Computational Cognitive Science

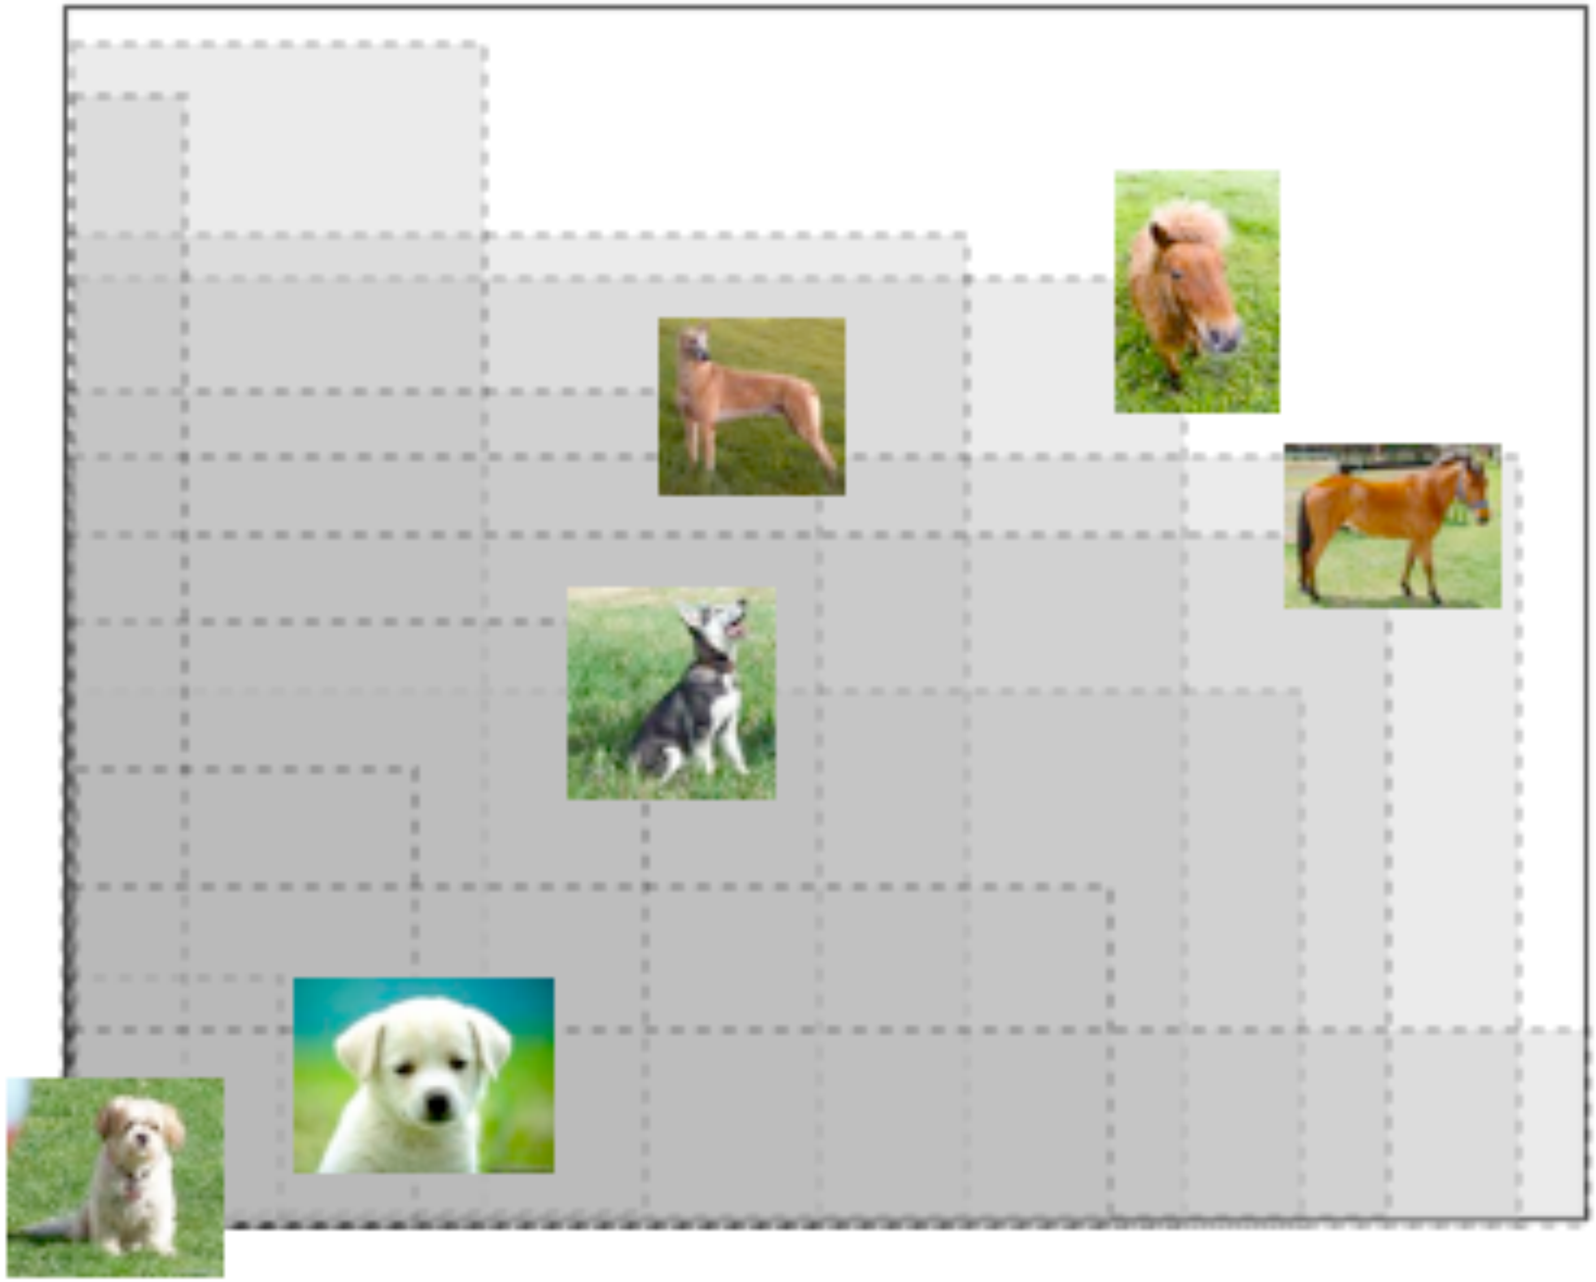

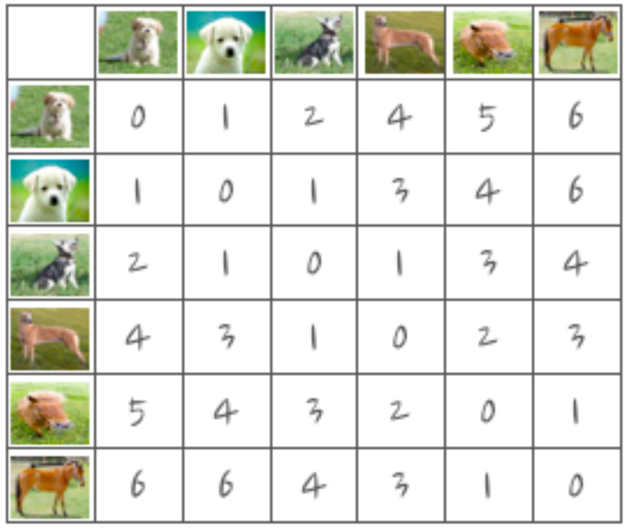

	0	1	2	4	5	6
	1	0	1	3	4	6
	2	1	0	1	3	4
	4	3	1	0	2	3
	5	4	3	2	0	1
	6	6	4	3	1	0

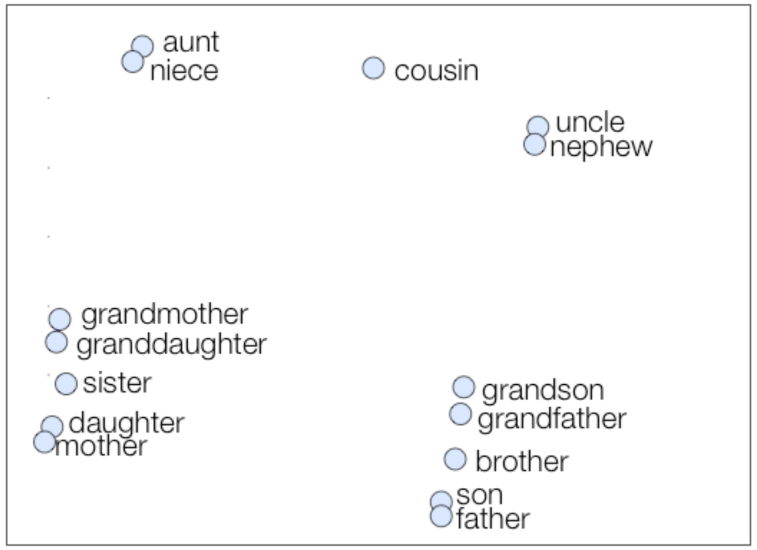

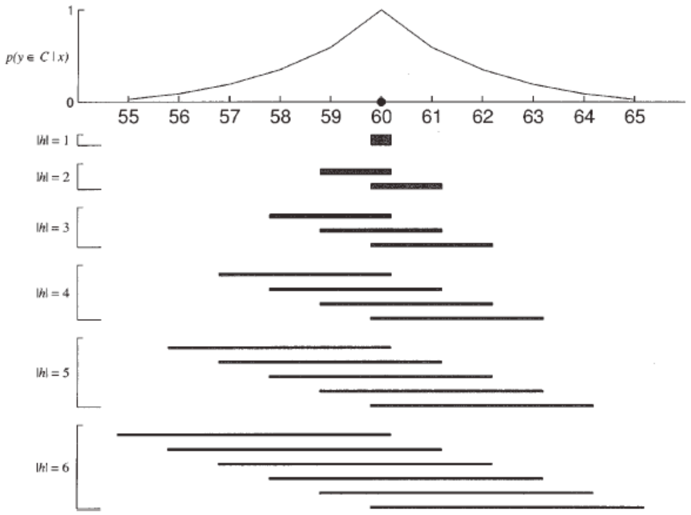

Lecture 4: Generalisation

Lecture outline

- The problem of generalisation
  - Defining the problem, and deriving a solution
  - Tangent: Multi-dimensional scaling
- Extending our solution to the problem of generalisation
  - A new framework
  - Application to more abstract representations
- Next: Other kinds of inductive generalisation

Lecture outline

- ➡ The problem of generalisation
  - Defining the problem, and deriving a solution
  - Tangent: Multi-dimensional scaling
- Extending our solution to the problem of generalisation
  - A new framework
  - Application to more abstract representations
- Next: Other kinds of inductive generalisation

A simple problem

Which of these things are the same?

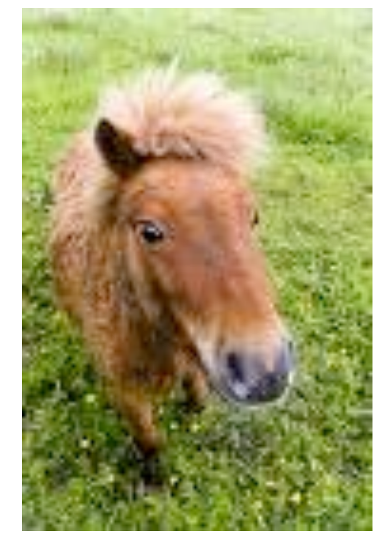

This is the problem of generalisation

You've seen one thing, and you need to figure out what other things in the world are "like" it in some way

- Note that this is not a *logical* or *deductive* problem: logically, everything could be, or everything could not be
- This is an *inductive* problem: you must go from limited data, just one or a few examples, to arrive at a general conclusion

This is the problem of generalisation

Most of the problems humans have to deal with are inductive problems.

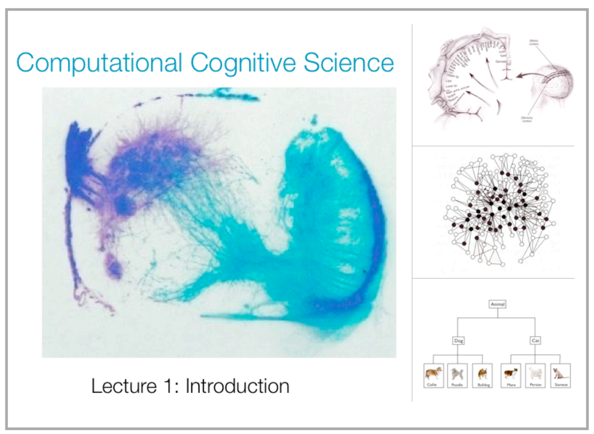

Will I like future lectures of this class?

Is this strange food edible and good?

Is my baby developing normally?

The problem of generalisation

How do people decide to generalise from one data point to another?

- Intuitively: the more similar it is, the more likely you are to say it is "the same kind of thing"
  - Vague: What is the functional form of the generalisation gradient?
  - Non-explanatory: Why would it be sensible to generalise in this way?

Can we formalise the problem of generalisation in such a way as to avoid these issues?

Defining the problem of generalisation

- Assume items can be located in some psychological "similarity space"
  - Things that are closer in that space are more similar
  - Later, we'll talk about how to capture this; for now assume it

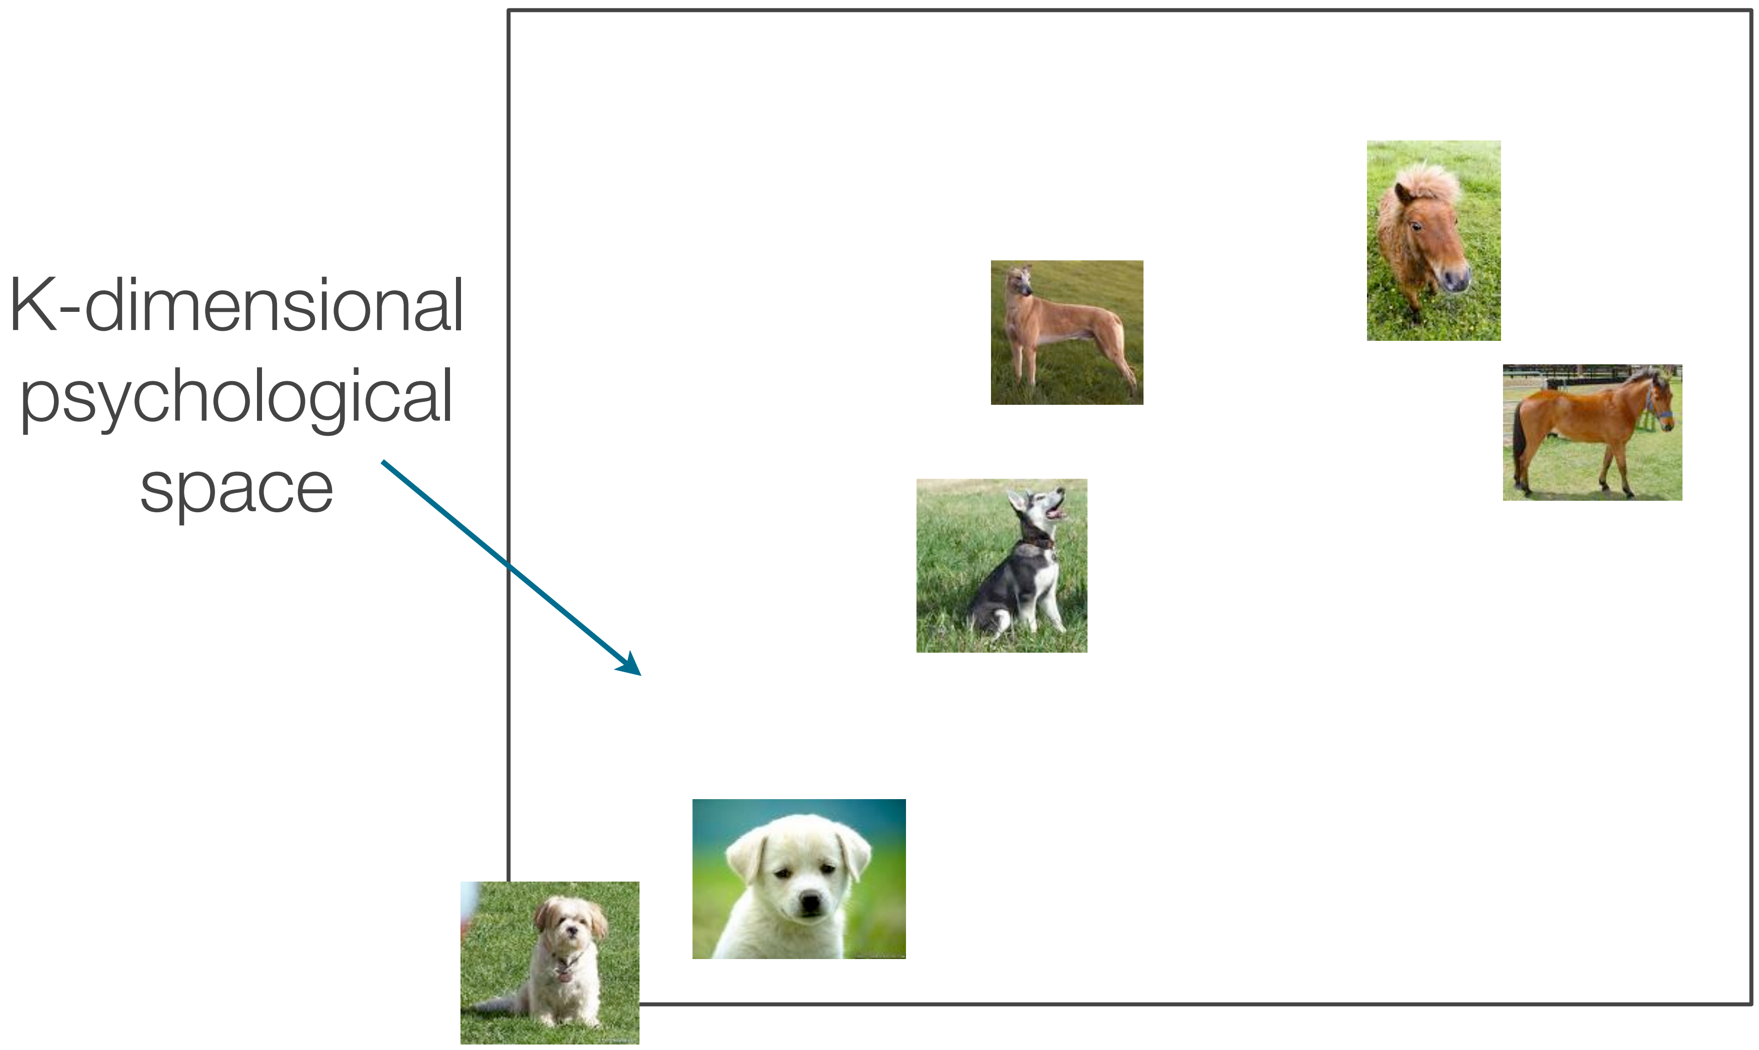

Defining the problem of generalisation

- Assume items can be located in some psychological "similarity space"
- Generalisation is the problem of identifying the *region* in that space where things are "the same kind of thing"

