BIO MOTION LAB

Summer School in Computational Sensory-Motor Neuroscience

Models: Some introductory thoughts

Niko Troje

Department of Psychology
School of Computing, Department of Biology
Queen's University

Queen's UNIVERSITY


Queens
Fleming Hall
Jemmett Wing
78 Fifth Field Company Lane
JEMMETT WING

BioMotion Lab

Looking at people: biological motion, face recognition
Sensorimotor control of bipedal gait: terrestrial locomotion in pigeons

Program for today

1. Reality, models, representations
2. What is a good model?
3. Looking at people: faces, biological motion

Matlab exercises in the afternoon

1. Reality, models, representations

Topics and questions

Parallels between perception and scientific modeling
What is a model?
What is a good model?

What is ‘Reality’?

- “ What is real? How do you define real? If you’re talking about what you can feel, what you can smell, what you can taste and see, then real is simply electrical signals interpreted by your brain. This is the world that you know.”—
- Morpheus’s answer to Neo in *The Matrix*, 1999

http://www.youtube.com/watch?v=WnEYHQ9dscY

What is ‘Reality’?

- “You assume it is what you see, hear, touch, taste and smell - but is this assumption justified?”

http://www.youtube.com/watch?v=qjGRySVyTDk

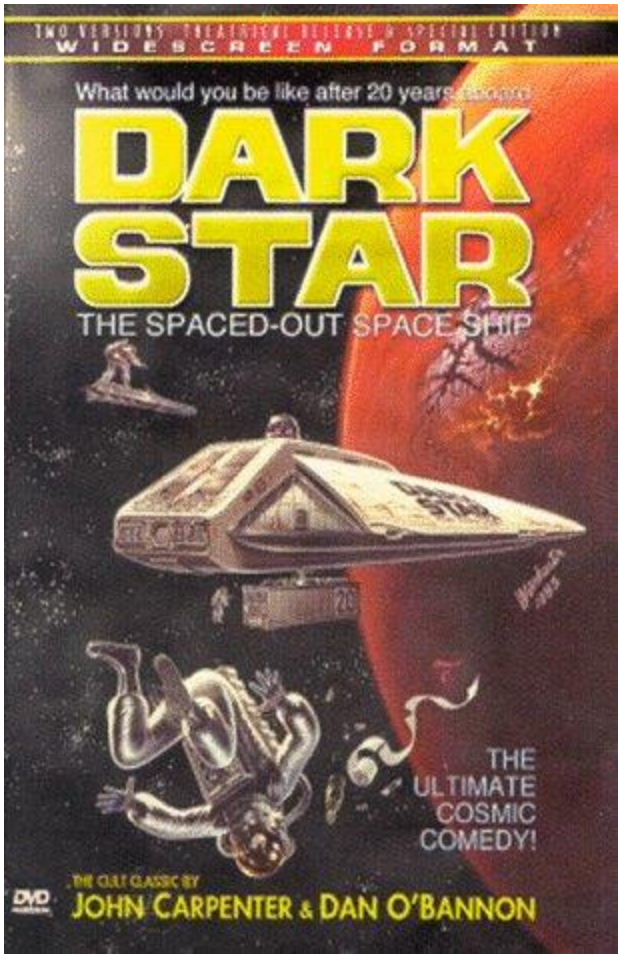

What is ‘Reality’?

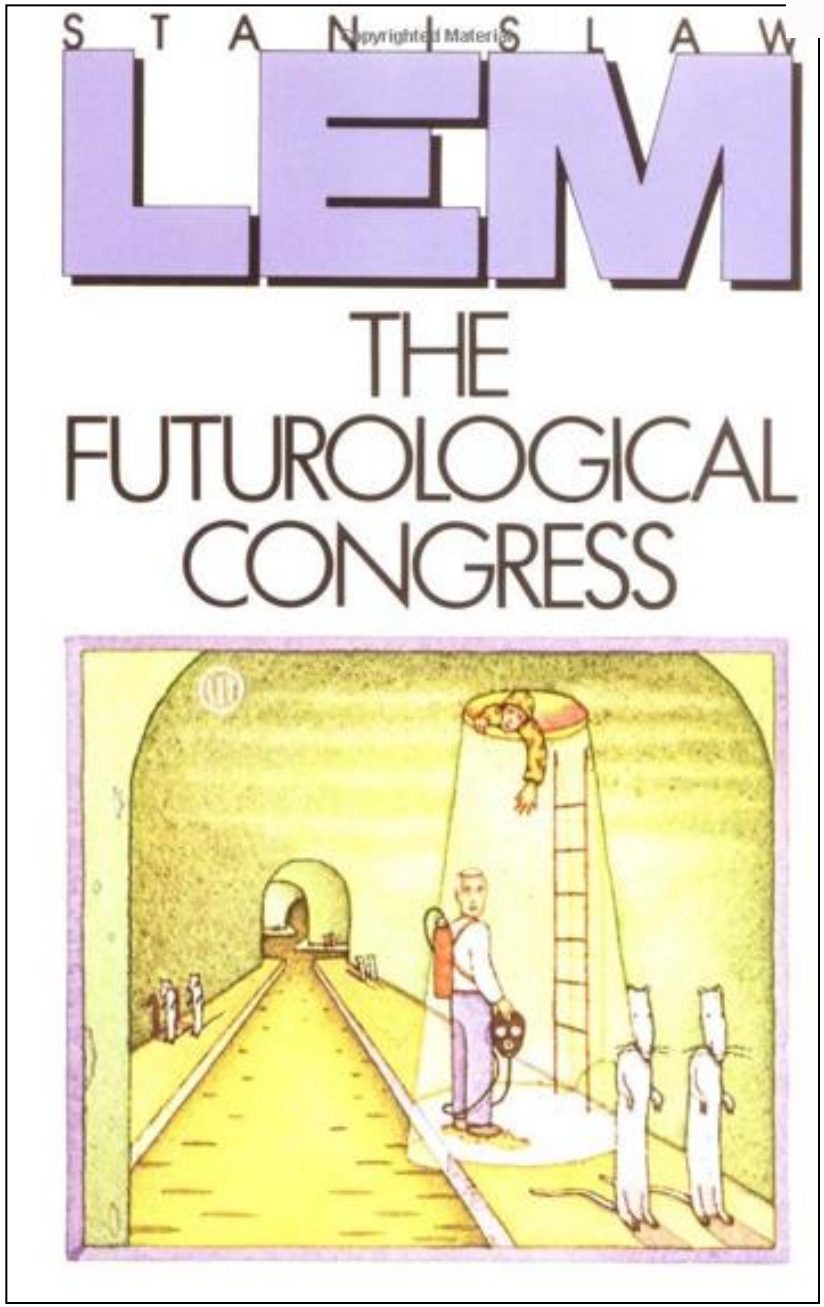

From Empiricism to Constructivism

David Hume (1711–1776) thought about the apparent discrepancy between the subjective experience of direct perception despite the fact that all we have is some limited measurements.

- Reliability refers to the consistency of measurements.
- Validity: refers to the relationship of the measurement to what is being measured

Hume's suggestion: Perceived reality is reliable and that is all that counts. Whether it is valid as well, we don't know.

And there is no way to find out!

Plato (427 – 347 BC)

Vision as reflection: The allegory of the cave

"The prison-house is the world of sight."

http://www.historyguide.org/intellect/allegory.html

Plato's Allegory of the Cave

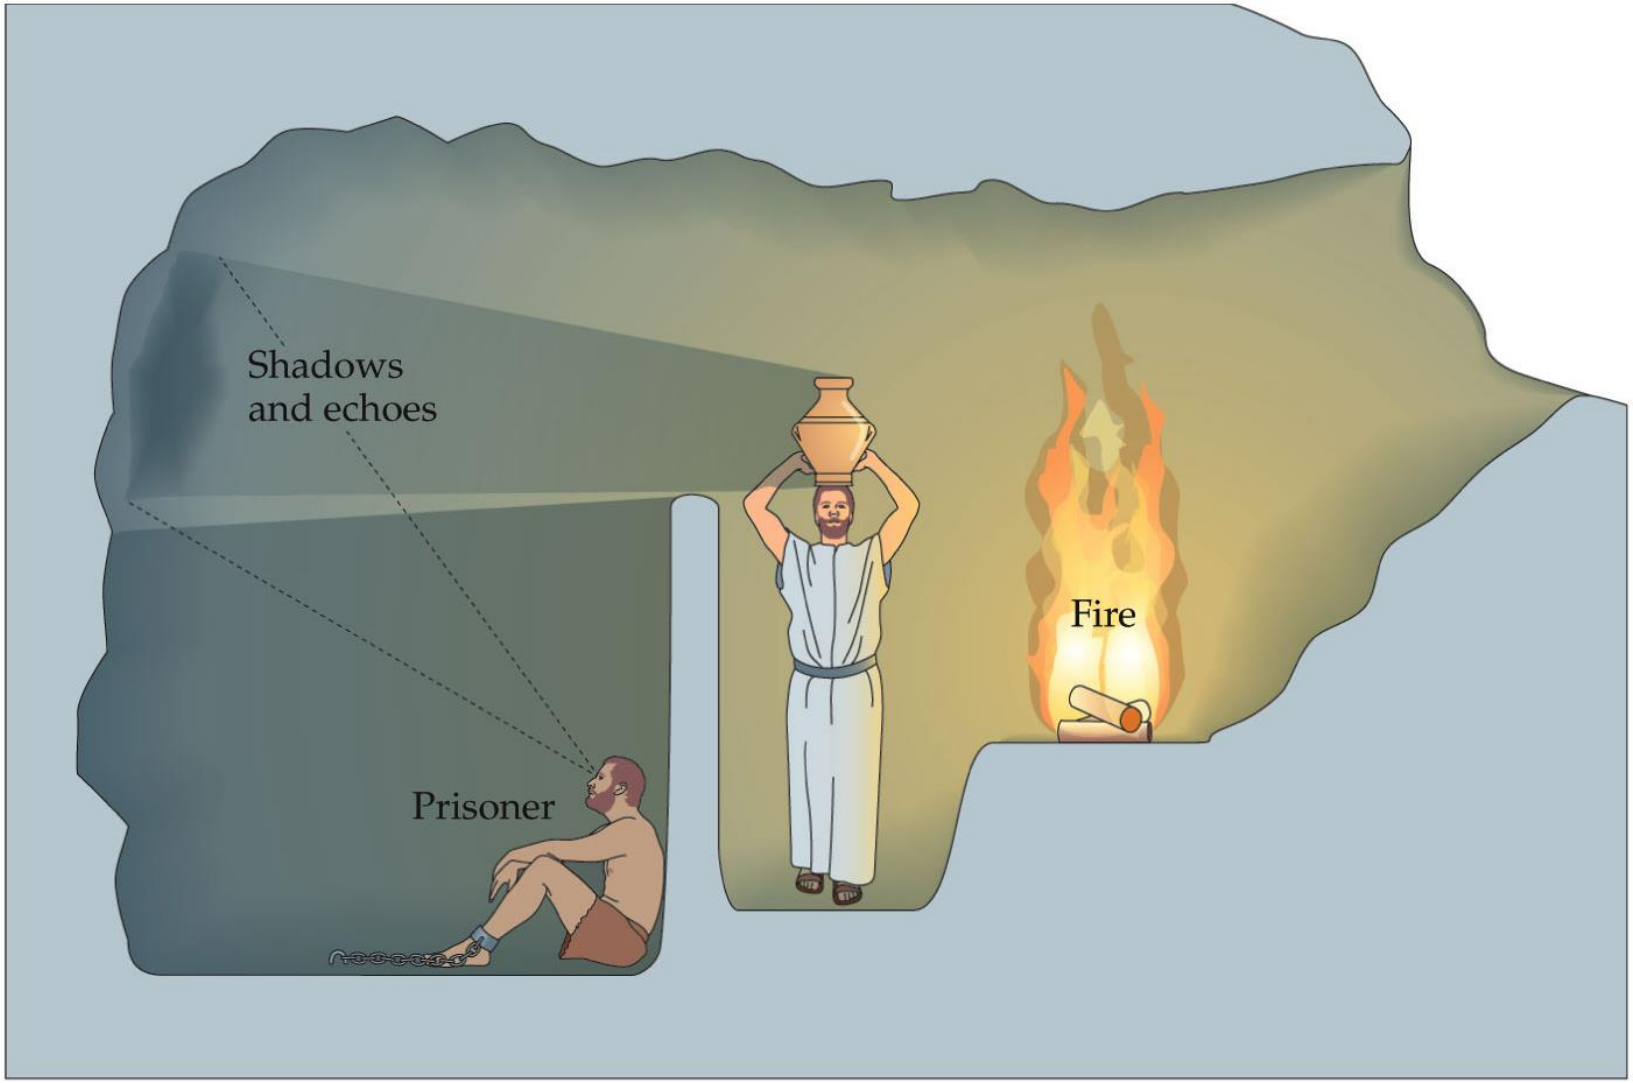

The perceiving brain as a modeler

The brain has to infer reality from sensory measurements!

