BSCS 2019 - Neural Computation

IV - Models of cognition

Mihály Bányai
banyai.mihaly@wigner.mta.hu
http://golab.wigner.mta.hu/people/mihaly-banyai/

- The mental model of the environment
- Prediction of behaviour
- Probing the mental representations

- The mental model of the environment
- Prediction of behaviour
- Probing the mental representations

The need for an internal model

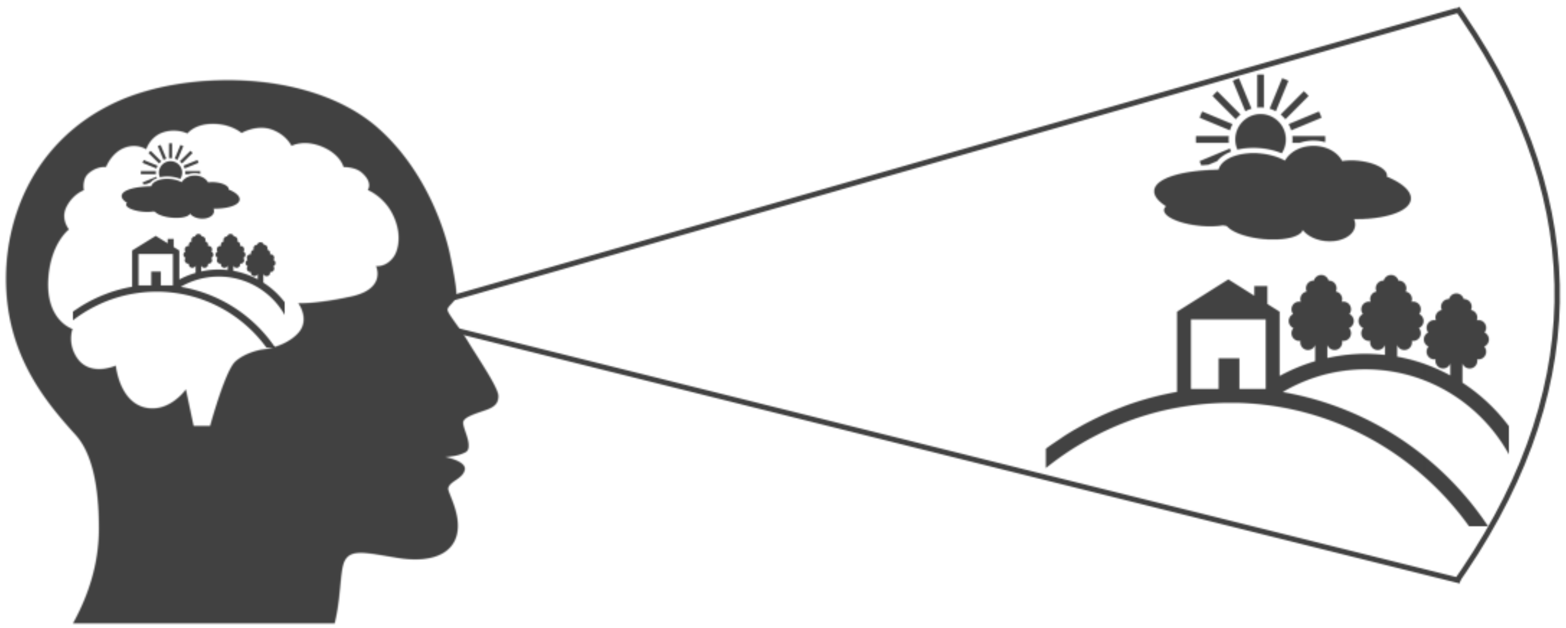

"To learn computer vision, first learn computer graphics."
Geoffrey Hinton

- In order to make decisions in complicated situations, we need to be able to predict the forthcoming events as well as the outcomes of our actions

- To do this, the brain needs to establish an internal model of the world

Probabilistic mental models

- we have seen that sensory experience is inherently ambiguous
- we need internal models that handle uncertainty in a consistent manner
- probabilistic models describe how unobserved variables effect the distribution of observations, thus they are ideally suited to be used as mental model candidates
- perception can be regarded as inference of the probability distribution of latent quantities conditioned on observations in a probabilistic model
    - e.g. what is the probability that there is a songbird or a howling gorilla in my environment given the visual and auditory input that I currently receive?
    - to answer this, as we have seen, we have to define a probabilistic model that tells us what is the probability distribution of these sensory values conditioned on the not directly observed presence of the animals
        - i.e. what kind of sounds does a bird and a gorilla typically make

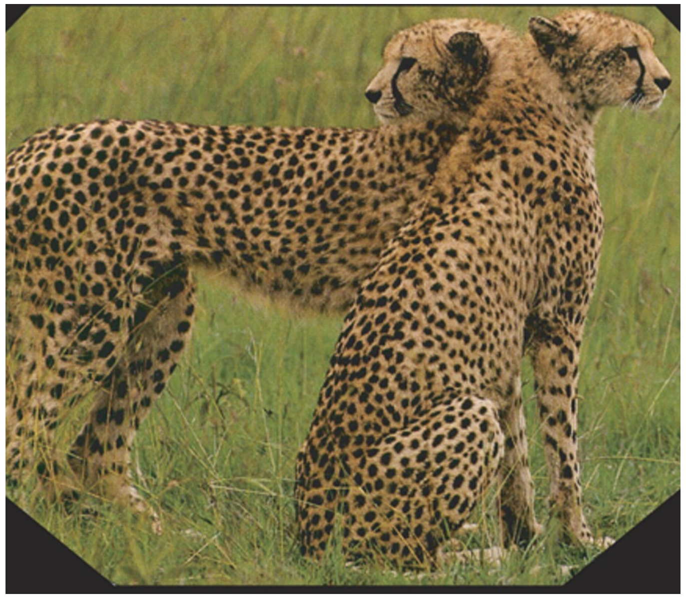

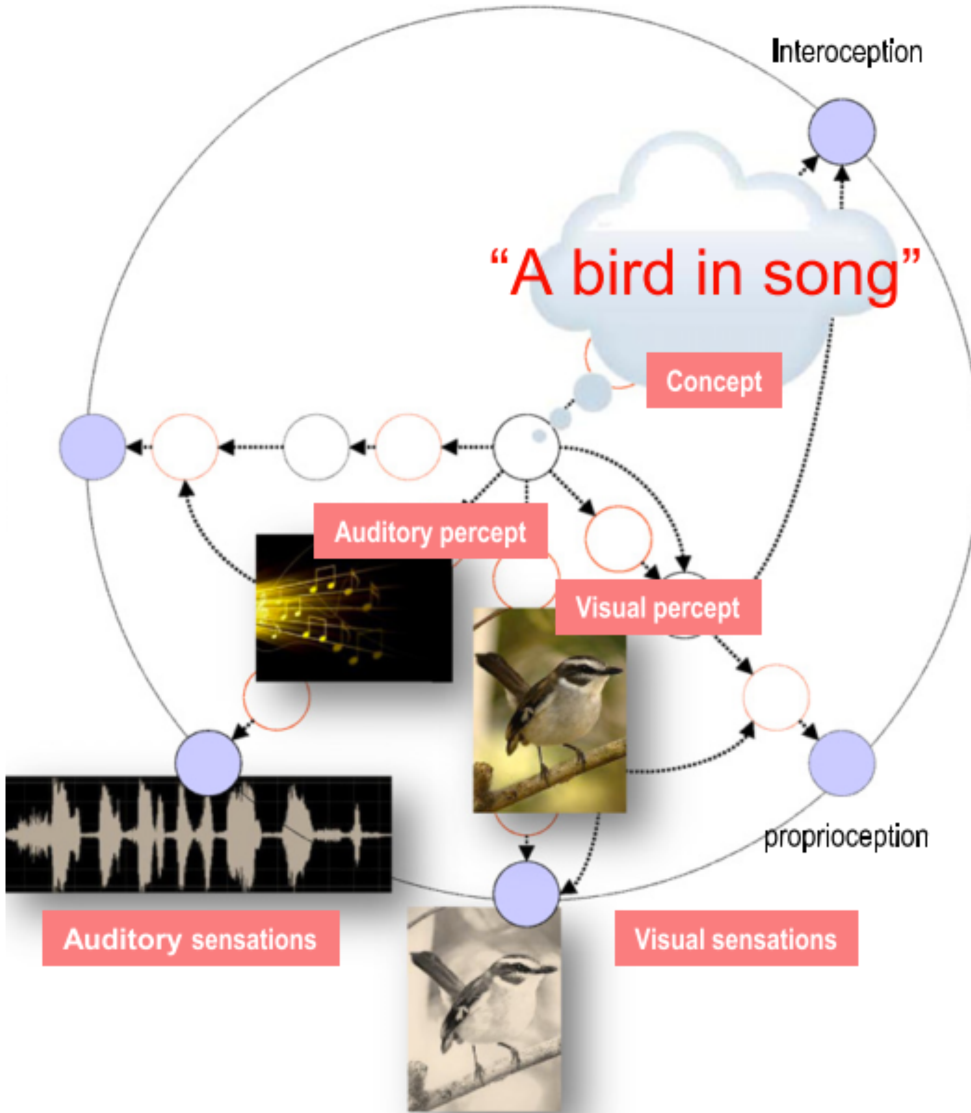

Adapting the model to stimulus statistics

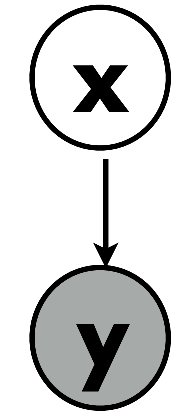

- as new and new observations arrive, we always add them to the axiom set
- this changes the inferred distributions of the variables of the model
- this can be achieved using parameter learning algorithms, that tune the PMFs or PDFs in the model to fit a set of observations as well as possible
    - choose P(x) and P(y|x) so that the probability of the observations $\{y_1 \ldots y_t\}$, that is maximal
        - we assume that observations are independent, so their probability distribution factorises
        - $\Pr(\{y_1 \ldots y_t\}) = \prod_i P(y_i) = \prod_i \sum_x P(y_i,x) = \prod_i \sum_x P(y_i|x)\, P(x)$, we have to maximalise this
    - in the case of continuous distributions, we can adapt the probabilities by adjusting the parameters of the PDFs
    - e.g. for a Gaussian, we need to figure out what is the mean value and the variance of the quantity that maximises the probability of the observations
    - artificial intelligence applications do the same - they fit models to data in order to predict new data.
        - the algorithms that are developed in machine vision, language processing, etc. can be used in computational cognitive science and neuroscience as well