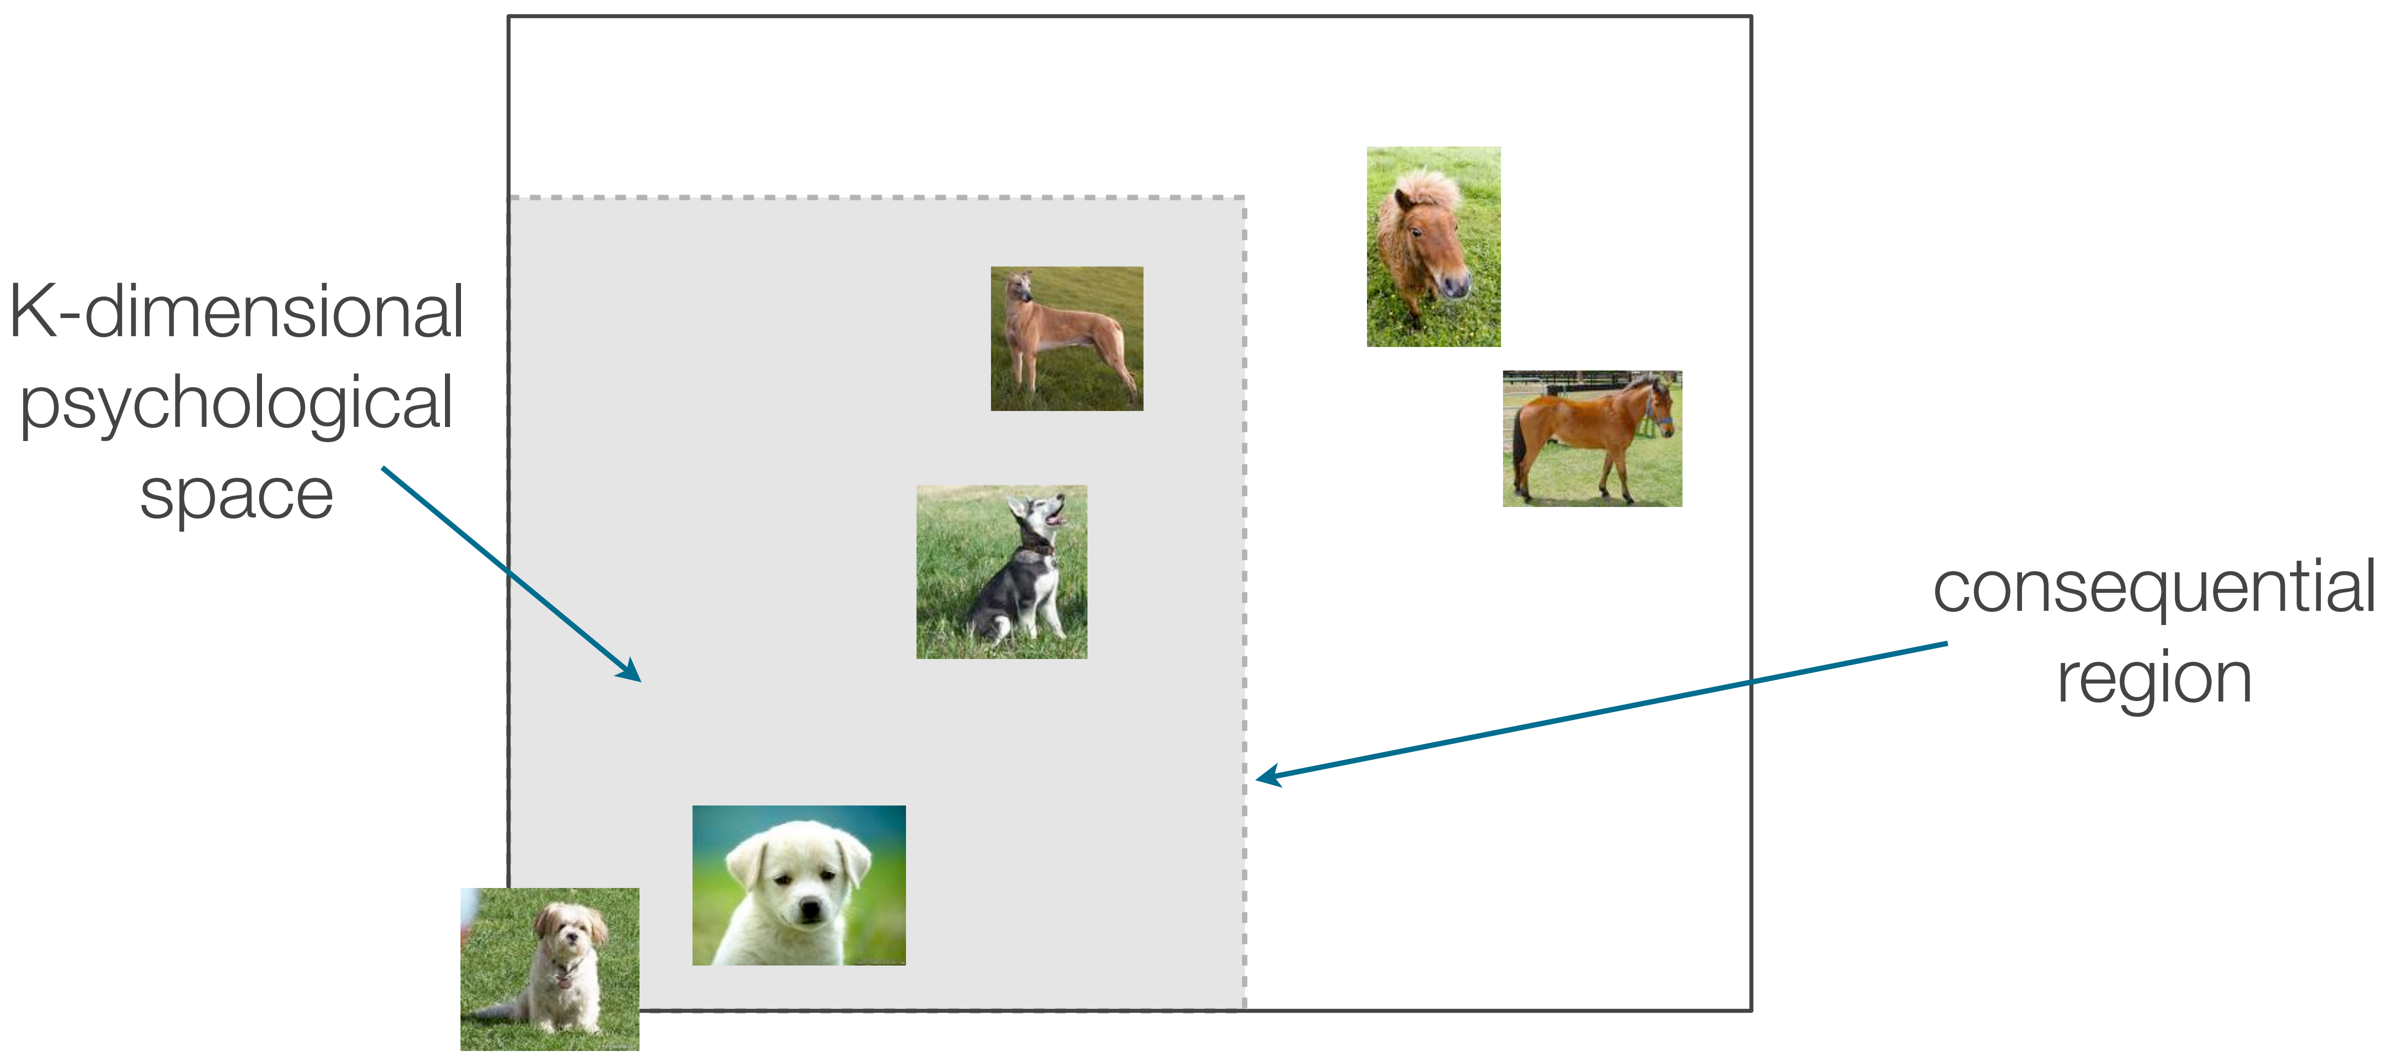

Two questions follow

- What is the consequential region?
- What is the probability that some point x lies within the consequential region?

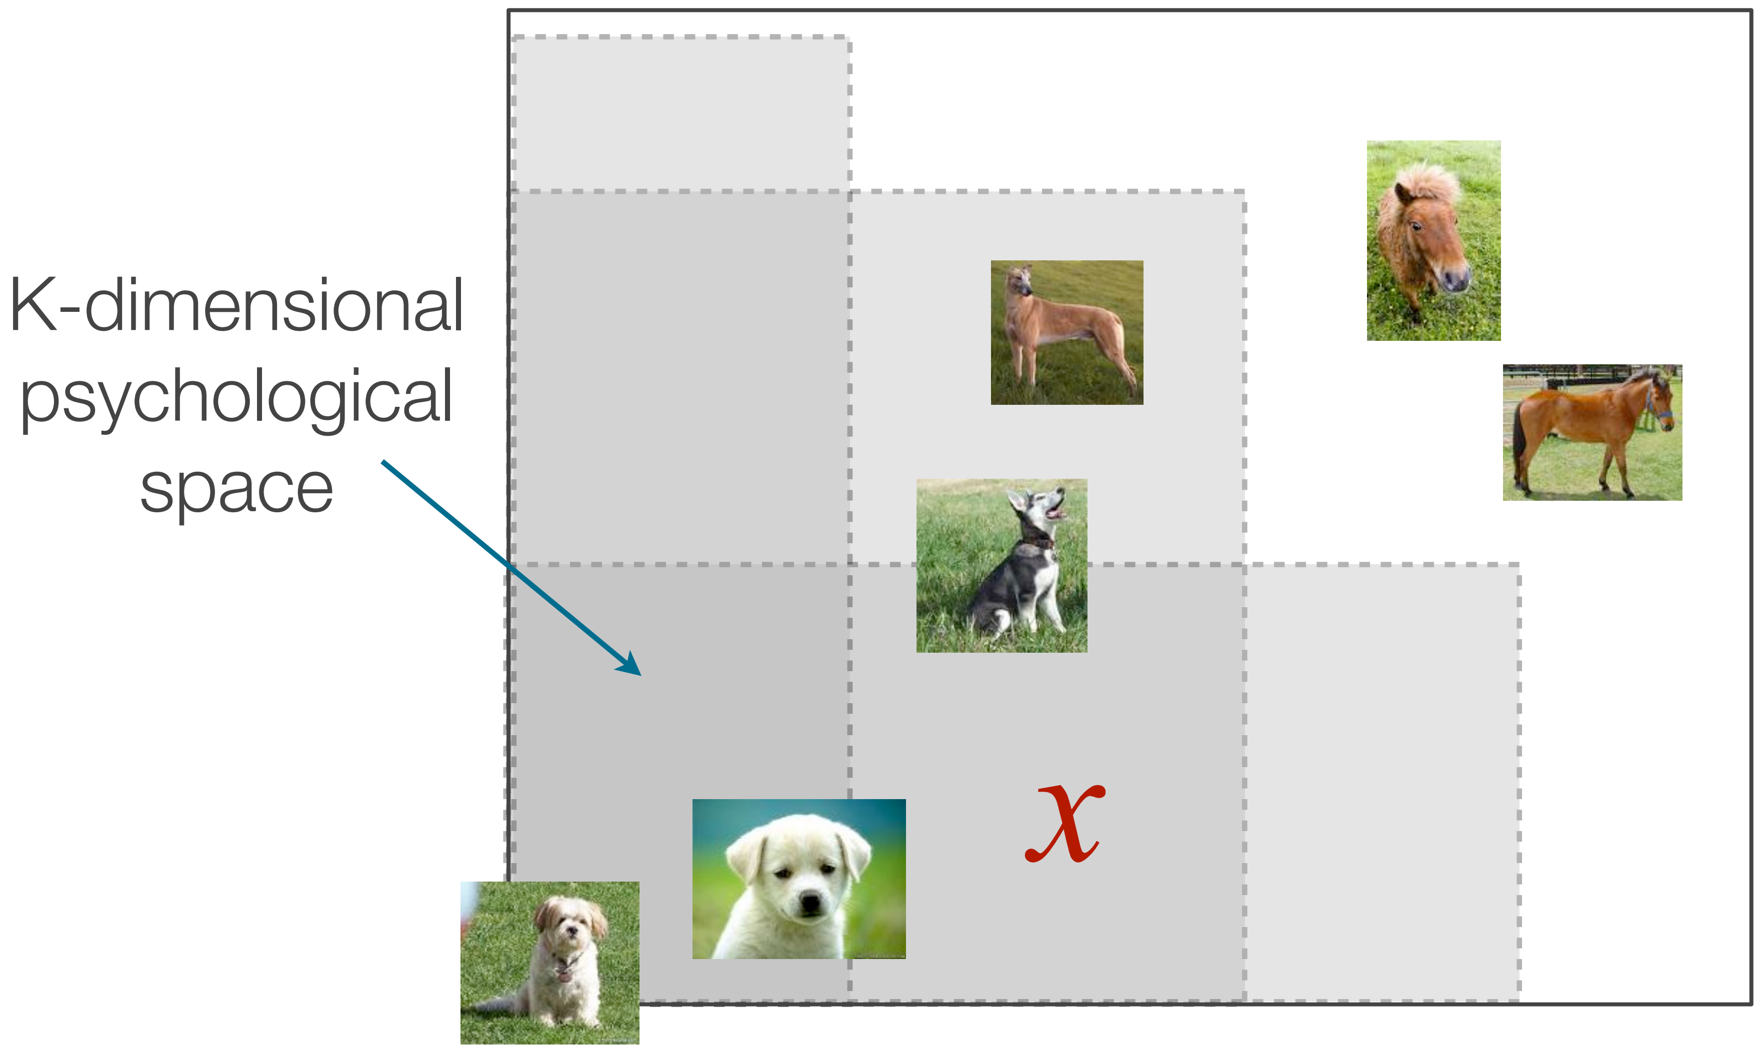

Our solution

- We can answer both of these questions by integrating over all possible consequential regions
- Good intuitions - but how do we put numbers to them?

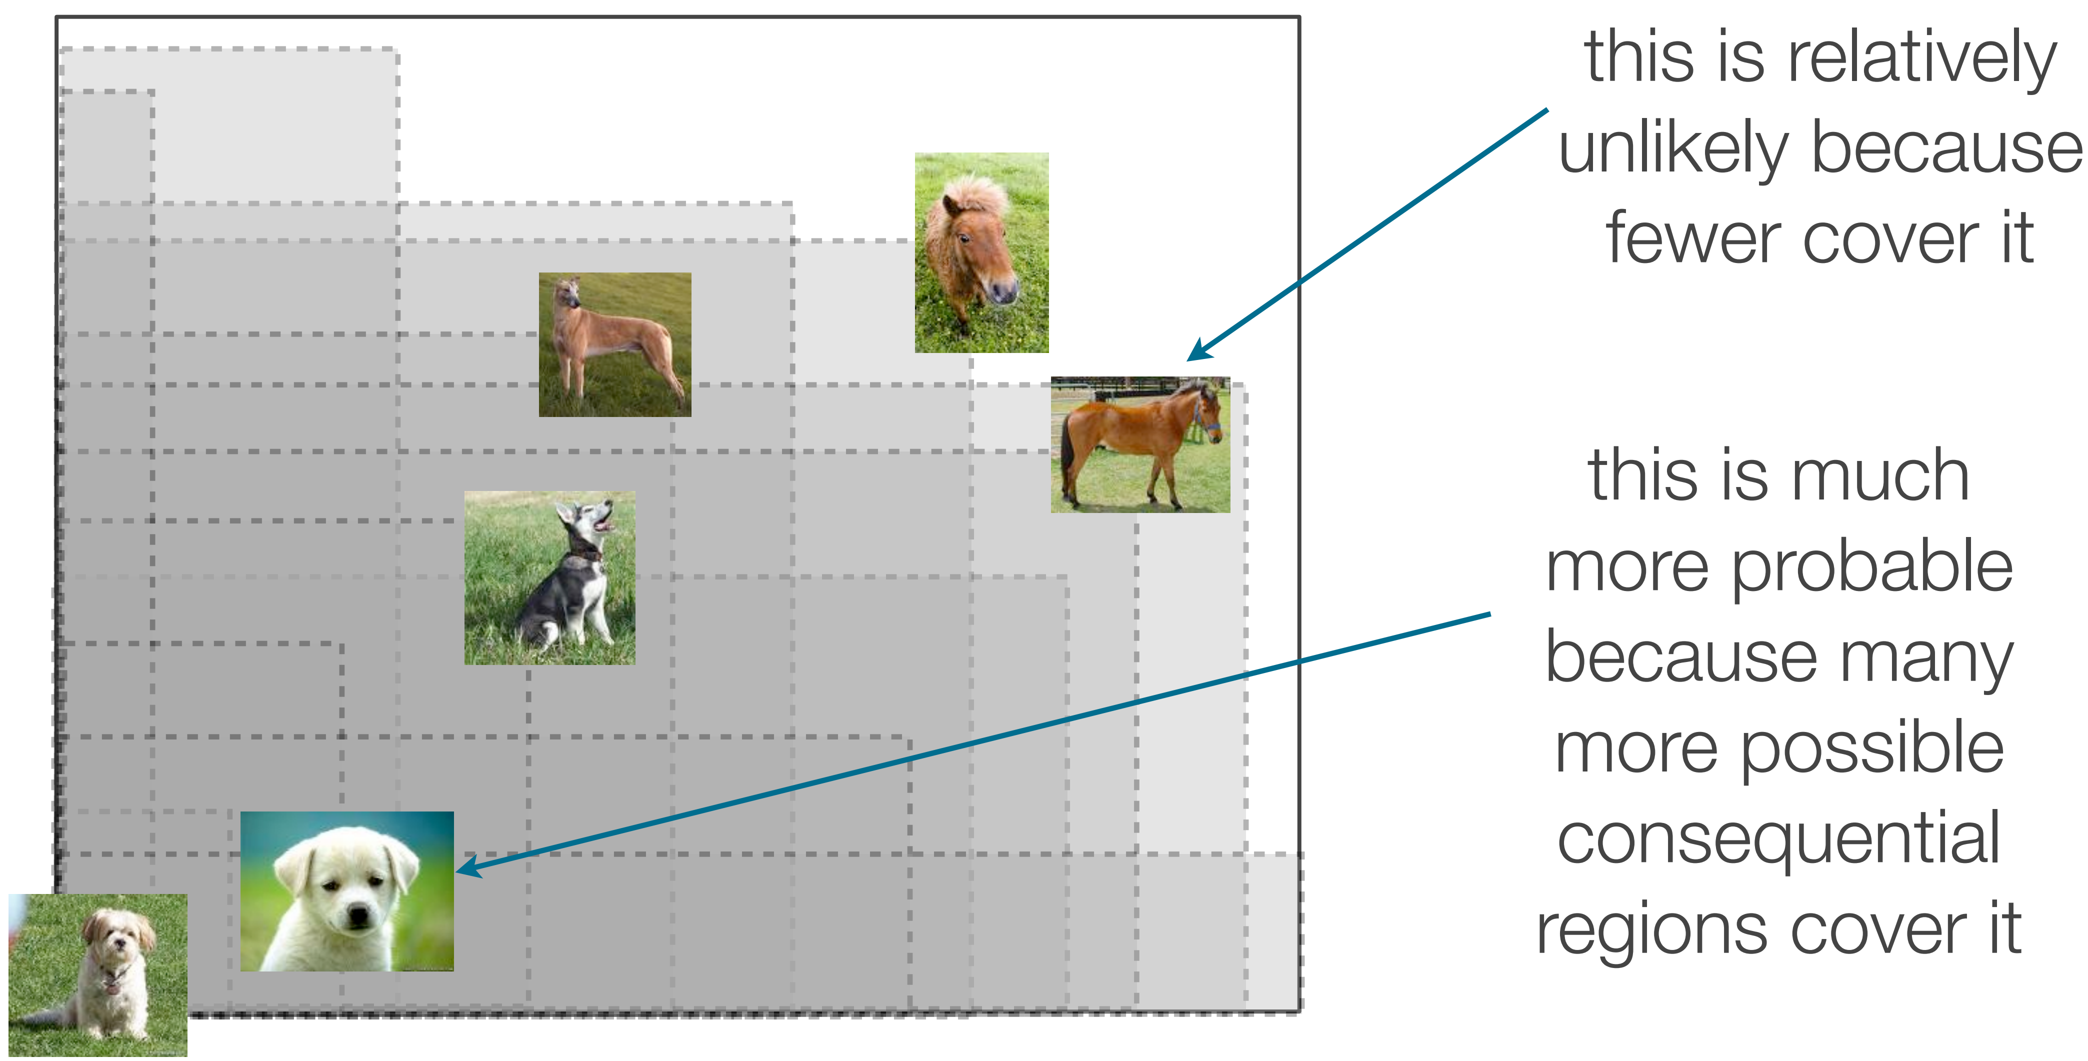

Our solution

- For some *single* consequential region of size s, the conditional probability that x is contained in the region is just the ratio $m(s,x)/m(s)$ of the measure of the overlap to the measure of the whole such region