What is ‘Reality’?

- Perception and your sense of reality are the products of evolution:
  - Sensory systems provide a survival advantage
  - Which senses are being developed is determined by the importance of type of energy in the environment for an animal
  - Human senses are limited to only certain kinds of energy in the environment
  - Therefore, humans’ sense of reality is also limited
  - For each sensory system, we always have to ask:
    - How does is work?
    - Why does it work as it works?

What is ‘Reality’?

- Some species sense energies that humans cannot:
  - Bees can sense the e-vector of polarized light
  - Dogs and cats can hear sounds of higher frequencies
  - Bats use ultrasound for echolocation
  - Birds, turtles, and amphibians use magnetic fields to navigate
  - Rattlesnakes can locate sources of infrared energy (e.g., heat)
  - Elephants can hear very low-frequency sounds

Reception and transduction ...

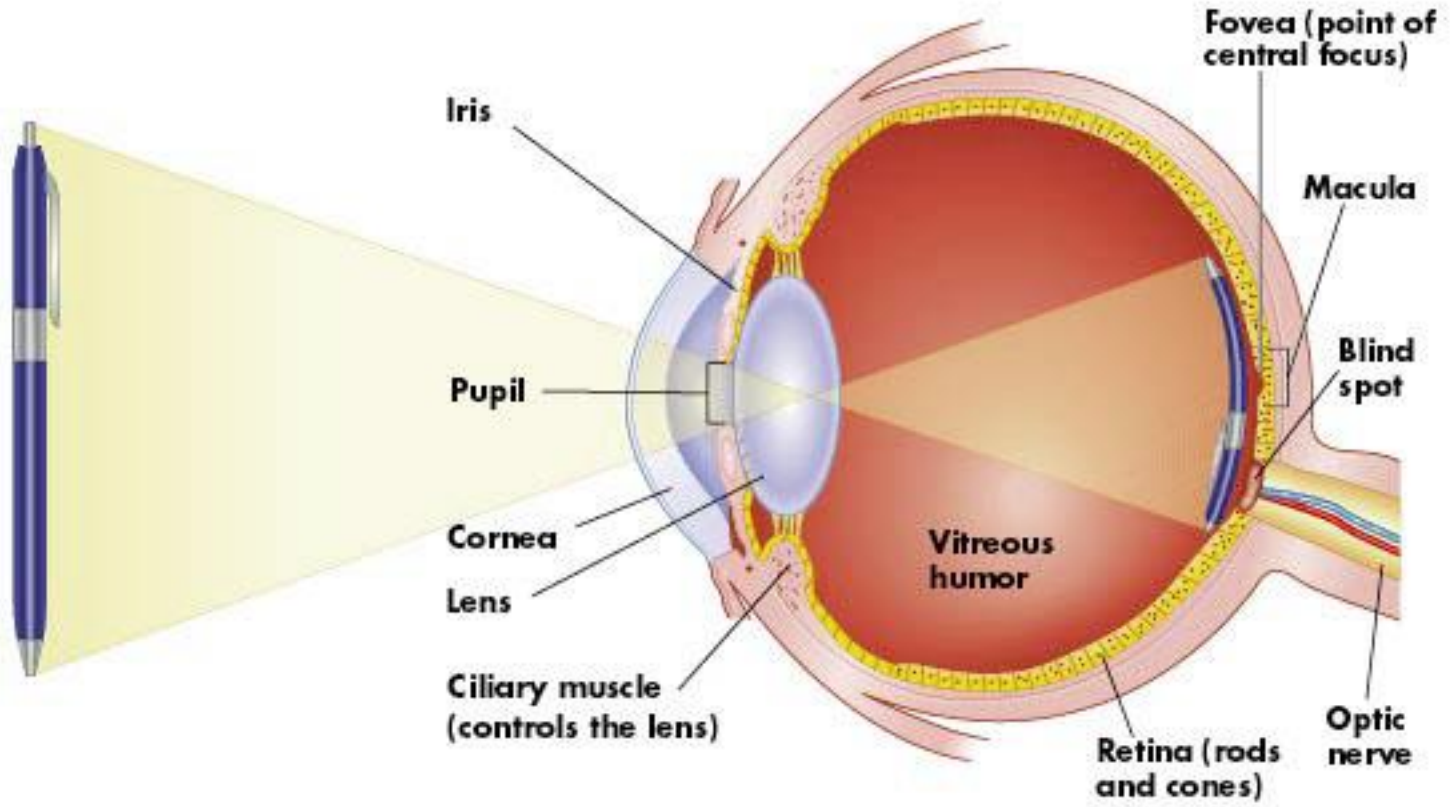

… transmission, propagation, coding

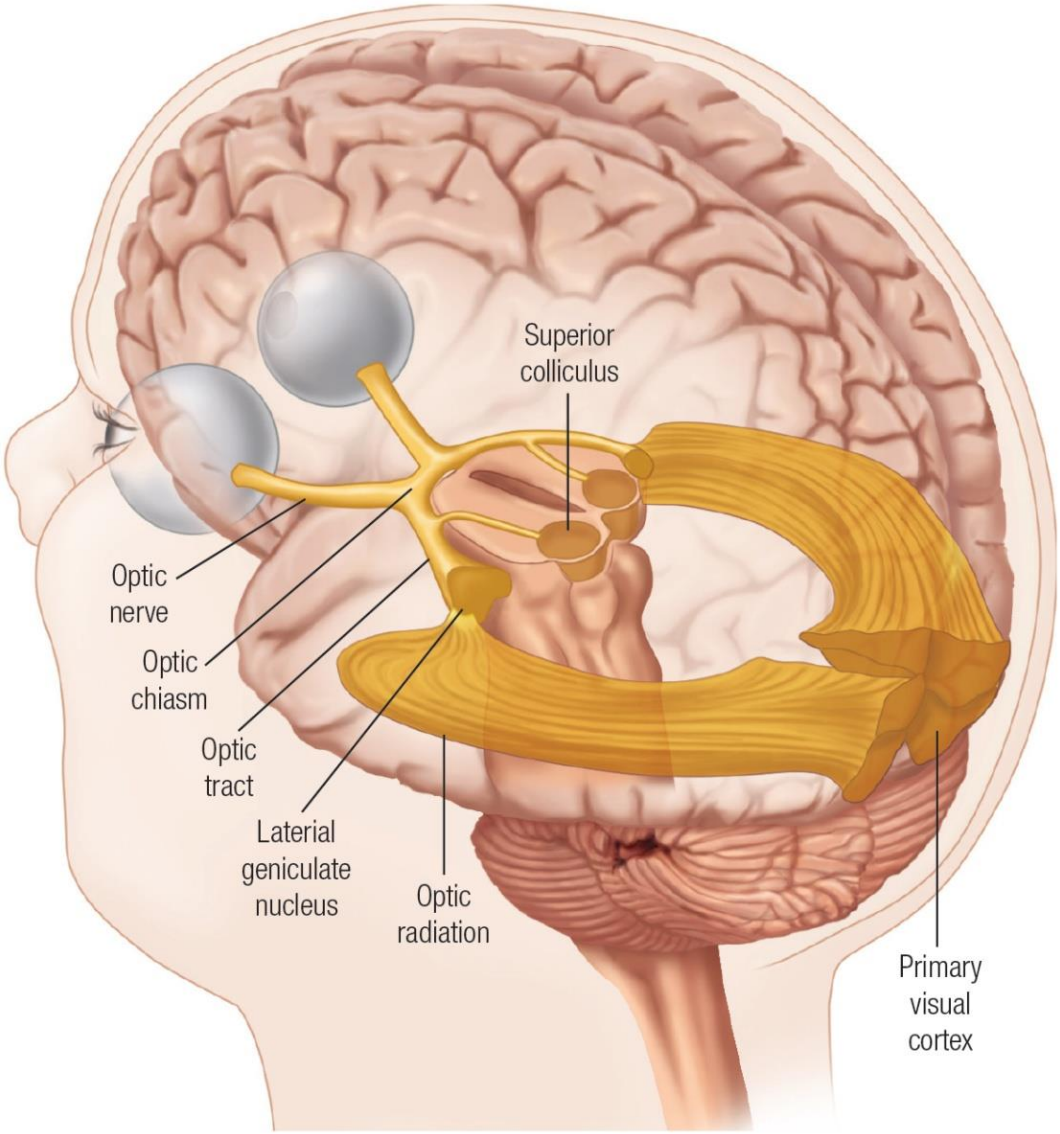

… transmission, propagation, coding

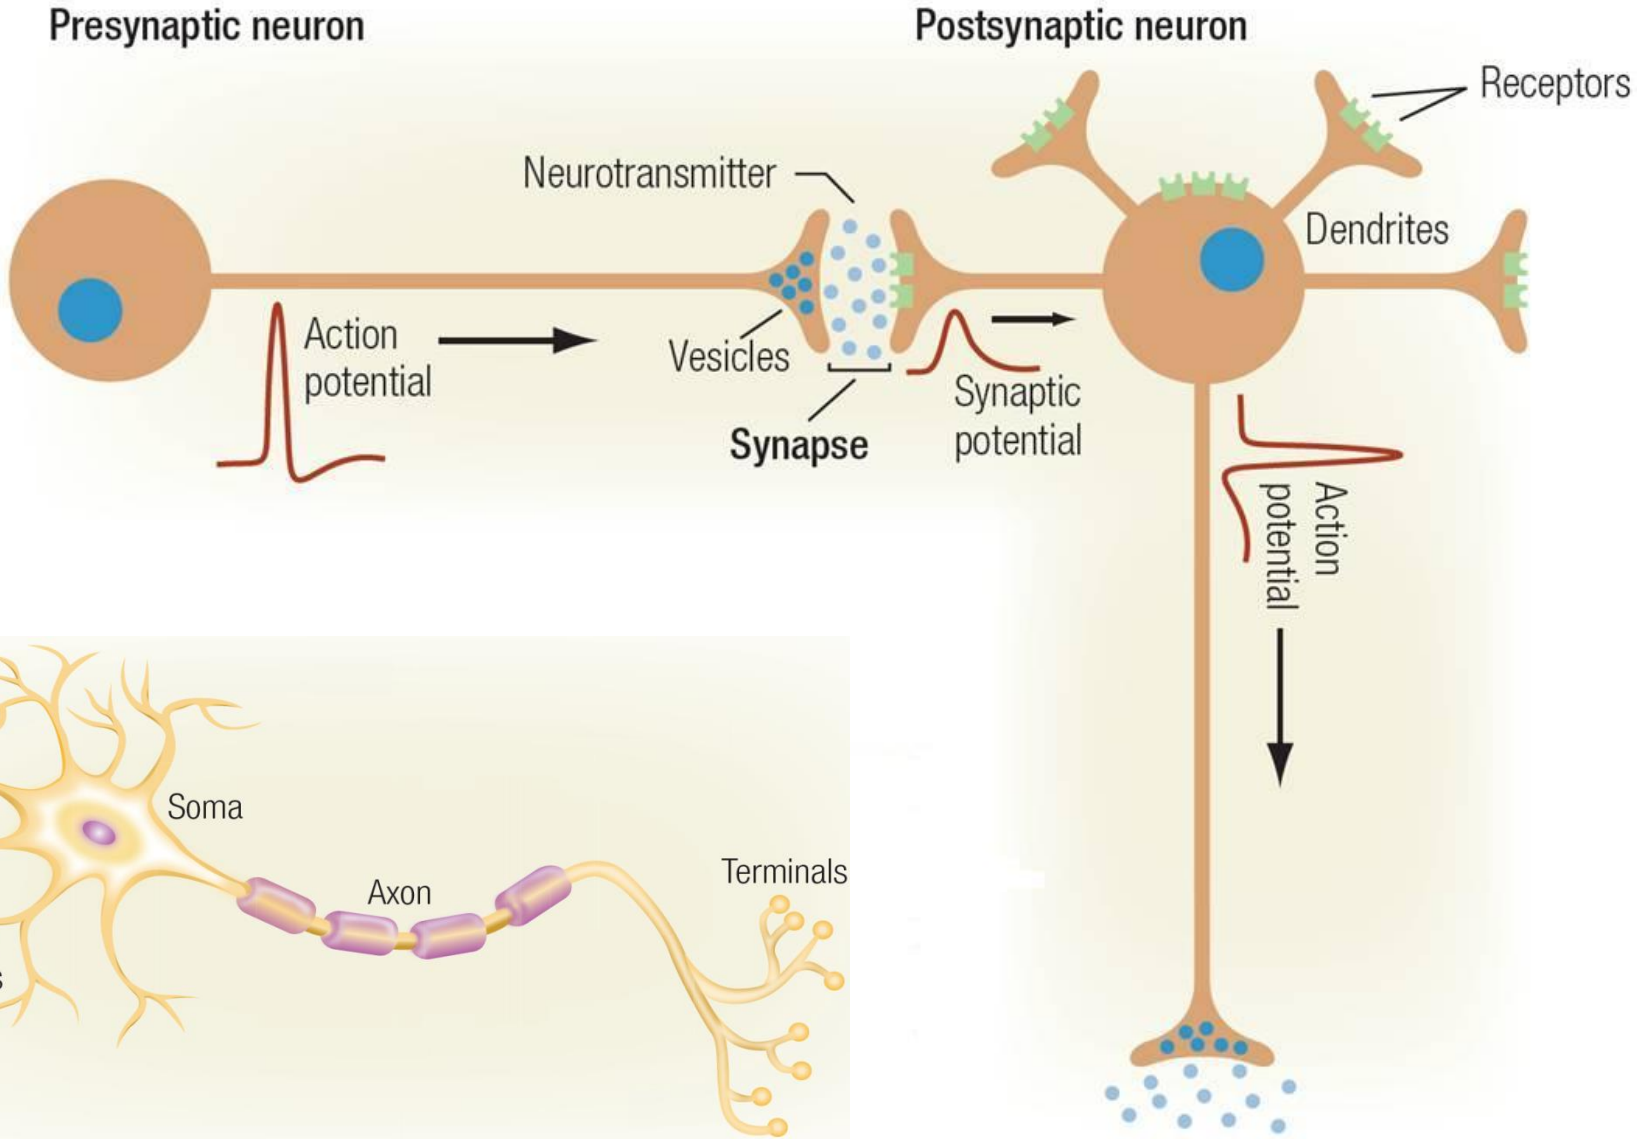

ascending connections
descending connections
lateral connections

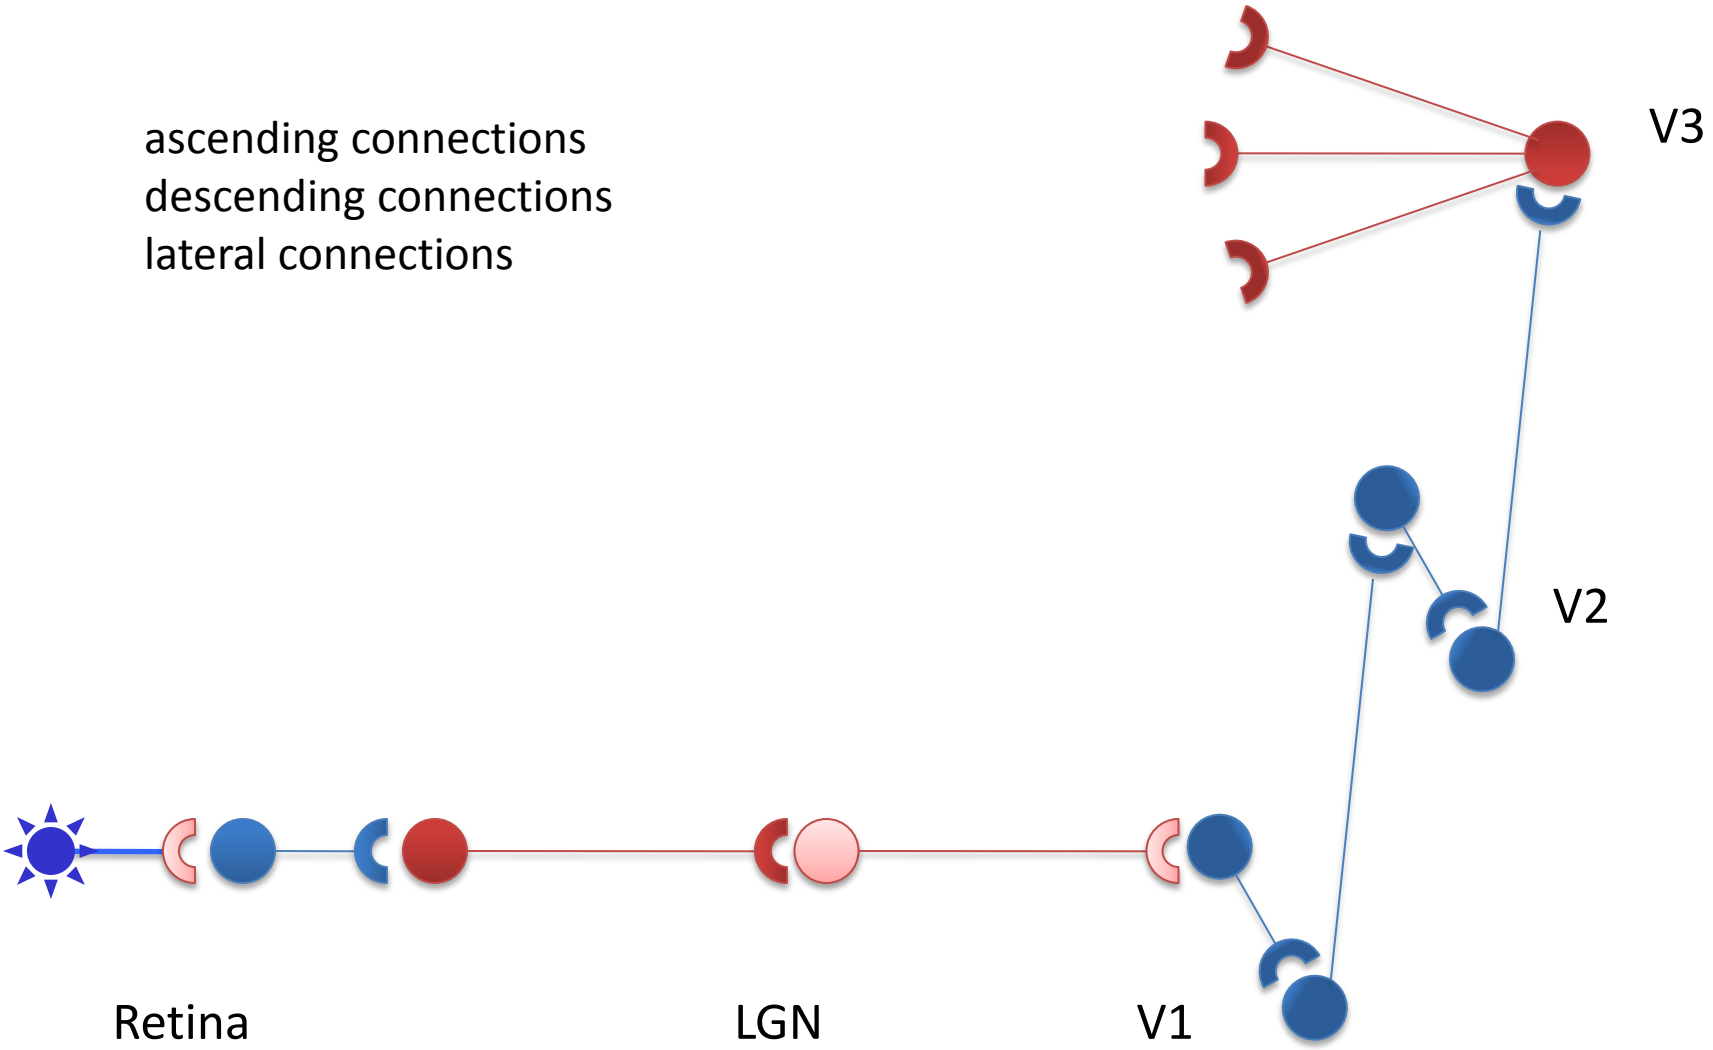

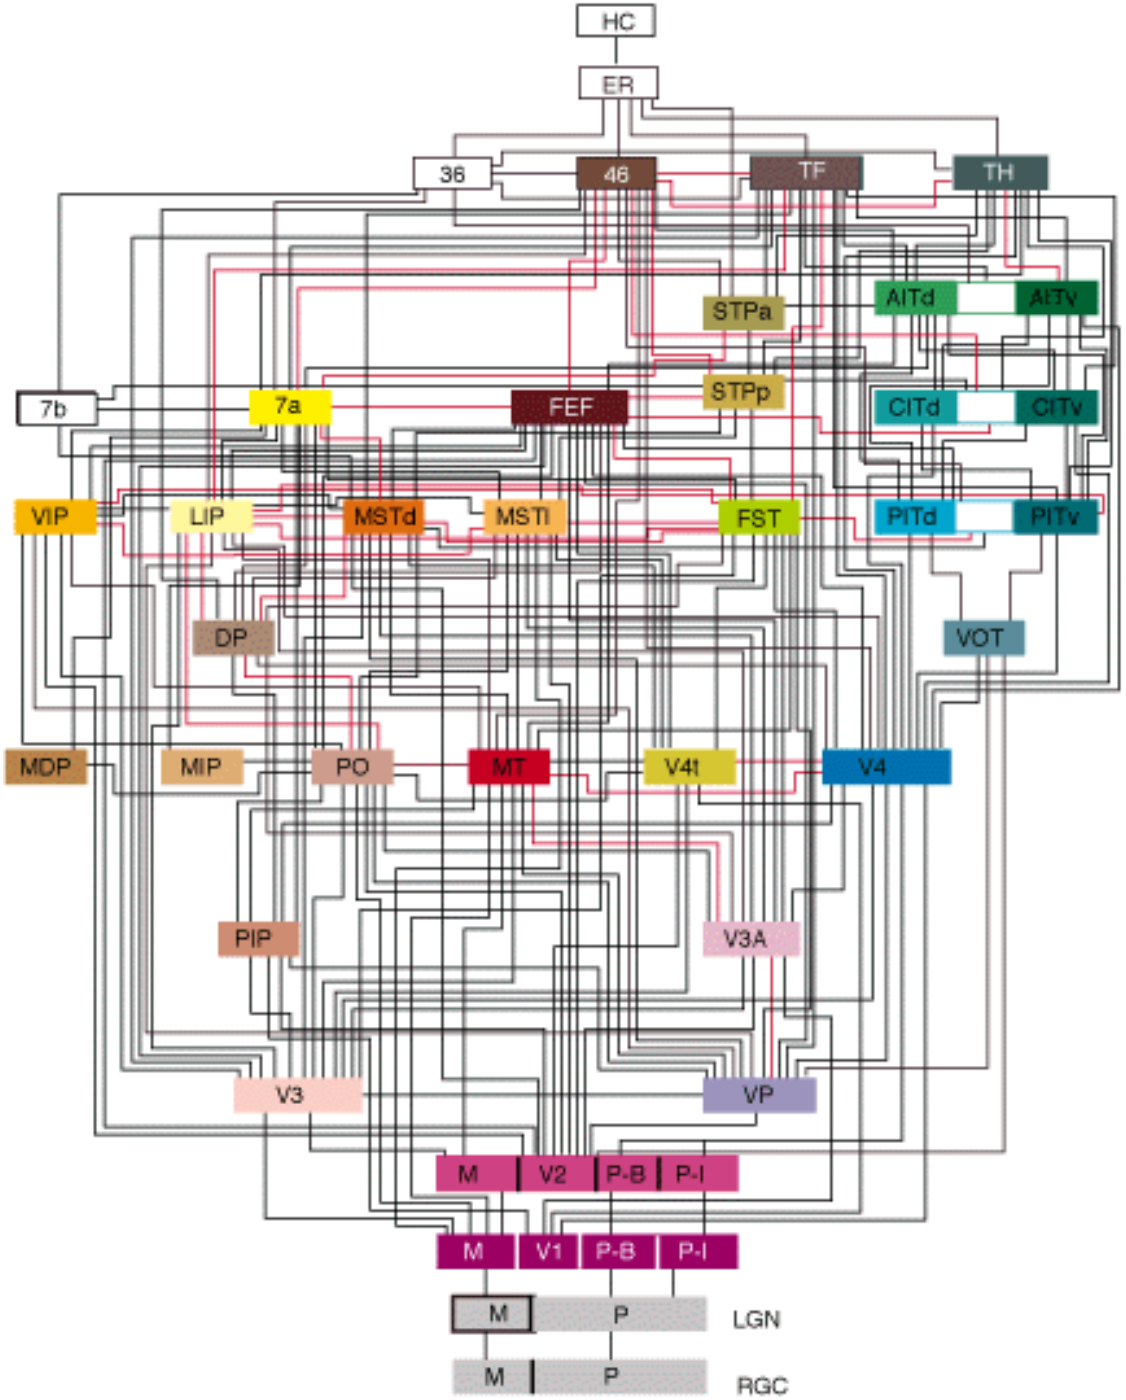

Van Essen et al. 1986

Johannes Müller (1801 – 1858)

