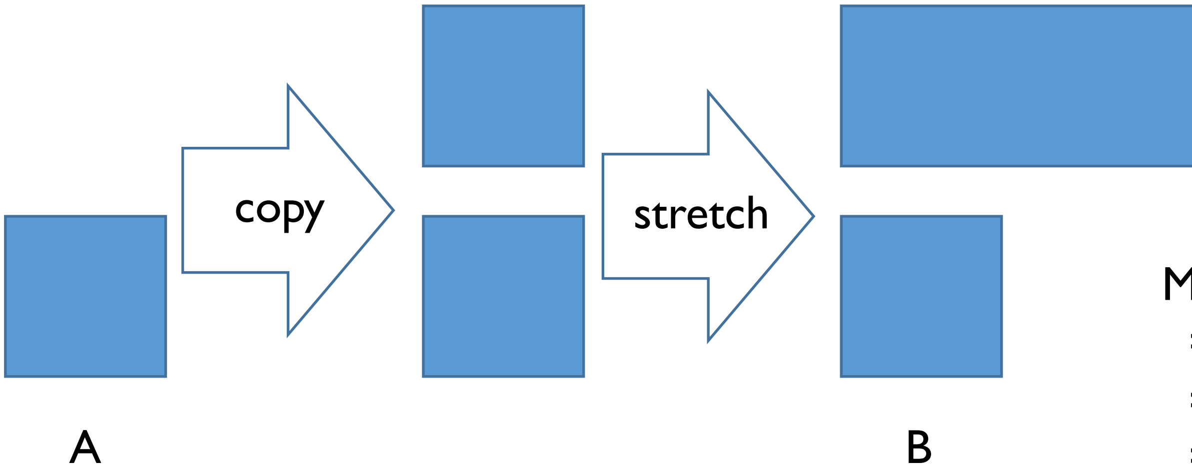

Many transformations
= very DISSIMILAR
= EASY to distinguish
= FAST (small) reaction time

Transform(B → A) = 1 step

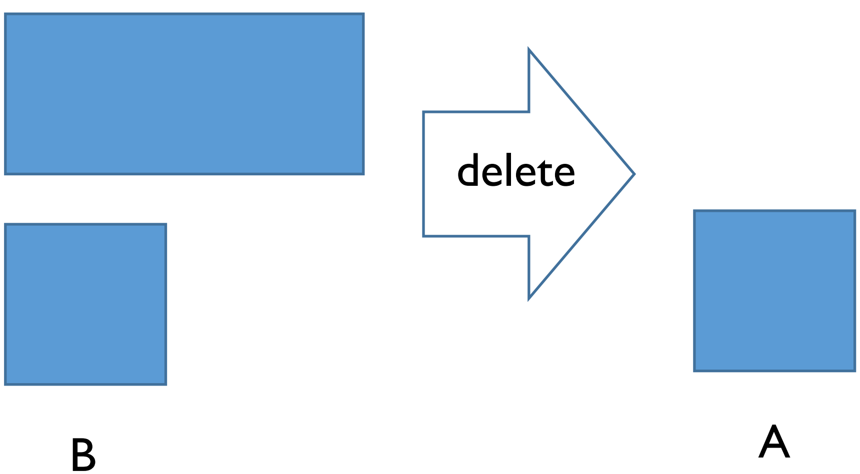

Few transformations
= very SIMILAR
= HARD to distinguish
= SLOW (large) reaction time

The four deductive reasoning scenarios (MP, MT, DA, AC)

Valid arguments:

(1) Modus ponens is when you "affirm the antecedent"….

If today is a Thursday, then it is a weekday

Today is a Thursday

Therefore today is a weekday

(2) Modus tollens is when you "deny the consequent"….

If today is a Thursday, then it is a weekday

Today is NOT a weekday

Therefore today is not Thursday

Invalid arguments:

(3) Affirmation of the consequent…

If today is a Monday, then it is a weekday Today is a weekday Therefore today is Monday

(4) Denial of the antecedent…

If today is a Monday, then it is a weekday Today is NOT a Monday Therefore today is not a weekday

Premise monotonicity vs premise non-monotonicity

Here is a colour I like

Do I like this colour?

Maybe??? Purple and red are a little bit similar
so it’s possible but hard to say for sure

Here is a colour I like

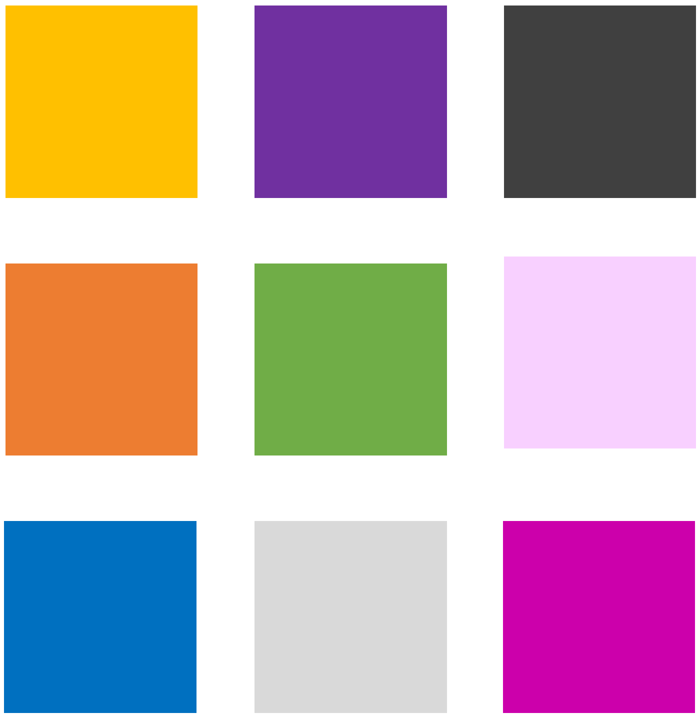

Do I like this colour?

I like all the colours?

Adding **<u>more examples</u>** of "things I like" **<u>increases</u>** your belief that I like "red"

PREMISE MONOTONICITY

Here is a colour I like

Do I like this colour?

Here is a colour I like

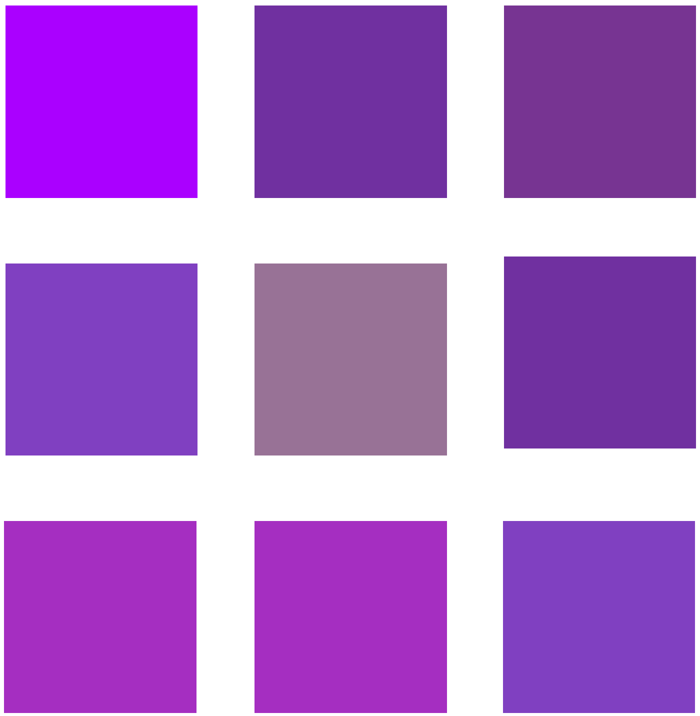

Do I like this colour?