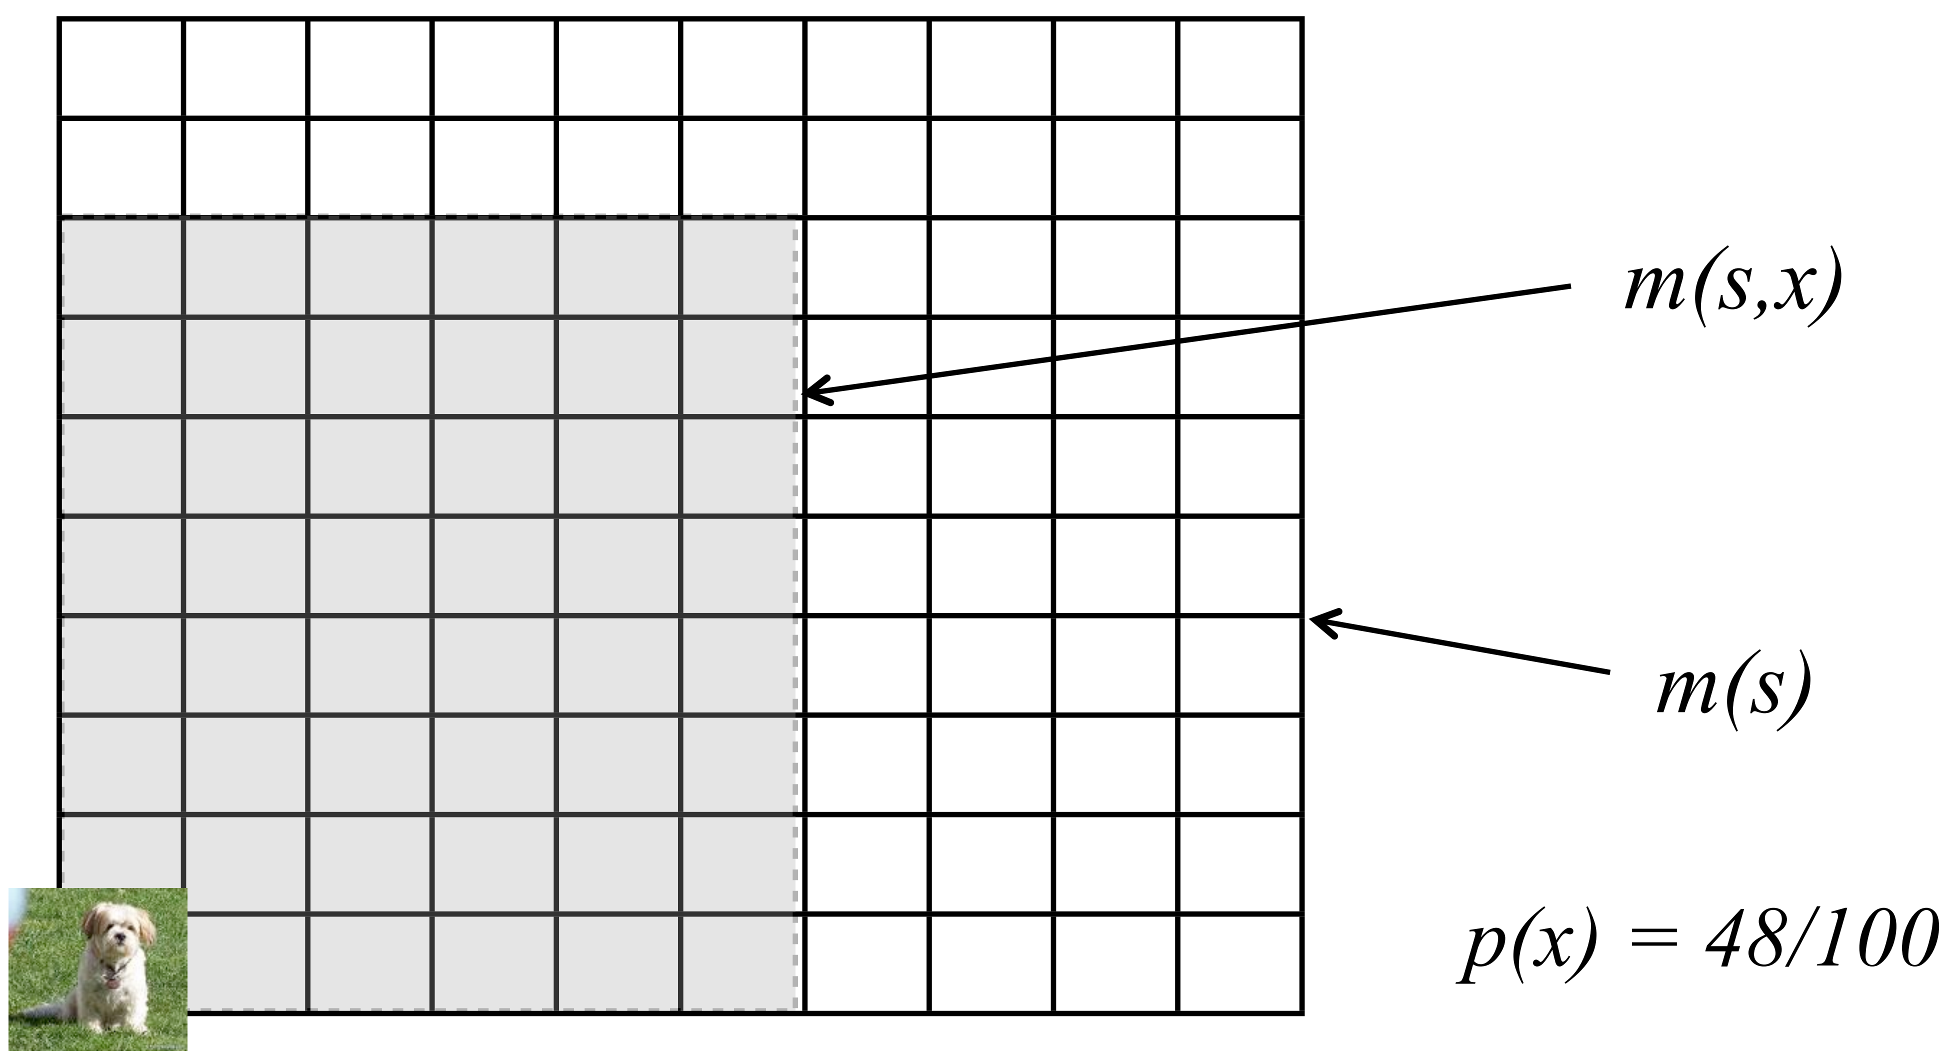

Our solution

- We calculate the probability that x is contained in any arbitrary region by integrating over all possible regions

$$g(\mathbf{x}) = \int_0^\infty p(s) \frac{m(s,\mathbf{x})}{m(s)} ds$$

$g(x)$ = probability that a response learned to stimulus 0 will generalise to x

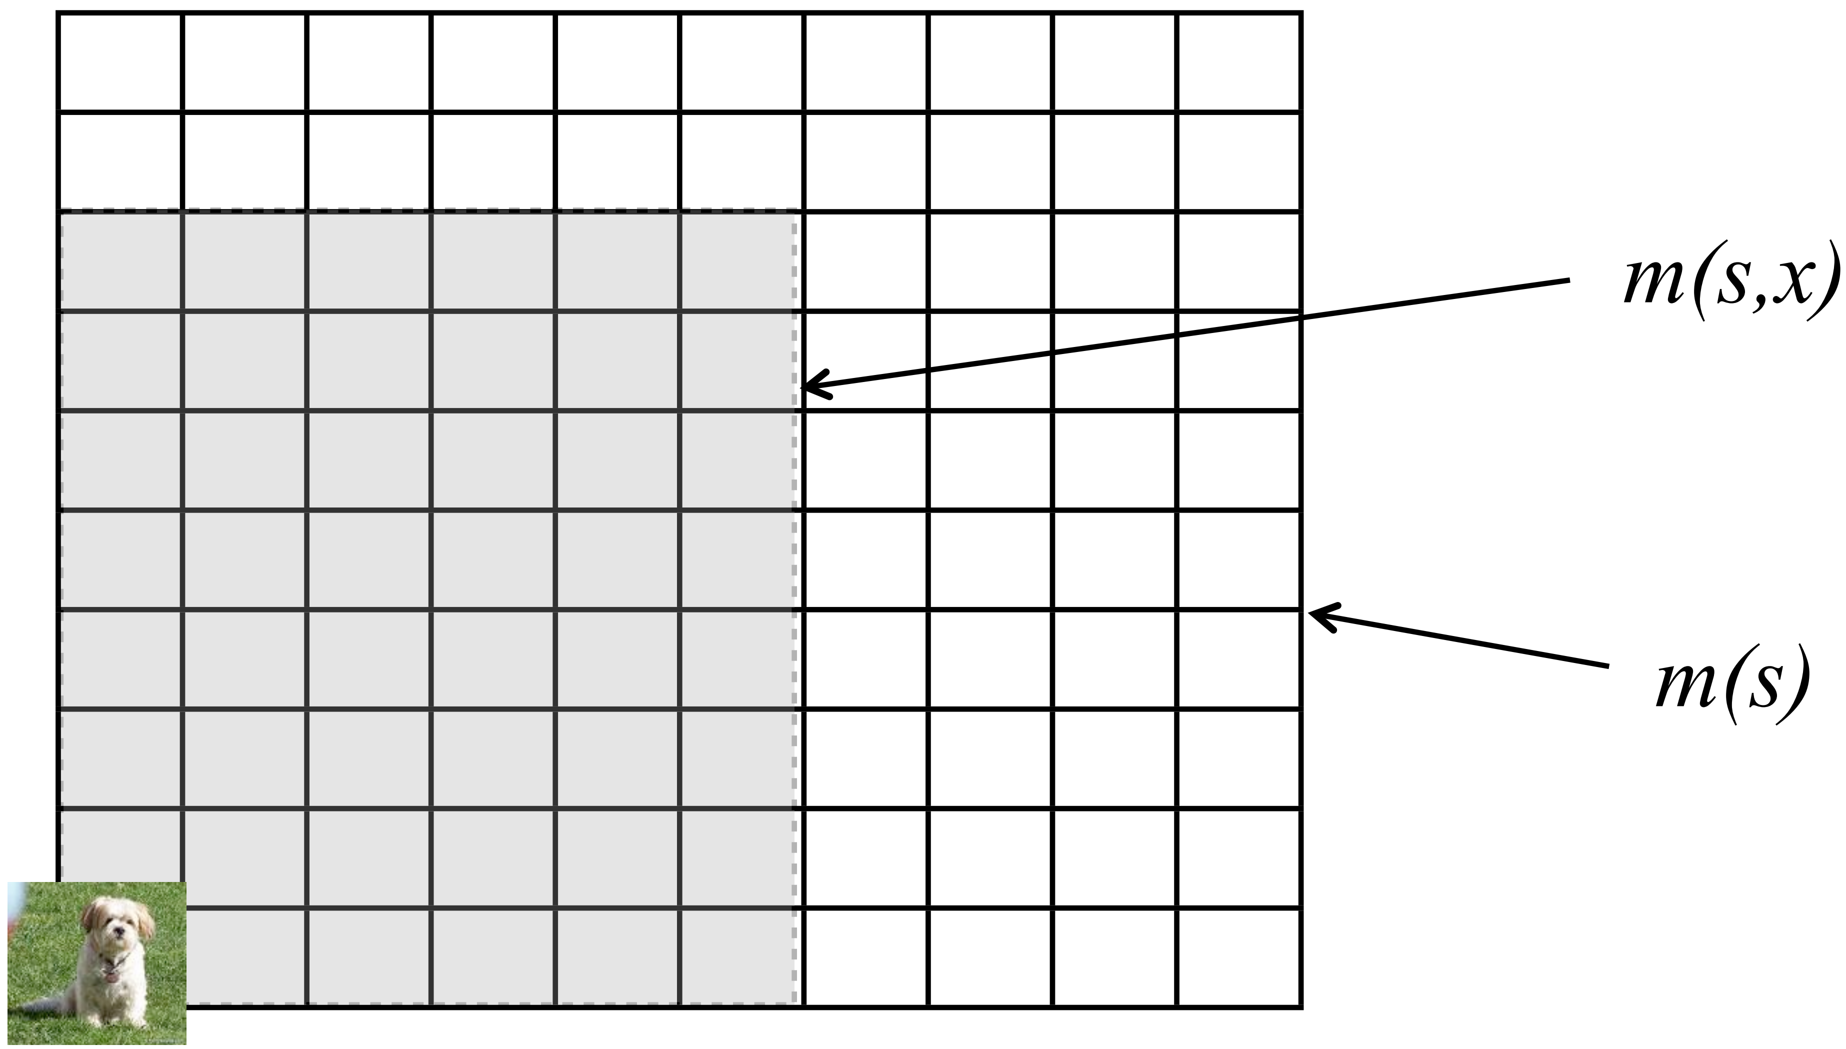

Our solution

- We calculate the probability that x is contained in any arbitrary region by integrating over all possible regions

$$g(\mathbf{x}) = \int_0^\infty p(s)\,\frac{m(s,\mathbf{x})}{m(s)}\,ds$$

$g(x)$ = probability that a response learned to stimulus 0 will generalise to x

- After some math, we find that generalisation is a function of the distance d in the psychological space in which:

$$g(d) = \int_d^\infty p(s)\,ds - d\int_d^\infty \frac{p(s)}{s}\,ds$$

Has unit value at $d=0$

$$g'(d) = -\int_d^\infty \frac{p(s)}{s}\,ds$$

Monotonically decreases with increasing d

$$g''(d) = \frac{p(d)}{d}$$

Is concave upward (mostly)

Our solution

- What does this mean for the *actual* shape of the generalisation gradient?

$$g(\mathbf{x}) = \int_0^\infty p(s) \frac{m(s,\mathbf{x})}{m(s)} ds$$

That depends on the nature of $p(s)$

Regardless of what $p(s)$ is, the resulting generalisation curve looks pretty much exponential

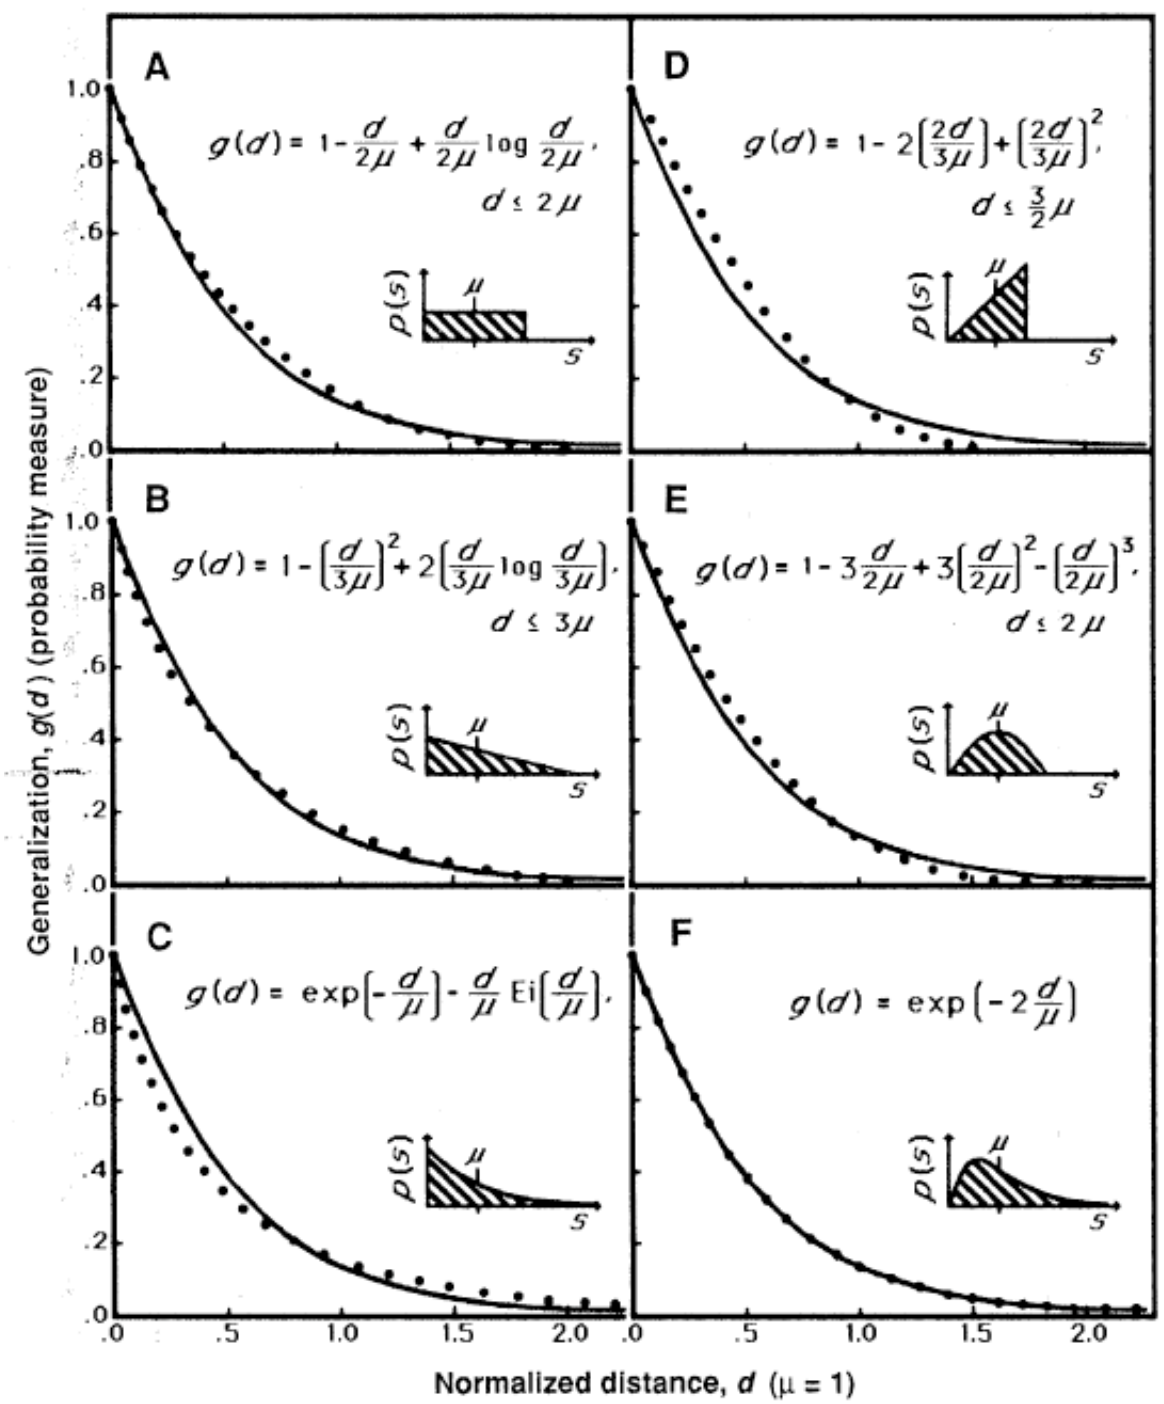

Prediction

- This predicts that generalisation should always be an exponentially decreasing function in psychological space

To test: We need a way to measure psychological space

Idea: Have people rate the similarity of items in the space

On a scale of 0 to 6, how similar are these two items to each other? (0 = very similar; 6 = not at all)

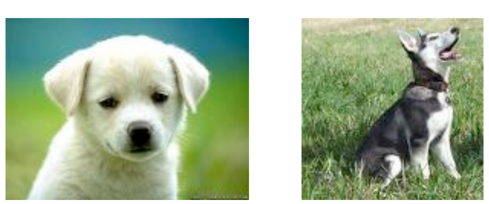

= 4

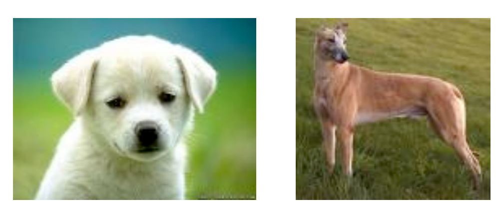

= 5

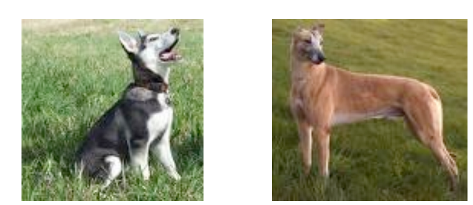

= 3

Prediction

- Based on the reported similarities, one can derive a *distance matrix* of all distances between all objects

Distance matrix Δ where $\delta_{i,j}$ is the distance between objects i and j

$$\Delta := \begin{pmatrix} \delta_{1,1} & \delta_{1,2} & \cdots & \delta_{1,I} \\ \delta_{2,1} & \delta_{2,2} & \cdots & \delta_{2,I} \\ \vdots & \vdots & & \vdots \\ \delta_{I,1} & \delta_{I,2} & \cdots & \delta_{I,I} \end{pmatrix}.$$

						
	0	1	2	4	5	6
	1	0	1	3	4	6
	2	1	0	1	3	4
	4	3	1	0	2	3
	5	4	3	2	0	1
	6	6	4	3	1	0

- This is the psychological space we assume!