The mental model is continuously updated

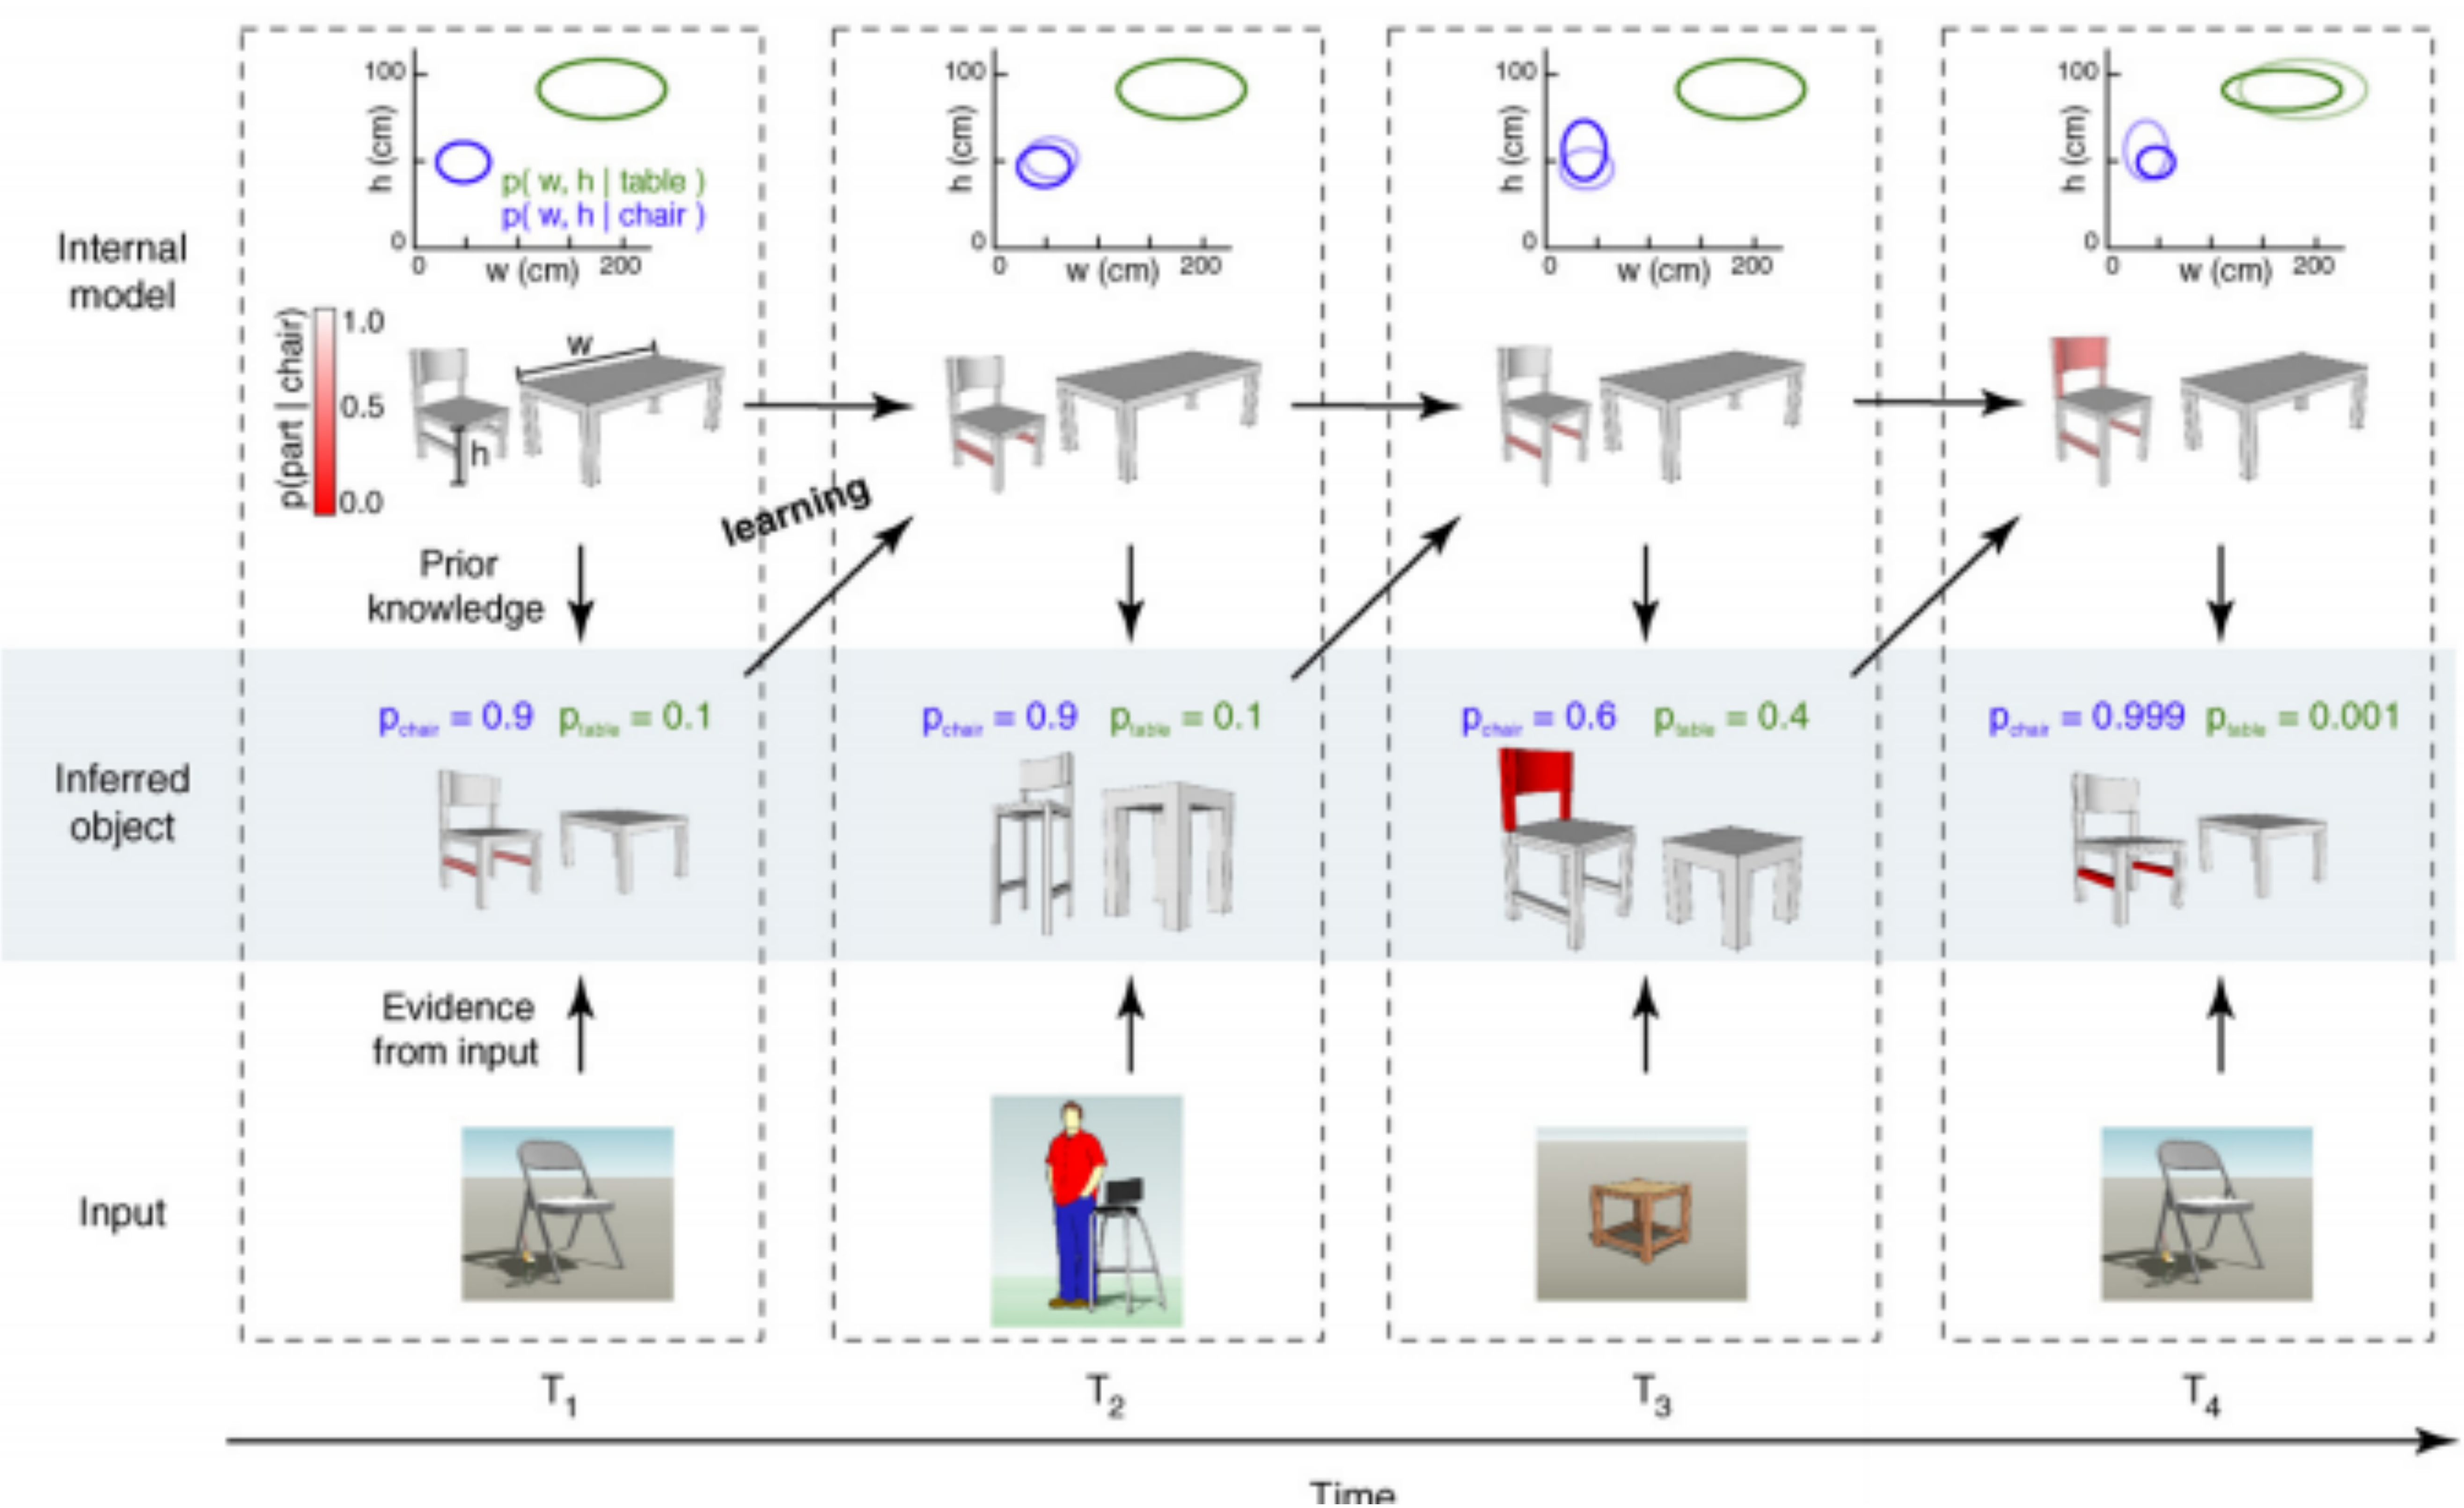

Compression of observations

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
    - it would be too much data (even in hyperthymesia)

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
    - it would be too much data (even in hyperthymesia)
    - it would be unnecessarily clumsy to access it

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
    - it would be too much data (even in hyperthymesia)
    - it would be unnecessarily clumsy to access it
    - we couldn’t generalise - wouldn’t know what to expect when a spotty dog shows up

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
    - it would be too much data (even in hyperthymesia)
    - it would be unnecessarily clumsy to access it
    - we couldn’t generalise - wouldn’t know what to expect when a spotty dog shows up
- the brain needs to use such models that reflect the property statistics of the environment

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
    - it would be too much data (even in hyperthymesia)
    - it would be unnecessarily clumsy to access it
    - we couldn't generalise - wouldn't know what to expect when a spotty dog shows up
- the brain needs to use such models that reflect the property statistics of the environment
    - object identities are **invariant** to a number of transformations, e.g. viewing angle or lighting differences

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
    - it would be too much data (even in hyperthymesia)
    - it would be unnecessarily clumsy to access it
    - we couldn't generalise - wouldn't know what to expect when a spotty dog shows up
- the brain needs to use such models that reflect the property statistics of the environment
    - object identities are **invariant** to a number of transformations, e.g. viewing angle or lighting differences
    - I can compress well when I'm aware of typical regularities

Pointer
http://www.xkcd.com/1155/

given: ~100000 Bytes

Compression of observations

- we cannot store every detail of all our memories - e.g. once I got bit by a brown dog, once by a white one
  - it would be too much data (even in hyperthymesia)
  - it would be unnecessarily clumsy to access it
  - we couldn't generalise - wouldn't know what to expect when a spotty dog shows up
- the brain needs to use such models that reflect the property statistics of the environment
  - object identities are **invariant** to a number of transformations, e.g. viewing angle or lighting differences
  - I can compress well when I'm aware of typical regularities

Pointer
http://www.xkcd.com/1155/

→ "untidy room with puma"

given: ~100000 Bytes

useful: ~40 Bytes

- The mental model of the environment
- Prediction of behaviour
- Probing the mental representations

Predictions derived from probabilistic inference

- we have to check whether the behaviour of humans and animals supports the idea that probabilistic inference takes place in the brain
- this is needed to justify any neural-level application of probabilistic models
- experiments: stimulus -> organism -> behaviour
- theory: stimulus -> perception model -> decision model -> predicted behaviour
- the simplest prediction of probabilistic models comes with the Bayes theorem: the posterior distribution is located between the prior and the likelihood
  - this means that the perception will be **biased** towards learned regularities from actual stimulus content s: sensory signal, x: perceived quantity

posterior likelihood prior

$$p(x|s) = p(s|x)\, p(x) / p(s)$$

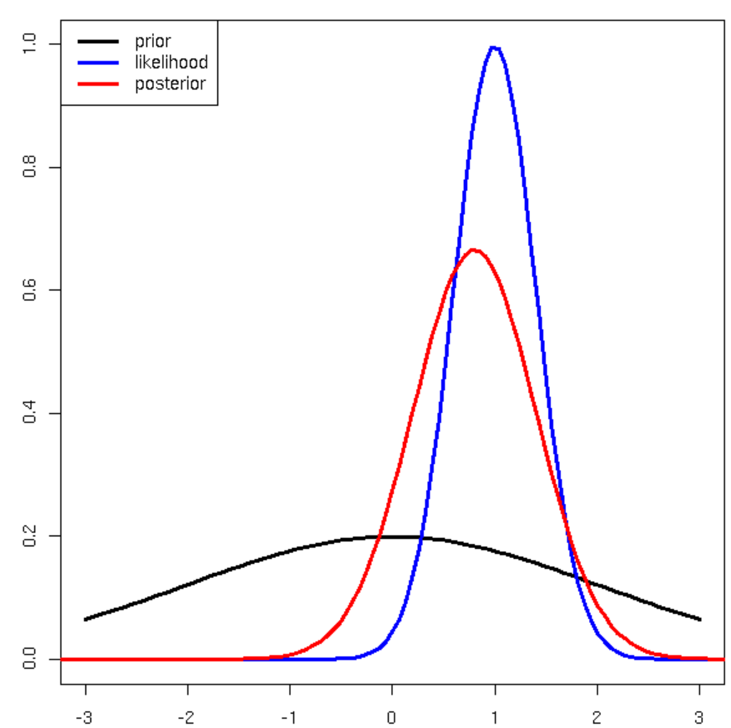

Illusions revisited

- What is the learned regularity of the environment that modulates sensory information in these images?
    - contour continuity
    - uniformity of colours and modulation by shadows

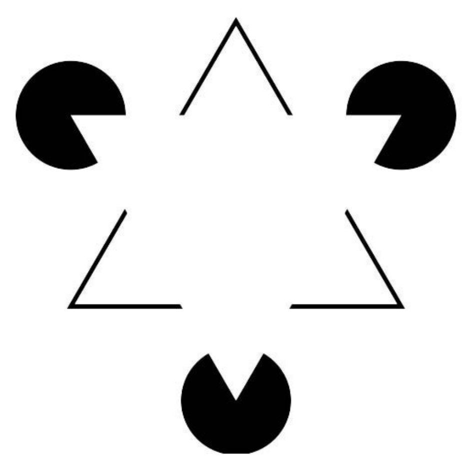

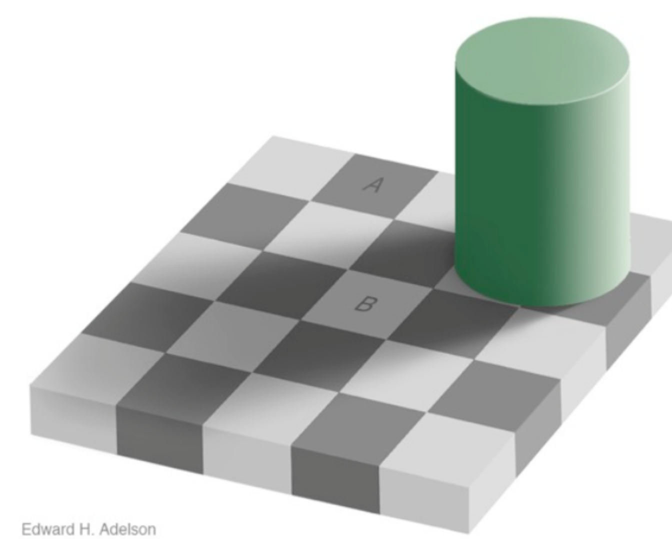

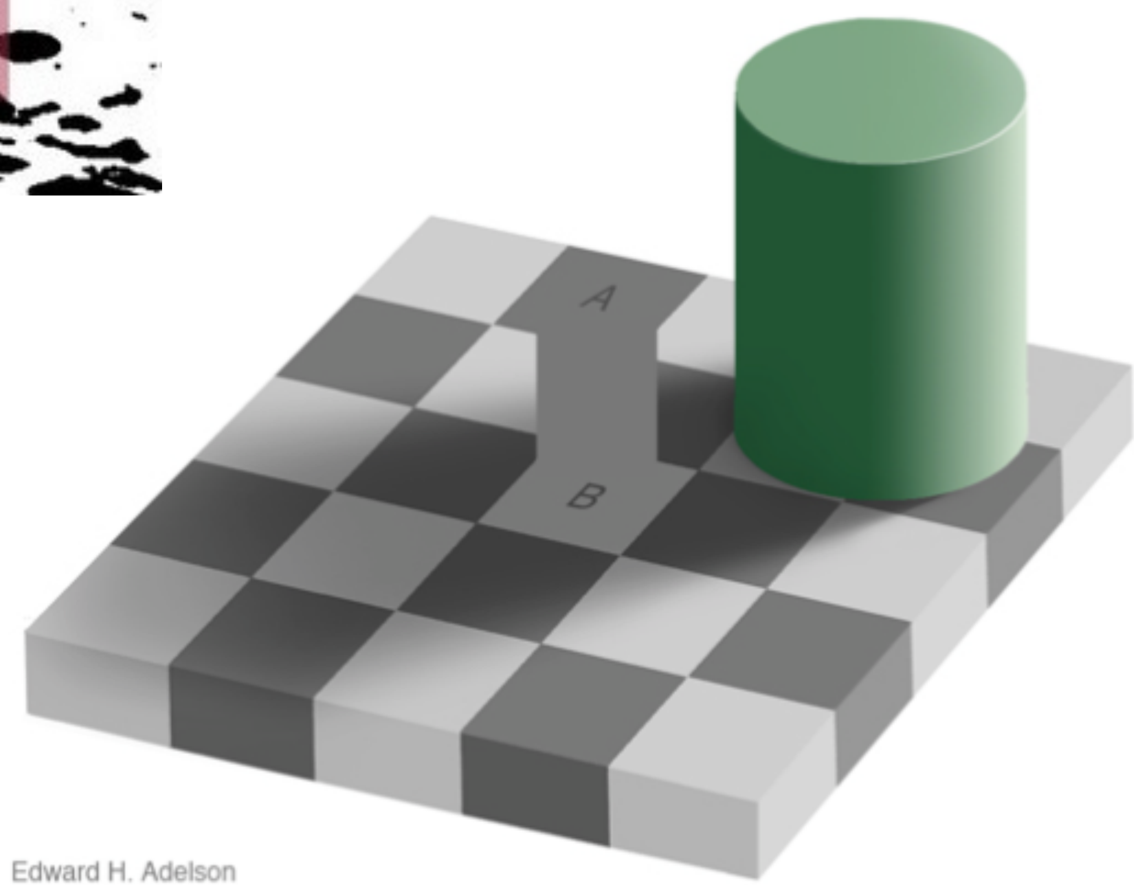

Edward H. Adelson

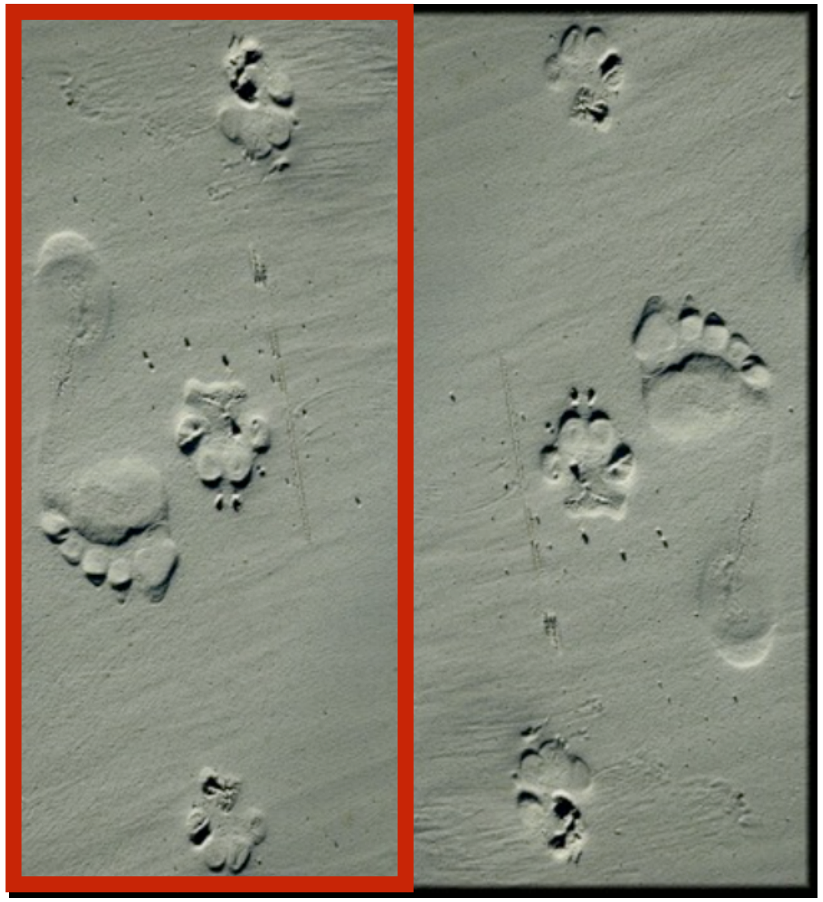

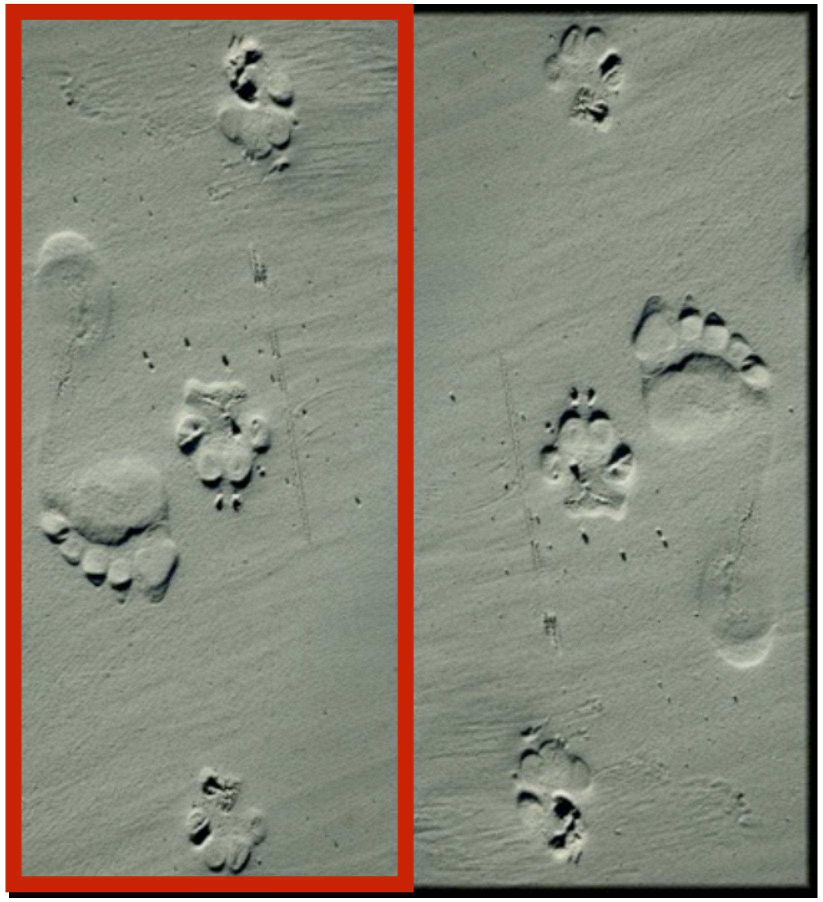

- Explanation 1: the sun is shining from below and footprints are hollow
- Explanation 2: the sun is shining from above and footprints are embossed