"Elements of Physiology"
Doctrine of specific nerve energies:
The nature of a sensation depends on which sensory fibers are stimulated, not on how the fibers are stimulated.

There is only one code and our brain has only access to this code, but never to the physical stimulus itself.

"The immediate objects in the perception of our senses are merely particular states induced in the nerves"

Edwin Boring (1886 - 1968)

The “Illusion of Direct Perception”

There is no direct perception. Every physical stimulus undergoes transduction, is propagated along axons, transmitted across synapses and somehow processed in the brain.

The amazing thing about perception is that despite this indirectness the world seems so real. And that it usually works quite well!

Ontology and Epistemology

Ontology
A branch of philosophy concerned with articulating the nature and structure of the world

Epistemology
The nature of human knowledge and the understanding that can possibly be acquired through different types of inquiry and investigation

The ontological question
What is the form and nature of reality and, therefore, what is there that can be known about it?

The epistemological question
What is the nature of the relationship between the knower and what can be known?

The methodological question
How can the inquirer go about finding out whatever he or she believes can be known?

Objectivistic vs. constructivistic worldview

Objectivistic worldview

- Reality is comprised of distinct objects with properties independent of the inquiring observer
- If two observers do not understand a phenomenon in the same way, it is due to human imperfection
- Disputes over the true nature of Reality can be resolved through additional inquiry

Constructivistic worldview

- Realities are local and specific in the sense that they vary between groups of individuals
- Constructions, being ontological elements of realities, are not absolutely true or correct in any sense, only more or less informed and sophisticated.
- Even meaningful constructions may be termed malconstructions, as they obviously are too simplistic or inconsistent.

The constructivistic worldview

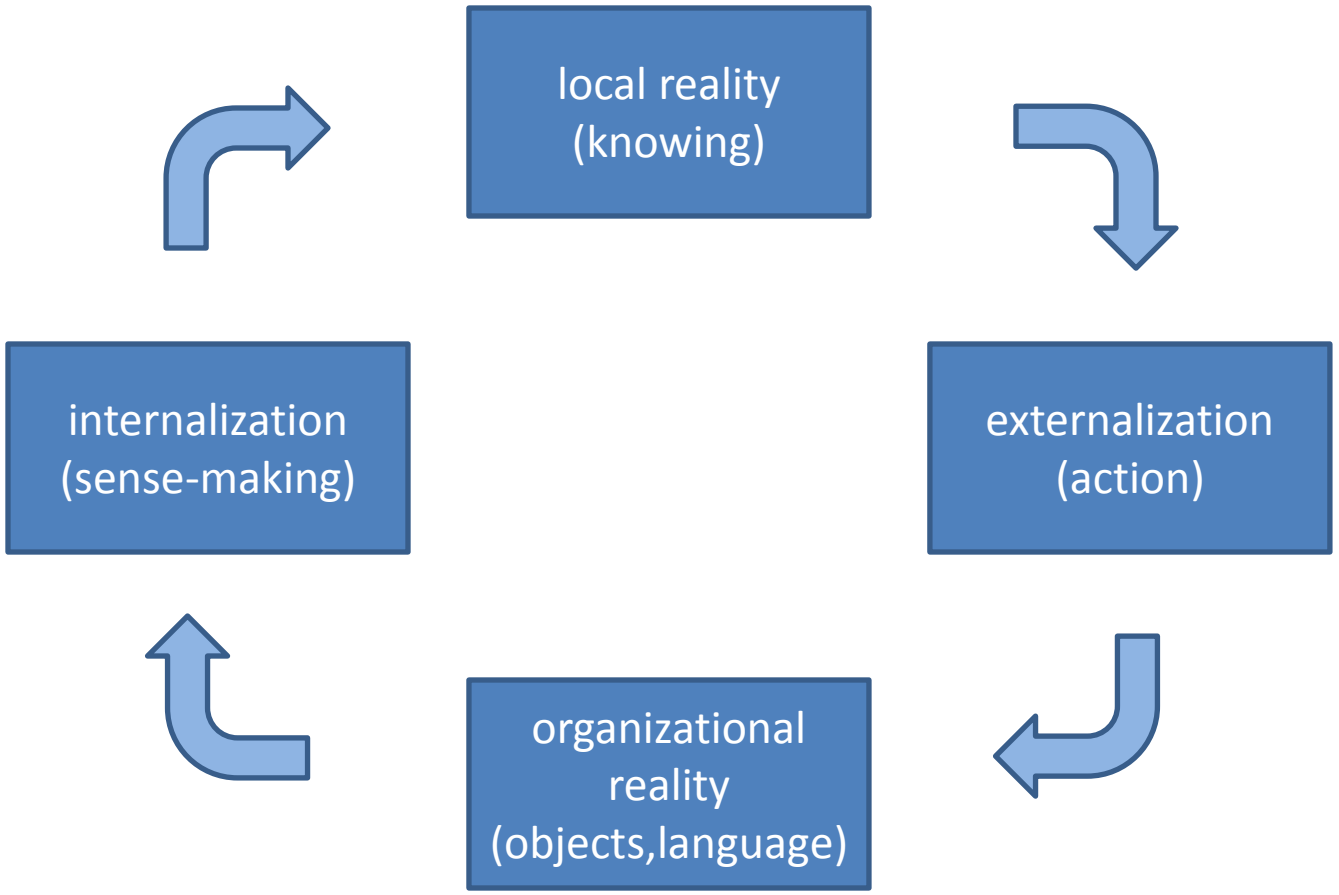

after Berger & Luckmann (1966) The Social Construction of Reality

The constructivistic worldview

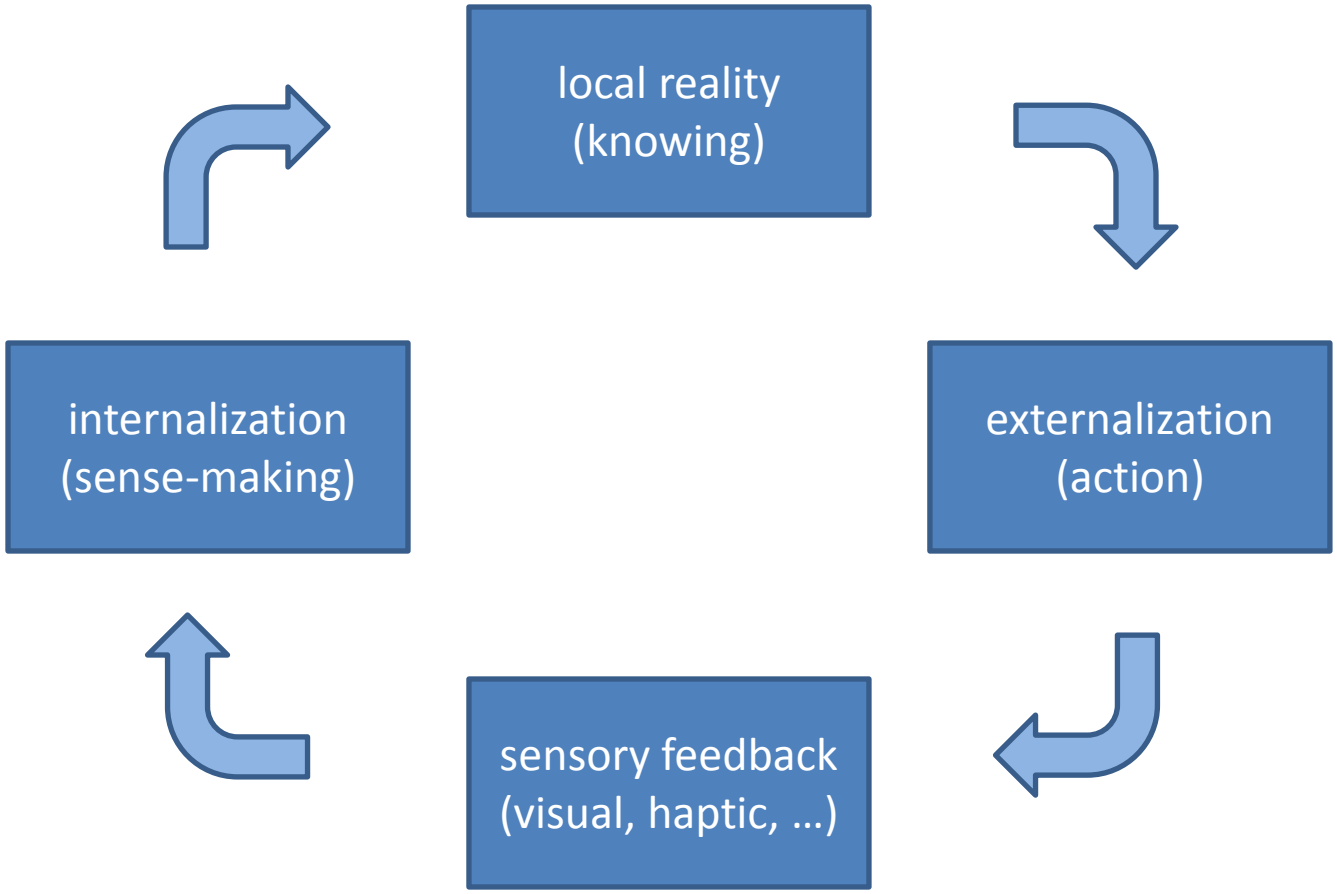

[task, state, context]
→ motor command

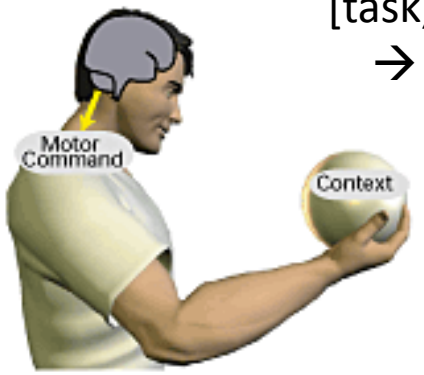

"inverse model"

[state, motor command, context]
→ sensory feedback

"forward dynamic model"

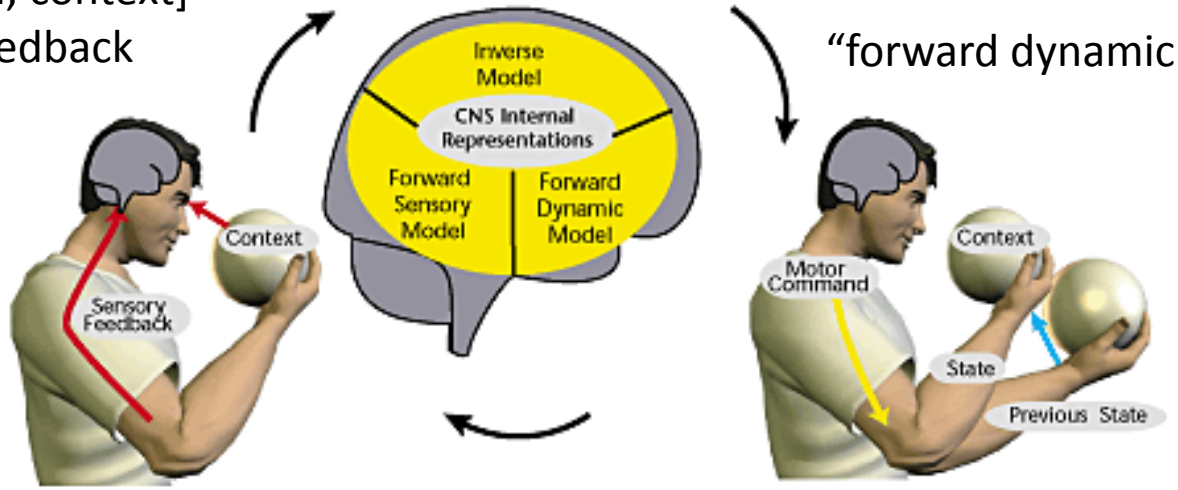

"forward sensory model"

[previous state, motor command, context]
→ state

Wolpert & Ghahramani, 2000

What is a model?