Characteristics of the psychological space

- Can we take this matrix of similarities and use it to visualise where the items are?
- Use a technique called *Multi-Dimensional Scaling (MDS)*

Basically uses numerical optimisation to find the points in a k-dimensional space that preserves the distances as well as possible - i.e., that minimises a function like the following

$$\min_{x_1,\ldots,x_I} \sum_{i<j} (\|x_i - x_j\| - \delta_{i,j})^2$$

In R: `cmdscale(d,eigs=TRUE)`

distance matrix (d)

returns eigenvalues for each dimension (eigs=TRUE)

```
> deltas
     [,1] [,2] [,3] [,4] [,5] [,6]
[1,]    0    1    2    4    5    6
[2,]    1    0    1    3    4    6
[3,]    2    1    0    1    3    4
[4,]    4    3    1    0    2    3
[5,]    5    4    3    2    0    1
[6,]    6    6    4    3    1    0
> mds <- cmdscale(d=deltas,eig=TRUE)
```

Result of MDS

- k=1

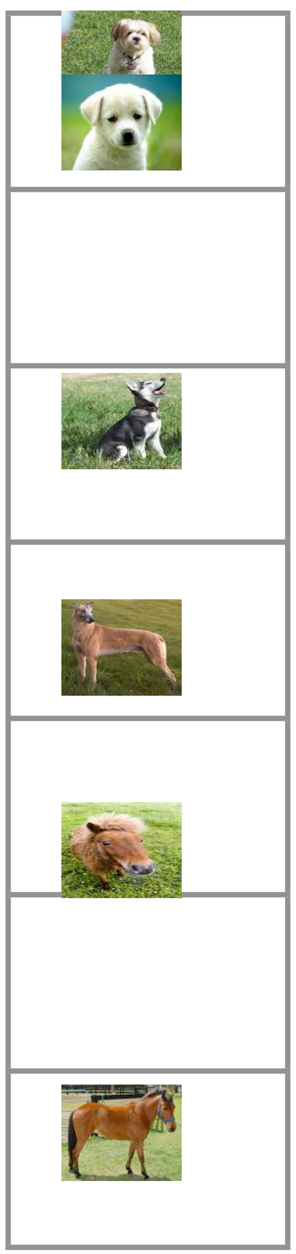

- k=2

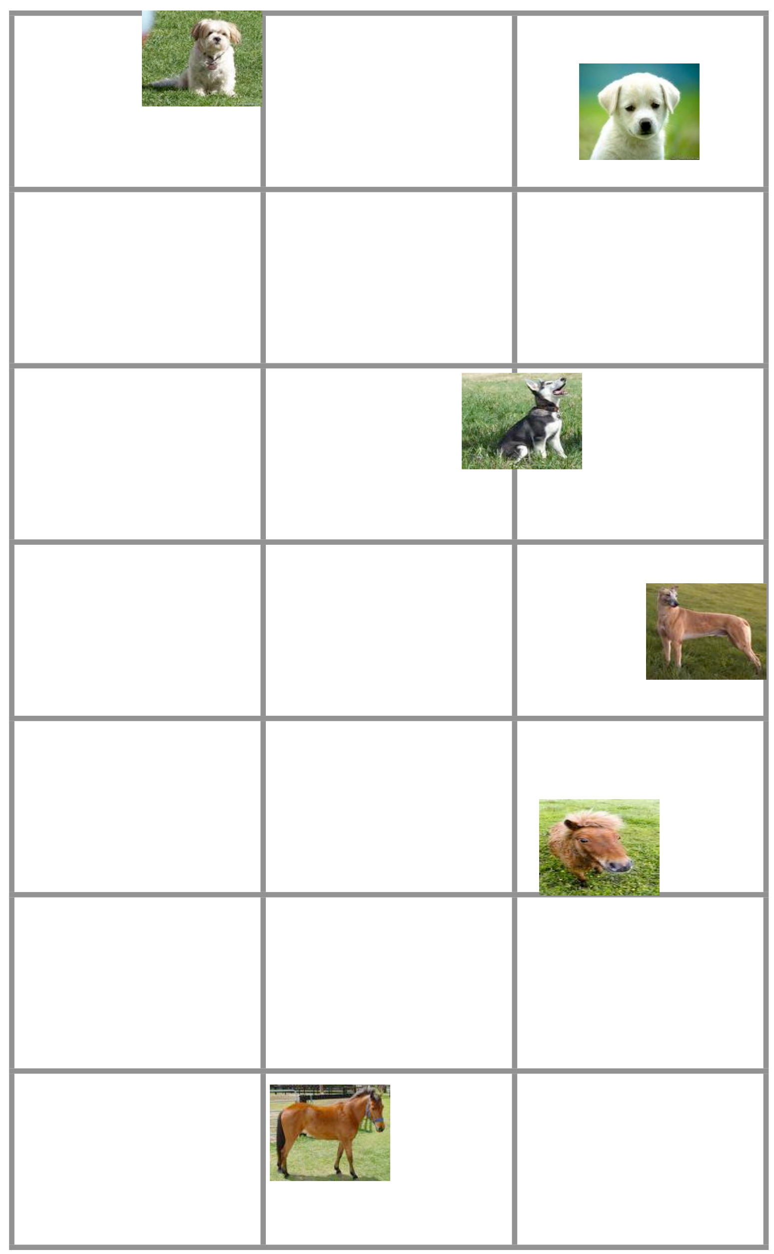

eigenvalues

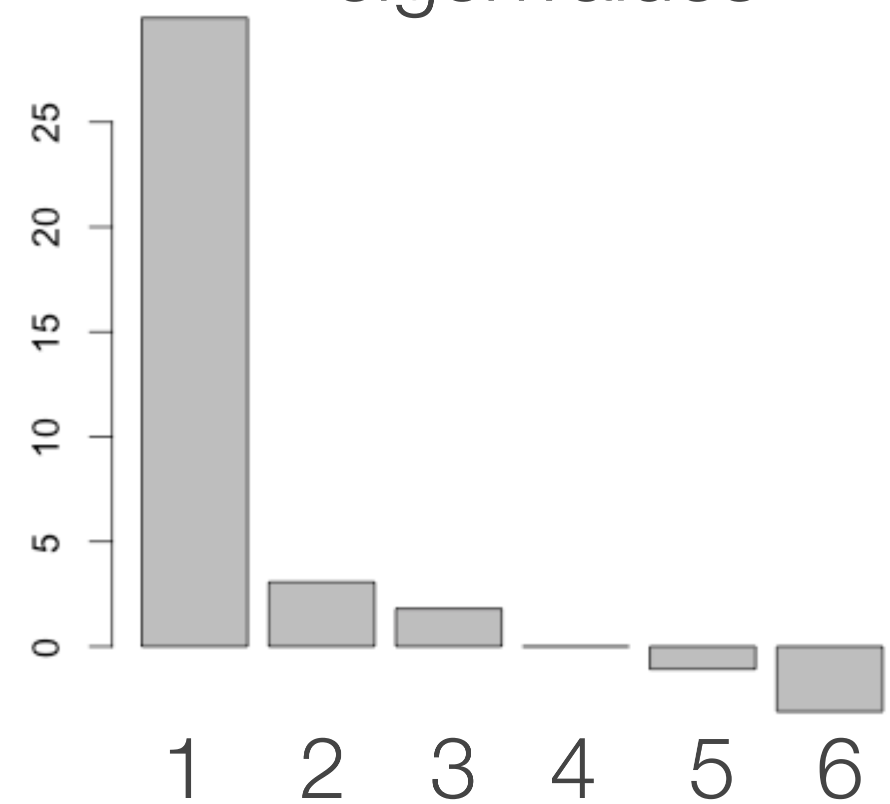

relative magnitudes of eigenvalues indicate the relative contribution of that dimension in reproducing the distance matrix: first dimension is doing most of the work

Some examples of MDS visualisations

- When the data aren't fake, MDS can be quite revealing about people's mental representations

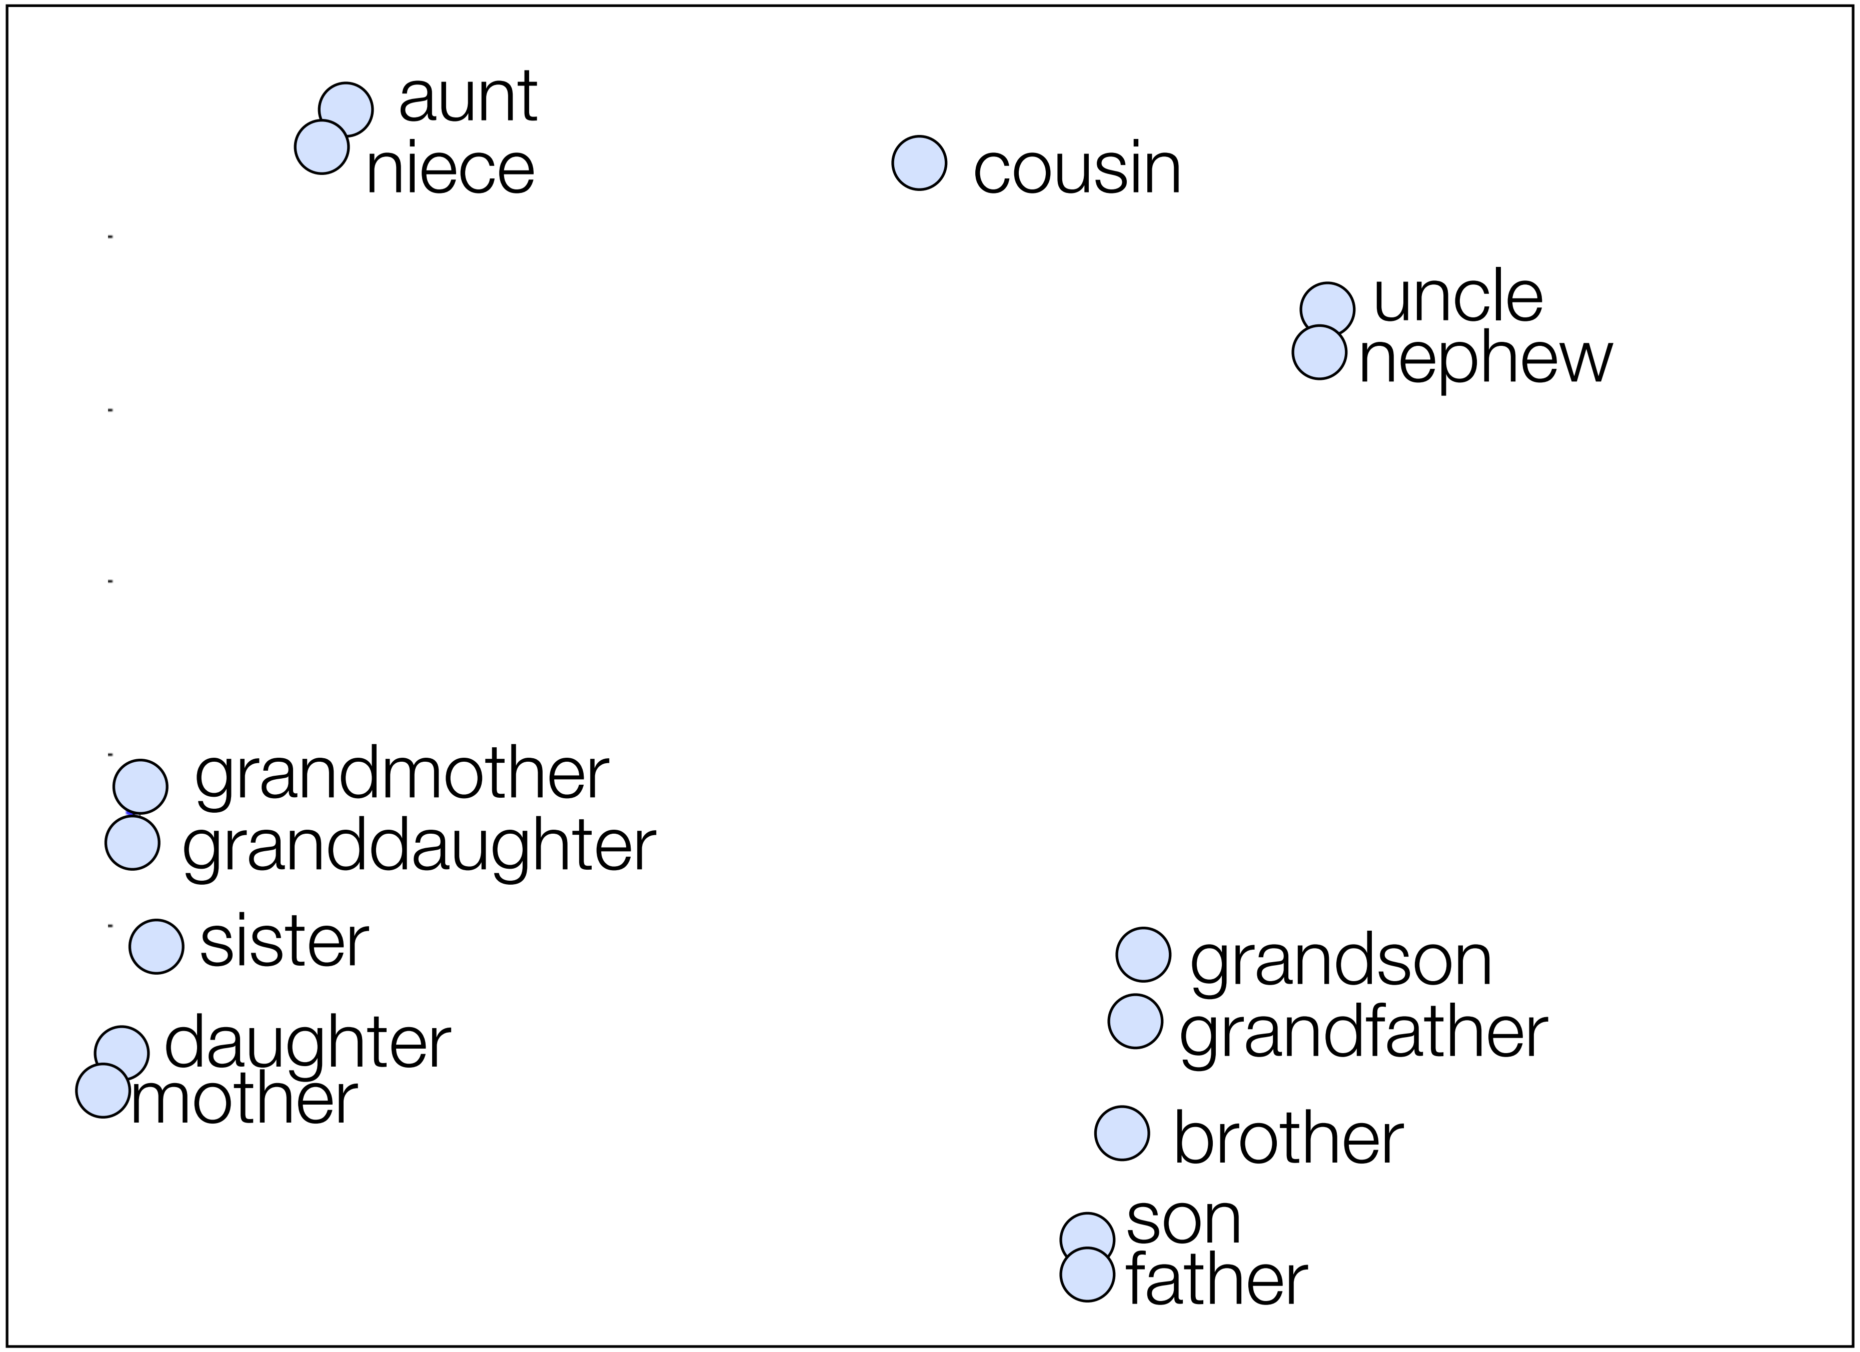

Some examples of MDS visualisations

- When the data aren't fake, MDS can be quite revealing about people's mental representations

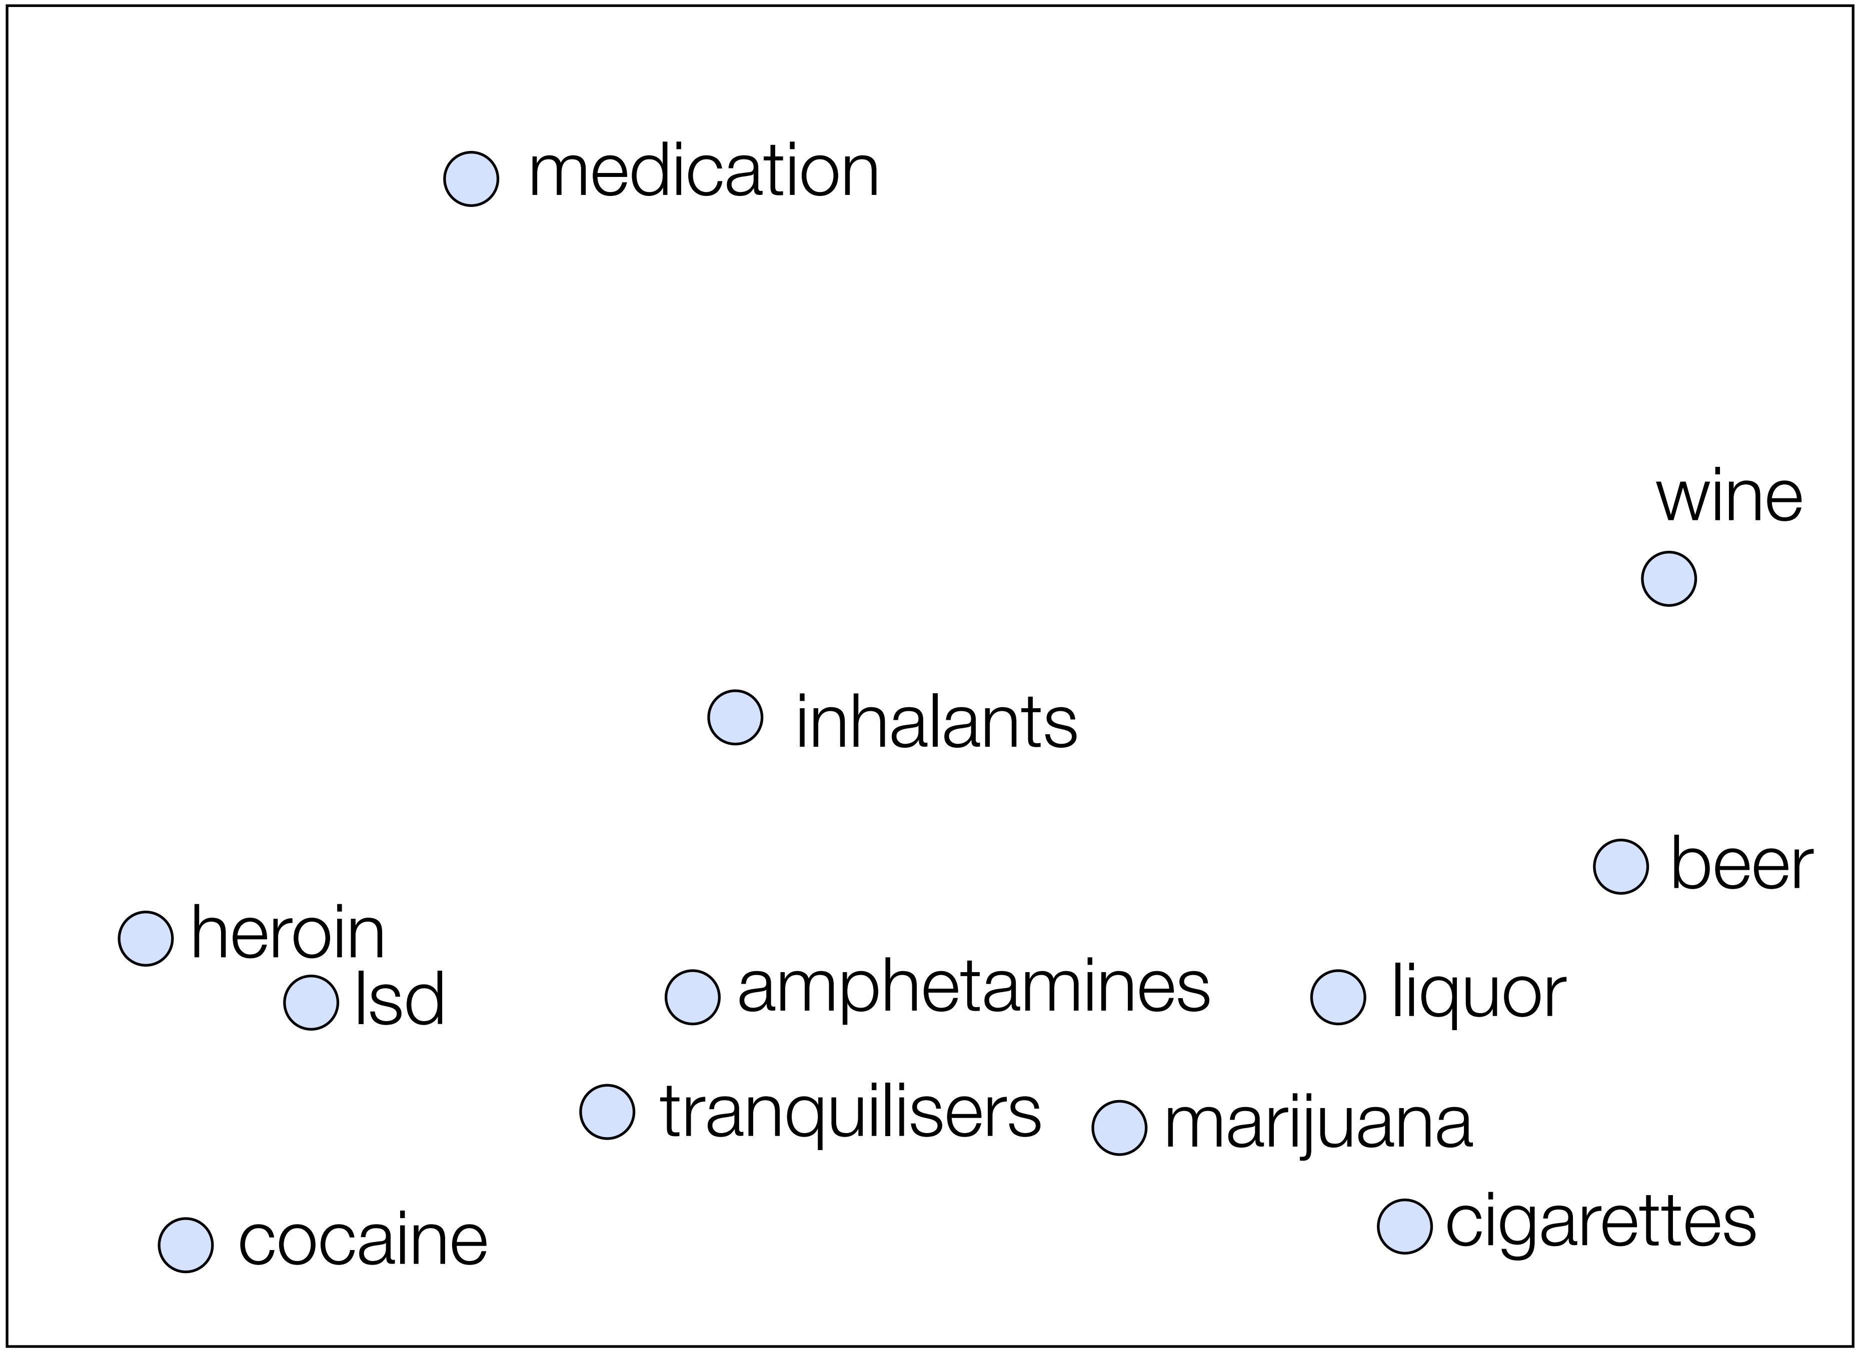

Prediction

- Key point is that we can derive a sensible measure of "psychological space" using similarity data, and visualise it using MDS
- Our earlier analysis predicted that generalisation should follow an exponential curve in psychological space

Does it?