Apparently I only like purple?

Adding **<u>more examples</u>** of "things I like" **<u>decreases</u>** your belief that I like "red"

PREMISE NON-MONOTONICITY

Category sampling and property sampling (tutorials)

Some parts of this explanation go beyond what was in the tutorials. The new content is <u>NOT</u> examinable

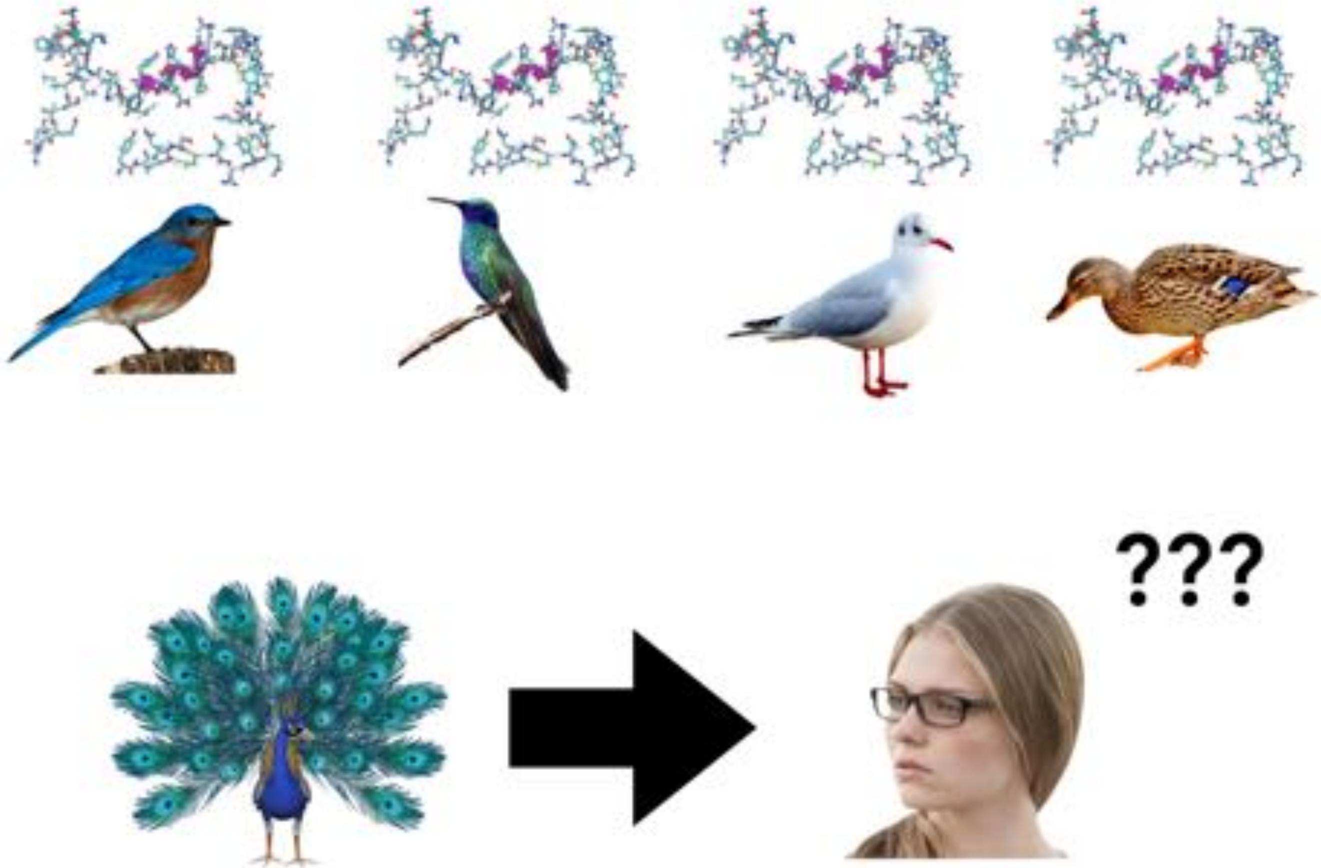

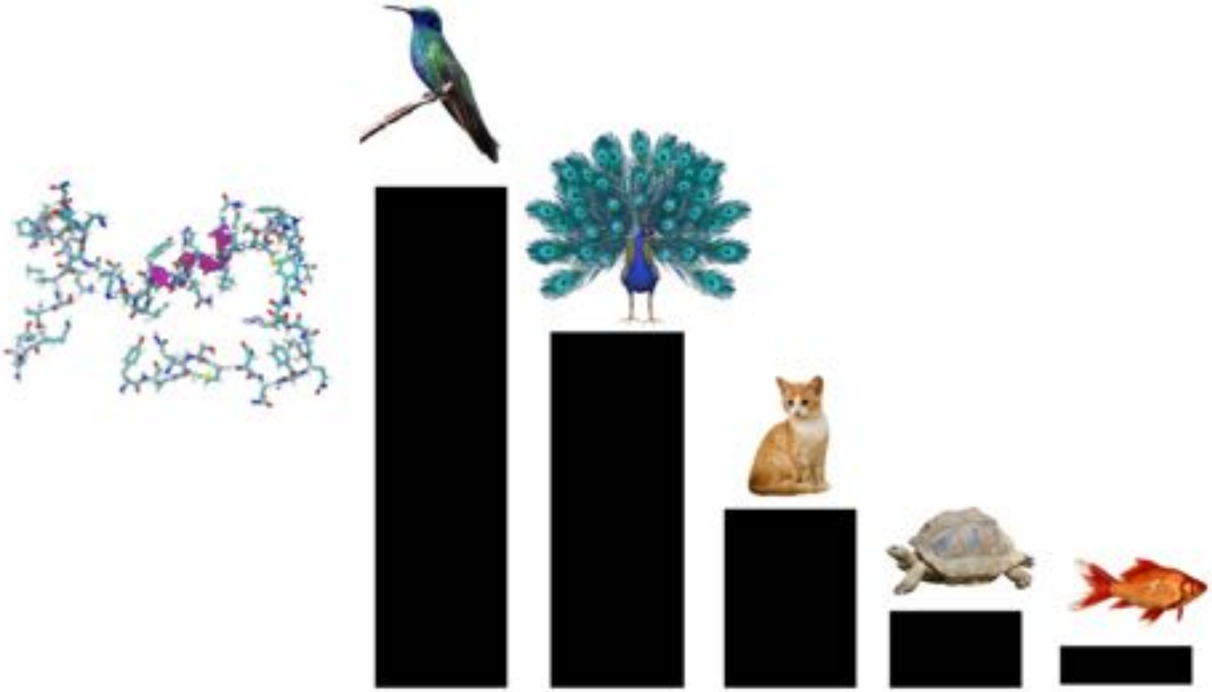

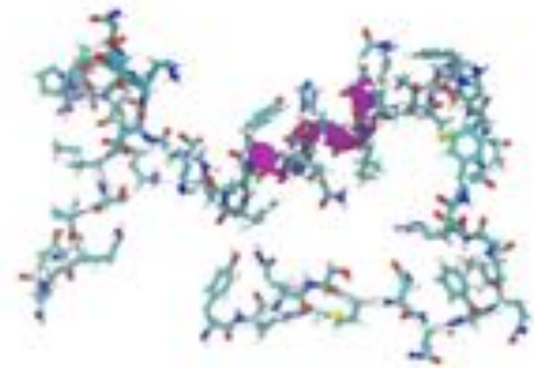

"Category sampling"