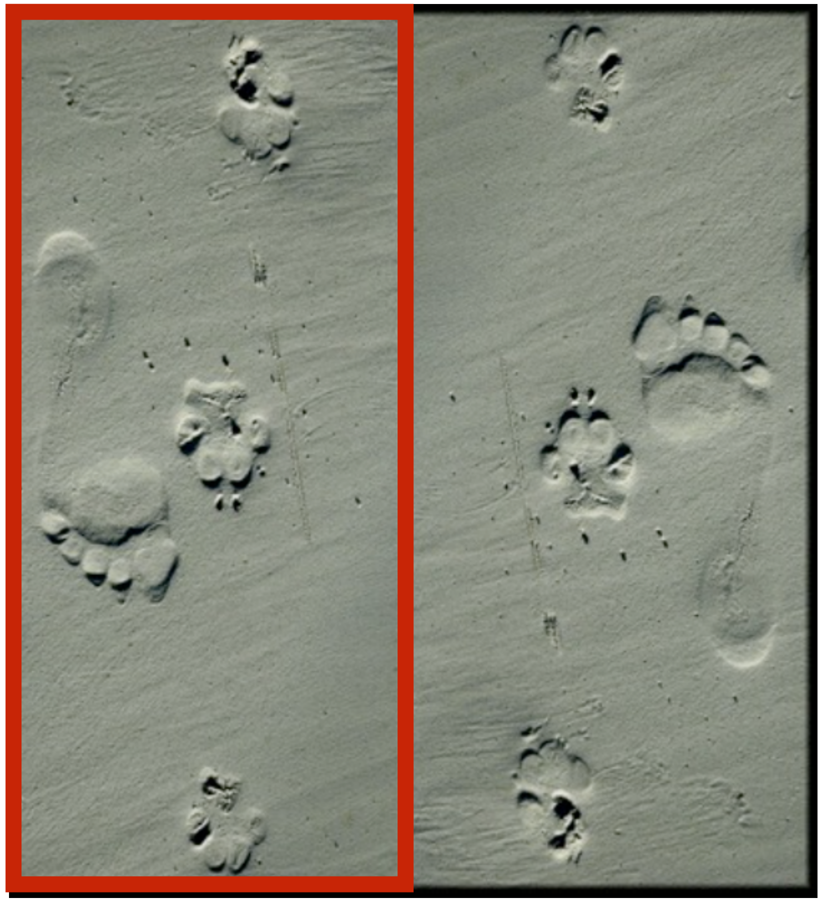

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] *California Institute of Technology 136-93, Pasadena, California 91125, USA*
[2] *Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy*
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature *neuroscience* • volume 1 no 3 • july 1998

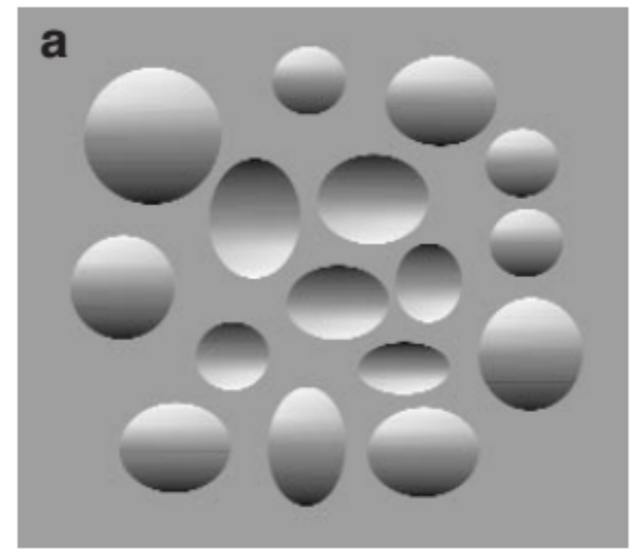

- Explanation 1: the sun is shining from below and footprints are hollow
- Explanation 2: the sun is shining from above and footprints are embossed

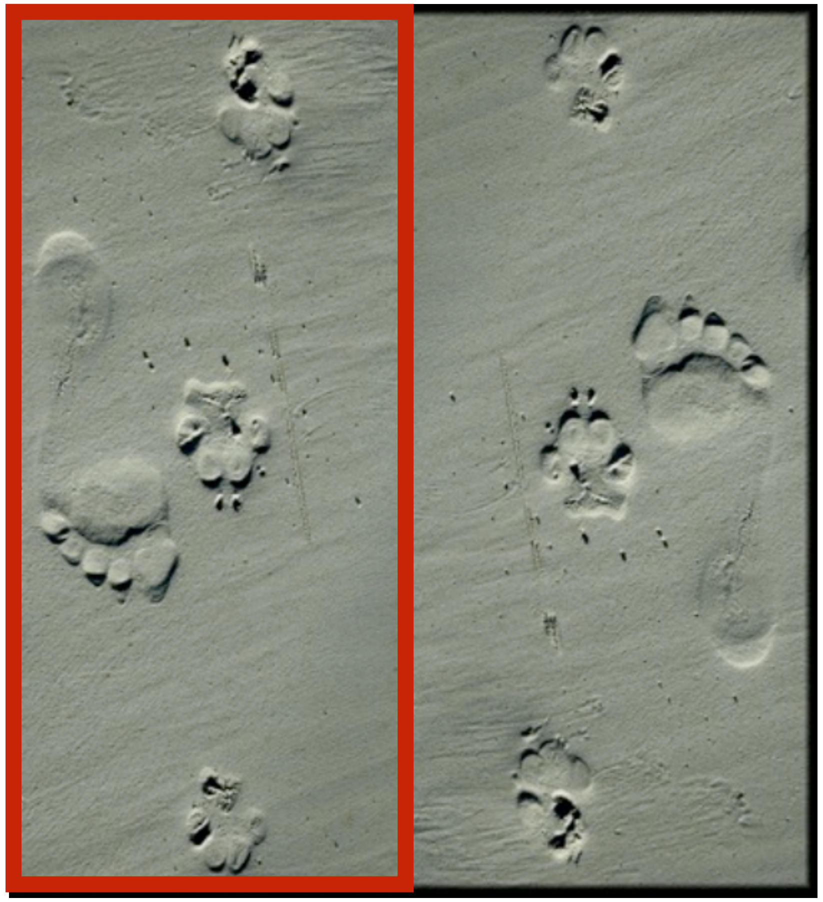

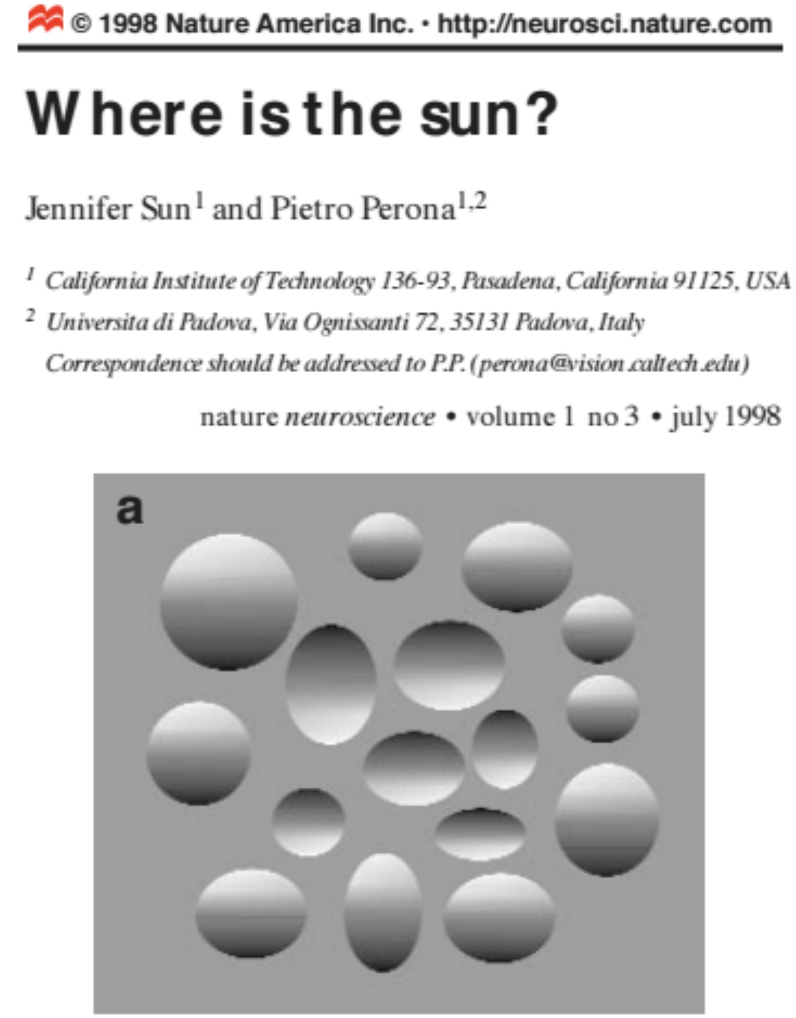
© 1998 Nature America Inc. • http://neurosci.nature.com

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] *California Institute of Technology 136-93, Pasadena, California 91125, USA*
[2] *Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy*
Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature *neuroscience* • volume 1 no 3 • july 1998

- Explanation 1: the sun is shining from below and footprints are hollow
- Explanation 2: the sun is shining from above and footprints are embossed

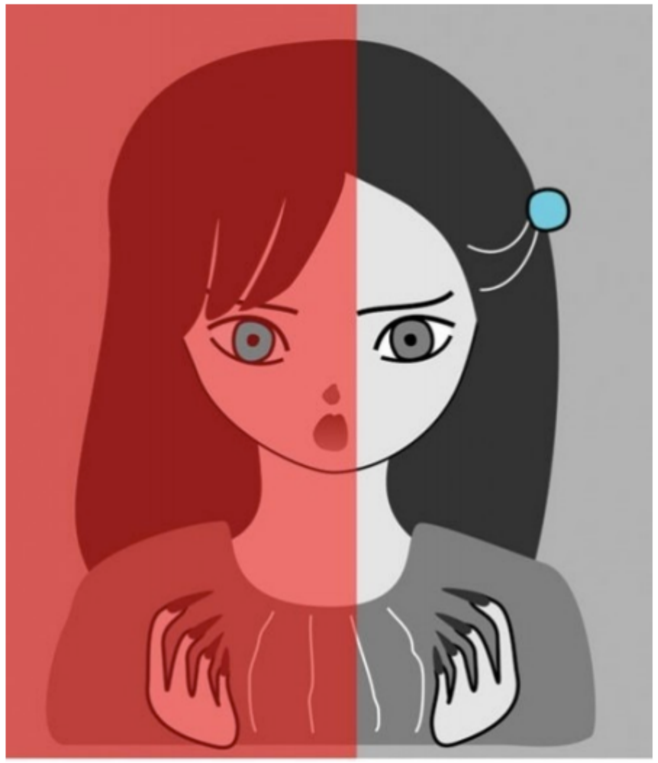

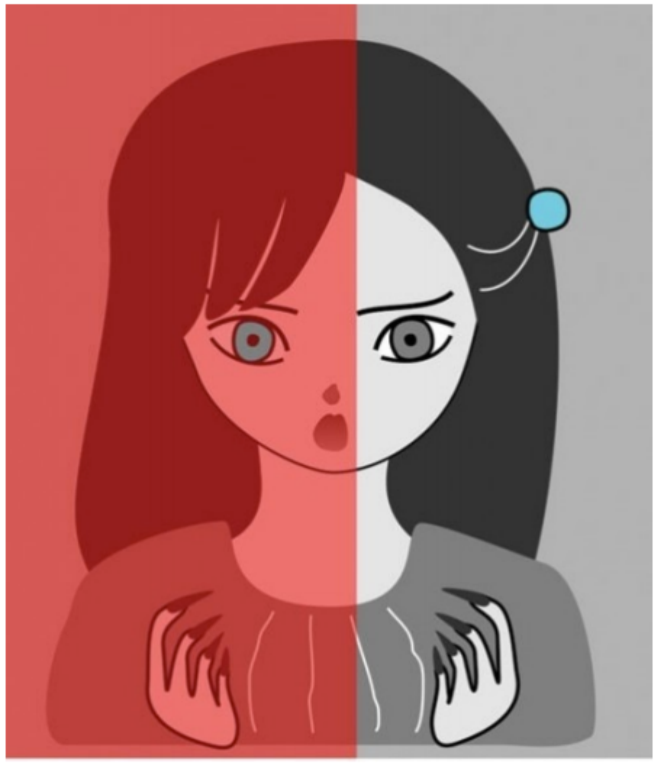

1.