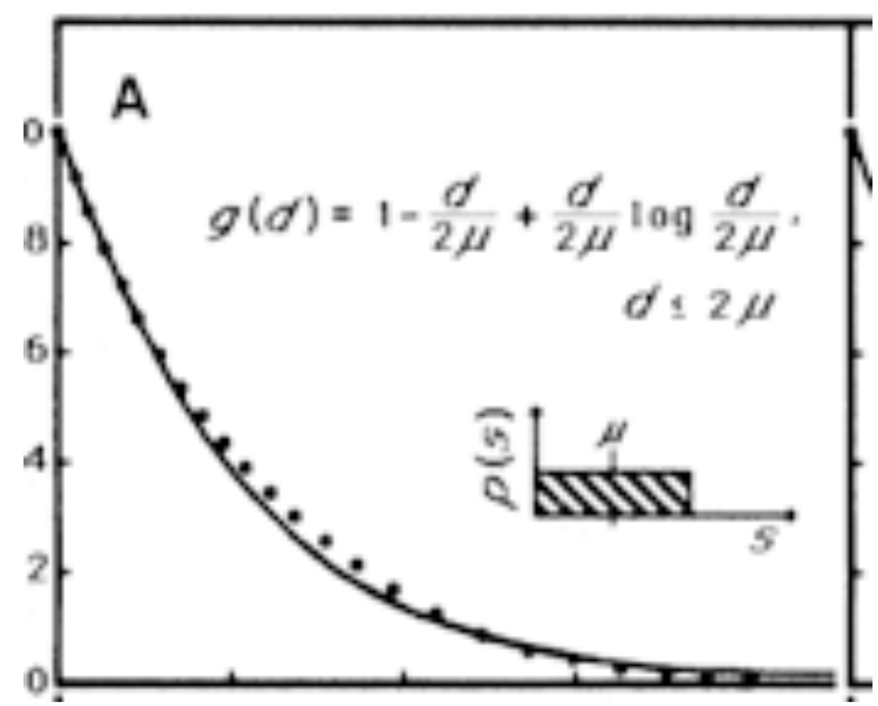

Testing the prediction

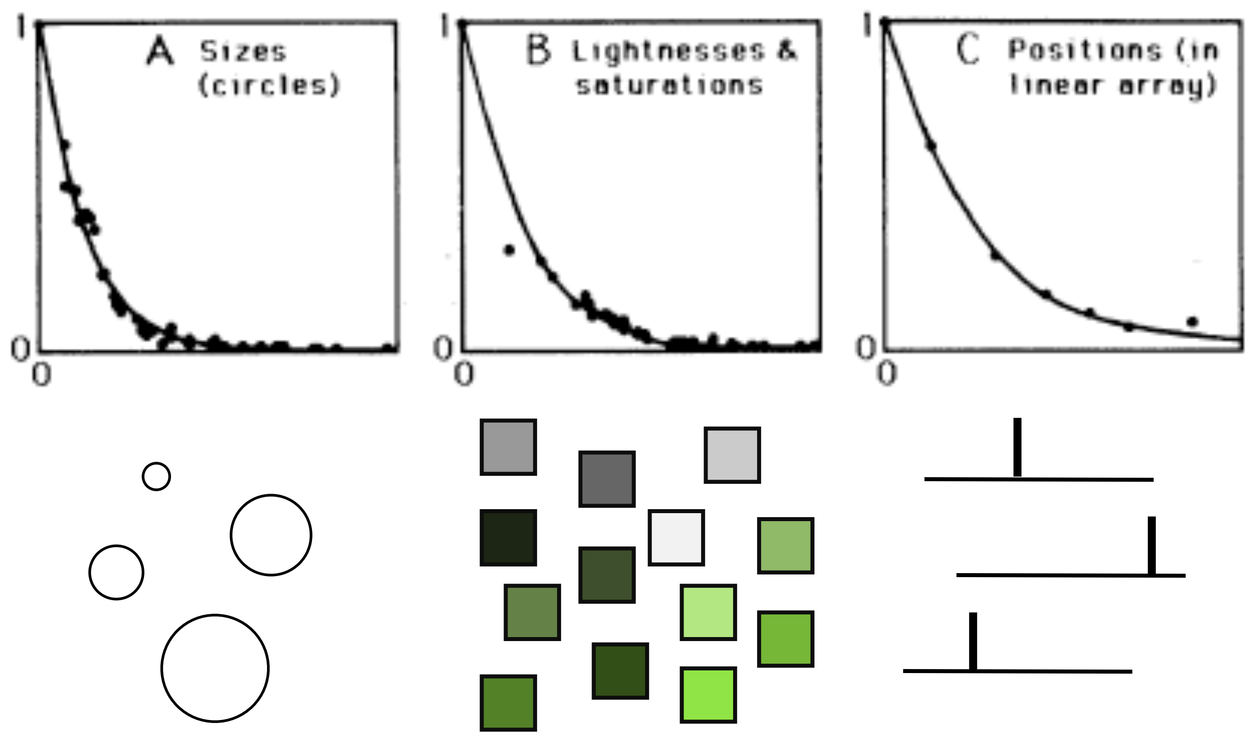

Indeed they do!

Datasets include fairly low-level visual ones

Testing the prediction

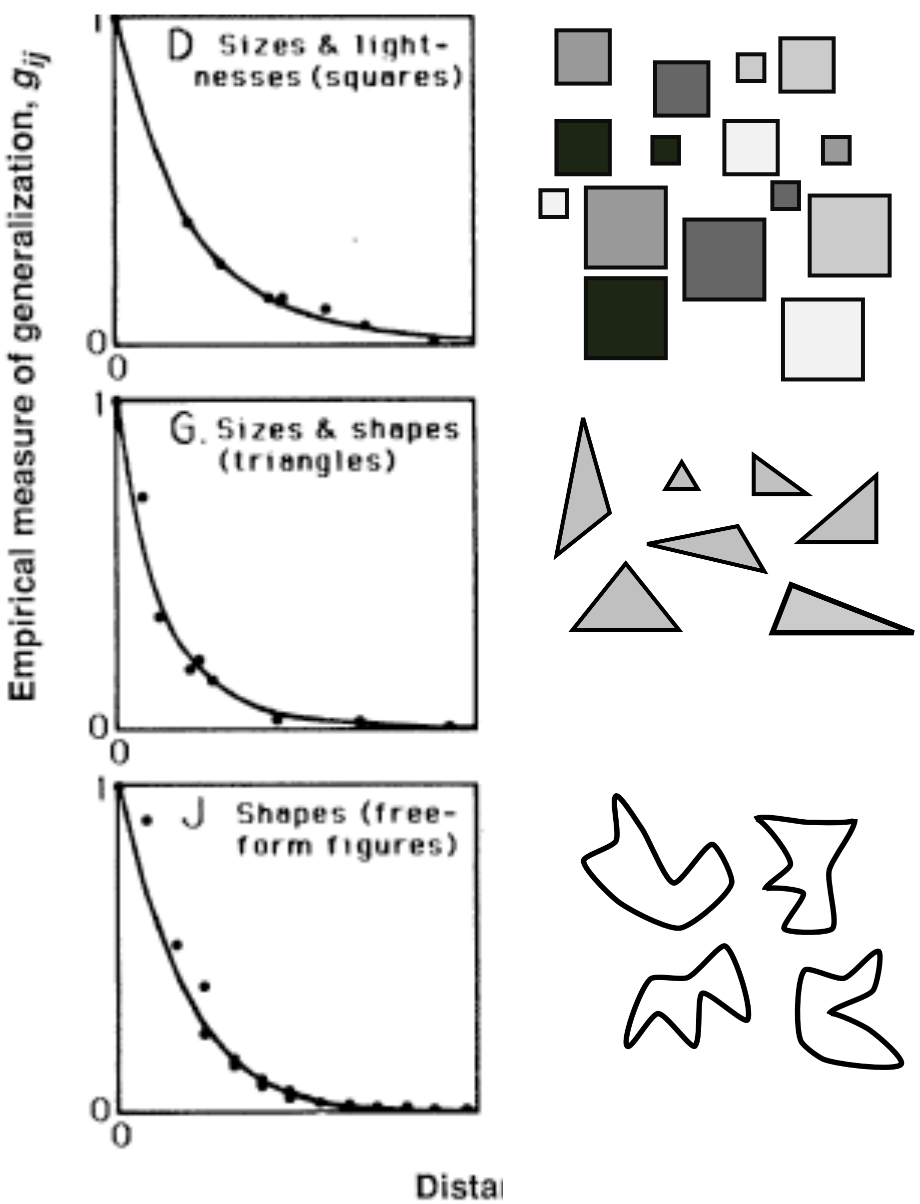

Indeed they do!

Datasets include fairly low-level visual ones

Slightly more multi-dimensional visual ones

Testing the prediction

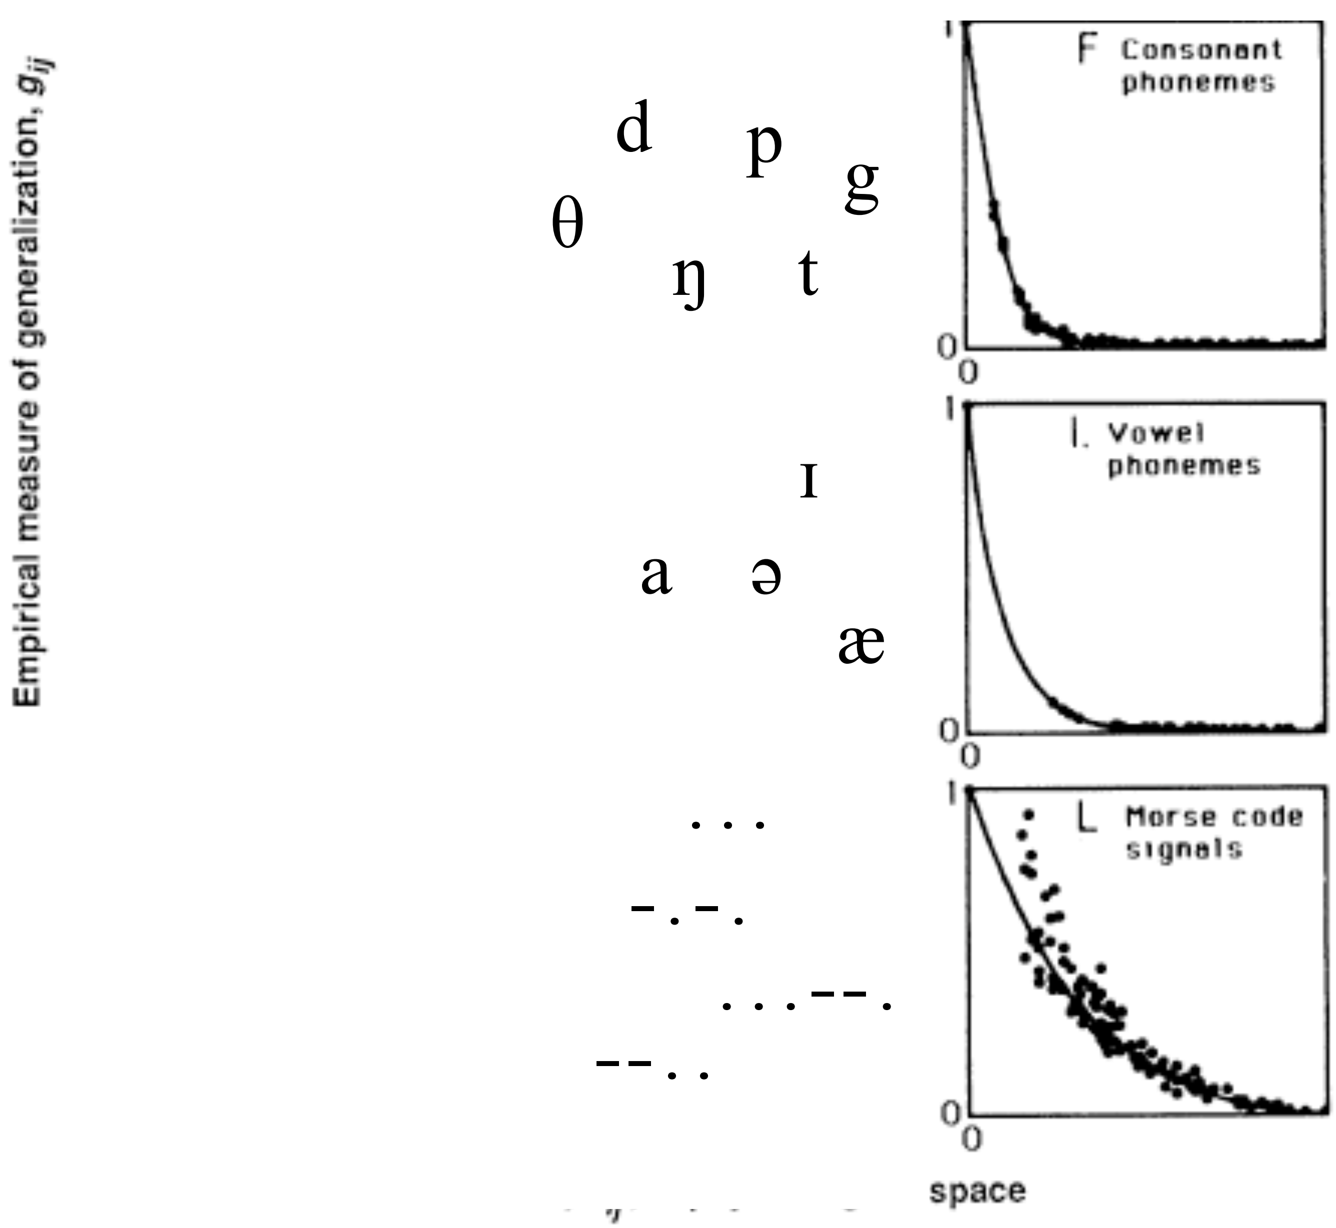

Indeed they do!

Datasets include fairly low-level visual ones

Slightly more multi-dimensional visual ones

Sounds

Testing the prediction

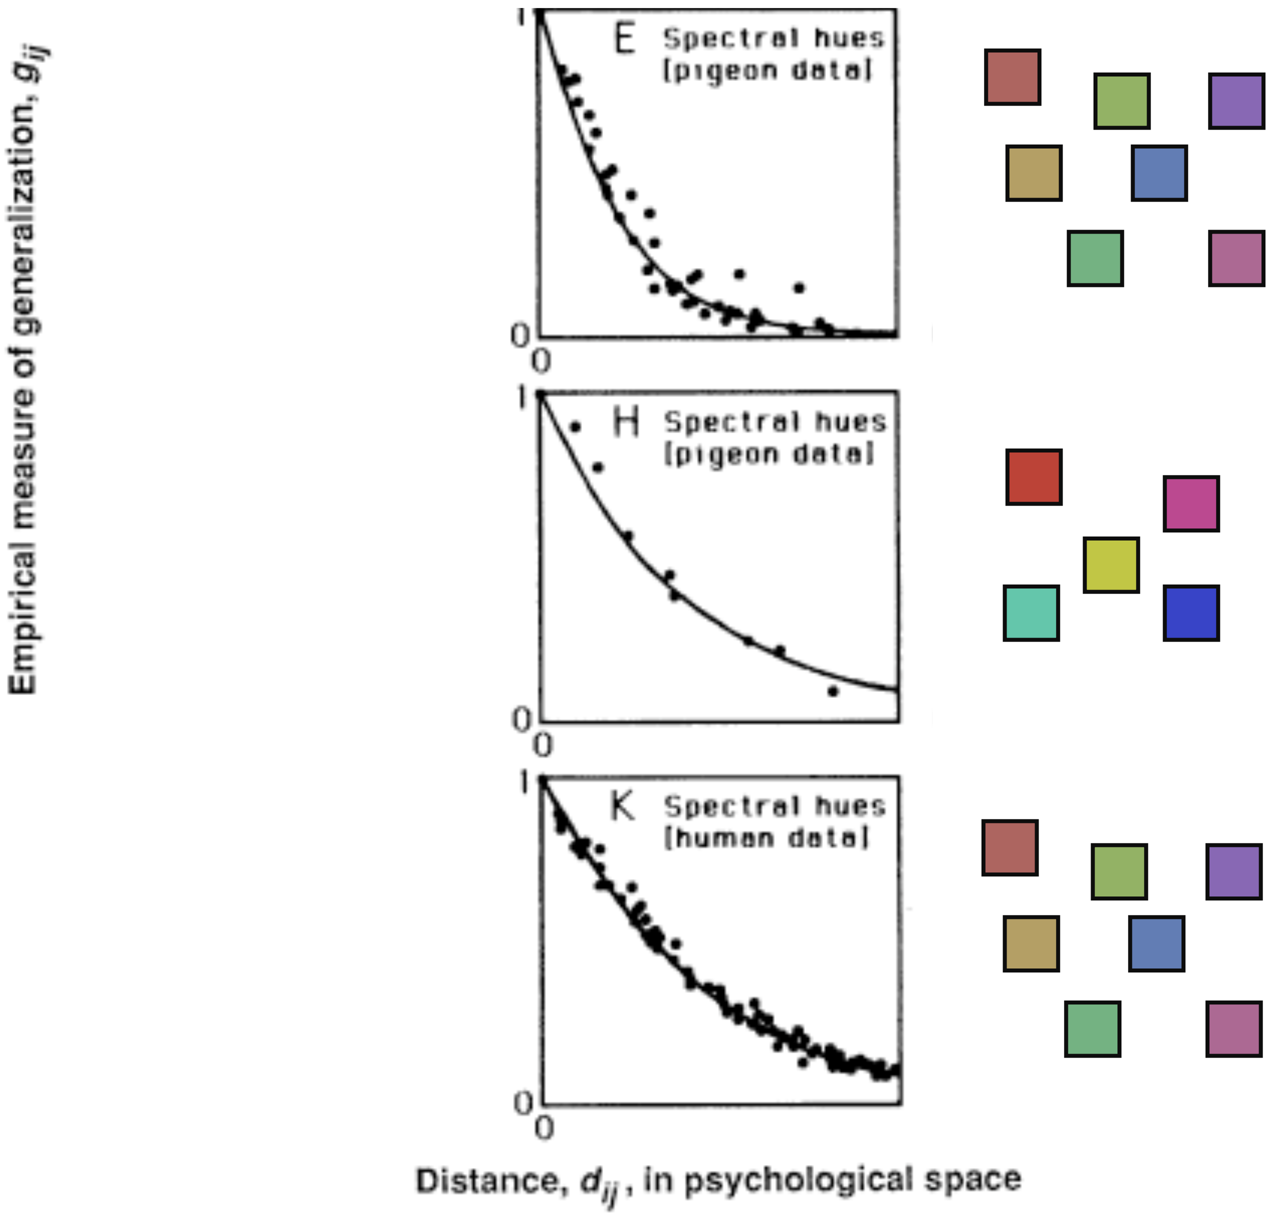

Indeed they do!