We **selected small birds** (e.g., because they could fit in the cage)

... when we tested them, it turned out that they all had plaxium blood

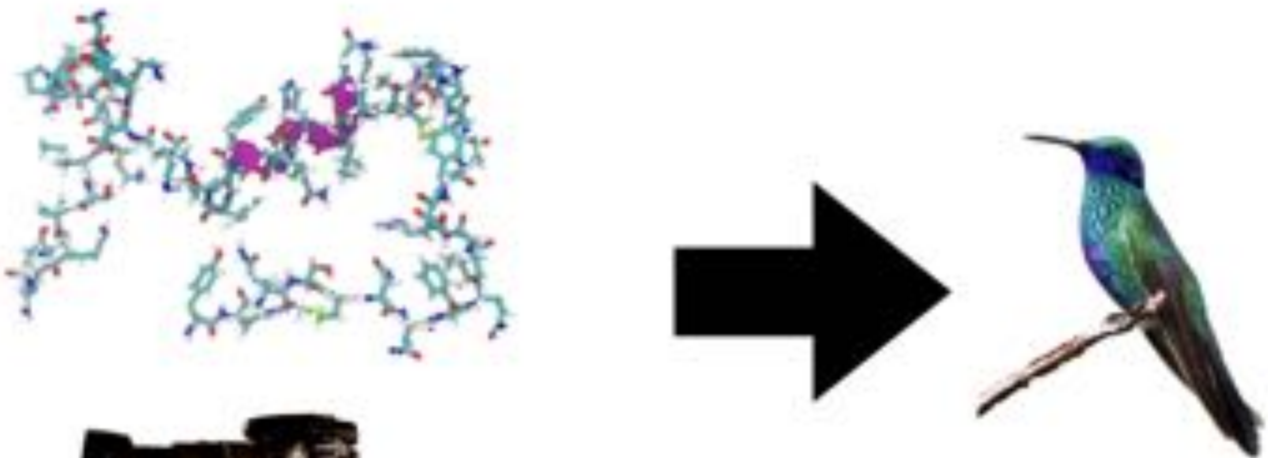

"Property sampling"

We detected **plaxium positive animals** (e.g., with a special camera)

... when we examined them, it turned out that they were all small birds

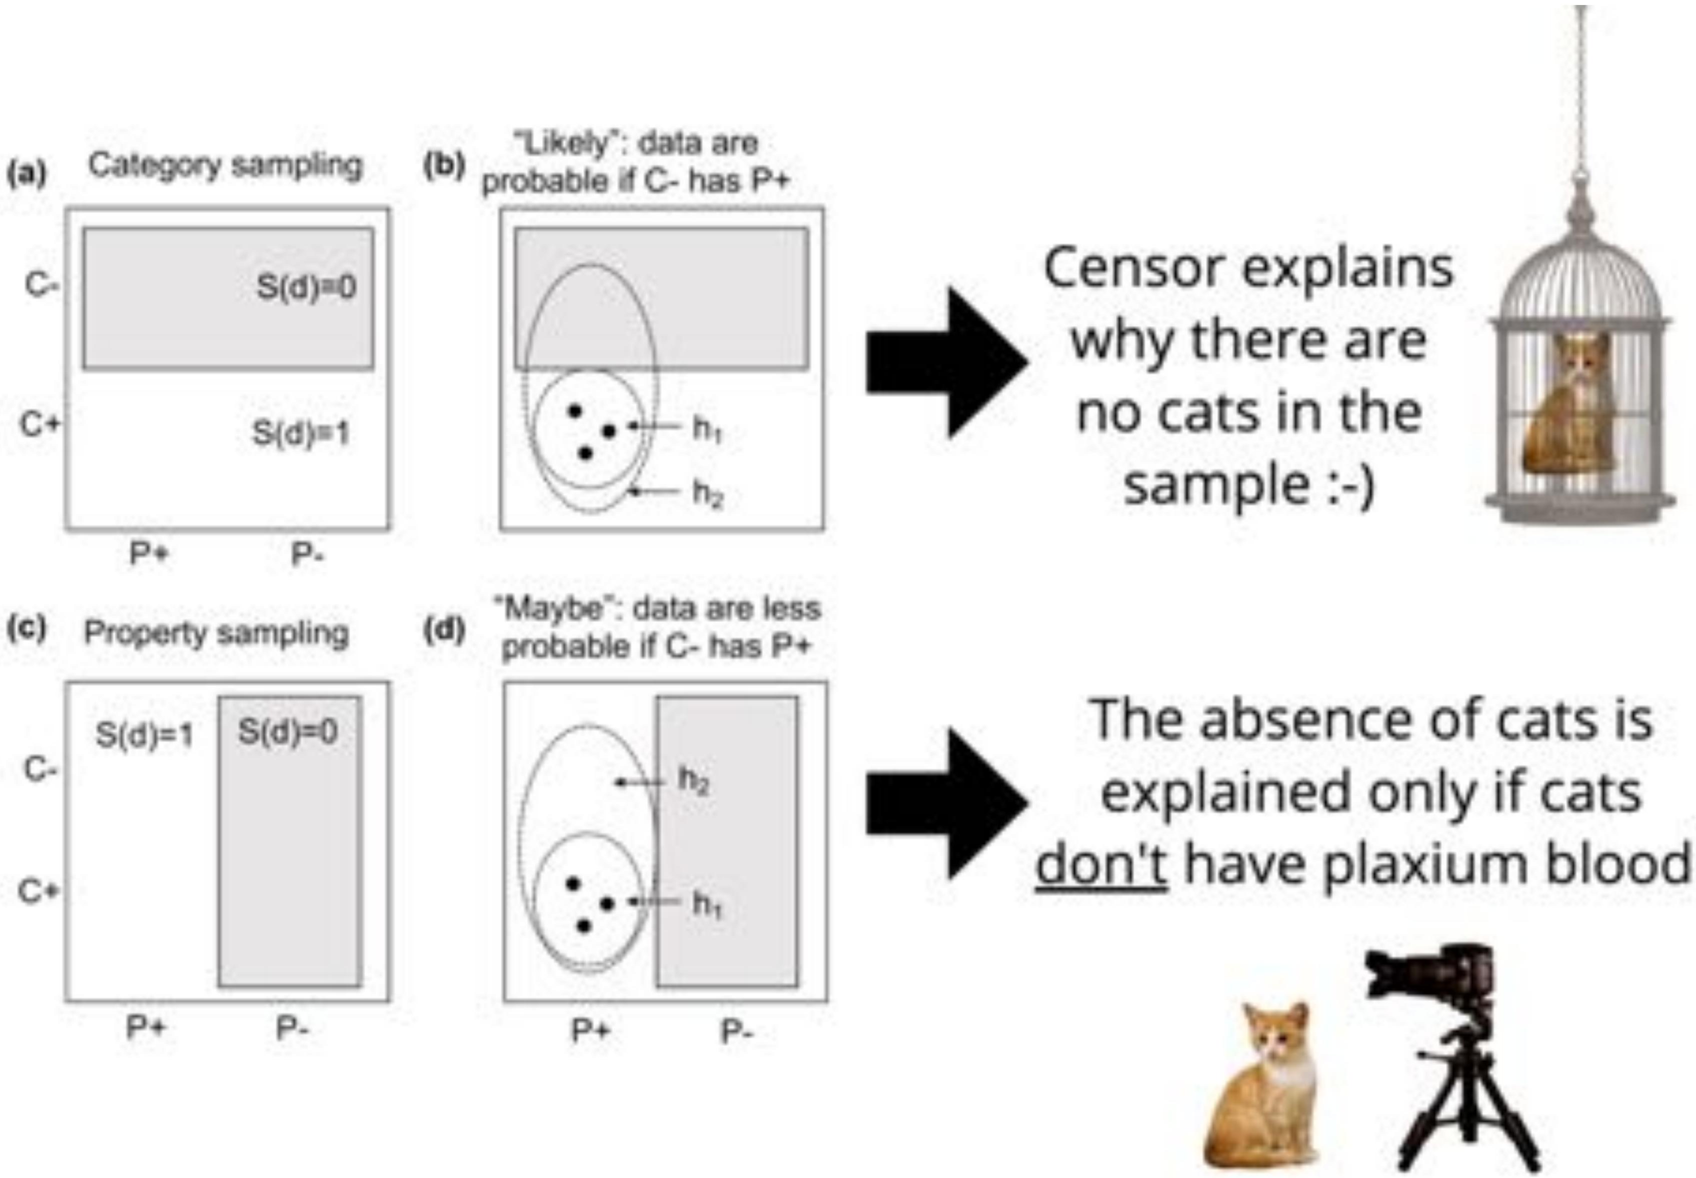
(a) Category sampling
C-
S(d)=0
C+
S(d)=1
P+
P-
(b) "Likely": data are probable if C- has P+
h_1
h_2
Censor explains why there are no cats in the sample :-)
(c) Property sampling
S(d)=1
S(d)=0
C-
C+
P+
P-
(d) "Maybe": data are less probable if C- has P+
h_2
h_1
P+
P-
The absence of cats is explained only if cats don't have plaxium blood