+

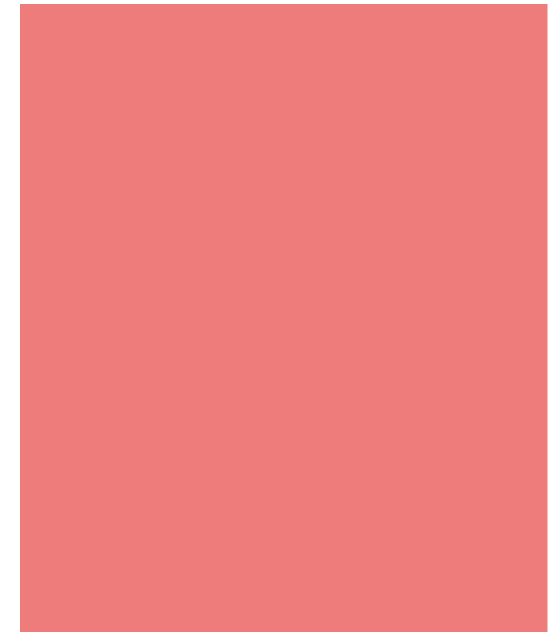

freely placed

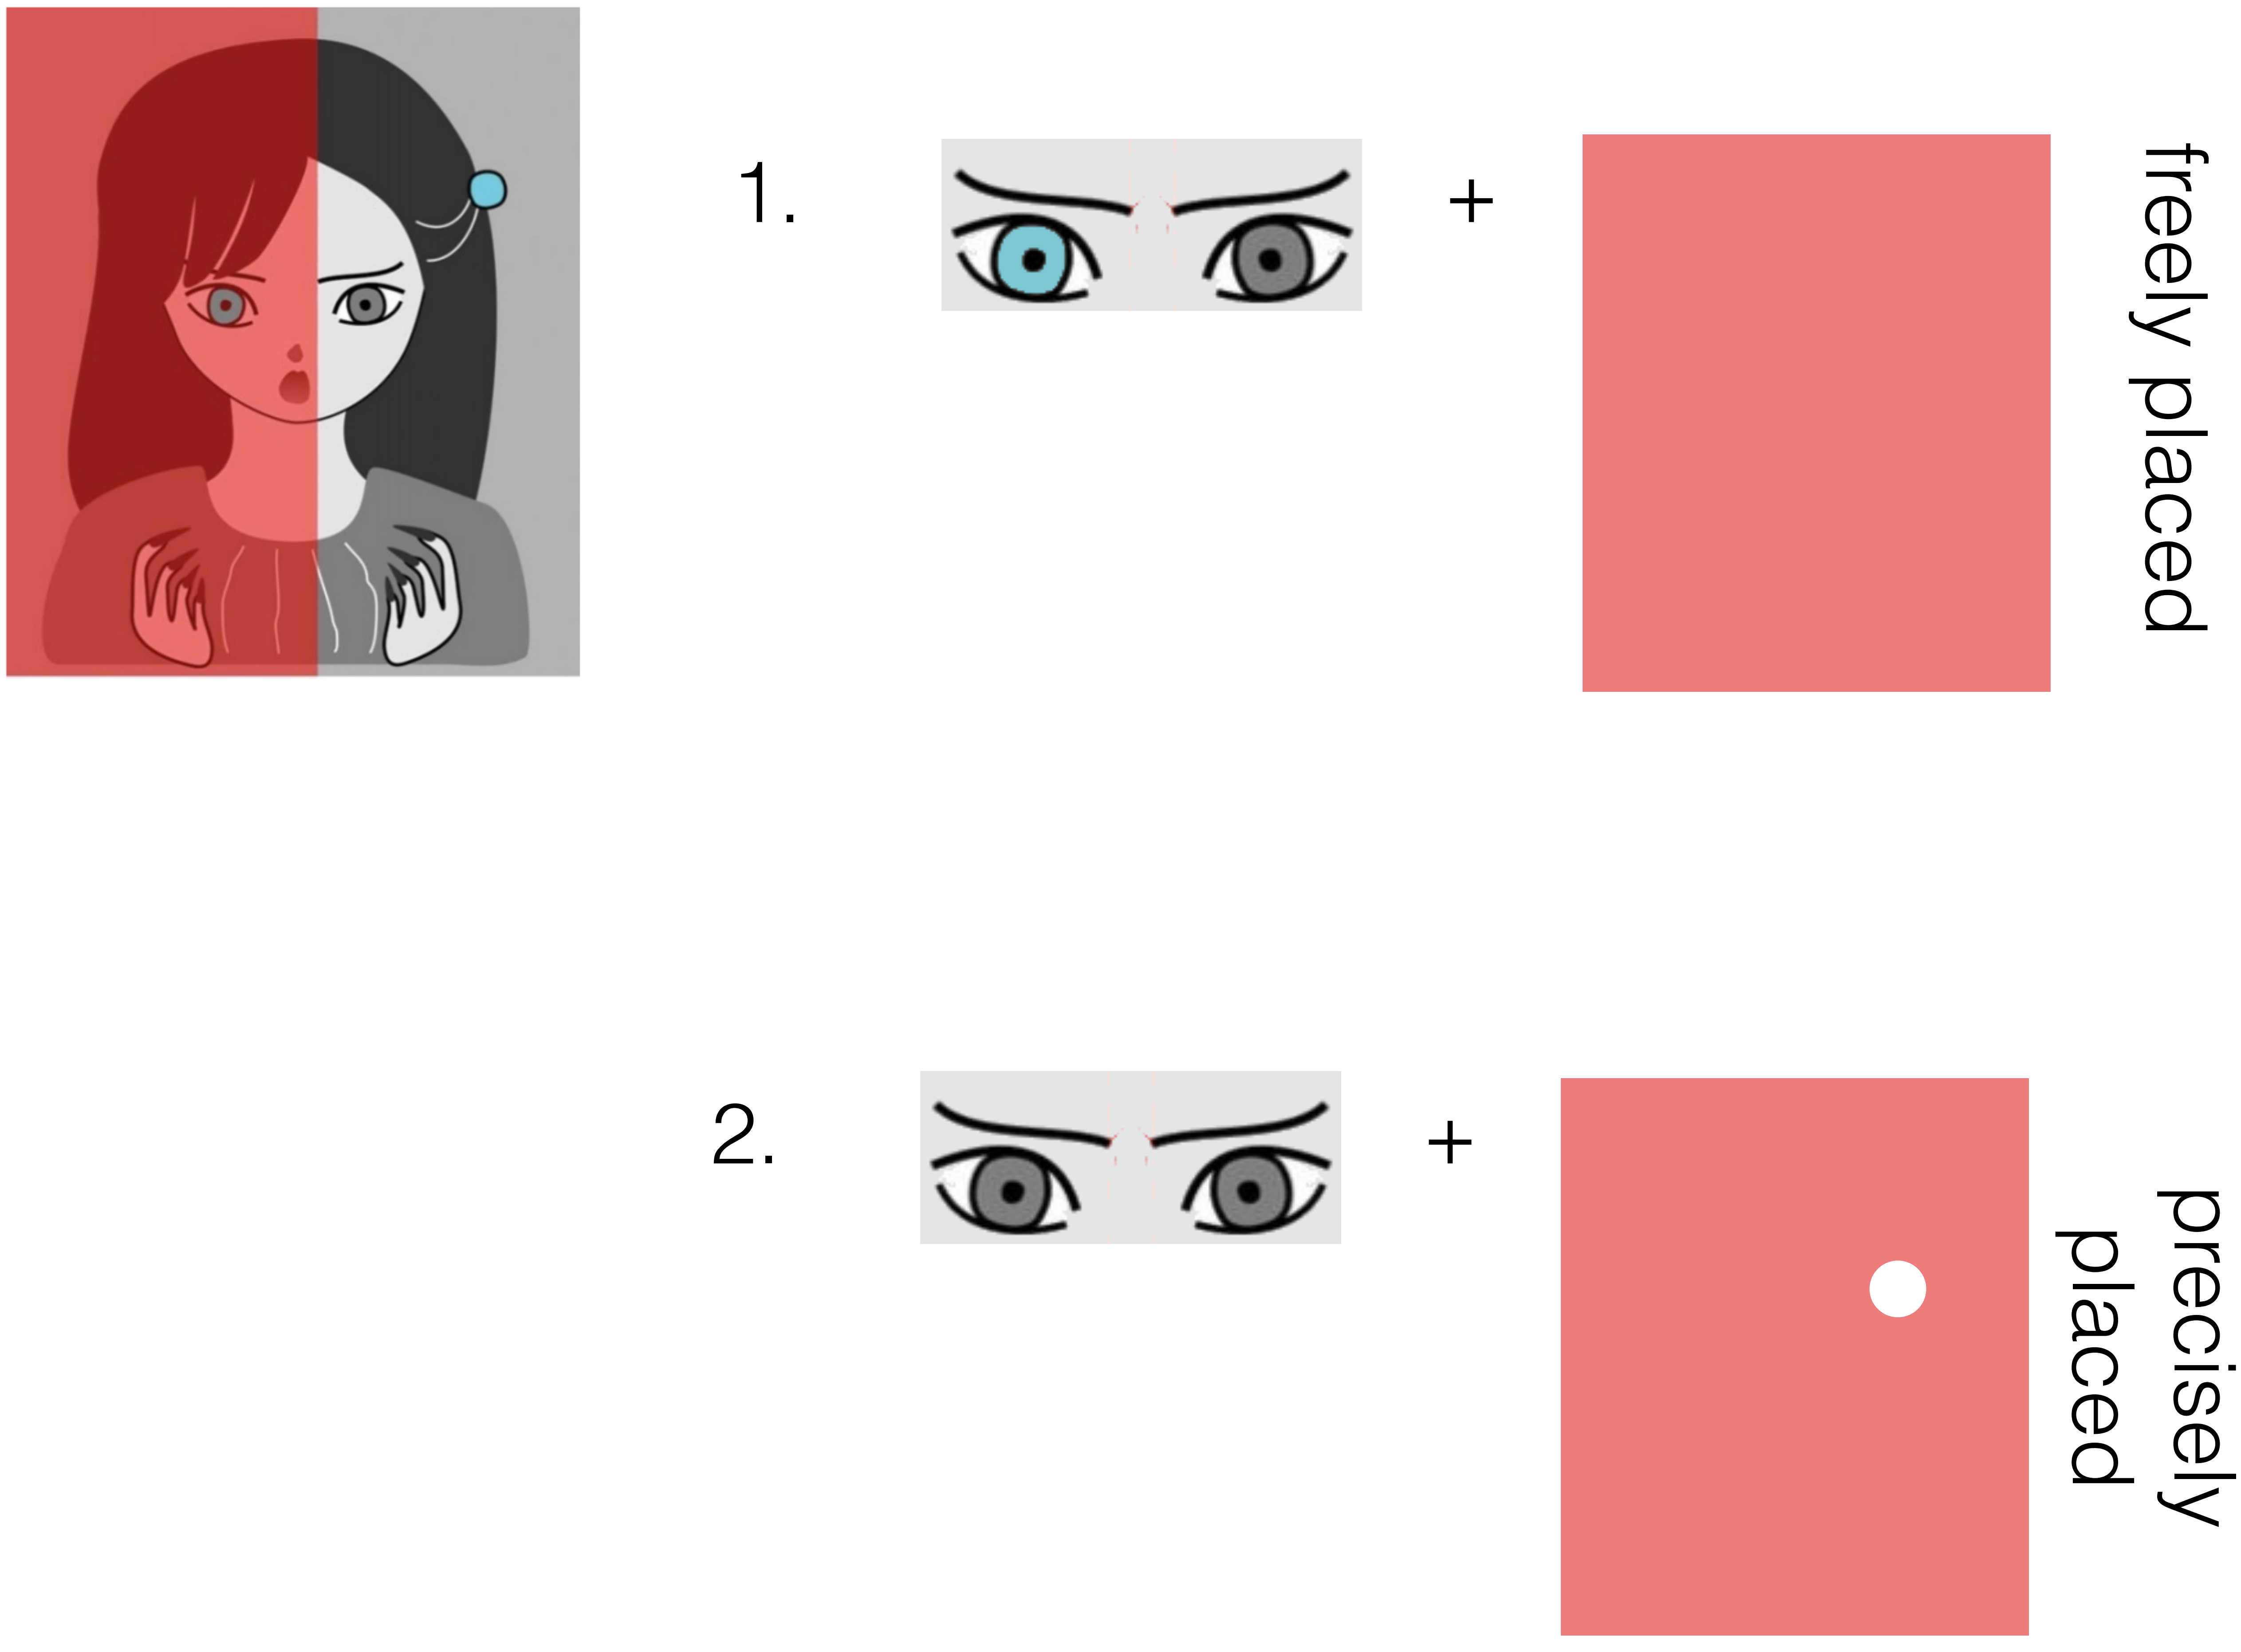
1.
+
freely placed
2.
+
precisely placed