Datasets include fairly low-level visual ones

Slightly more multi-dimensional visual ones

Sounds

Non-humans

Interim summary

- Derived a result showing that generalisation should follow an exponential function in psychological space

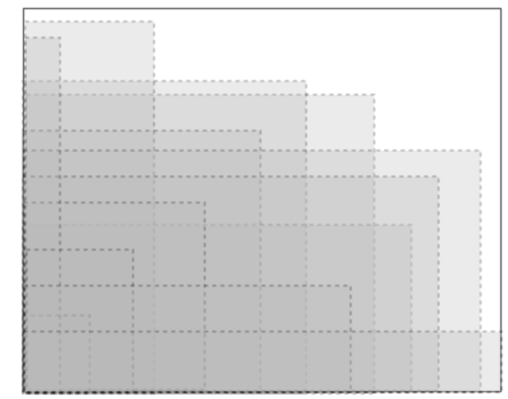

- Figured out how to measure that space, and how to visualise it

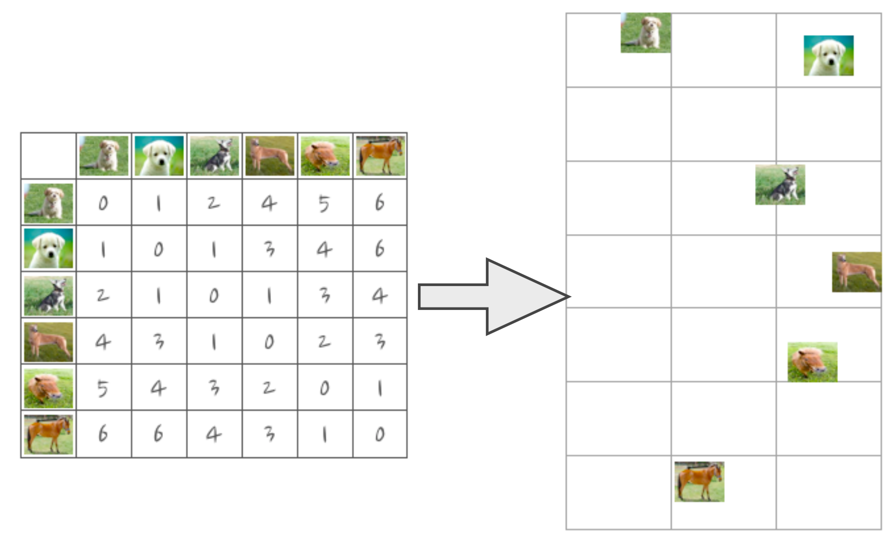

- Tested this prediction on multiple datasets

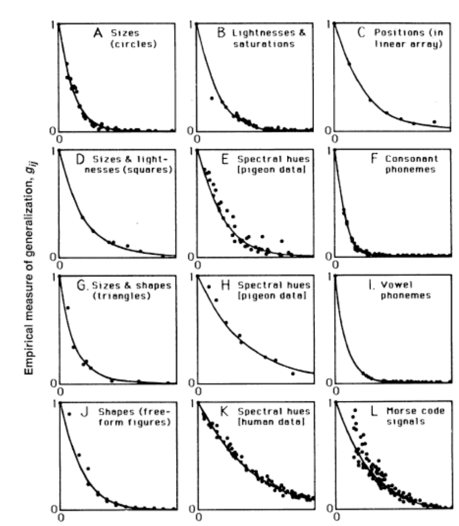

Limitations / assumptions

- This question is always important for figuring out if the result is psychologically interesting / applicable, and determining where to go next
- Assumptions are:
  - Consequential regions are convex, of finite extension, and centrally symmetric
  - Psychological space exists and we can approximate it
  - It has a metric (i.e., distance measure)
  - It is continuous
  - We only care about generalising from one point

Can we extend the analysis so these are not true?

Lecture outline

- The problem of generalisation
  - Defining the problem, and deriving a solution
  - Tangent: Multi-dimensional scaling
- ➡ Extending our solution to the problem of generalisation
  - A new framework
  - Application to more abstract representations
- Next: Other kinds of inductive generalisation

Extending the problem of generalisation

- Arbitrary representational structure, not just a metric space: Bayesian inference over hypotheses

Hypothesis space $\mathcal{H}$ is the set of possible consequential regions

$p(h)$ is the prior probability
of each hypothesis in the set

$p(x|h)$ is the probability
of that hypothesis given data x

by Bayes' Rule, we can calculate
the posterior probability as
$p(h|x) \propto p(x|h)\, p(h)$

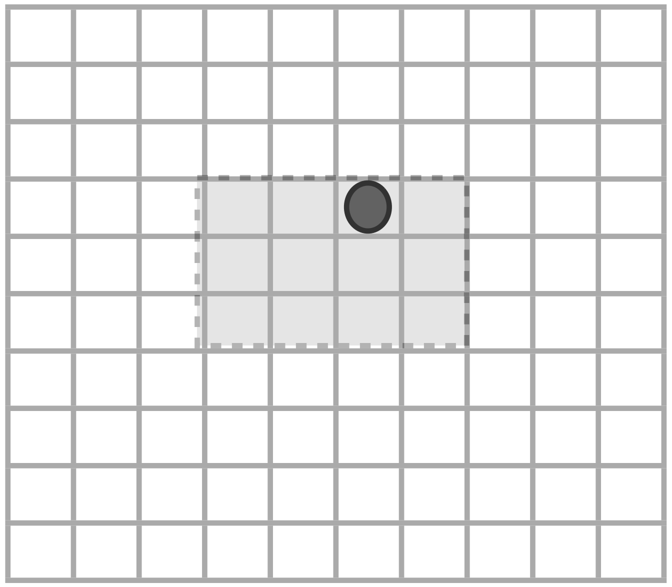

$p(h) = 1/|\mathcal{H}|$

$p(x|h) = 1/12$

Extending the problem of generalisation

Probability of generalising to point y is given by summing the probabilities of all hypothesised consequential regions that contain y

$$p(y \in C|x) = \sum_{h:y \in h} p(h|x).$$

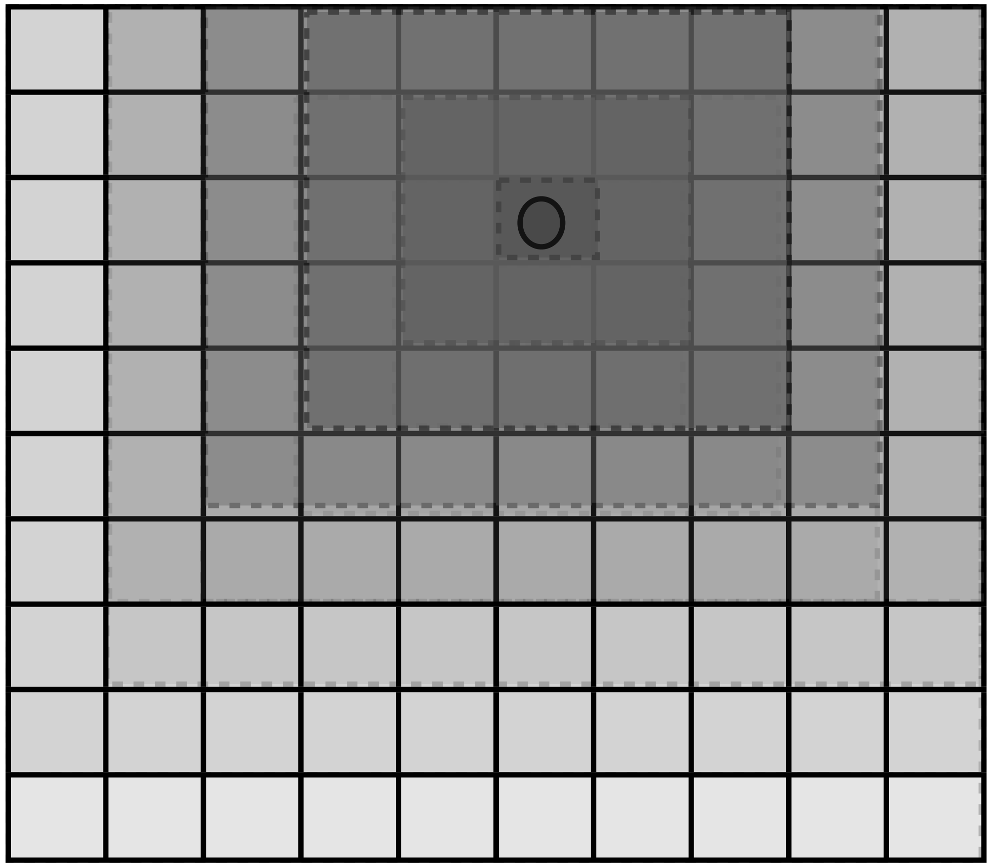

Extending the problem of generalisation

Probability of generalising to point y is given by summing the probabilities of all hypothesised consequential regions that contain y

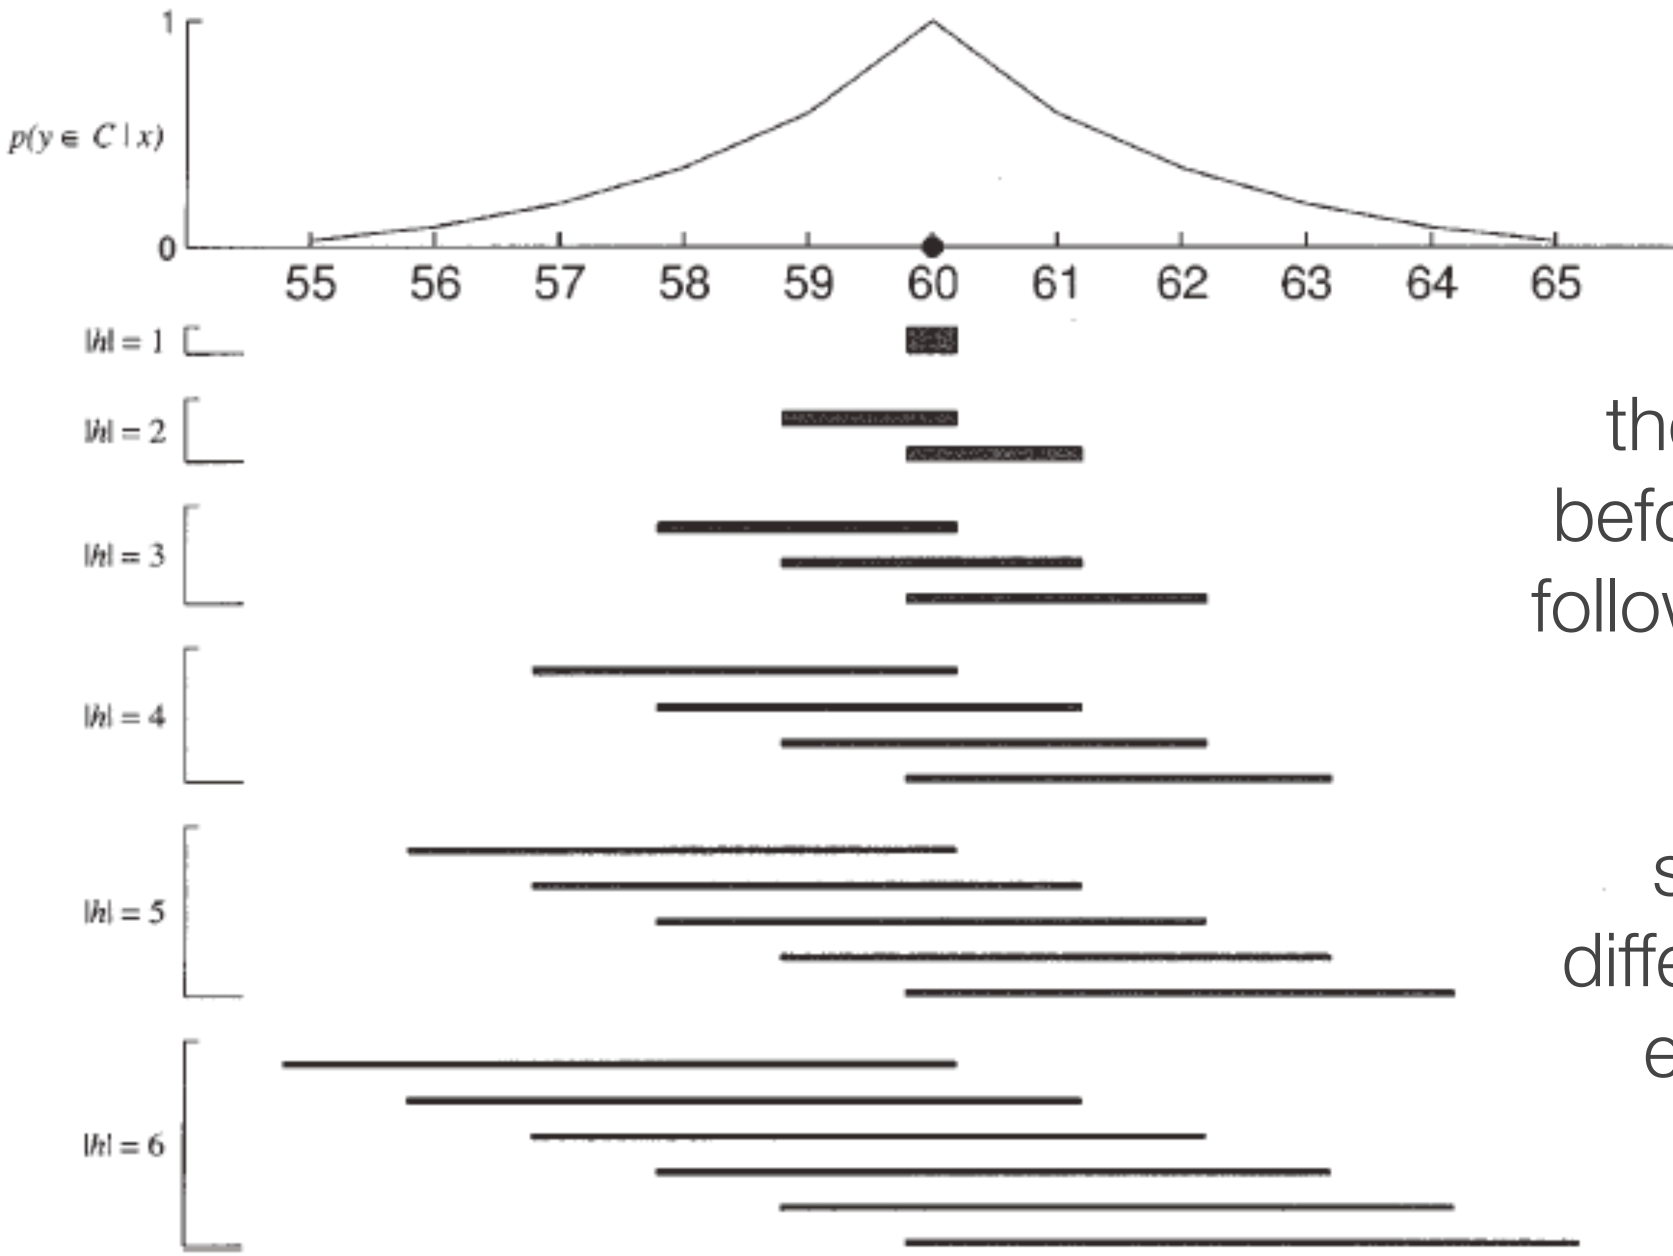

the logic is the same as before -- generalisation still follows an exponential curve

so far nothing is very different... but this is easily extendible to multiple datapoints

Extending the problem of generalisation

For multiple datapoints just recalculate $p(h|x)$ for the new datapoint, and redo the sum over all hypotheses to get the generalisation gradient!