A model formulates a certain input (cause) – output (effect) relationship.

A model is a theory.

A model can make predictions. It generates hypotheses.

A model structures (organizes) something else.

A model is a purposeful representation of something else.

A model can be used to simulate or anticipate the behavior of something else.

Models and representations

Discriminative models, also called conditional models, are a class of models for modeling the dependence of an unobserved variable *y* on an observed variable *x*.

Discriminative models can be

- deterministic: f: *x* → *y*
- probabilistic (they learn conditional probabilities): P(y|x)

Examples:

- linear regression, generalized linear regression
- support vector machine
- linear discriminant analysis

Models and representations

Generative models are models for randomly generating observable data, typically given some hidden parameters.

A generative model specifies a joint probability distribution over observations and associated attributes (labels).

Generative models learn joint probabilities: P(x,y)

Examples:

- Correlations, principle components analysis, factor analysis
- Hidden Markov models
- Gaussian mixture models
- Restricted Boltzmann machine

Generative models can be used as discriminative models (but not vice versa).

Generative models are “purposeful representations of something else”.

Discriminative vs. generative models

Generative models learn joint probabilities

Discriminative models learn conditional probabilities

$$p(y|x) = \frac{p(x,y)}{p(x)}$$

Example

Assume you are given data in the form (*x,y*):
((1,0), (1,0), (2,0), (2,1)

The joint probabilities are:

	y=0	y=1
x=1	1/2	0
x=2	1/4	1/4

The conditional probabilities are:

	y=0	y=1
x=1	1	0
x=2	1/2	1/2

Bayes' formula

$$p(x, y) = p(y|x)\, p(x)$$
$$= p(x|y)\, p(y)$$

→ $$p(y|x) = \frac{p(x|y)\, p(y)}{p(x)}$$

$$p(y|x) \sim p(x|y)\, p(y)$$

posterior likelihood prior

$$\hat{y} = \arg\max_y\ p(y|x) = \arg\max_y\ p(x|y)\, p(y)$$

Thomas Bayes 1702 - 1761

$$p(y|x) = \frac{p(x|y)p(y)}{p(x)}$$

Short Bio

University of Edinburgh (logic, theology)

English mathematician

Presbyterian minister

Fellow of the Royal Society

Publications

Bayes, T. (1731) Divine Benevolence, or an Attempt to Prove That the Principal End of the Divine Providence and Government is the Happiness of His Creatures.

$$p(B|I) < 1$$

Anonymous (1736) An Introduction to the Doctrine of Fluxions [1], and a Defence of the Mathematicians Against the Objections of the Author of the Analyst [2]

[1] Fluxions: The logical foundations of Isaac Newton's calculus

[2] George Berkeley

2. What is a good model?

… a good representation?
… a good decision?

Probability: Follow Bayes to maximize posterior probability.

Accuracy: Ability to closely reproduce data for which ground truth is available.

Predictability: Ability to generalize to new data.

Dimensionality: Ability to make predictions based on as few parameters as possible.

A farmer, interested in the construction of farm gates, investigates the relation between length (l) of the diagonal strut and the lengths (x and y) of the rails and stiles of the farm gate. Although he is familiar with the theory and practice of length measurement and with Euclid, the law of Pythagoras is utterly beyond his deductive powers.

He invents a model, a linear one, in which the lengths (l, x, y) are related by:
$l = x+y$

The model has several properties analogous to the behaviour of farm gates. However, when the farmer carries out detailed and precise measurements, he finds discrepencies. In order to explain them, he introduces a new effect: the Farm-Gate Contraction.

The Farm-Gate Contraction becomes part of his adjusted model:
$l = \lambda\ (x+y)$

The cardinal question about perception

How does the brain transform a stream of continuous, noisy, incomplete, unstructured, ambiguous sensory data into a meaningful, discrete, reliable and predictable model of the world?


Stimulus
Sensation
Perception

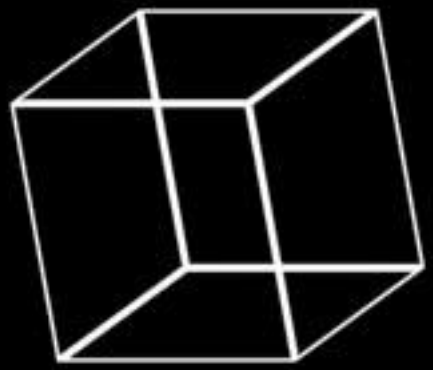

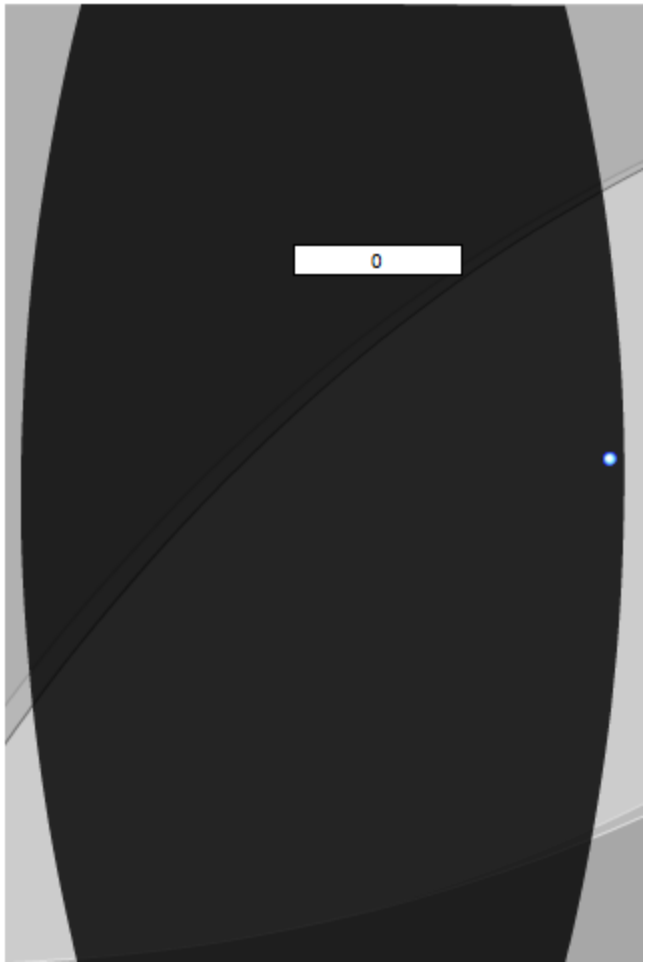
0

Perception (coming up with a model of the world) is more than bottom-up processing of sensory information.

Where is the additional information coming from?
What constrains the possible solutions?

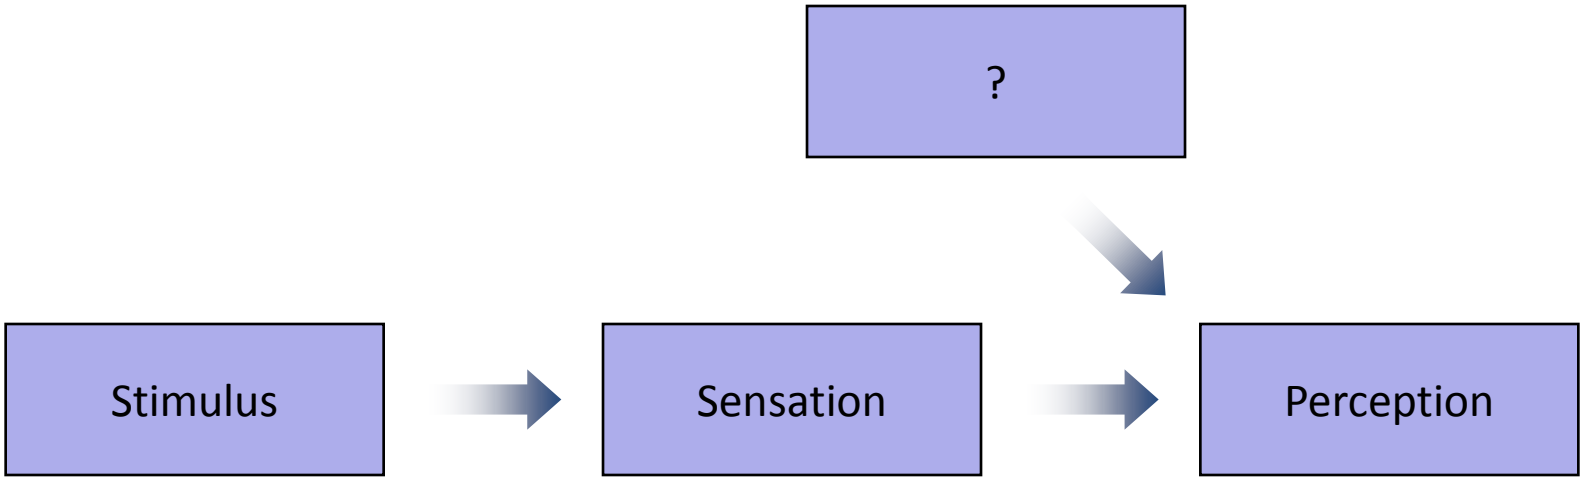

Gestalt (German): form, shape, figure

Kurt Koffka
(1886 – 1941)

Wolfgang Köhler
(1887 – 1967)

Max Wertheimer
(1880 – 1943)

The whole is more than the sum of its parts!