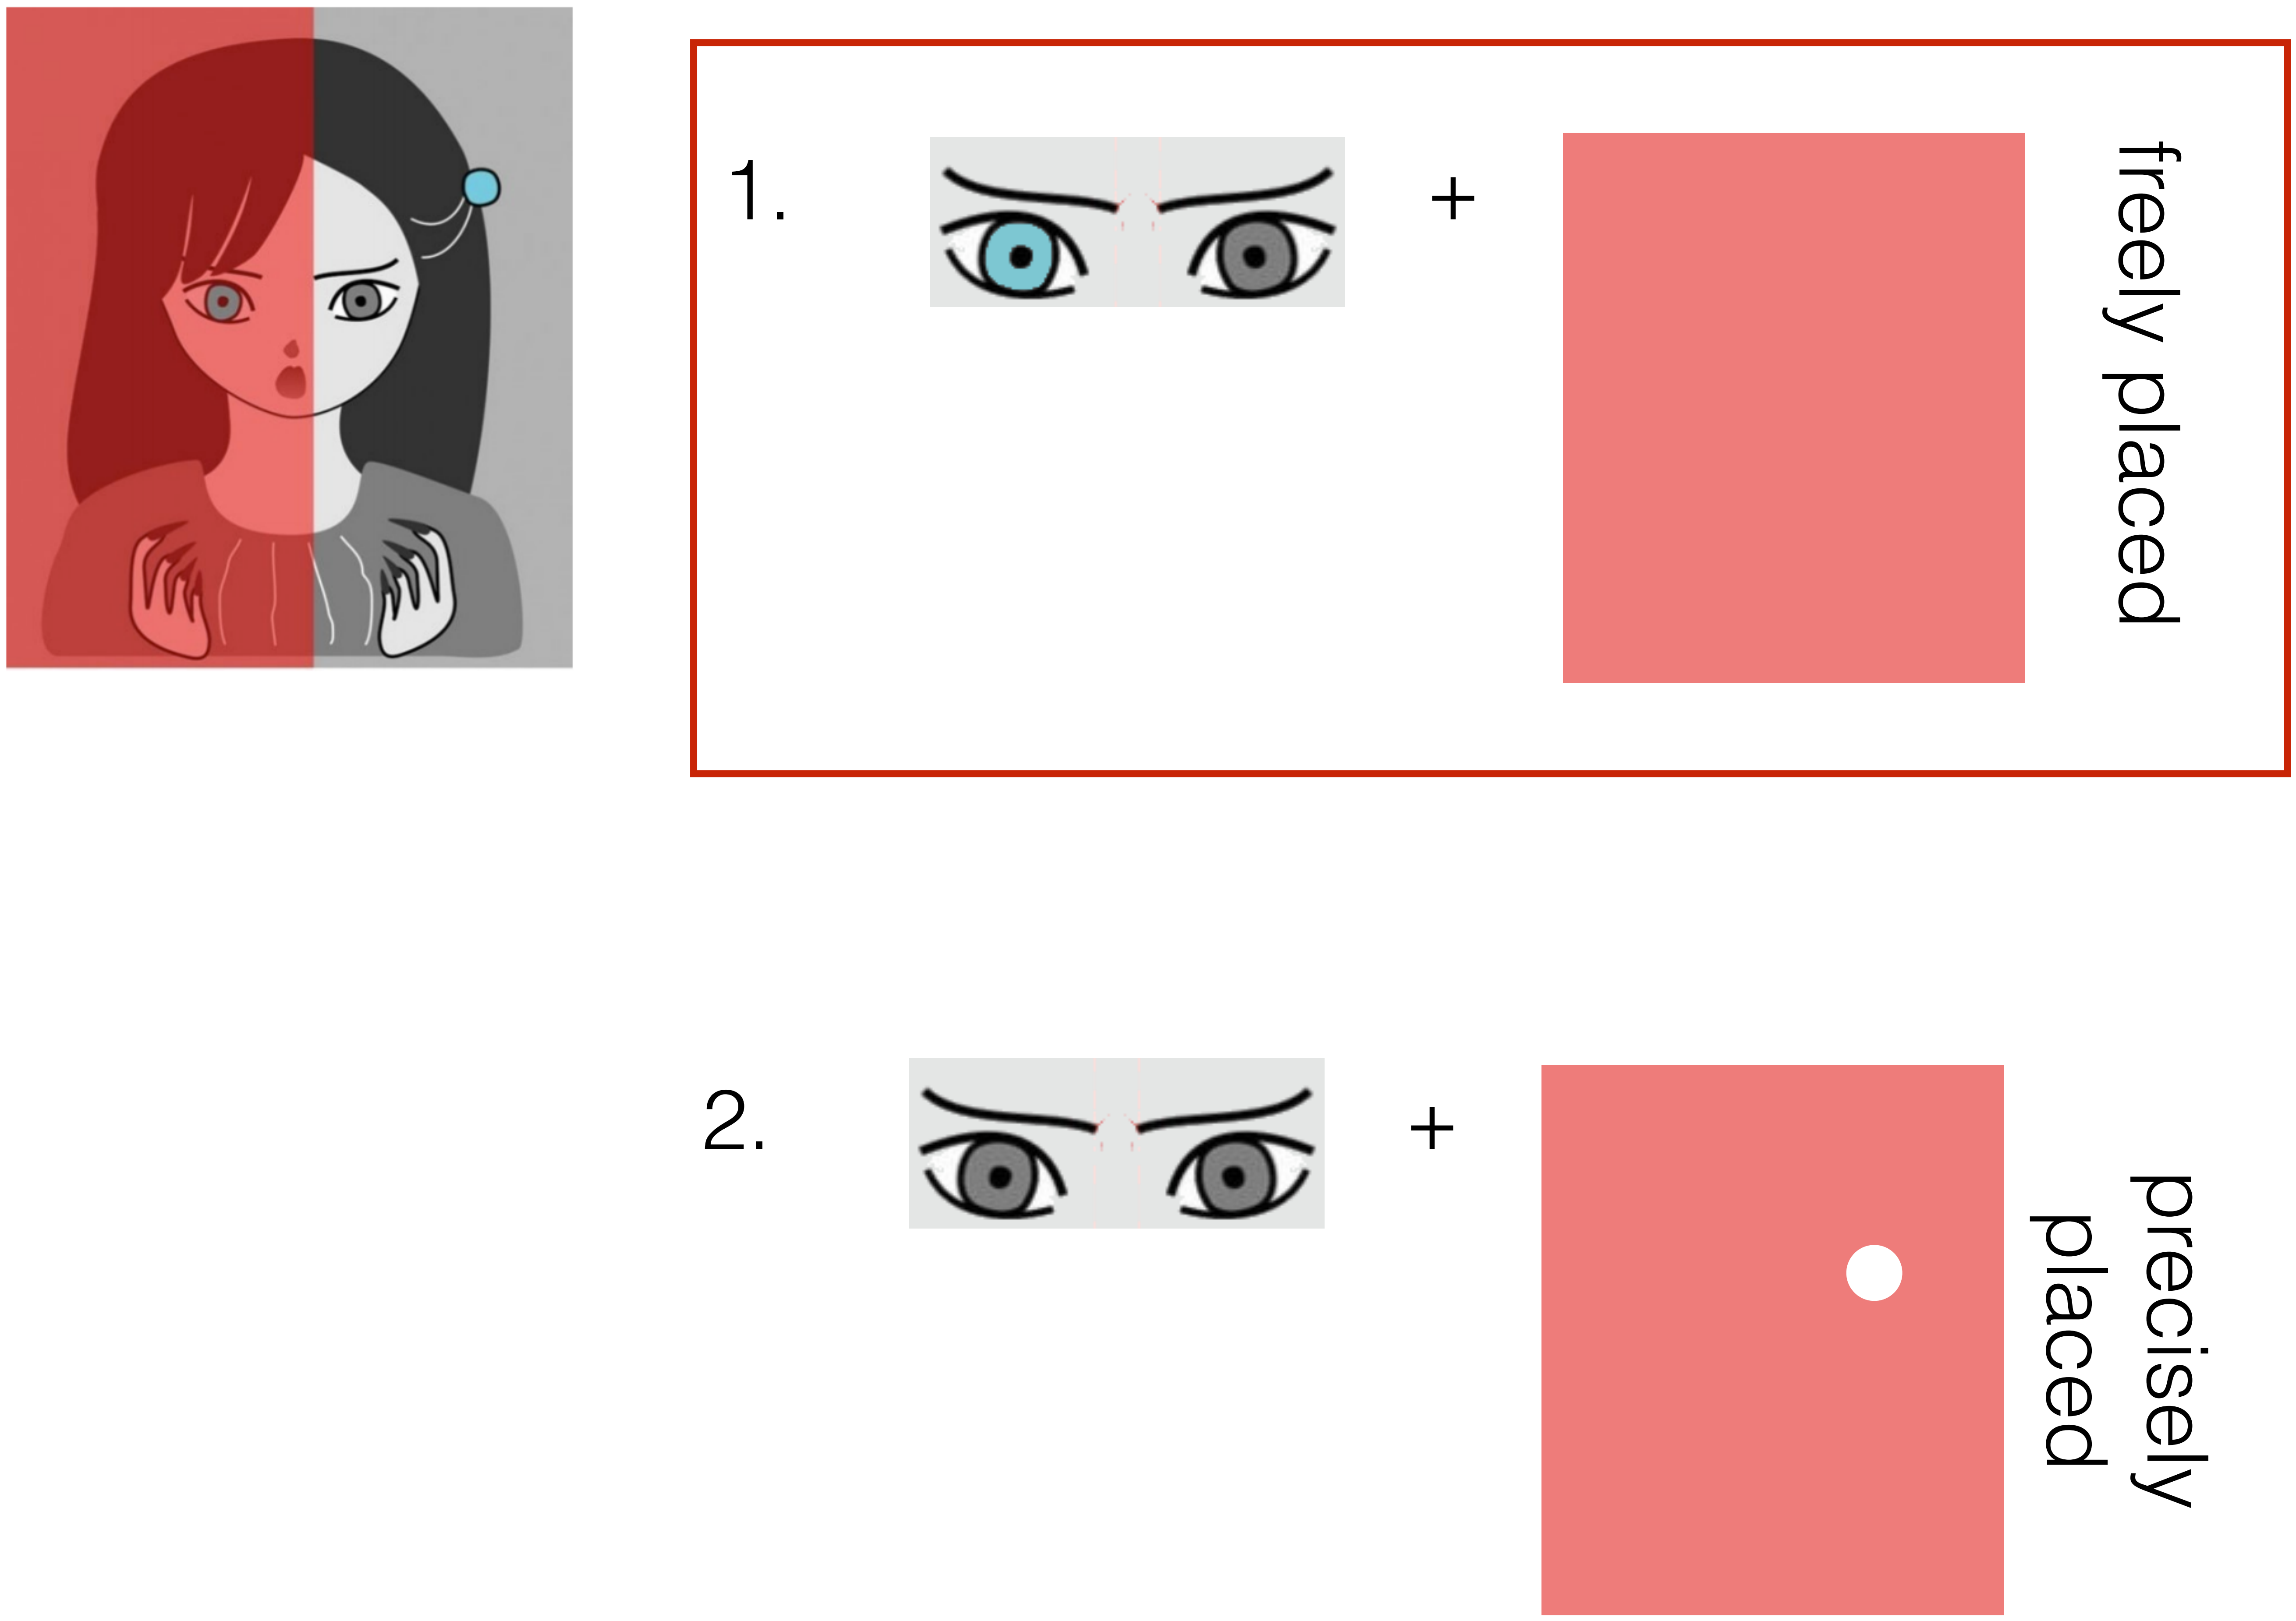

1.
+
freely placed
2.
+
precisely placed

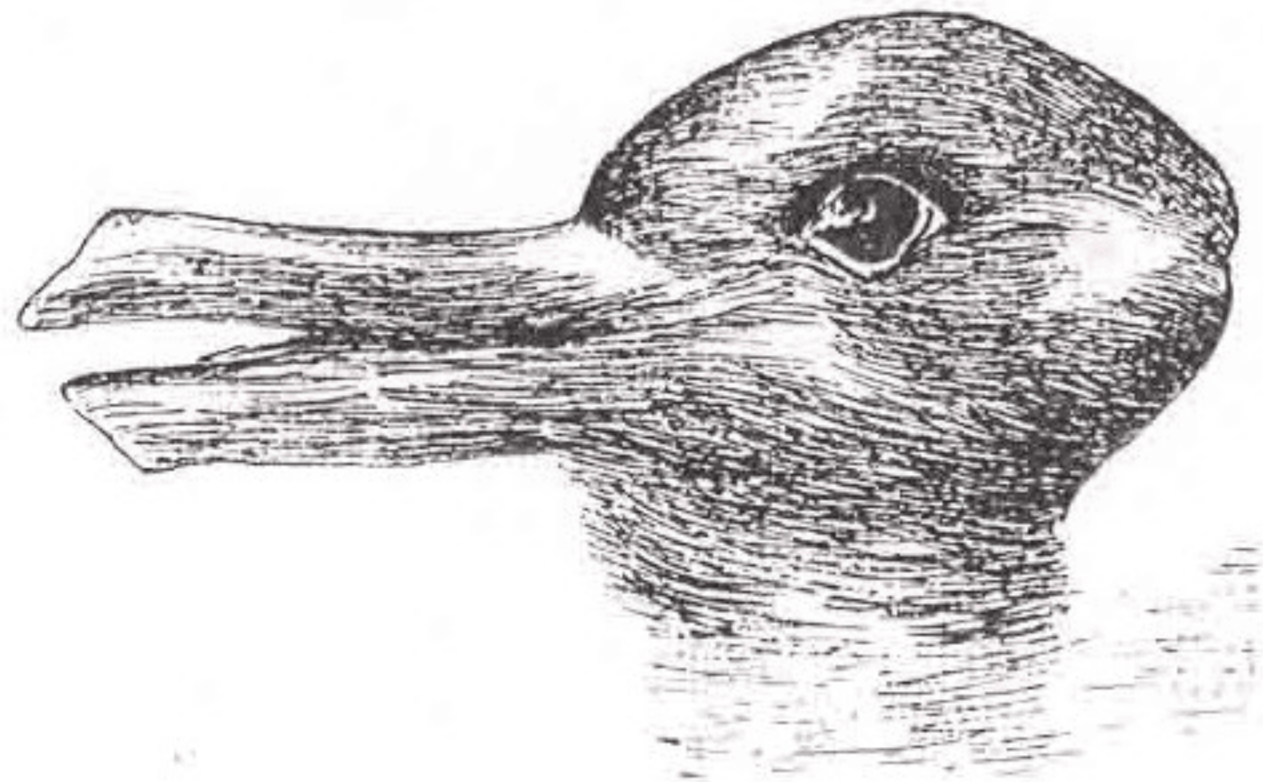

Percept Mot Skills. 1993 Apr;76(2):577-8.

The Easter bunny in October: is it disguised as a duck?

Brugger P[1], Brugger S.

Author information

Abstract

To study the influence of motivational expectancy on perception, the ambiguous drawing of a duck/rabbit was shown to 265 subjects on Easter and to 276 subjects in October. The ambiguous drawing, though perceived as a bird by a majority of subjects in October, was most frequently named a bunny on Easter. This biasing effect of expectancy upon perception was observed for young children (2 to 10 years) as well as for older subjects (11 to 93 years).

PMID: 8483671 [PubMed - indexed for MEDLINE]

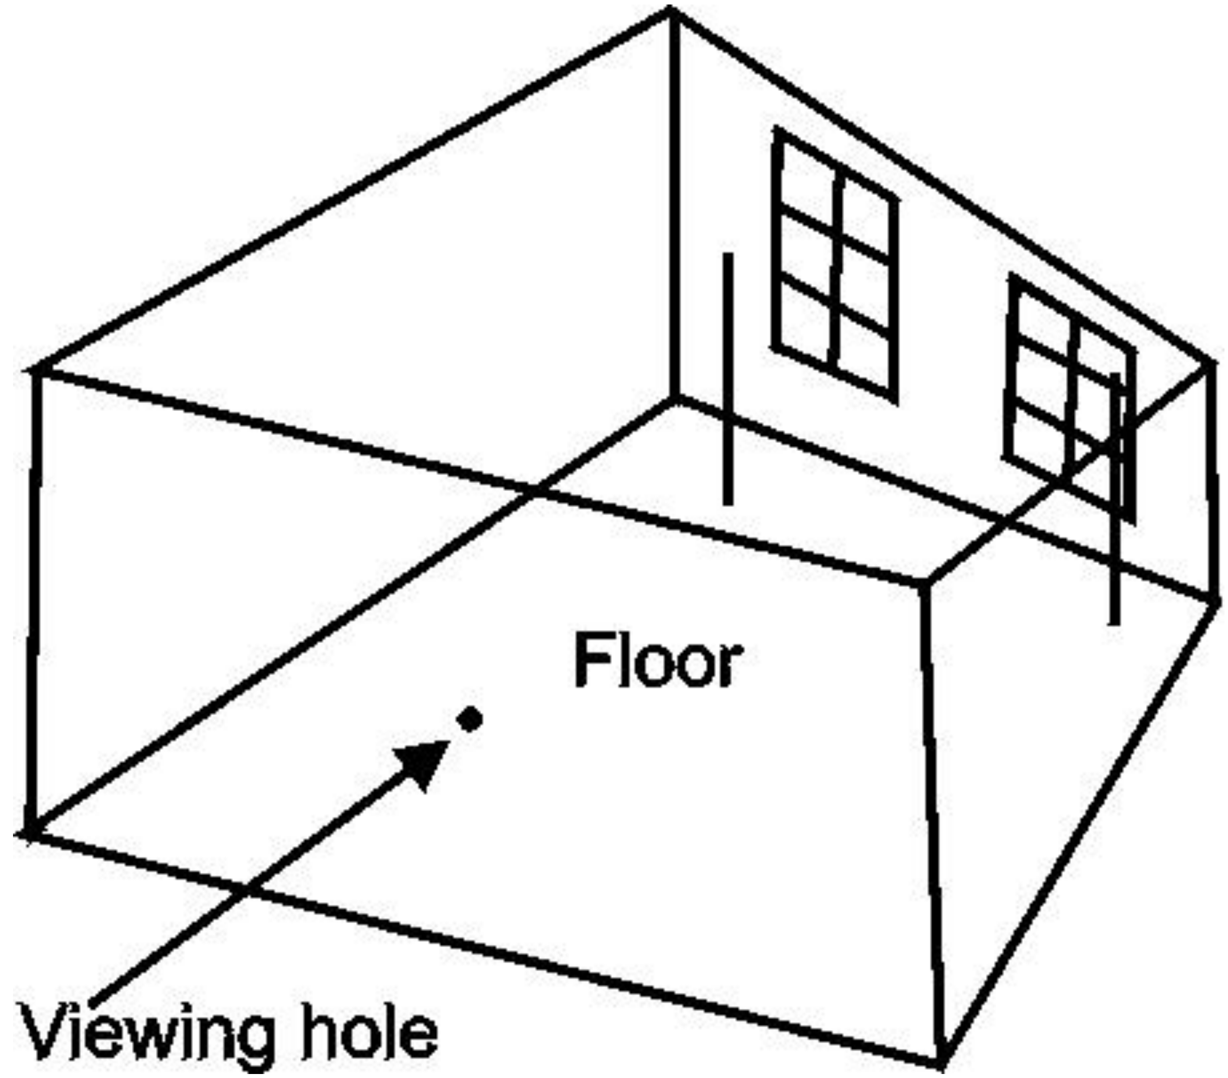
Floor
Viewing hole

Modelling everyday estimations of people

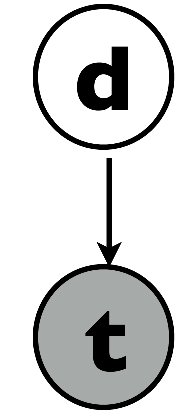

- subjects had to guess the duration of different phenomena conditioned on an observation
  - e.g. how long a representative will be in office if right now he/she has been for 3 years?
- we can assume that for these simple phenomena people have a reasonably good internal model about the distribution of durations
- if they do, their estimates should be consistent with a probabilistic model in which
  - the hidden variable is total duration d, and P(d) is the true distribution of durations
  - the observed variable is time from beginning to observation, t, and P(t|d) is U(t; 0,d), that is, it is equally likely e.g. to meet a representative at any point during his/her office time

real distribution of durations

line - prediction of prob. model using real distributions as priors
dots - estimations of people

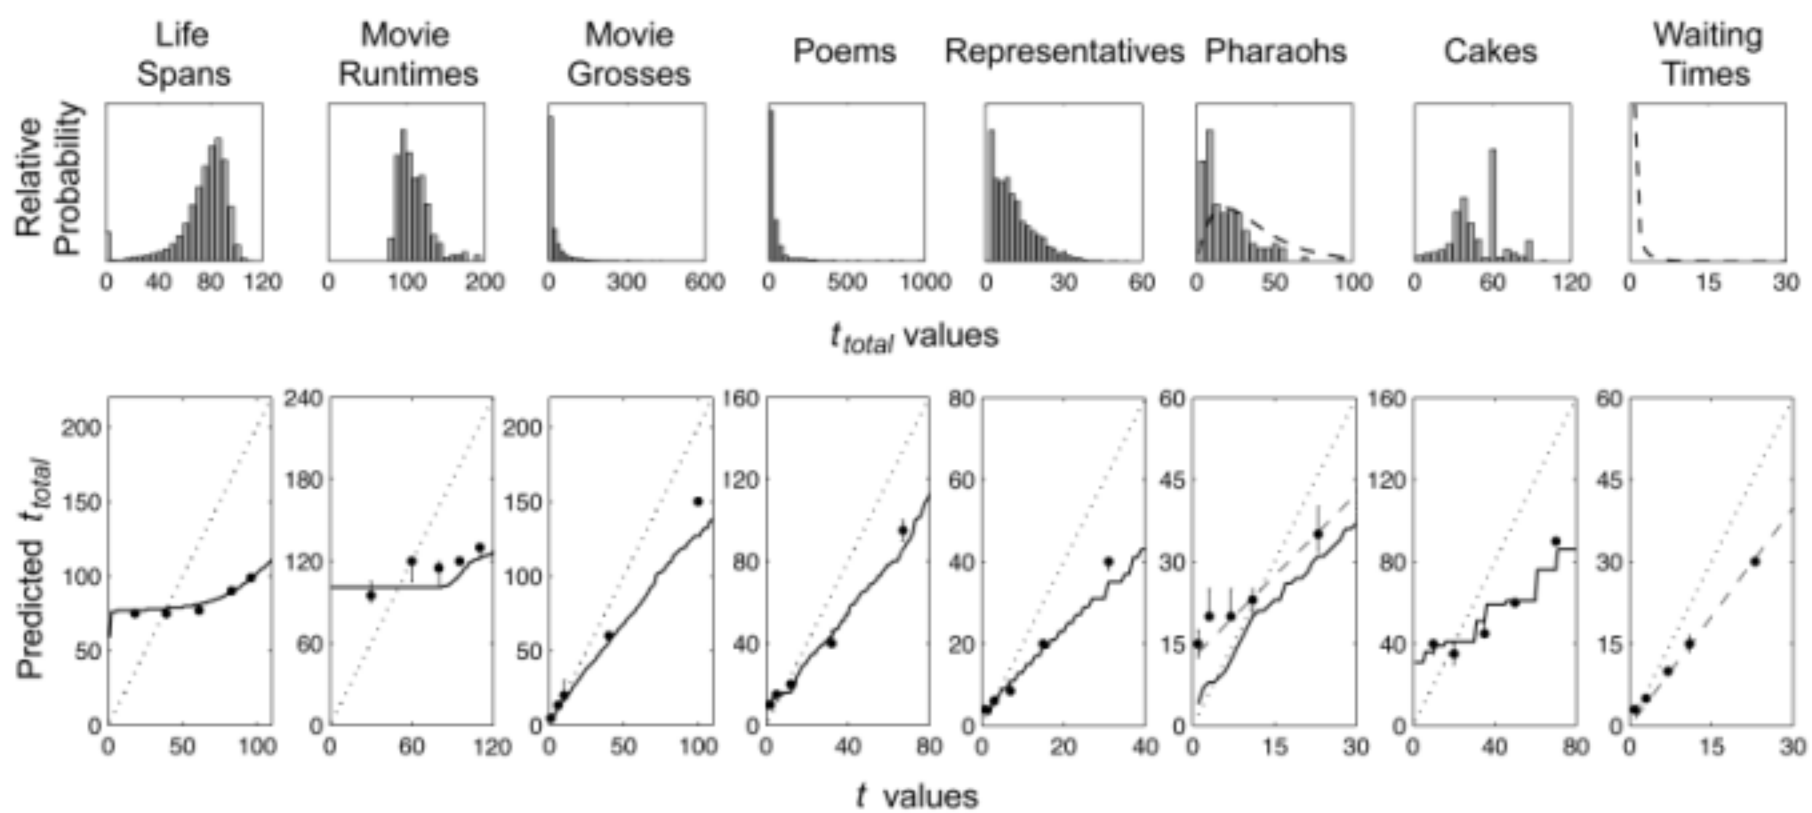

Griffiths, T. L., & Tenenbaum, J. B. (2006). Optimal predictions in everyday cognition. Psychological science, 17(9), 767-773.