Core prediction

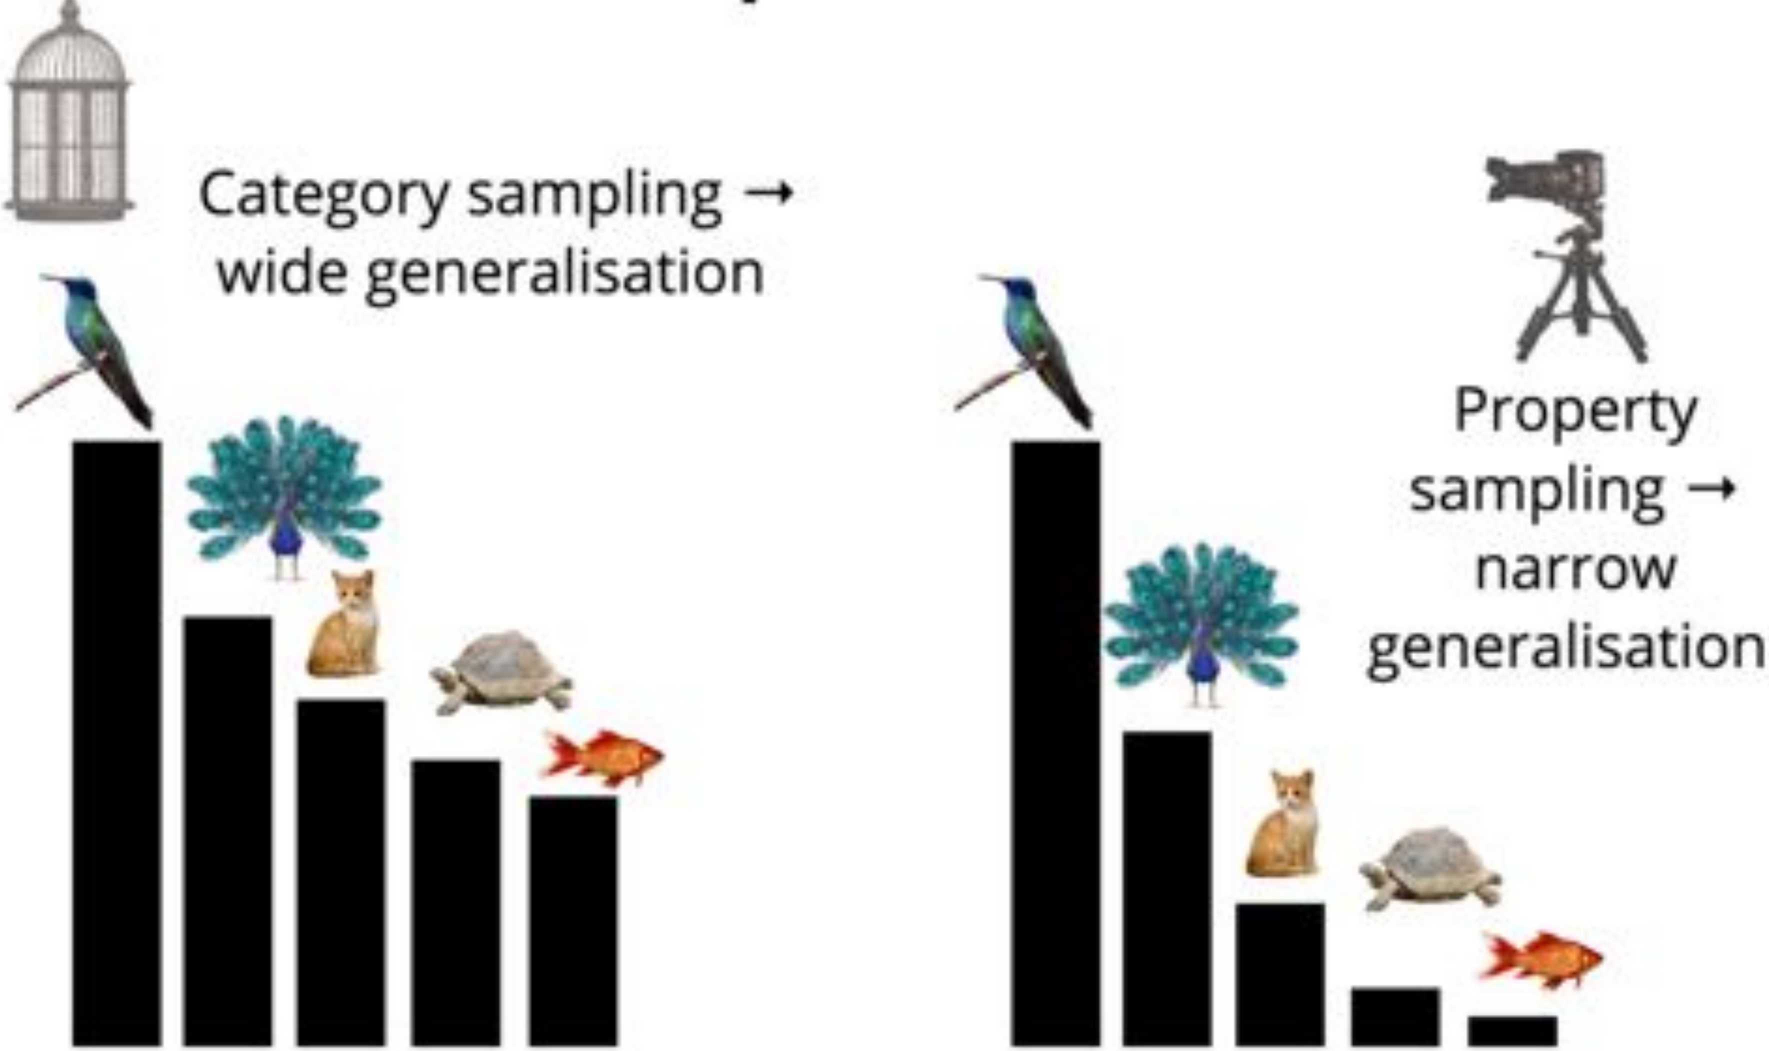

Does it work?