Law of proximity

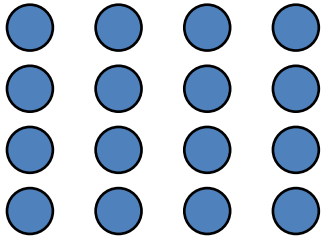

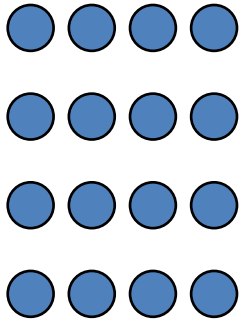

Law of similarity

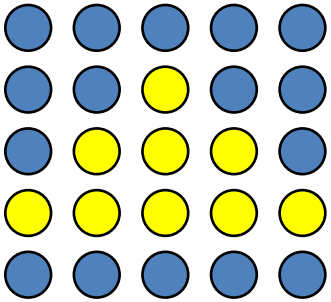

Law of good continuation

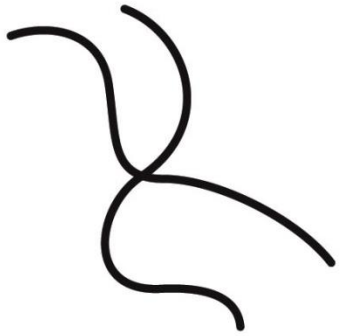

Law of proximity

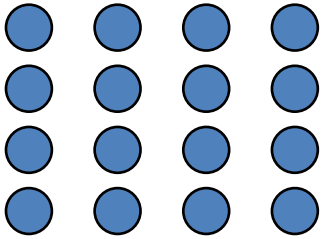

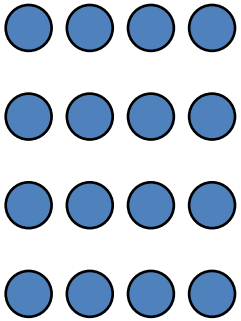

Law of similarity

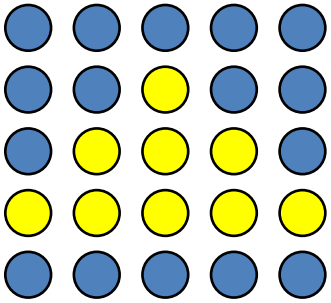

Law of good continuation

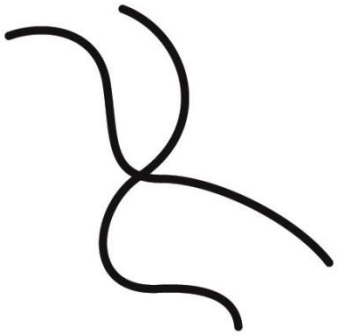

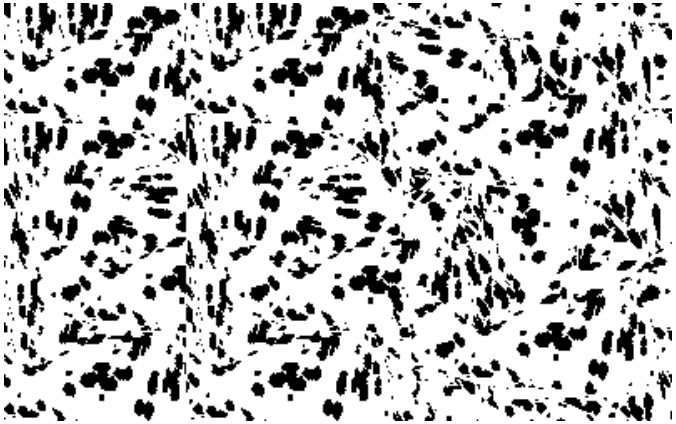

Law of common fate

Law of pragnanz

- Prägnanz: German for good figure, design, carrying meaning
- A stimulus pattern is organized in such a way that the resulting structure is as simple as possible.
- “The Simplicity Principle” (e.g. Pomerantz & Kubovy 1986)

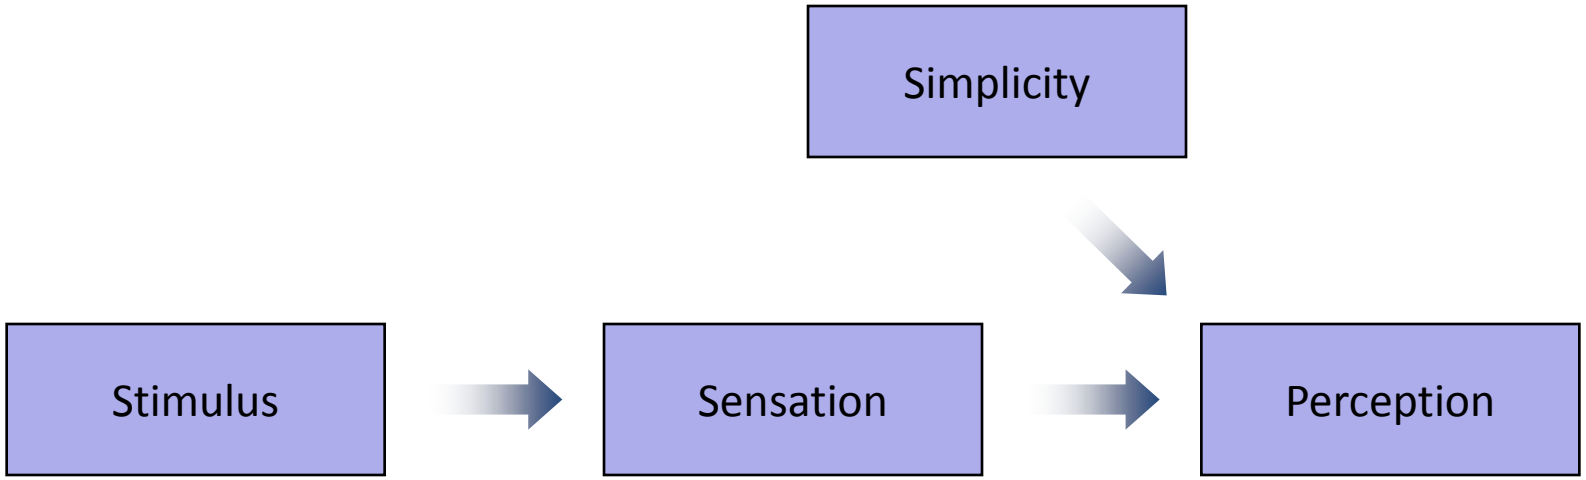

The Simplicity principle

- Occam's razor (William of Ockham, 14th century):
  "entities must not be multiplied beyond necessity"
  (entia non sunt multiplicanda praeter necessitatem)
- Information theory, channel capacity, redundancy (Shannon and Weaver 1949)
- Coding theory (Hochberg & McAlister 1953, Leeuwenberg 1969, 1971; Restle 1979)
- Attneave 1954, 1981: "What the perceptual system likes is short descriptions"
- Minimum description length (Rissanen 1989, Hinton & Zemel 1994)

Hermann von Helmholtz (1821 – 1894)

- Measured the speed of a nerve impulse
- Invented the ophthalmoscope
- Important publications
  - "Handbook of Physiological Optics"
  - "The Sensation of Tone as a Physiological Basis for the Theory of Music"
  - "On the Conservation of Force"

Very good reading is a speech he gave in 1878 "The Facts of Perception"
(Die Tatsachen in der Wahrnehmung)

http://www.marxists.org/reference/subject/philosophy/works/ge/helmholt.htm

The Likelihood principle

Helmholtz advocated the “Likelihood Principle”

- Perception as “unconscious inference”
- Sensory input will be organized into the most probable distal object or event consistent with that input
- (be aware that terminology is inconsistent with Bayes)

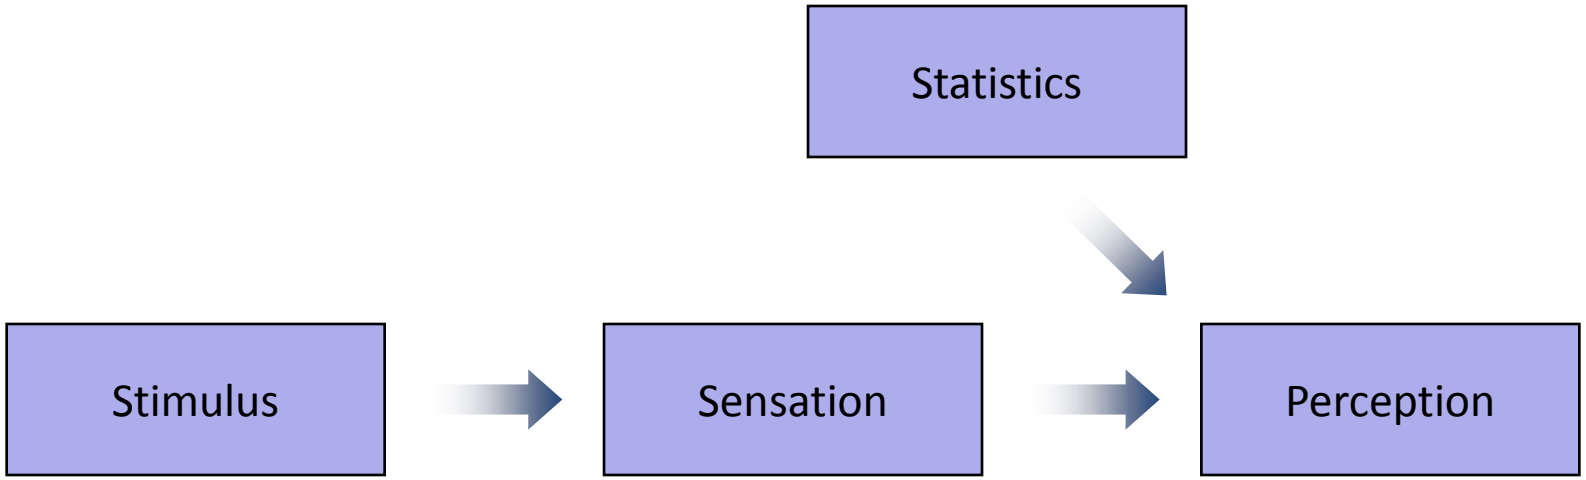

Likelihood

- Helmholtz's "unconscious inference"
- Oliver Selfridge's (1959) "pandemonium" model of letter recognition. Committee rules: Honor physics, avoid accidents
- Bayesian statistical decision theory:

$$p(w_i|s) \approx p(s|w_i)\, p(w_i)$$

posterior probability

likelihood

prior probability

"prior"