Multisensory integration

- if two different sensors provide conflicting information, we can measure how they are weighted against each other

- if the brain uses probability theory, then cues have to be combined according to the prior variance of the latent variables

  - the more prior variance, the less the modality will determine where the maximum of the posterior distribution of the source location is

- this is found in human and monkey experiments

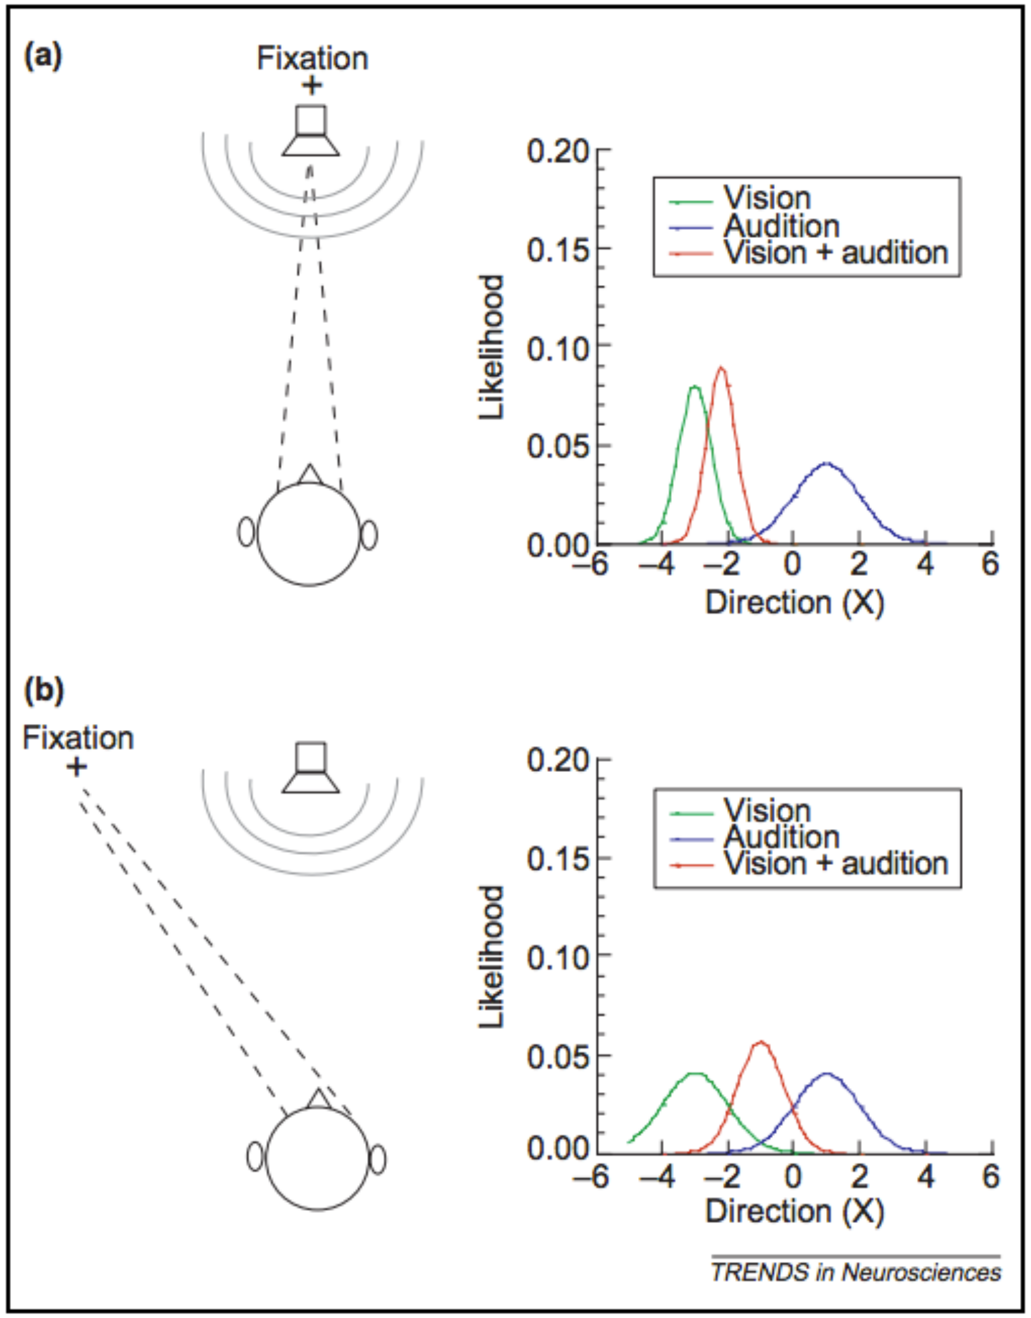

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. TRENDS in Neurosciences, 27(12), 712-719.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature, 415(6870), 429-433.

Pattern recognition

- grid patterns of symbols are assembled from building blocks
- the subjects view a lot of such patterns, but are not told what the building blocks are
- then they view pairs of patterns, one built from the building blocks, one randomly assembled
- they have to tell which one is more similar to the previously seen patterns
- human performance in identifying the patterns with similar statistics as seen before is well predicted by a probabilistic model, but not a model that only encodes pairwise symbol associations

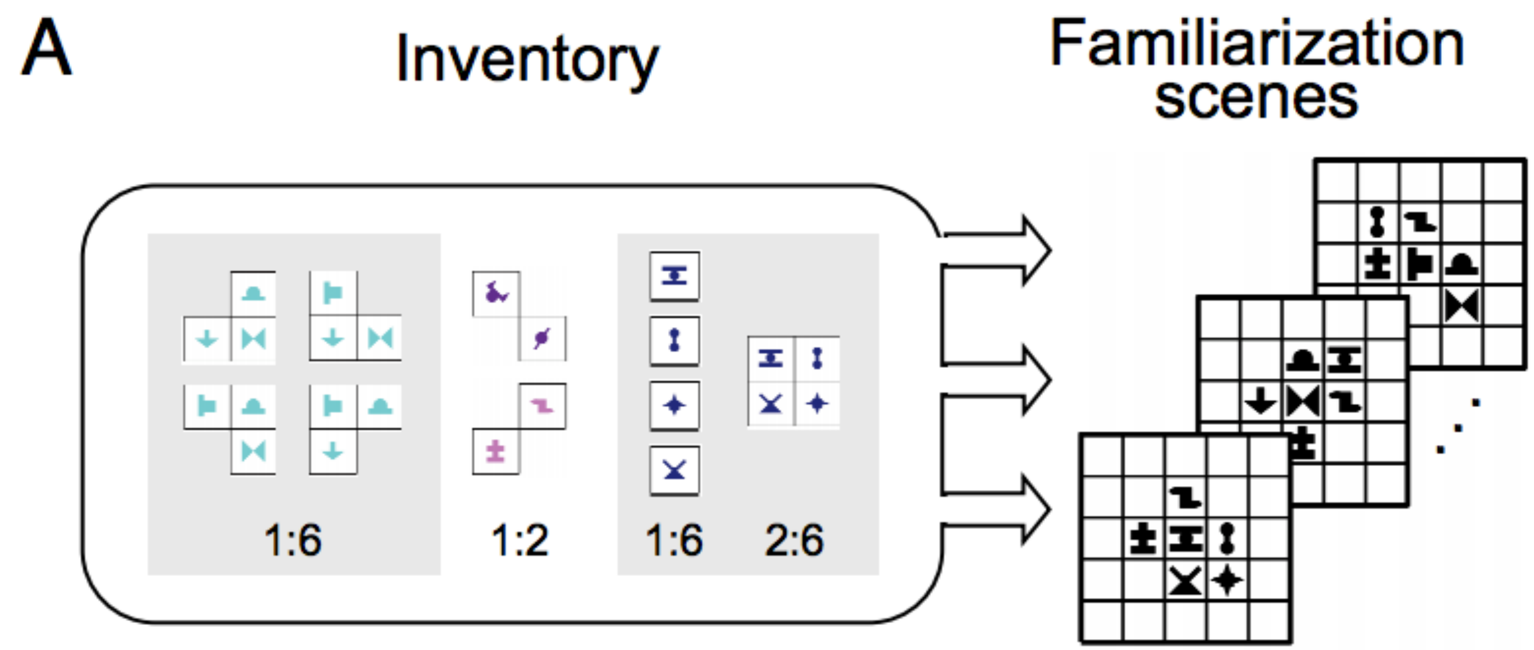

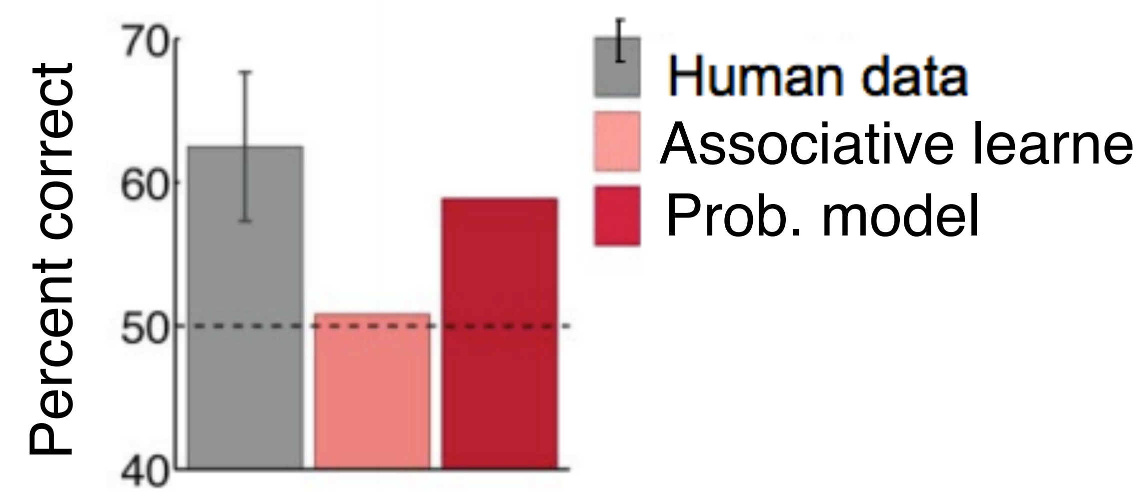

Orbán, G., Fiser, J., Aslin, R. N., & Lengyel, M. (2008). Bayesian learning of visual chunks by human observers. Proceedings of the National Academy of Sciences, 105(7), 2745-2750.

Intuitive physics

- Animals and humans need to predict the outcome of physical processes in order to make decisions about what action to take
- We can test this by building block towers in 3D simulation software and asking people whether they will fall over, and if yes, to which direction
- We can build a probabilistic model in which the probability of fall and its direction is determined by a physics simulation software used in computer games
- the predictions of such a model agrees with human estimates well

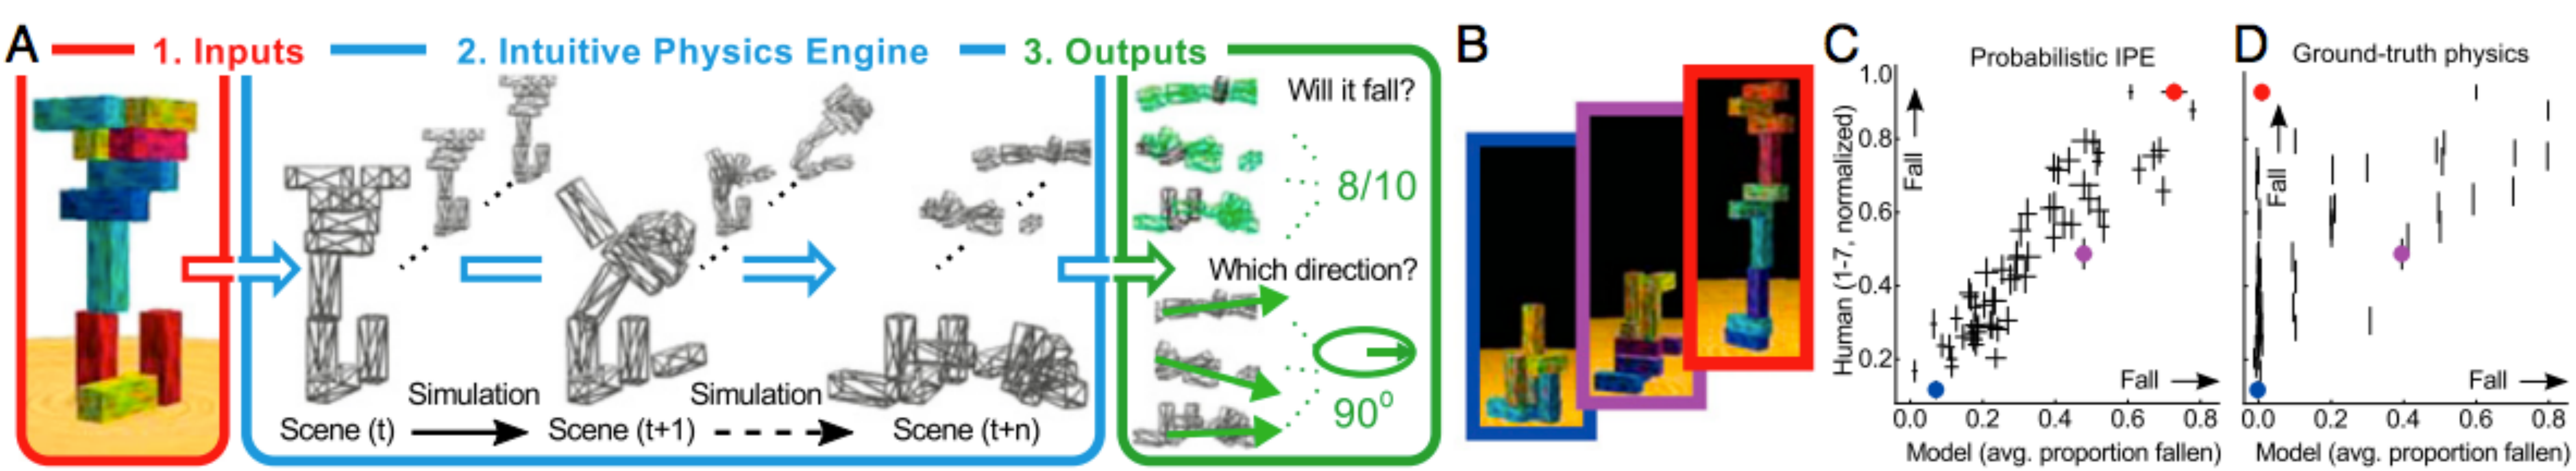

Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. Proceedings of the National Academy of Sciences, 110(45), 18327-18332.

Learning the structure of the mental model

- If we assume that probabilistic models are used in the brain, we have seen that the PMFs and PDFs can be adjusted based on observations
- But how to decide what kind of (latent) variables to use in the first place - i.e. what are the useful concepts?
- This can also be learned from observations
- We regard the structure of the model, the directed graph, as the object that has to be learned