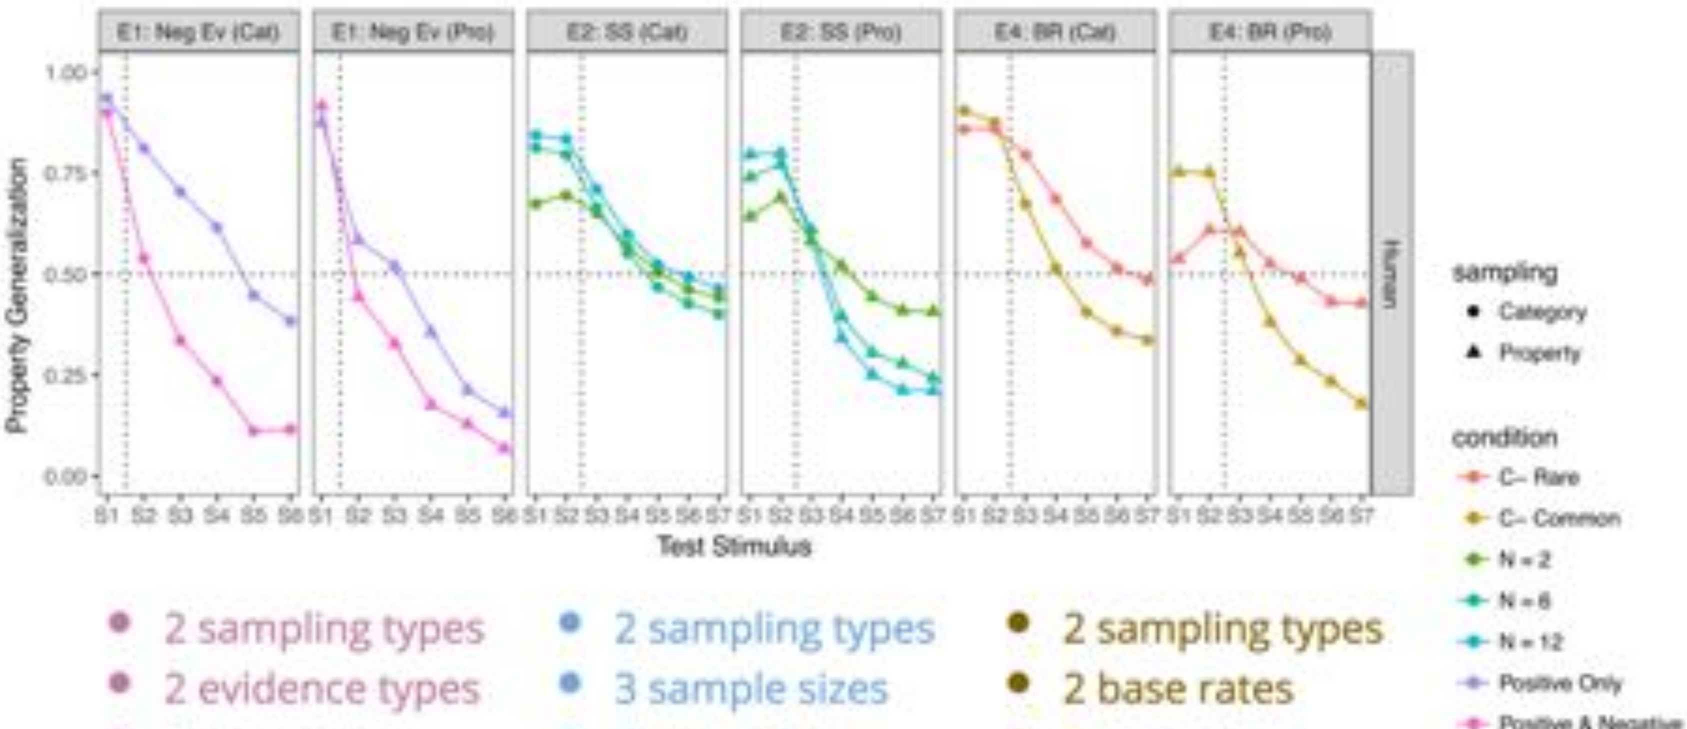

- 2 sampling types
- 2 evidence types
- 6 test points

- 2 sampling types
- 3 sample sizes
- 7 test points

- 2 sampling types
- 2 base rates
- 7 test points

Some fancy-pants modelling in which I show off…
BLAH BLAH BLAH… no-one cares ☺…

```
model {
  # mean and covariance matrix defining the Gaussian process
  for(i in 1:ncat) {
    mean_gp[i] <- m
    cov_gp[i,i] <- (sigma^2) + (tau^2)
    for(j in (i+1):ncat) {
      cov_gp[i,j] <- (tau^2) * exp(-rho * (test[i] - test[j])^2)
      cov_gp[j,i] <- cov_gp[i,j]
    }
  }

  # sample a function from the Gaussian process
  cov_gp_inv <- inverse(cov_gp)
  f ~ dmnorm(mean_gp, cov_gp_inv)

  # pass f through logit function to get a function on (0,1)
  for(i in 1:ncat) {
    phi[i] <- 1/(1+exp(-f[i]))
  }

  #
  # [SNIP]
  #
}
```