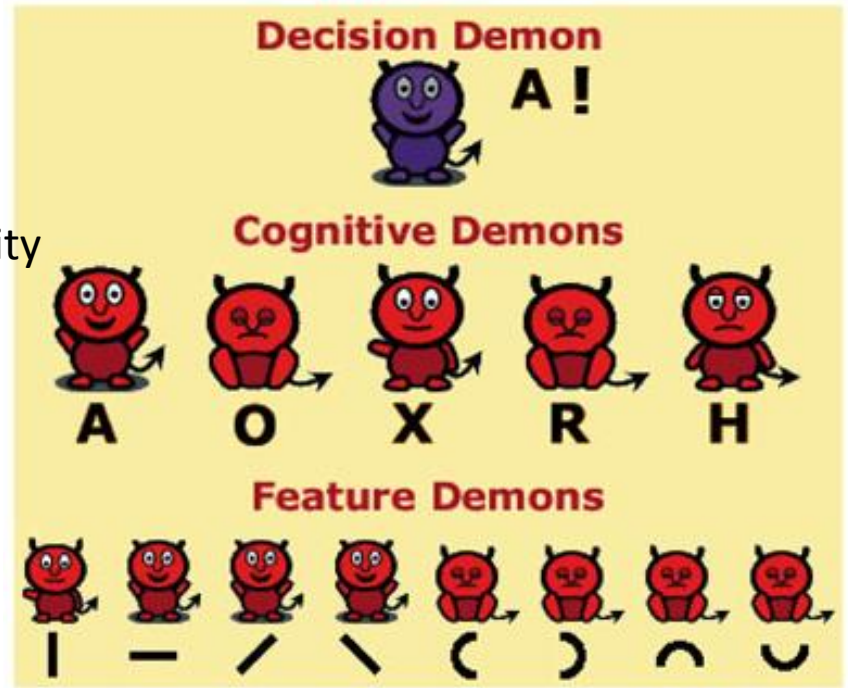

Bayes'ian Inference

$$p(w_i|s) \approx p(s|w_i)\, p(w_i)$$

posterior probability

likelihood

prior probability

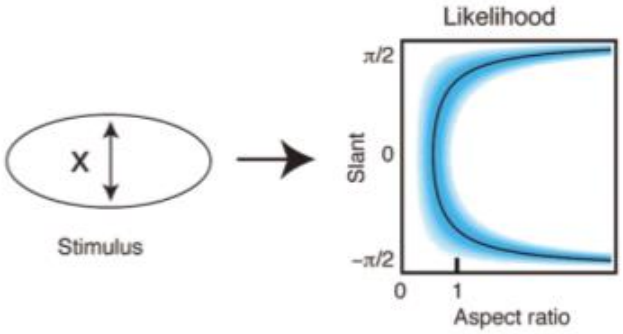

$$\widehat{w_i} = \arg\max_{w_i} p(w_i|s)$$

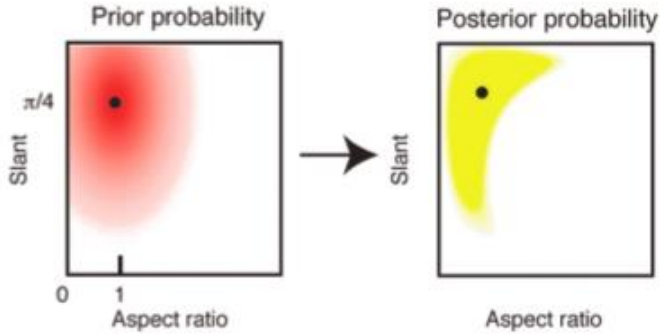

Bayes'ian Inference

$$p(w_i|s) \approx p(s|w_i)\, p(w_i)$$

posterior probability

likelihood

prior probability

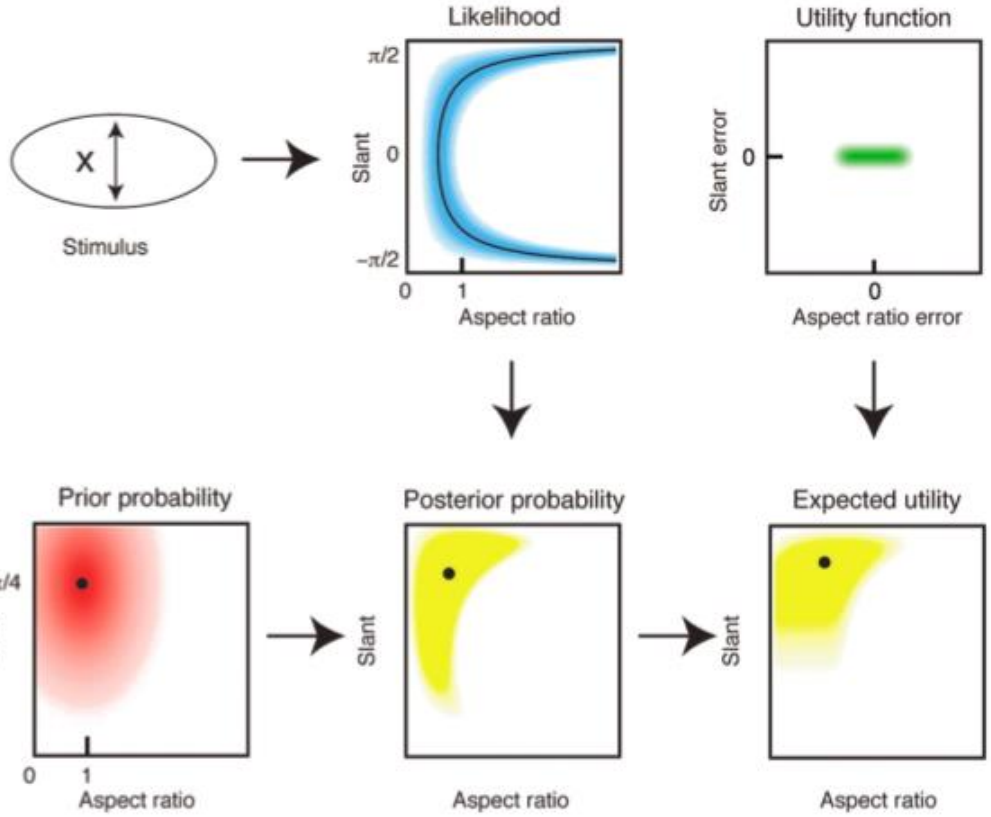

$$R = \arg\max_{r} [u(r,W)\, p(W|s)]$$

More examples for ambiguities in visual perception

what we want to know		what we can measure	context
aspect ratio	slant (of an ellipse)	aspect ratio of the retinal image	our example