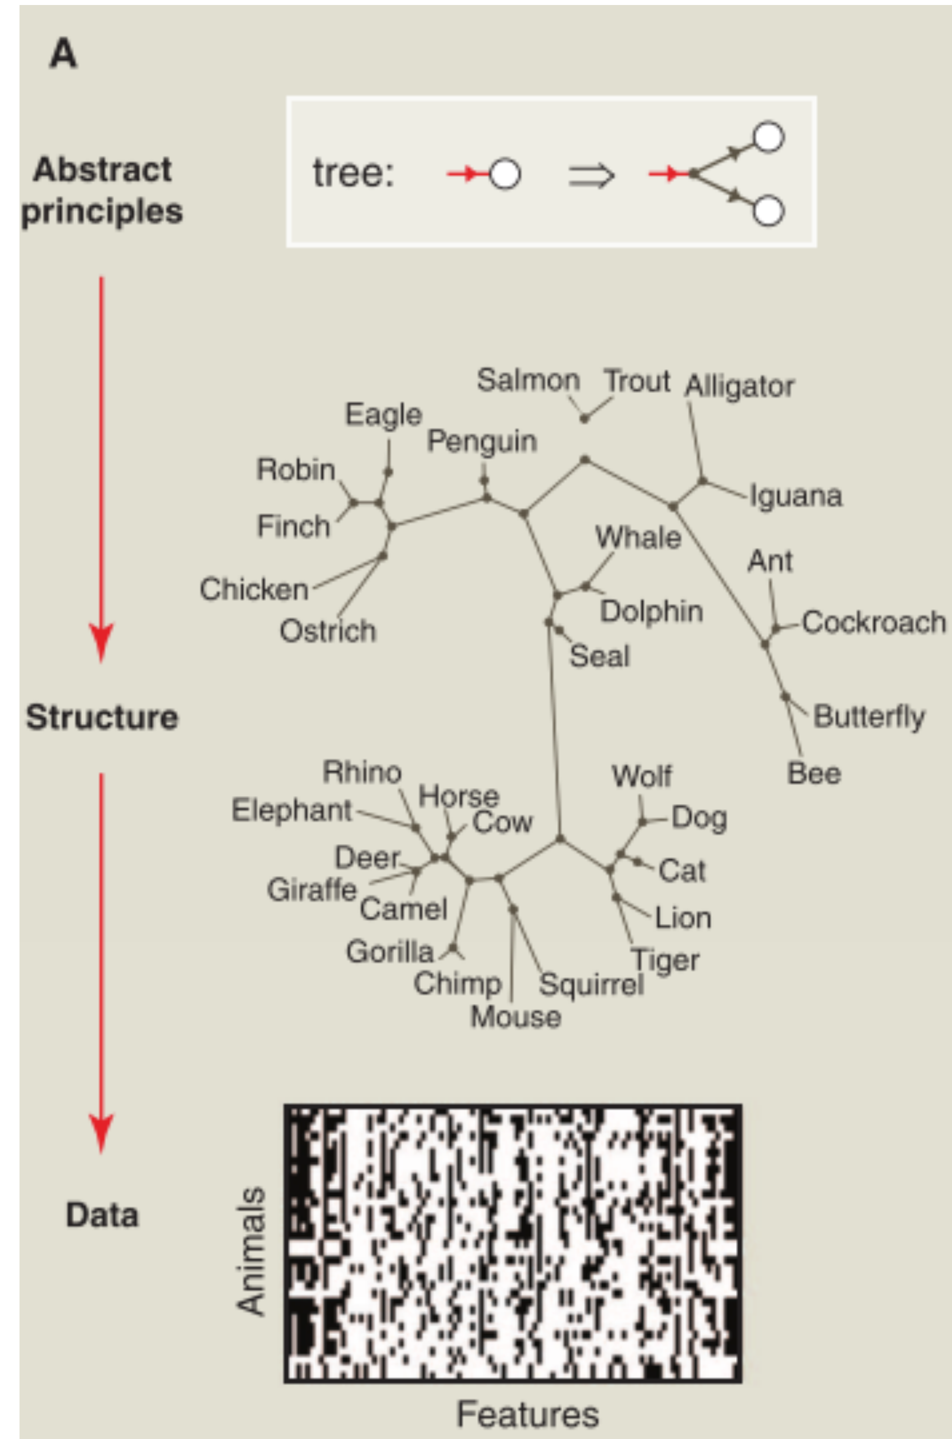

Recommended watching
http://videolectures.net/aaai2012_tenenbaum_grow_mind/

Pointer
https://probmods.org/

Predicting human taxonomy building

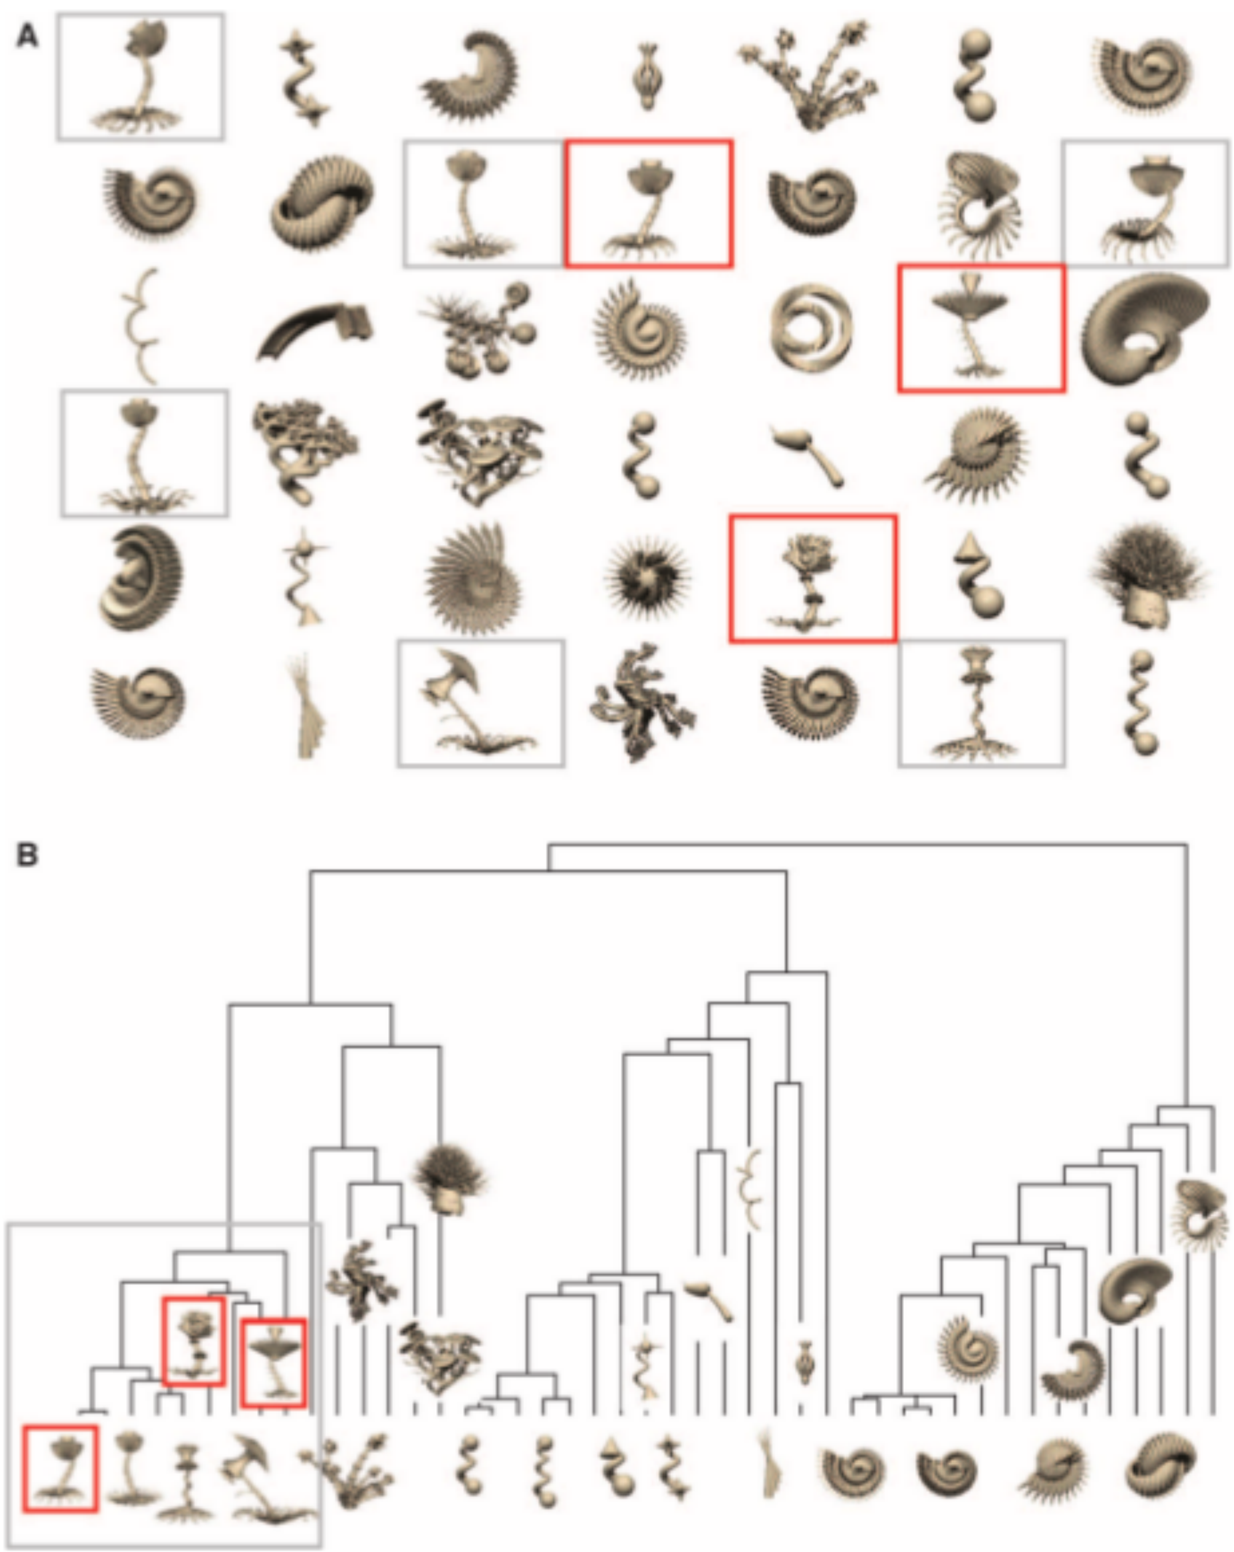

Tenenbaum et al, 2011

- Using observations about specific objects, we have to infer what the useful categories are
- Children learn languages similarly
- From a very few instances, they are able to build a useful structure
- A tree-like model structure is typical for many kind of stimuli

- The mental model of the environment
- Prediction of behaviour
- Probing the mental representations

Accessibility of the mental model to the experimenter

- how can we discover something about the distributions in the mental model instead of assuming it?
- can we say something about the exact probabilities the brain assigns to real-world quantities?
- if the mental model hypothesis is correct, then we have to be able to find a task-invariant representation of some real-world quantities
- the question is whether a signature of such a representation can be extracted from behavioural data, without probing the electrical activity of neurons

Face stimuli in psychophysics

- we need a stimulus type that our experimental subjects are experts in - that is, they have a detailed representation of the different possible stimuli

- everyone would have a slightly different learned statistics of faces based on their family&friends

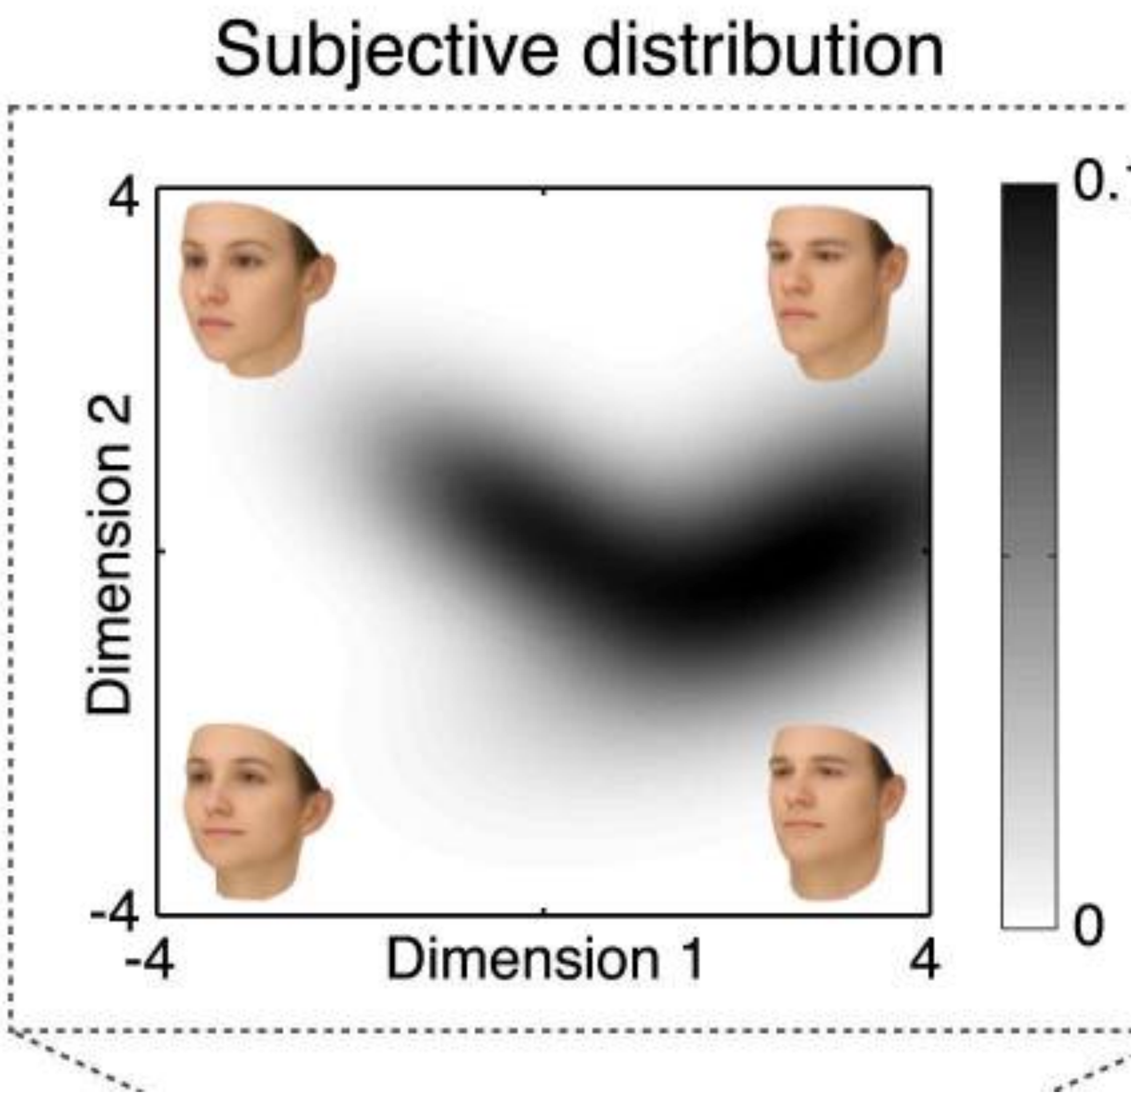

Task-independent mental representations

- An important question about mental models is whether we learn the information that is needed for each task independently, or we build a common body of knowledge that we can reuse in every new task
- the biases subjects show reveal what they think is typical and what isn't

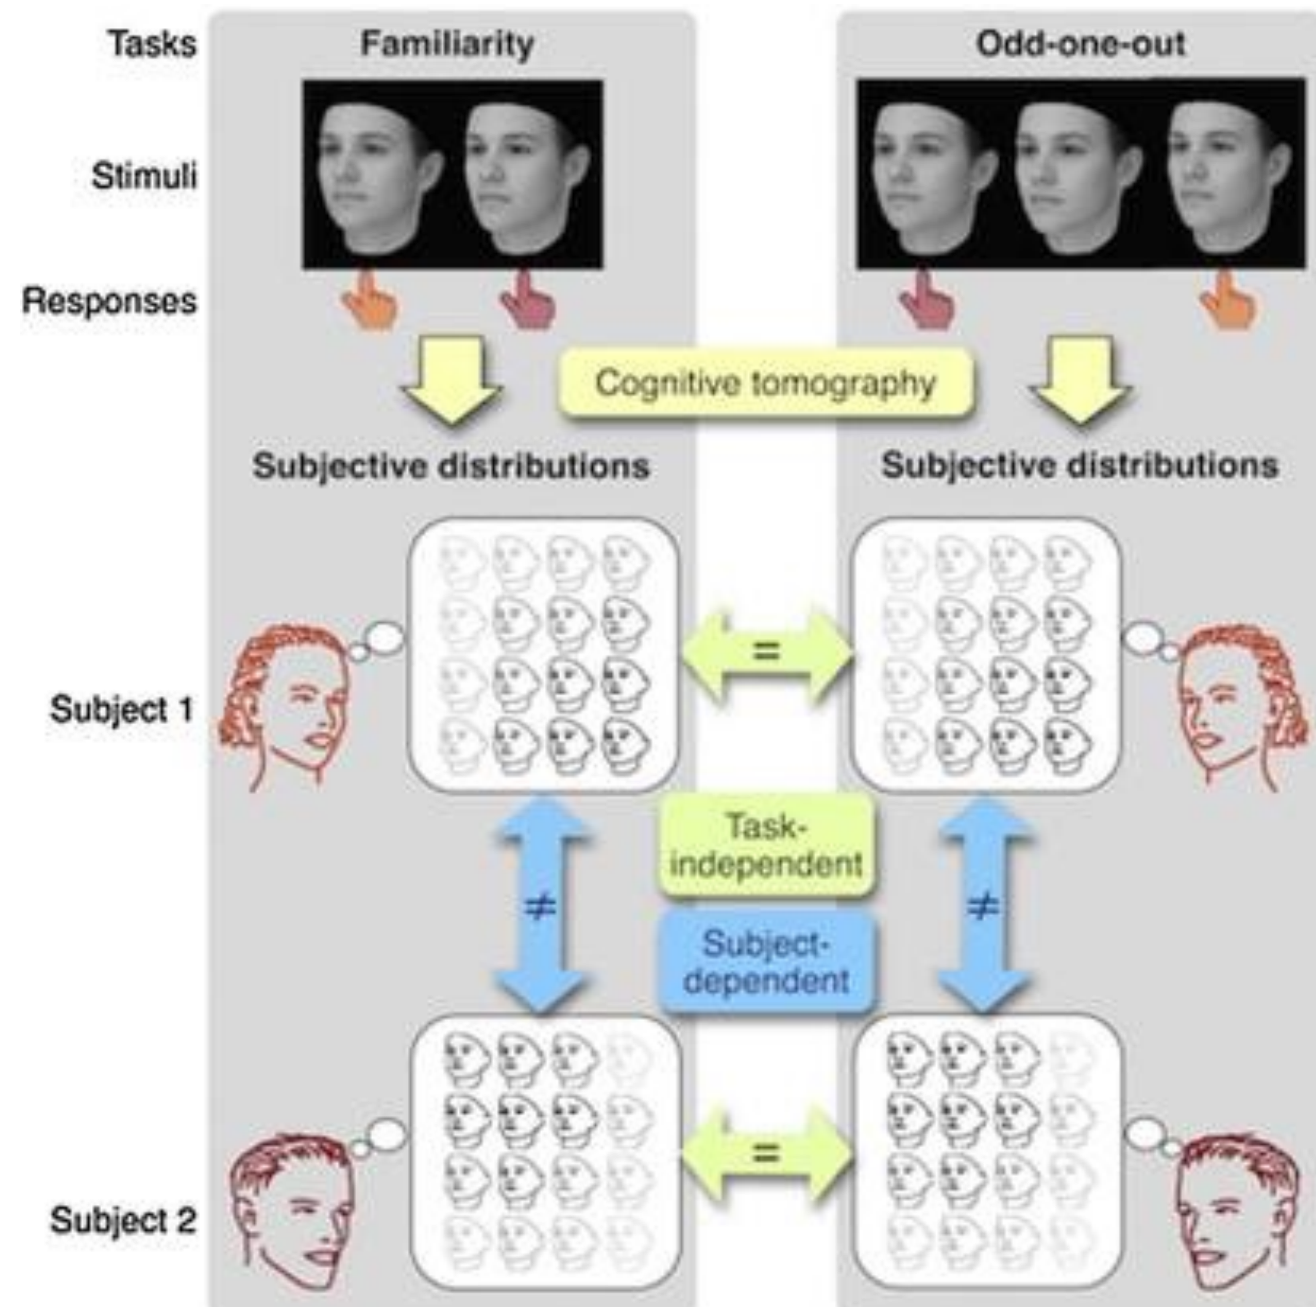

Inferring the face feature distribution of subjects

- First we have to assume a probabilistic model for face perception, that uses the internal model to assess the a priori probability of face features
- then we have to augment it with a simple decision making model to connect the perceptual results to the measured answers from the experiment
- One can invert such a model, so that given the measured answers to a lot of different stimuli, the shape of the mental PDF over facial features can be inferred