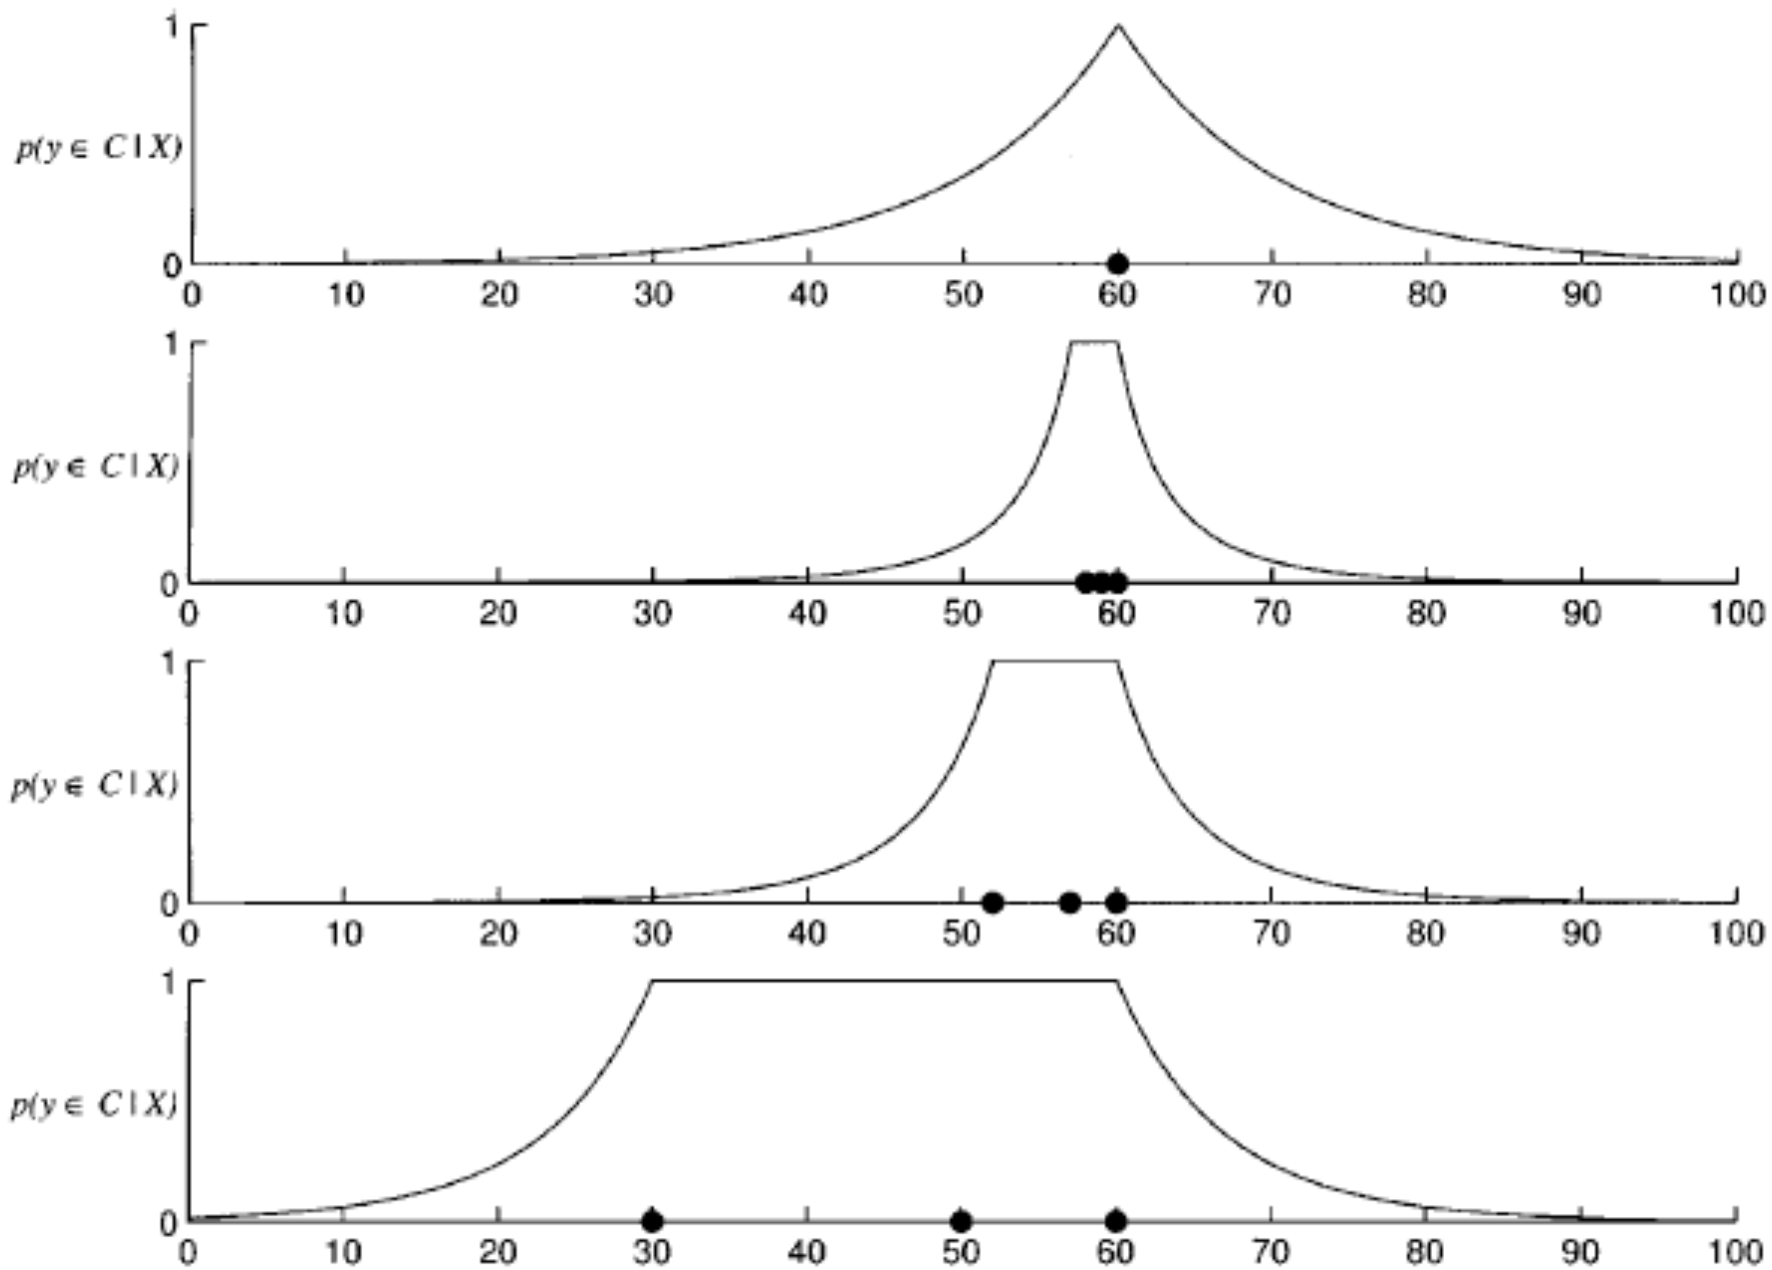

Extending the problem of generalisation

The generalisation gradient is always exponential,
but it tightens with additional data

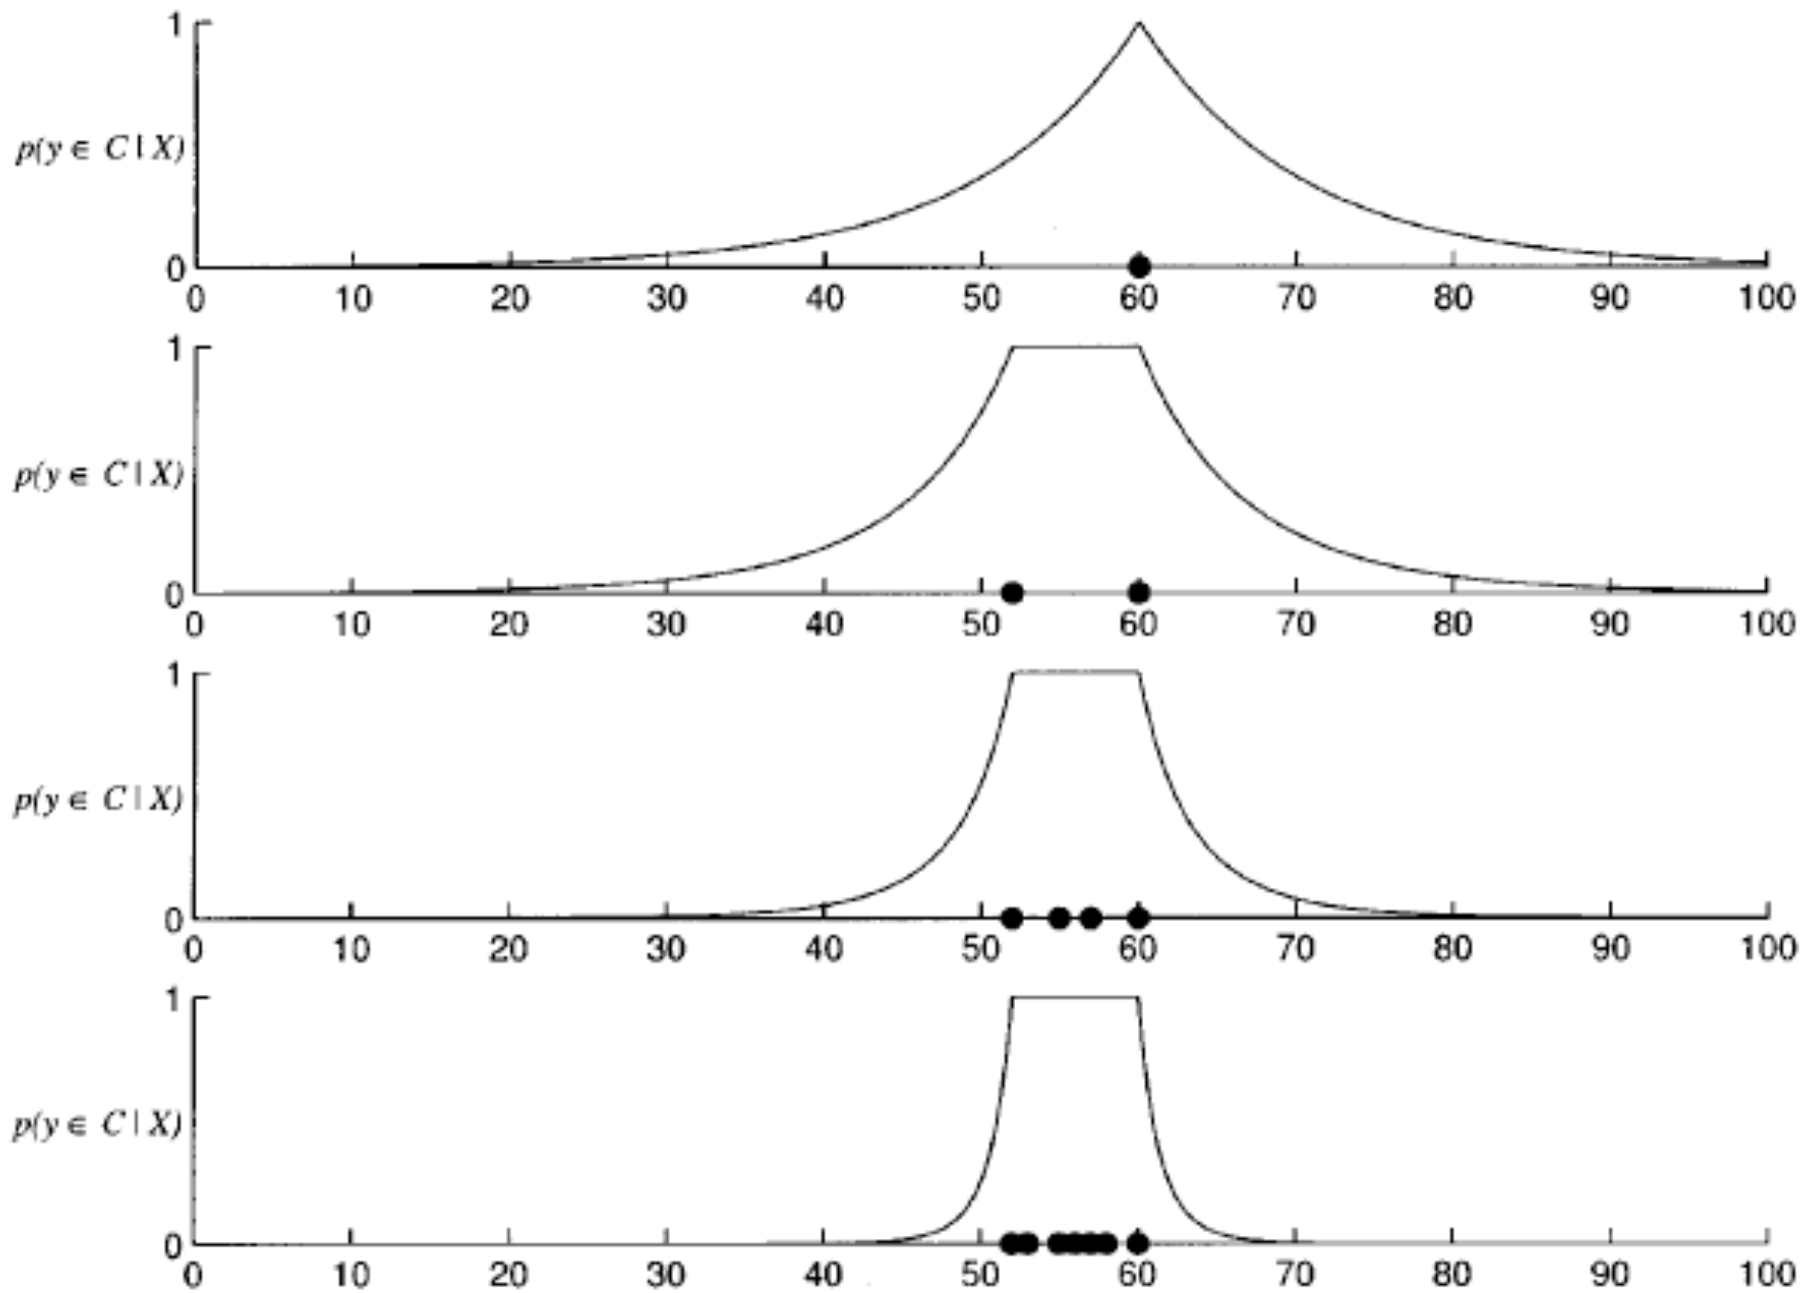

Extending the problem of generalisation

We can generalise it further, though, to include arbitrary hypothesis spaces

Look familiar? This is just the Lotto Problem from Friday!