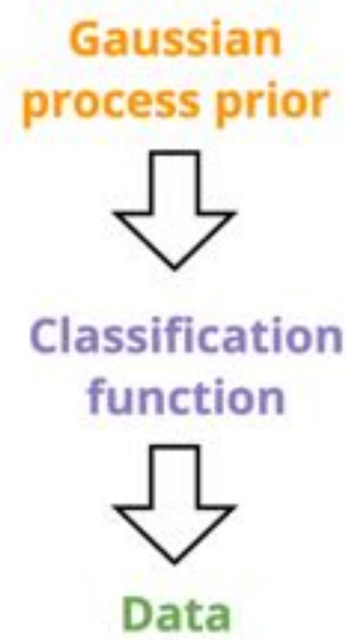

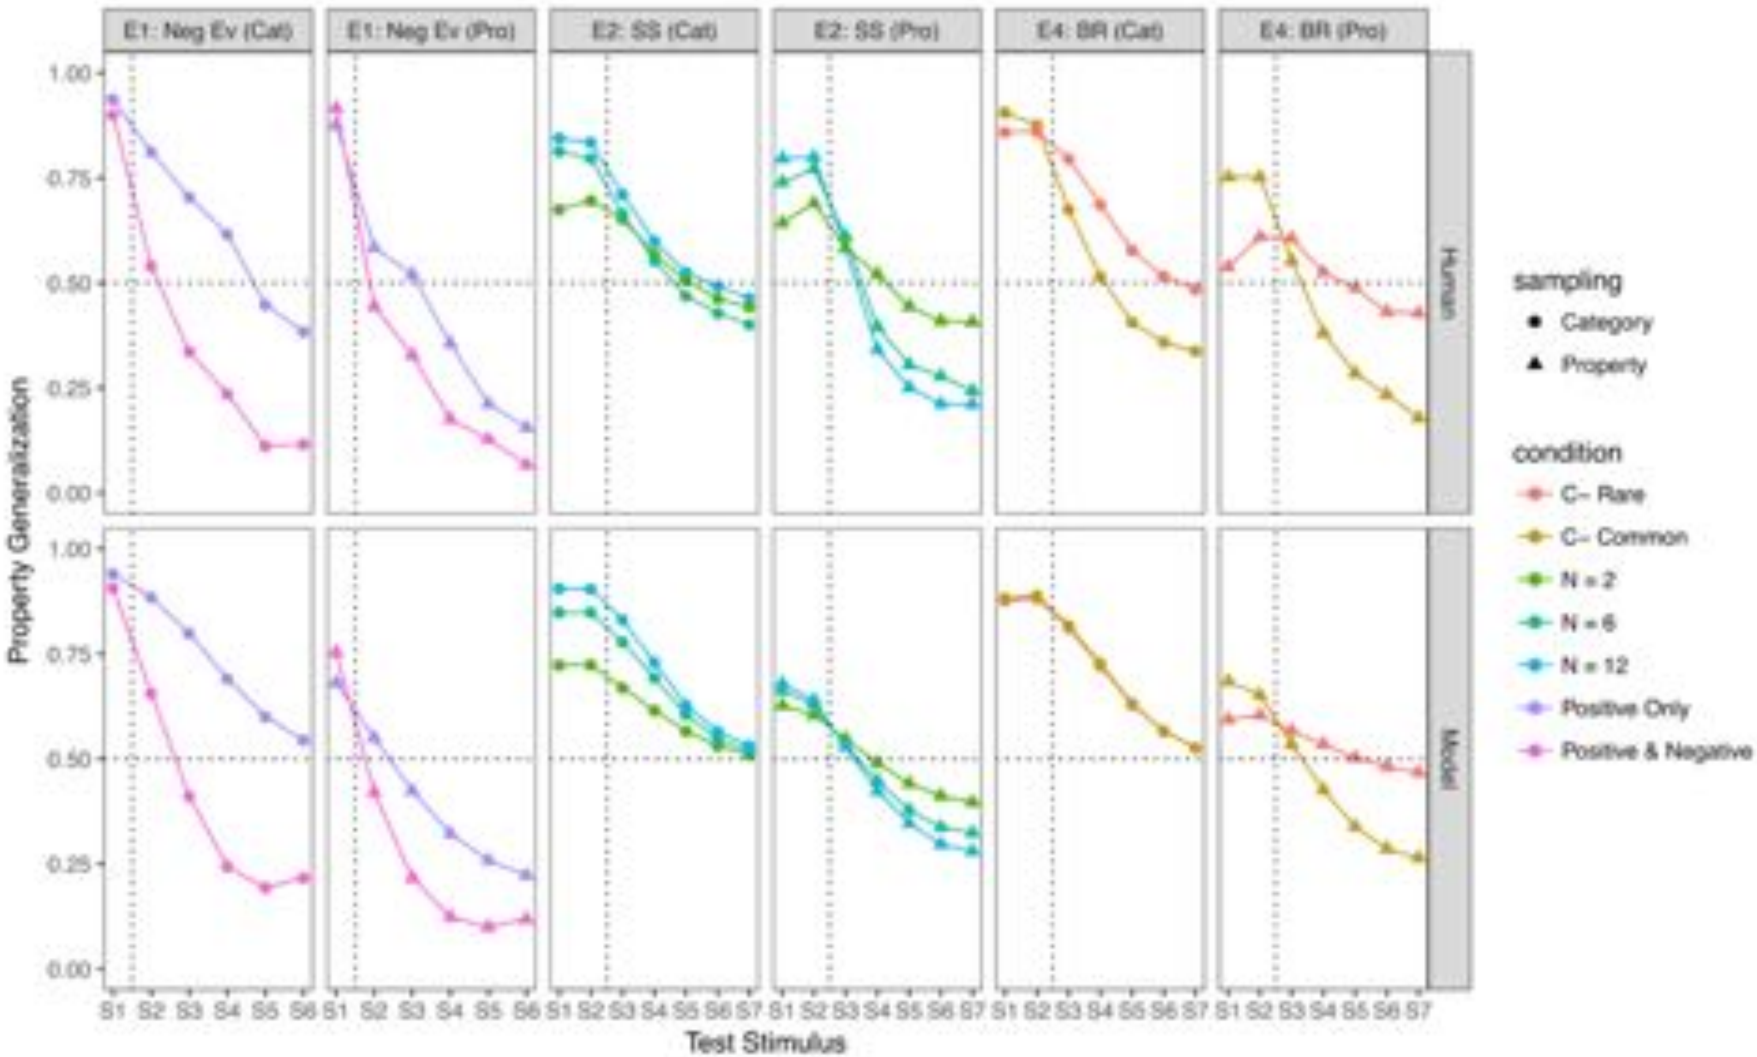
E1: Neg Ev (Cat)
E1: Neg Ev (Pro)
E2: SS (Cat)
E2: SS (Pro)
E4: BR (Cat)
E4: BR (Pro)
Human
Model
Property Generalization
1.00
0.75
0.50
0.25
0.00
S1 S2 S3 S4 S5 S6
S1 S2 S3 S4 S5 S6 S7
Test Stimulus
sampling
Category
Property
condition
C– Rare
C– Common
N = 2
N = 6
N = 12
Positive Only
Positive & Negative

“Do computational models in cognitive science and neuroscience *really* help us build intelligent machines?

(surprisingly, yes!)

Hm… I wonder what Google are up to these days?