Luminance depends on illuminance and reflectance

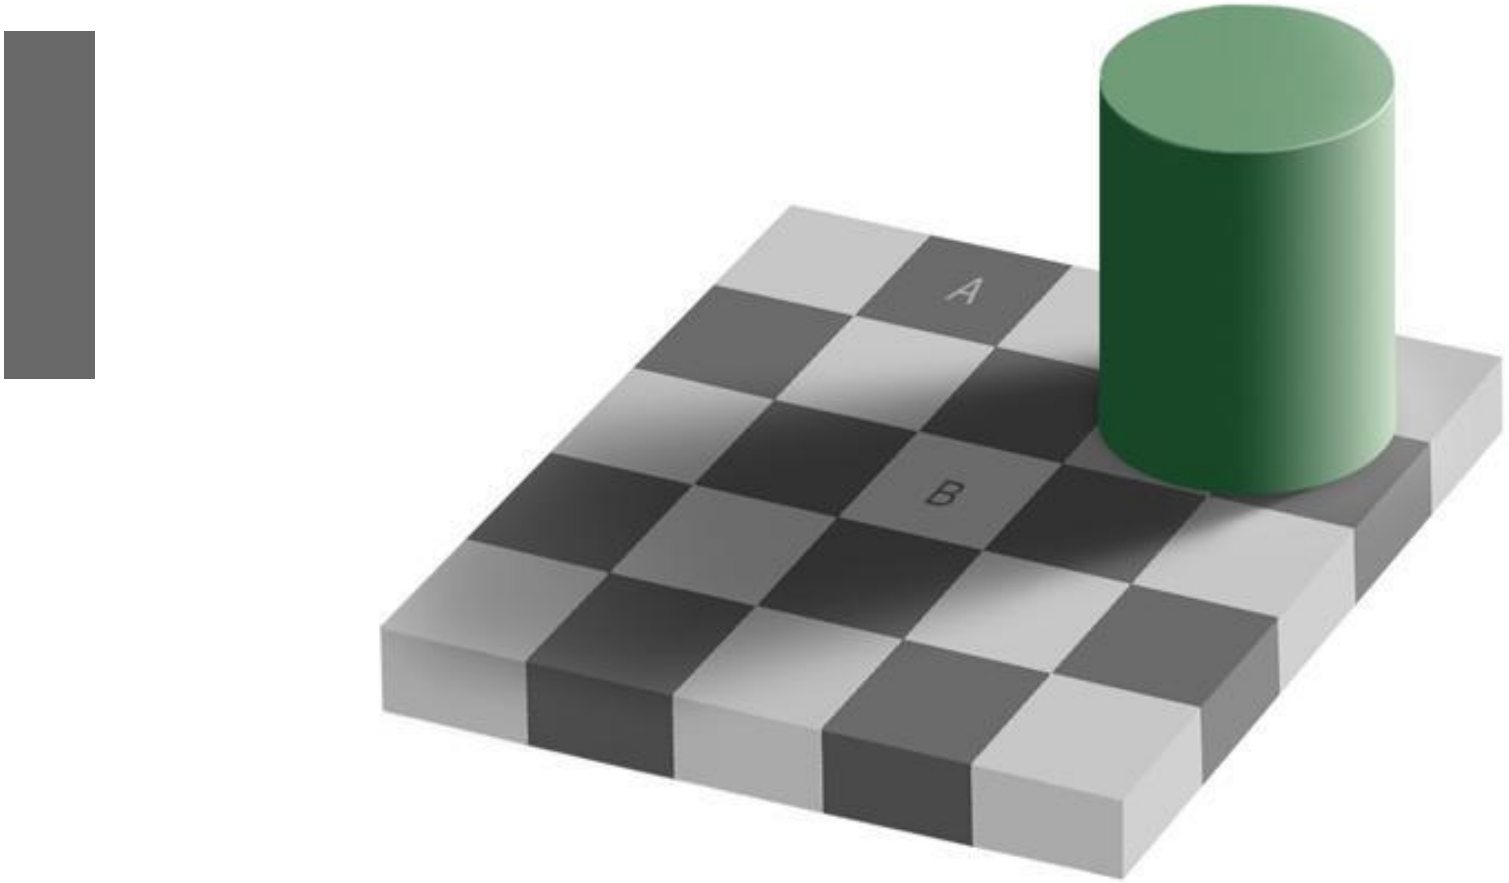

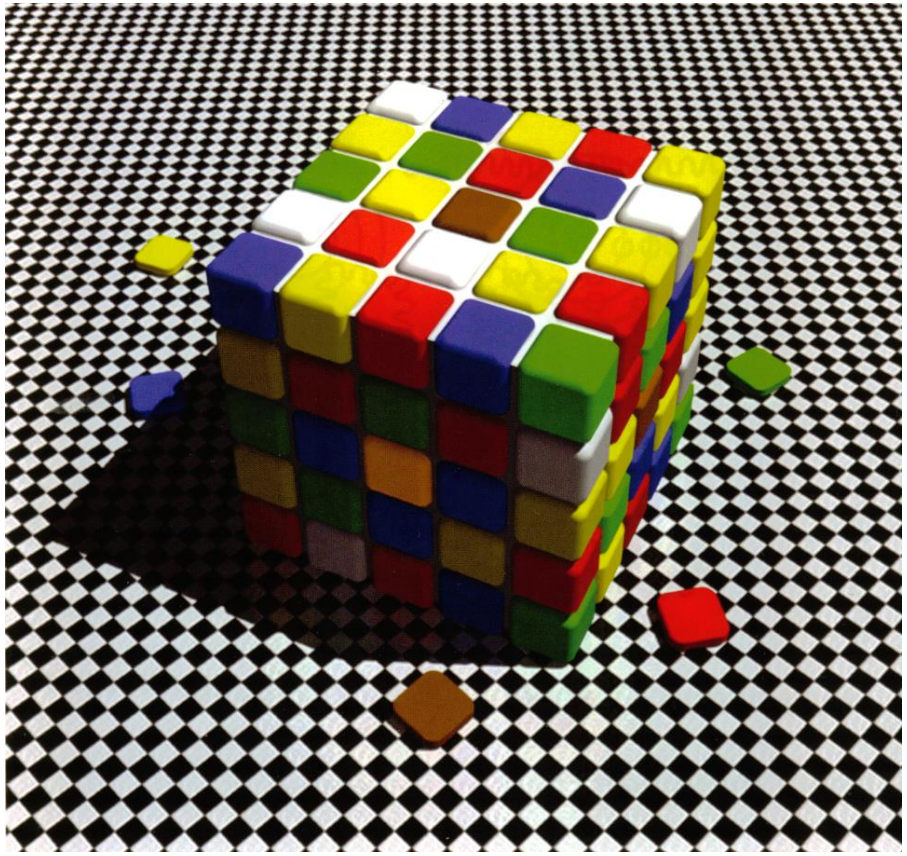

luminance = illuminance * reflectance


ECVP
2005
ECVP
2005

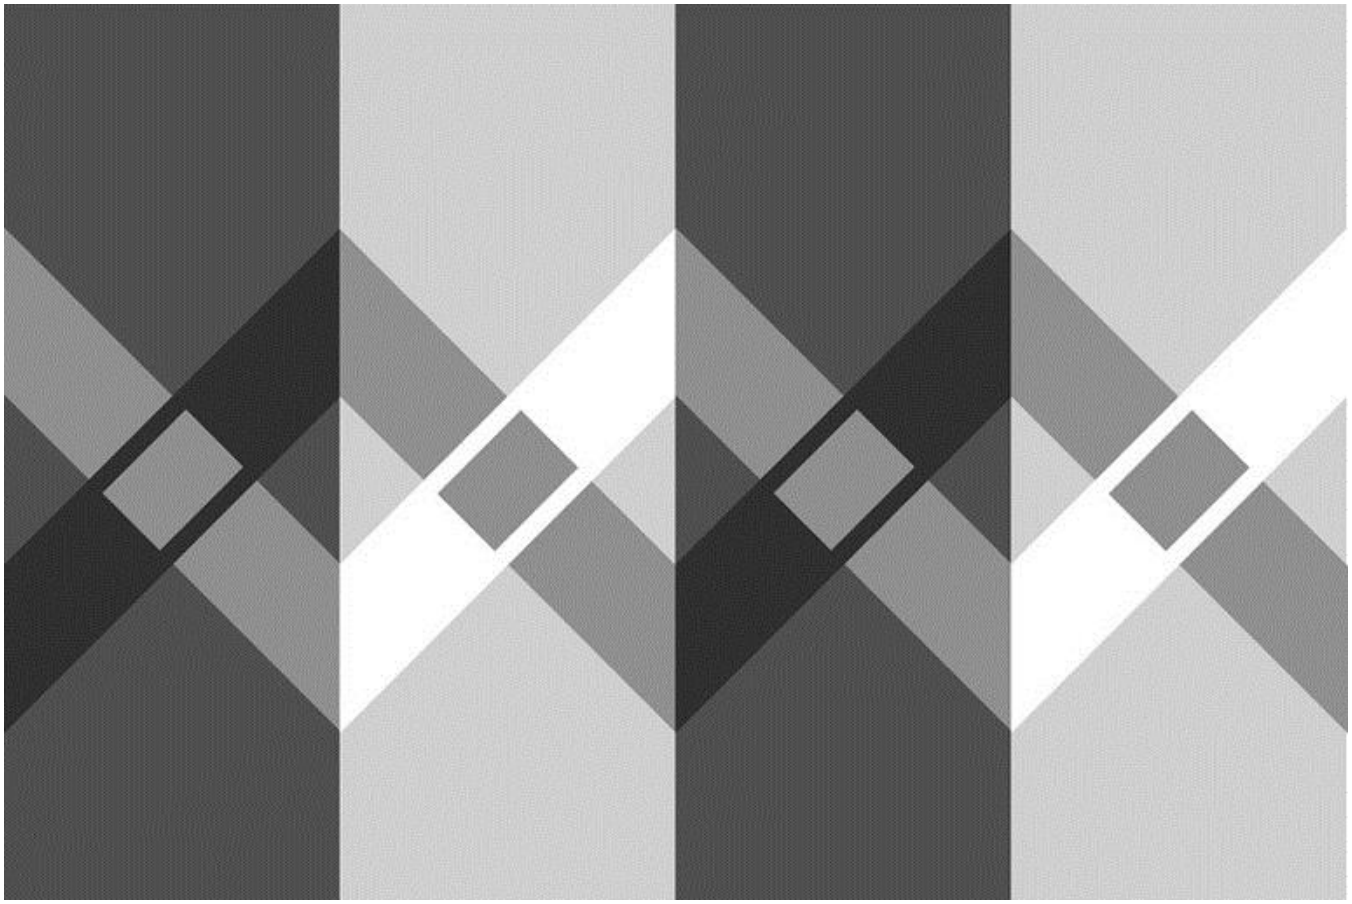

E. H. Adelson, 2000, www-bcs.mit.edu/gaz/

More examples for ambiguities in visual perception

what we want to know		what we can measure	context
aspect ratio	slant (of an ellipse)	aspect ratio of the retinal image	our example
reflectance	illuminance	luminance	lightness constancy
size	distance	retinal size	size – distance ambiguity
intensity	wavelength	absorption	principle of univariance
object motion	motion of the retina	retinal motion	efference copy theory
speed	direction of motion	image motion	aperture problem
3D shape	direction of light	2D retinal image	light-from-above bias
distance	motion of observer relative to object	retinal motion	motion parallax

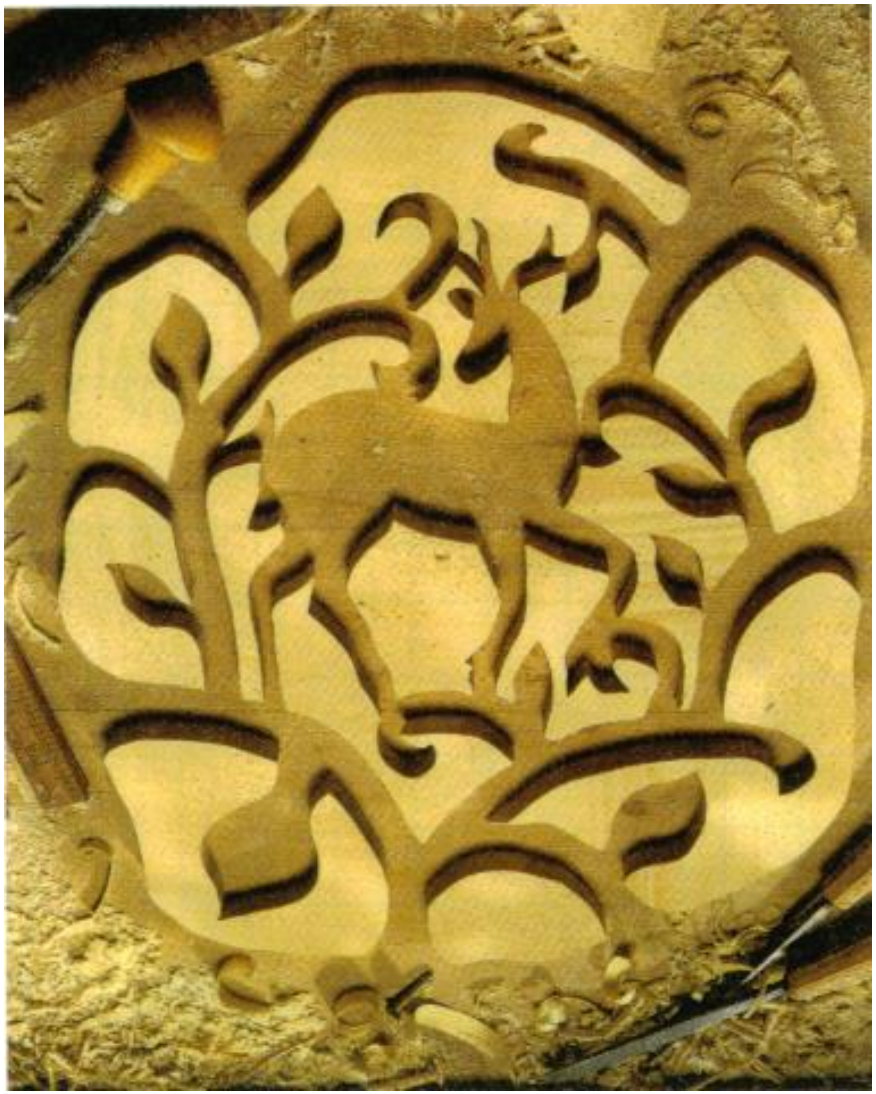

Light-from-above bias

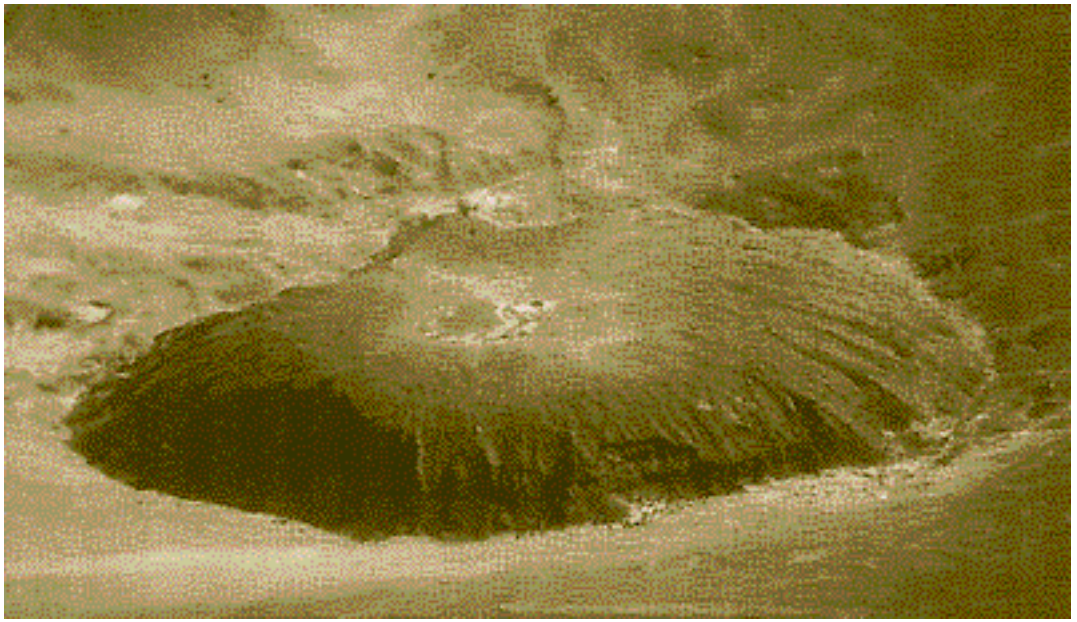

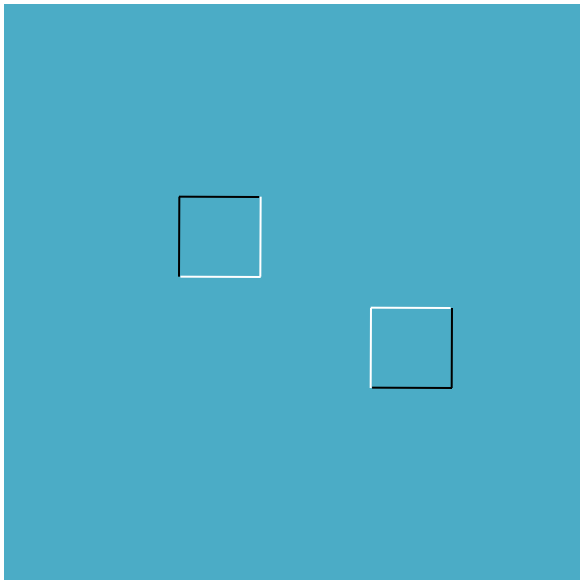

Rock, 1975

Leeuwenberg & Boselie 1988

Simplicity vs Likelihood

Simplicity or Likelihood as an organizing principle?

- What is role of redunancy?
- Simplicity: "entities must not be multiplied beyond necessity" (Occam's razor)
- Likelihood: Correlations in the data are necessary to make predictions

Simplicity vs Likelihood

It can be shown that there is no conflict and that the principles are in fact equivalent

- Mumford 1992, Chater 1996, Kolmogorov 1965
- Simplicity is quantified in the context of information theory: What is the minimum length required to encode data s, using a particular hypothesis w?
- The code will have two parts:
  (a) the encoding of the chosen hypothesis w_i
  (b) the encoding of the data s, given w_i
- Using Shannon's *noiseless coding theorem*, it can be shown that the minimal code length required to encode an item A_i which has a probability of occurrence of $p(A_i)$ is

$$l(A_i) = \log \frac{1}{p(A_i)}$$

Simplicity vs Likelihood

$$\widehat{w_i} = \arg\min_{w_i} \left[l_{w_i} + l_{s|w_i} \right]$$

$$= \arg\min_{w_i} \left[\log \frac{1}{p(s|w_i)} + \log \frac{1}{p(w_i)} \right]$$

$$= \arg\max_{w_i} \left[\log\, p(s|w_i) + \log\, p(w_i) \right]$$

$$= \arg\max_{w_i} \left[\log\, \left(p(s|w_i)\, p(w_i) \right) \right]$$

$$= \arg\max_{w_i} \left[p(s|w_i)\, p(w_i) \right]$$

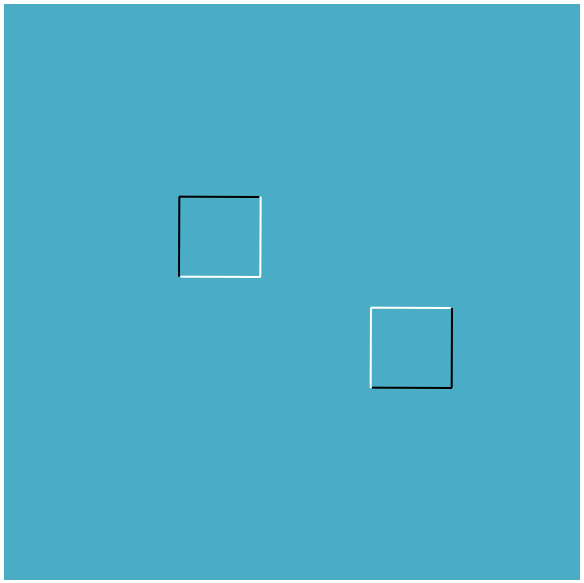

Horace Barlow (1921 -)

What is efficient coding?

- Efficient coding should convert hidden redundancy into a manifest, explicit, immediately recognizable form, rather than reduce it or eliminate it.