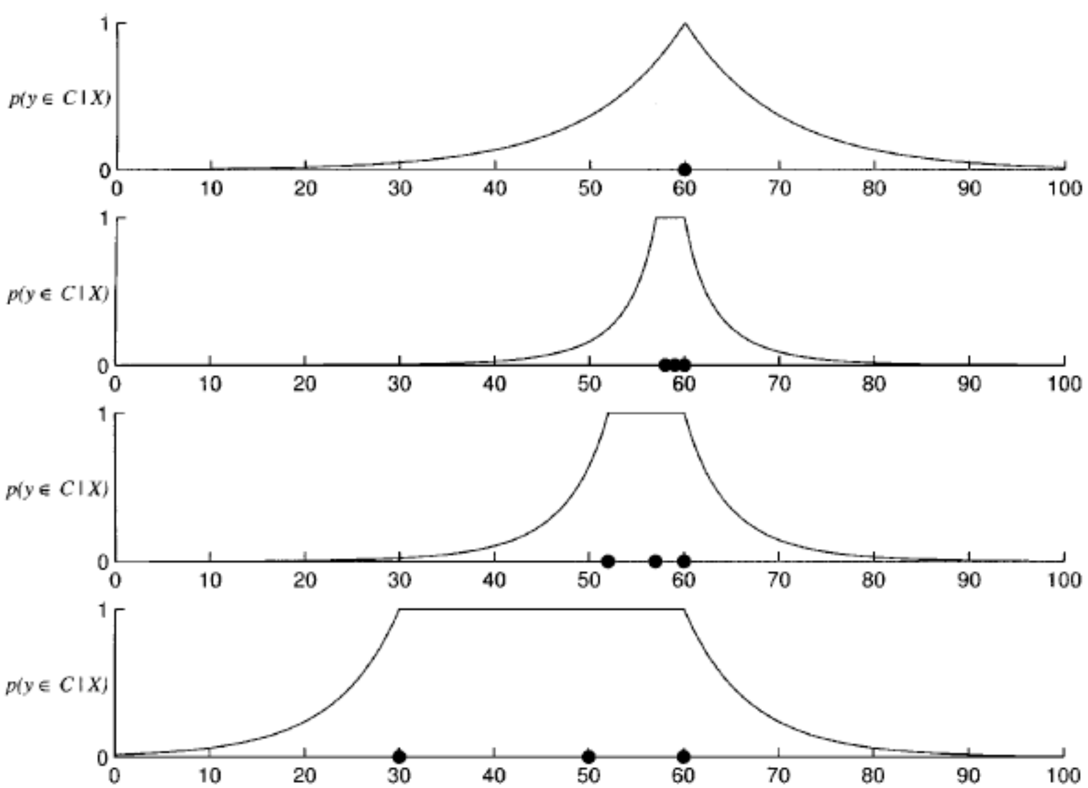

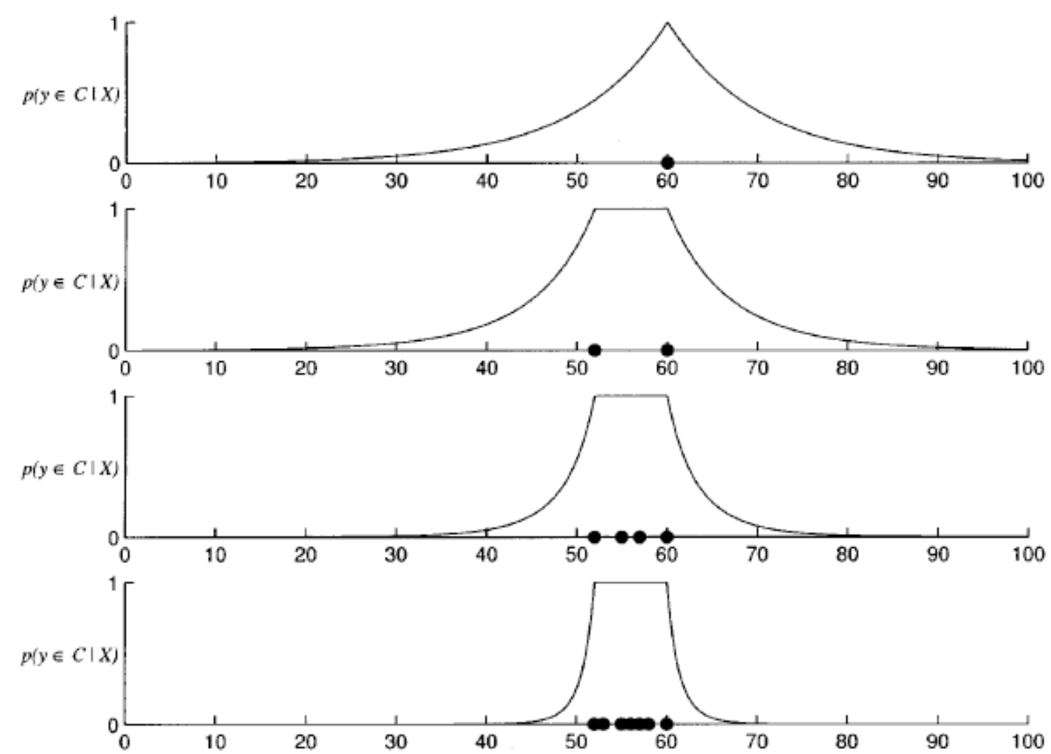

Extending the problem of generalisation

We can generalise it further, though, to include arbitrary hypothesis spaces

The number game:
Guess which number rule I'm thinking of

9 36 25 49 1

Extending the problem of generalisation

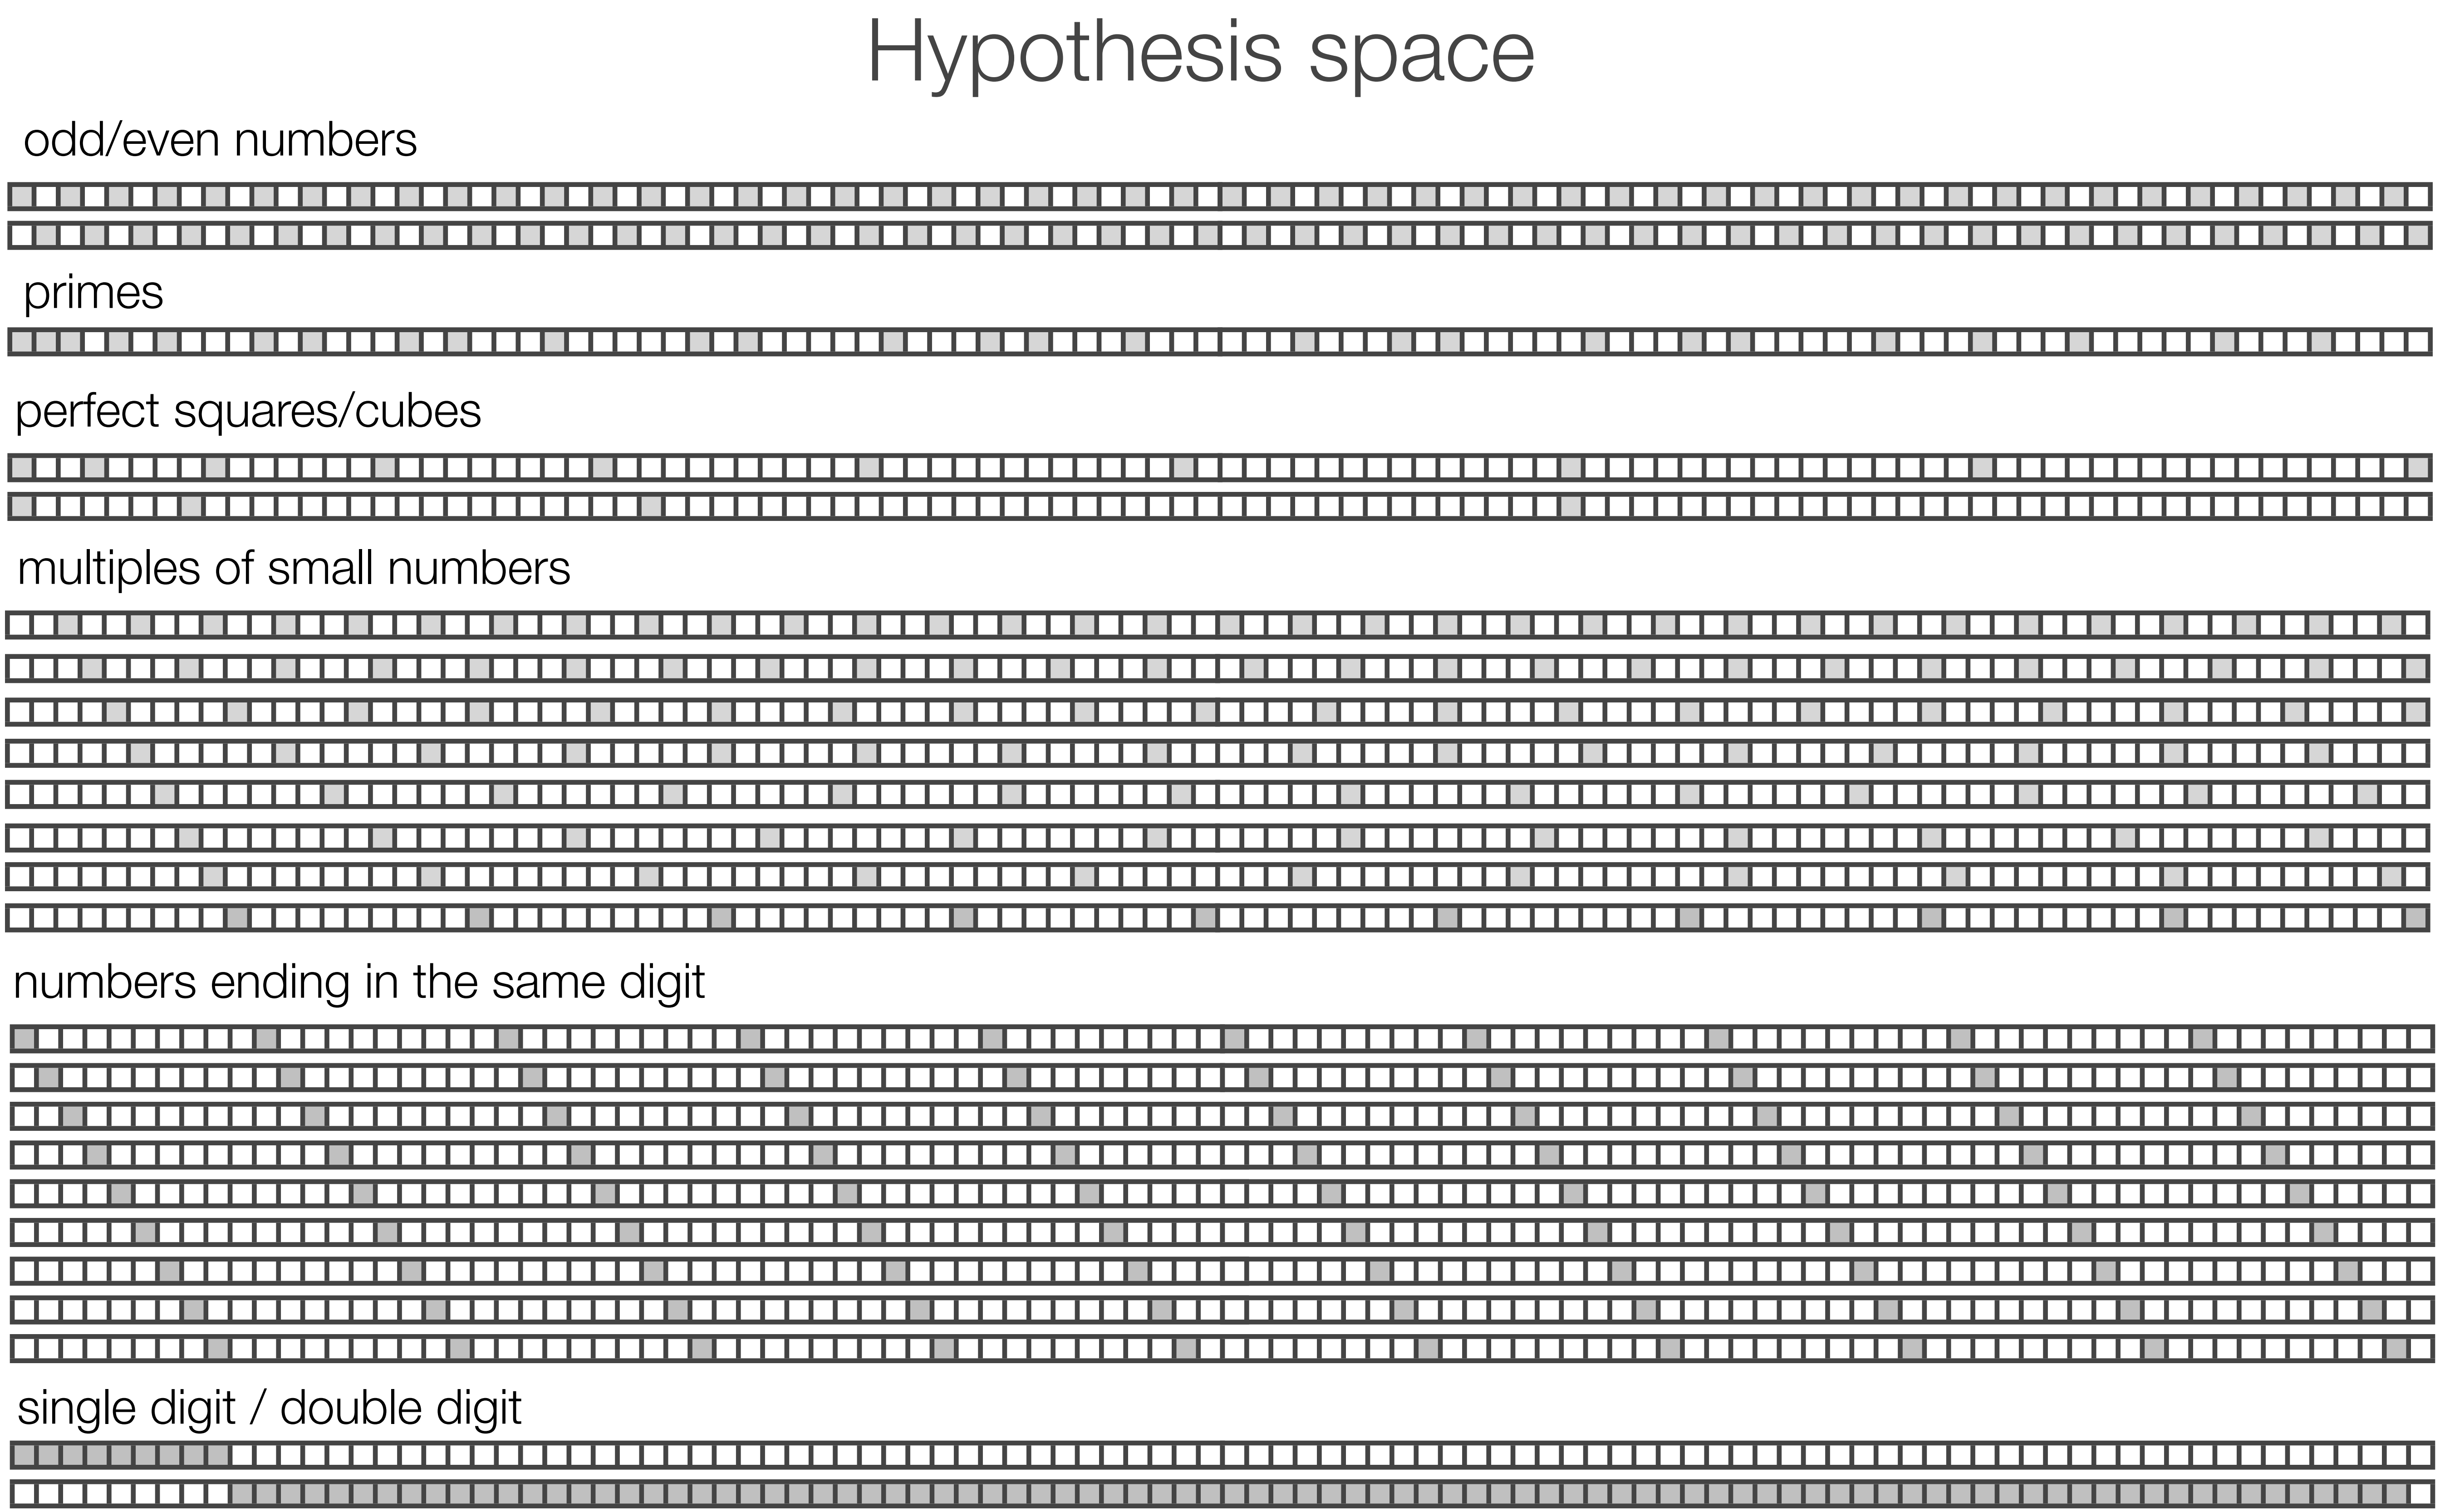

Extending the problem of generalisation

9

odd/even numbers

perfect squares/cubes

multiples of small numbers

numbers ending in the same digit

single digit / double digit

Extending the problem of generalisation

9 36

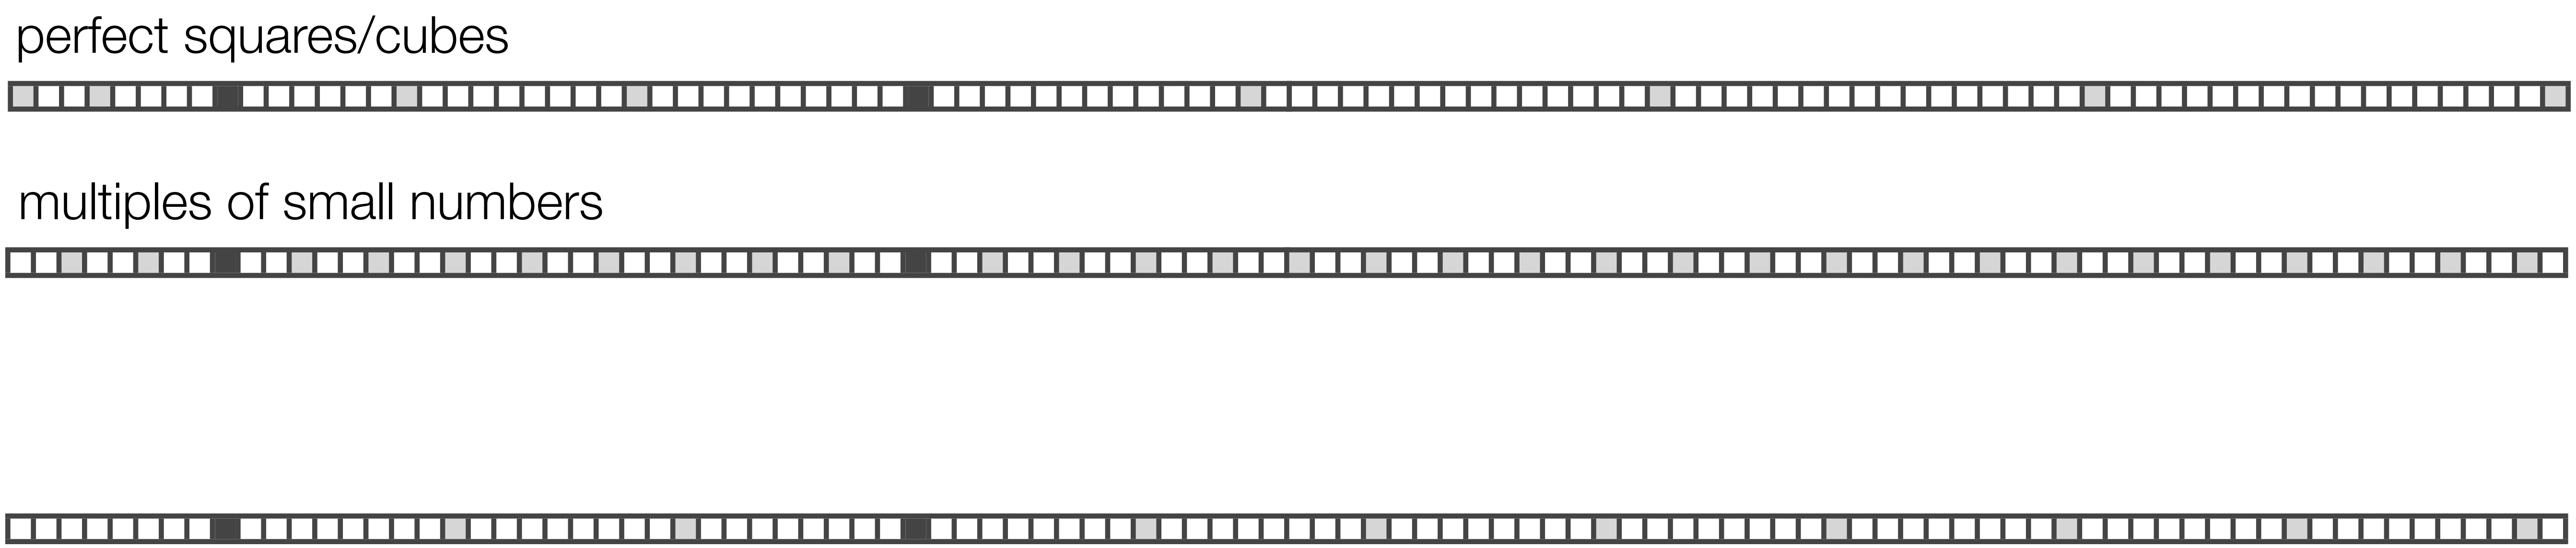

Extending the problem of generalisation

9 36 25

perfect squares/cubes

Extending the problem of generalisation

Summing over all hypotheses, with datapoint at 60

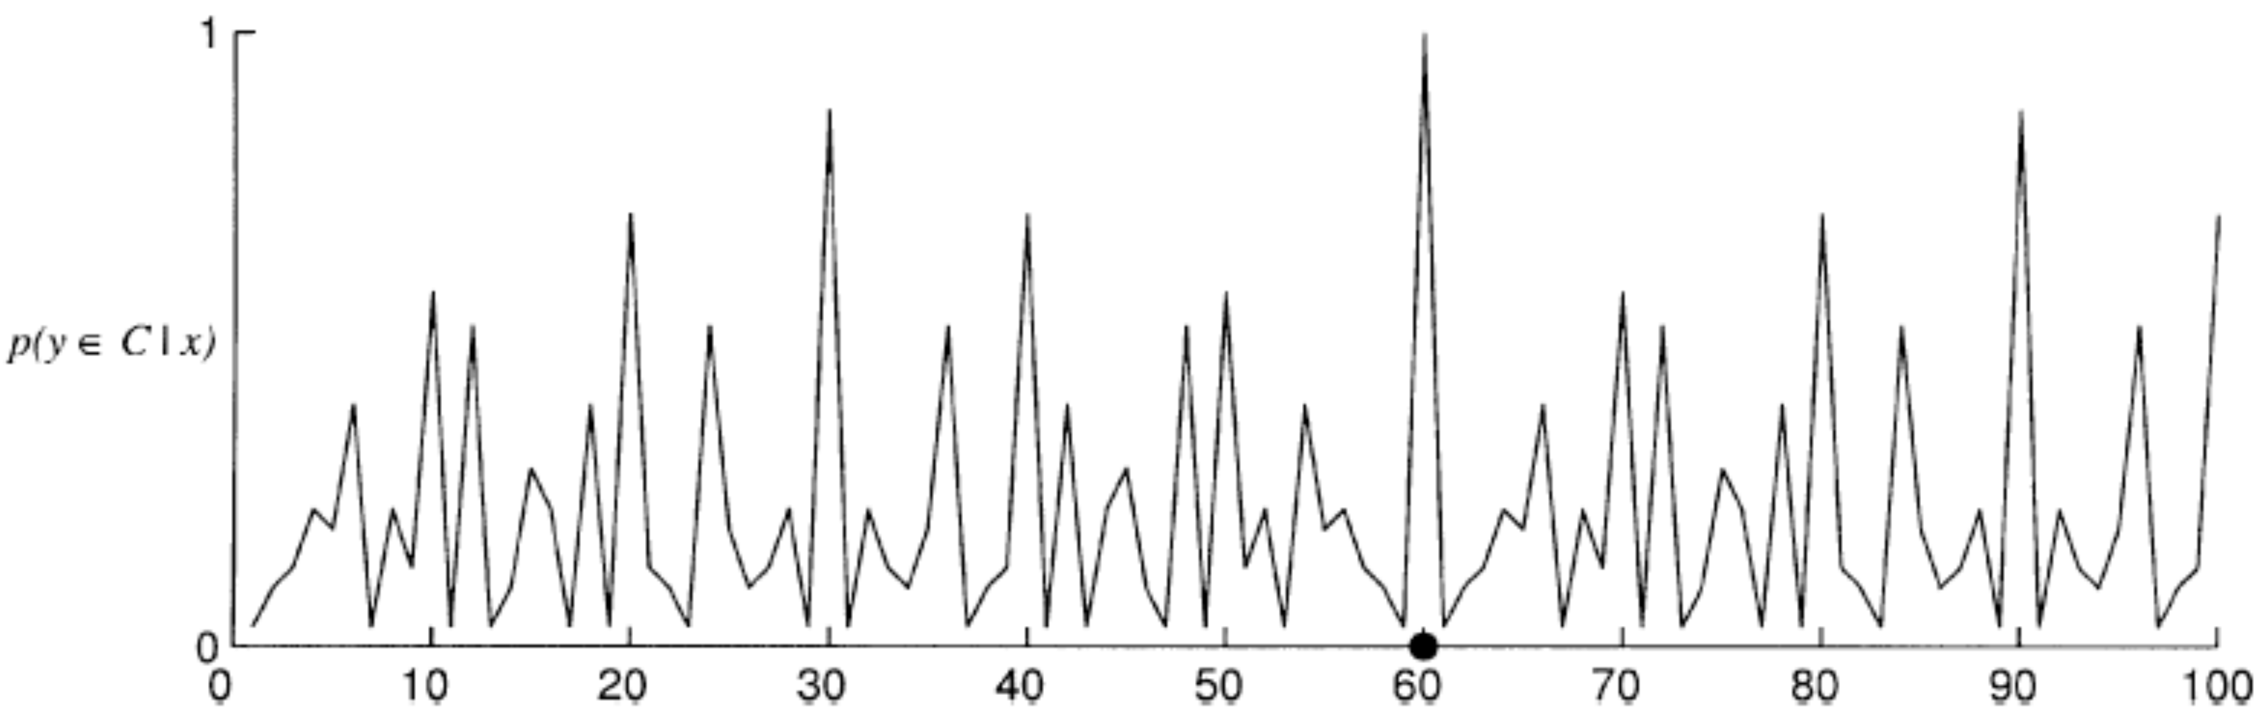

Generalisation gradients are still exponential, but in the relevant psychological space -- which we are not showing here (it would be derived from similarities over numbers, in which case 90 is much more similar to 60 than 47 is)

Final summary

- Adapted previous analysis to work within a Bayesian framework

$$p(h|x) \propto p(x|h)\, p(h)$$

- Showed that generalisation should still follow an exponential curve even if underlying space is not continuous or there are multiple datapoints

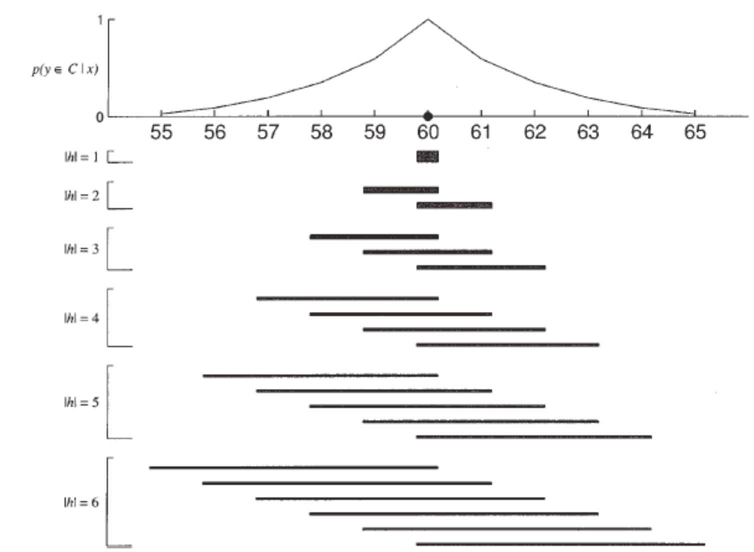

- Saw that this is just the lotto problem in another guise!

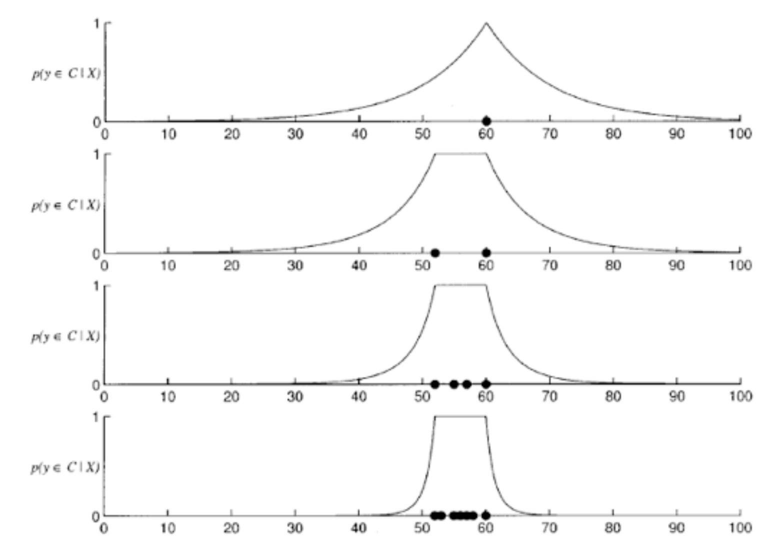

Lecture outline

- The problem of generalisation
  - Defining the problem, and deriving a solution
  - Tangent: Multi-dimensional scaling
- Extending our solution to the problem of generalisation
  - A new framework
  - Application to more abstract representations

➡ Next: Other kinds of inductive generalisation

Additional references (not required)

Deriving the exponential law

- Shepard, R. (1987). Toward a universal law of generalization for psychological science. *Science 237*(4820): 1317-1323.

Extension to more datapoints / abstract space

- Tenenbaum, J., & Griffiths, T. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences 24*(2): 629-641.