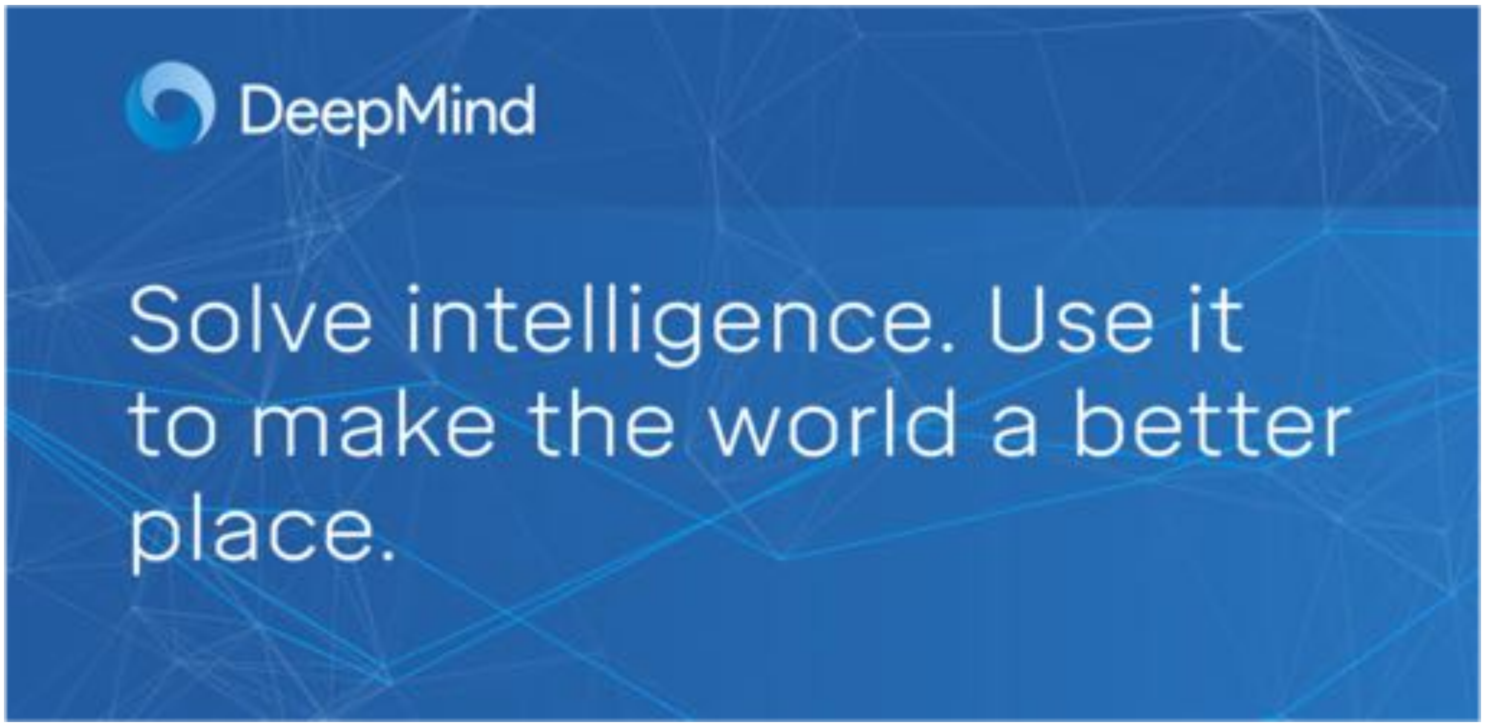

Oh, okay, teaching machines to play Atari games using… reinforcement learning

Better than human-level control of classic Atari games through Deep Reinforcement Learning.

LETTER

doi:10.1038/nature14236

Human-level control through deep reinforcement learning

Volodymyr Mnih[1*], Koray Kavukcuoglu[1*], David Silver[1*], Andrei A. Rusu[1], Joel Veness[1], Marc G. Bellemare[1], Alex Graves[1], Martin Riedmiller[1], Andreas K. Fidjeland[1], Georg Ostrovski[1], Stig Petersen[1], Charles Beattie[1], Amir Sadik[1], Ioannis Antonoglou[1], Helen King[1], Dharshan Kumaran[1], Daan Wierstra[1], Shane Legg[1] & Demis Hassabis[1]

The theory of reinforcement learning provides a normative account[1], deeply rooted in psychological[2] and neuroscientific[3] perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems[4,5], the former evidenced by a wealth of neural data revealing notable parallels between the phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms[3]. While reinforcement learning agents have achieved some successes in a variety of domains[6–8], their applicability has previously been limited to domains in which useful features can be handcrafted, or to domains with fully observed, low-dimensional state spaces. Here we use recent advances in training deep neural networks[9–11] to develop a novel artificial agent, termed a deep Q-network, that can learn successful policies directly from high-dimensional sensory inputs using end-to-end reinforcement learning. We tested this agent on the challenging domain of classic Atari 2600 games[12]. We demonstrate that the deep Q-network agent, receiving only the pixels and the game score as inputs, was able to surpass the performance of all previous algorithms and achieve a level comparable to that of a professional human games tester across a set of 49 games, using the same algorithm, network architecture and hyperparameters. This work bridges the divide between high-dimensional sensory inputs and actions, resulting in the first artificial agent that is capable of learning to excel at a diverse array of challenging tasks.

agent is to select actions in a fashion that maximizes cumulative future reward. More formally, we use a deep convolutional neural network to approximate the optimal action-value function

$$Q^*(s,a) = \max_{\pi} \mathbb{E}\left[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \ldots \mid s_t = s, a_t = a, \pi\right],$$

which is the maximum sum of rewards r_t discounted by γ at each time-step t, achievable by a behaviour policy $\pi = P(a|s)$, after making an observation (s) and taking an action (a) (see Methods)[19].

Reinforcement learning is known to be unstable or even to diverge when a nonlinear function approximator such as a neural network is used to represent the action-value (also known as Q) function[20]. This instability has several causes: the correlations present in the sequence of observations, the fact that small updates to Q may significantly change the policy and therefore change the data distribution, and the correlations between the action-values (Q) and the target values $r + \gamma \max_{a'} Q(s', a')$. We address these instabilities with a novel variant of Q-learning, which uses two key ideas. First, we used a biologically inspired mechanism termed experience replay[21–23] that randomizes over the data, thereby removing correlations in the observation sequence and smoothing over changes in the data distribution (see below for details). Second, we used an iterative update that adjusts the action-values (Q) towards target values that are only periodically updated, thereby reducing correlations with the target.

While other stable methods exist for training neural networks in the reinforcement learning setting, such as neural fitted Q-iteration[24], these methods involve the repeated training of networks *de novo* on hundreds of iterations. Consequently, these methods, unlike our algorithm, are too inefficient to be used successfully with large neural networks. We parameterize an approximate value function $Q(s,a;\theta_i)$ using the deep convolutional neural network shown in Fig. 1, in which θ_i are the param-

Better than human-level control of classic Atari games through Deep Reinforcement Learning.

https://storage.googleapis.com/deepmind-media/dqn/DQNNaturePaper.pdf

The theory of reinforcement learning provides a normative account, deeply rooted in psychological and neuroscientific perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations

Better than human-level control of classic Atari games through Deep Reinforcement Learning.