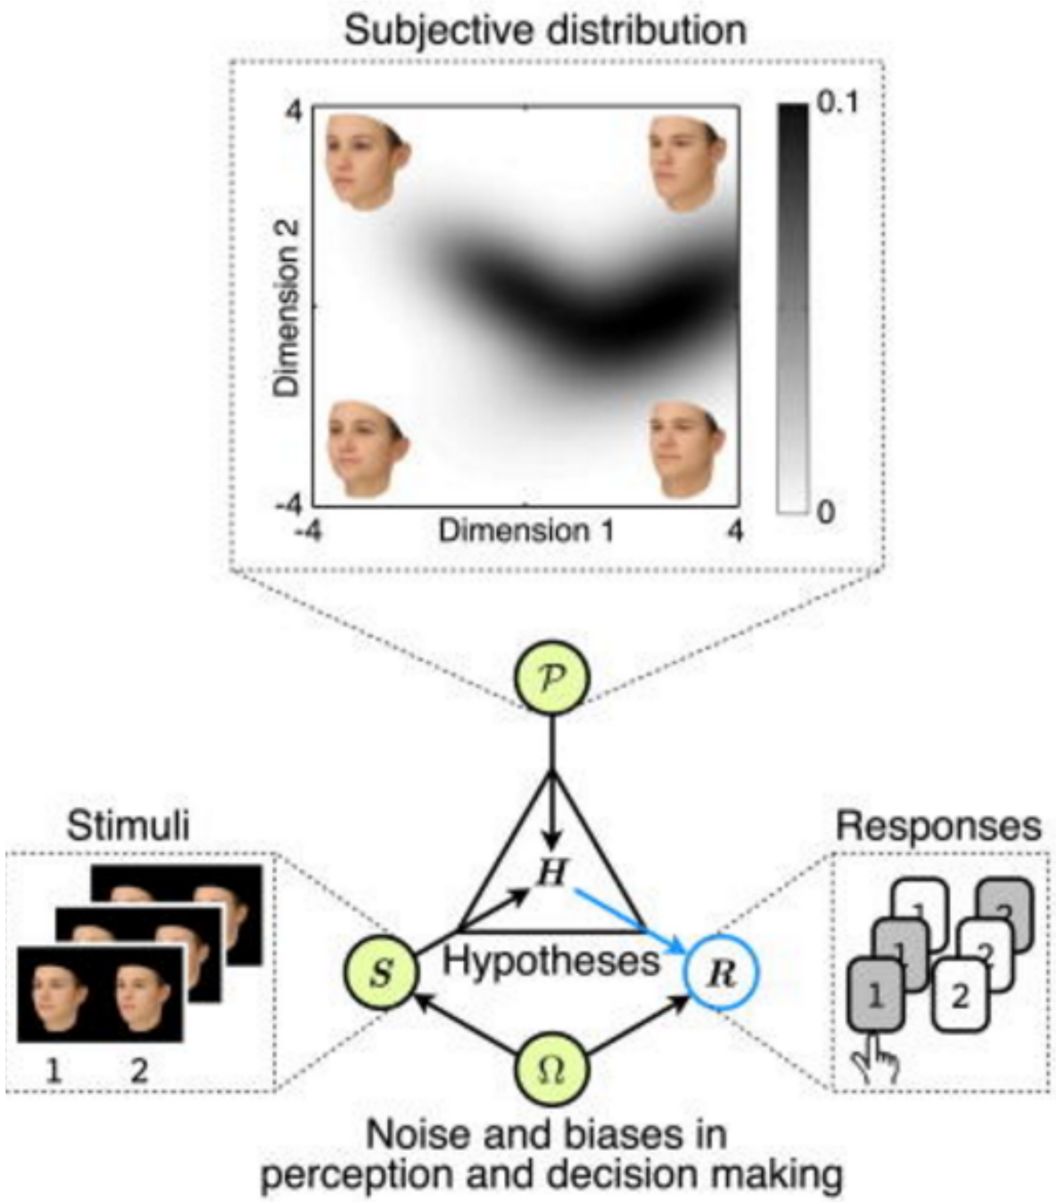

inference of the mental model

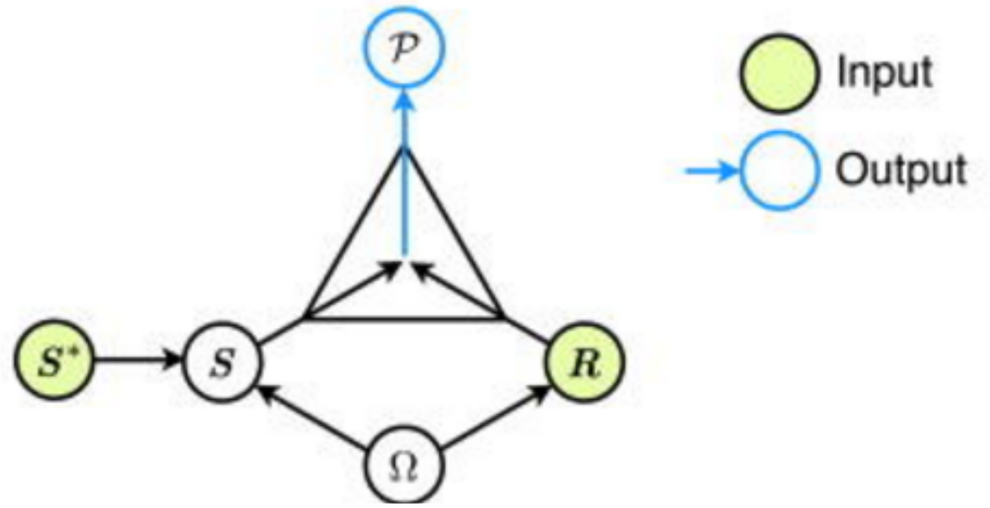

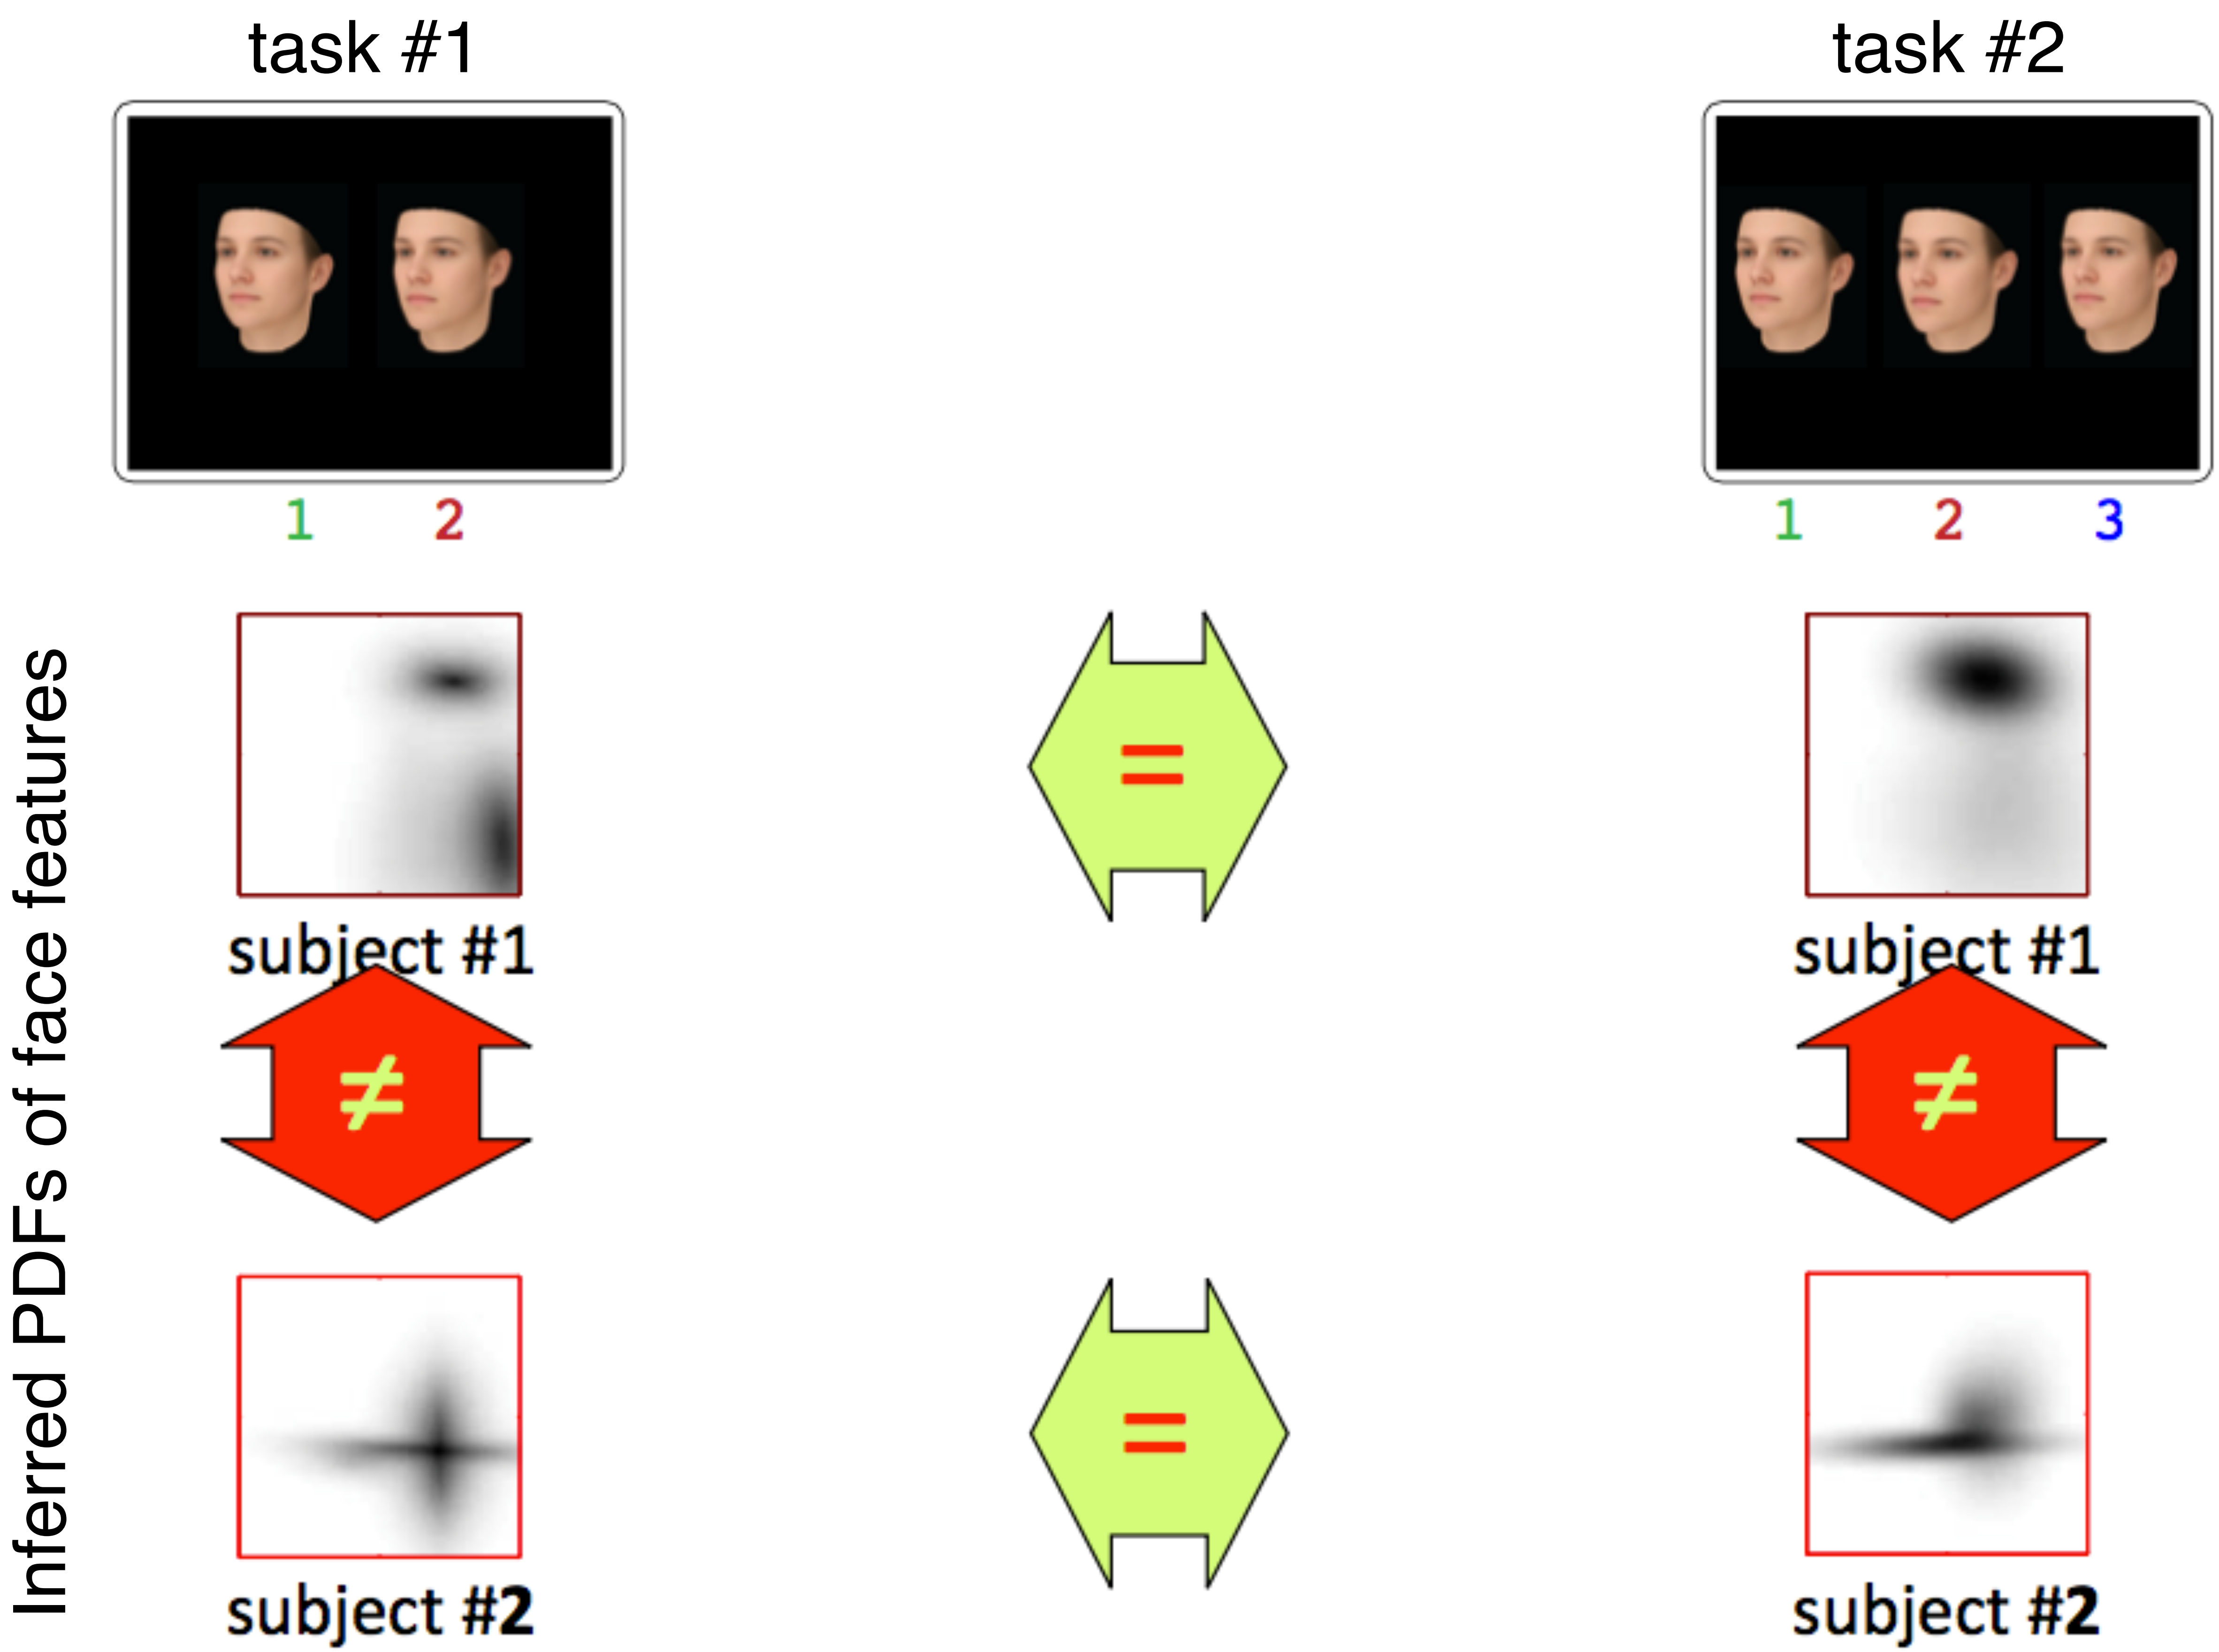

Houlsby, N. M., Huszár, F., Ghassemi, M. M., Orbán, G., Wolpert, D. M., & Lengyel, M. (2013). Cognitive tomography reveals complex, task-independent mental representations. Current Biology, 23(21), 2169-2175.

The way forward

- we have seen that probabilistic models can be used to predict behavioural data from psychophysics experiments

- we have also seen that we can infer the stability of mental representation in certain tasks using probabilistic models

- now we have to tie the model variables and inference algorithms to neurons