Don't try to defeat an AI at space invaders

You'll do much better at Ms Pac-Man

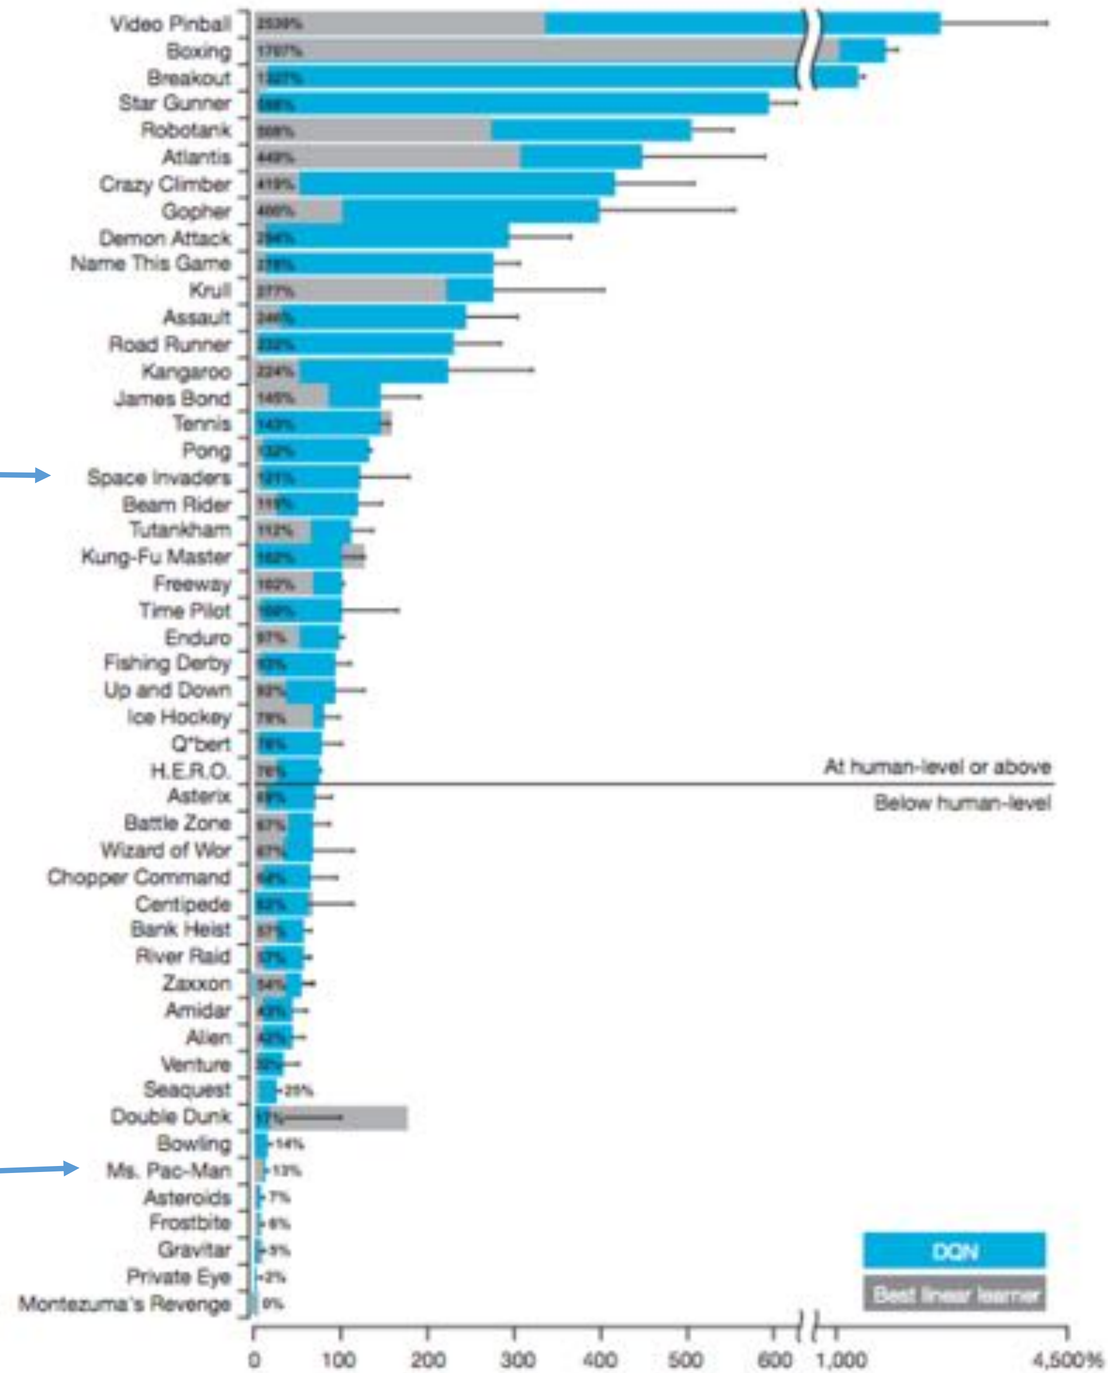

Cognitive science researchers interested in working out why humans are better than AI at some games, and worse at others…

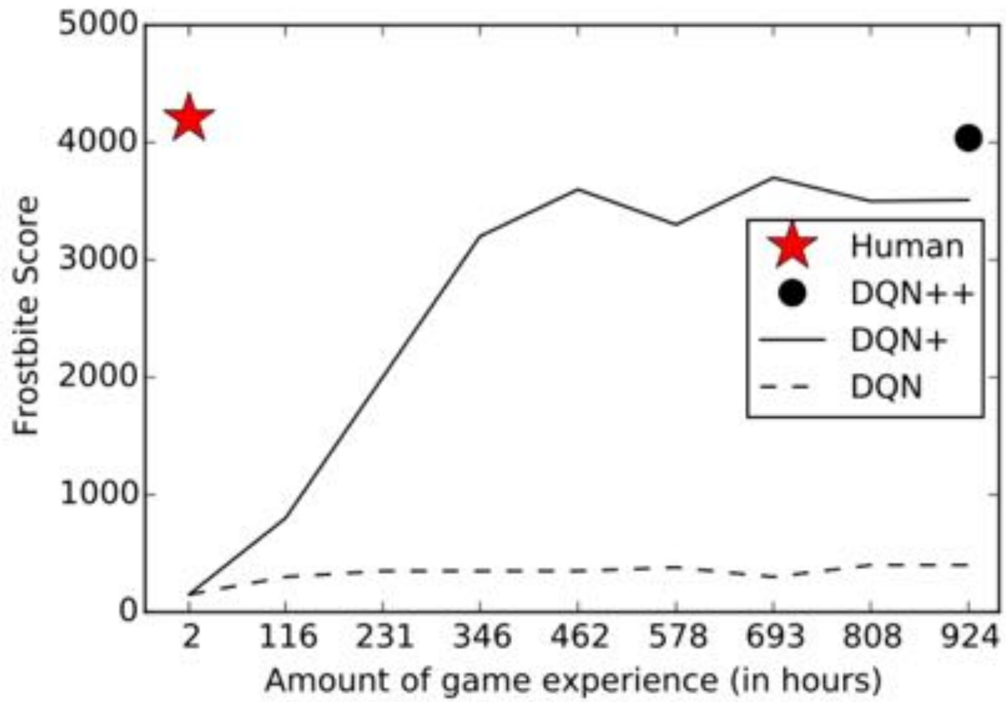

http://web.stanford.edu/class/psych209/Readings/LakeEtAlBBS.pdf

Cognitive science researchers interested in working out why humans are better than AI at some games, and worse at others…

a woman riding a horse on a dirt road

an airplane is parked on the tarmac at an airport

a group of people standing on top of a beach

Figure 6: Perceiving scenes without intuitive physics, intuitive psychology, compositionality, and causality. Image captions are generated by a deep neural network (Karpathy & Fei-Fei, 2015) using code from `github.com/karpathy/neuraltalk2`. Image credits: Gabriel Villena Fernández (left), TVBS Taiwan / Agence France-Presse (middle) and AP Photo / Dave Martin (right). Similar examples using images from Reuters news can be found at `twitter.com/interesting_jpg`.

http://web.stanford.edu/class/psych209/Readings/LakeEtAlBBS.pdf

Any other questions???

A/Prof Danielle Navarro
d.navarro@unsw.edu.au
compcogscisydney.org

Learning, reasoning, induction, decision making, computational modelling, statistics