From Raw Data to Abstract Concepts

Harri Valpola

Computational neuroscience group

Helsinki University of Technology

Overview

- Starting point: the brain needs to support intelligent behaviour
  - Control and decision making
  - Predict the consequences of what-if scenarios
- Q1: What kind of structures need to be represented?
  - Conditional probabilities
  - Hierarchy of abstractions
- Q2: How to select relevant information?
  - Distributed selection
  - Two timescales: attention and learning

How to behave intelligently

- Evaluate the rewards for each potential action and choose the one that maximises expected reward
- Problem: overwhelmingly difficult mapping

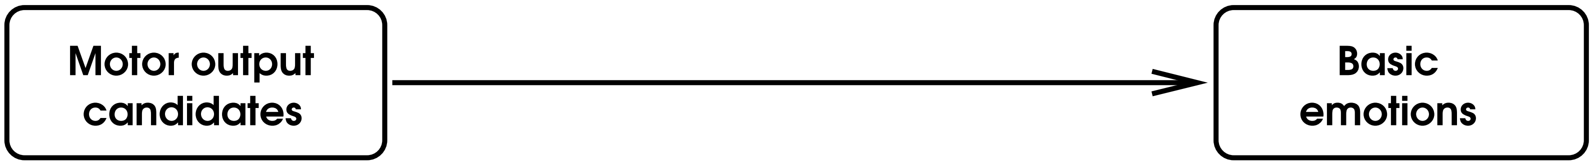

How to behave intelligently

- Solution: add sensory consequences as an intermediate step
- Problem: the mappings are still too complex, depend on context

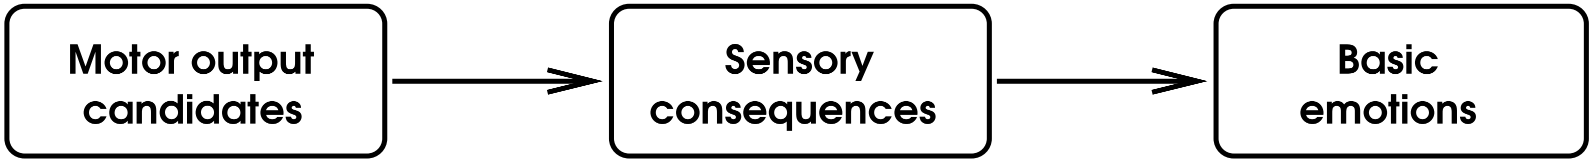

How to behave intelligently

- Solution: mappings modulated by context
- Problem: what is the context?

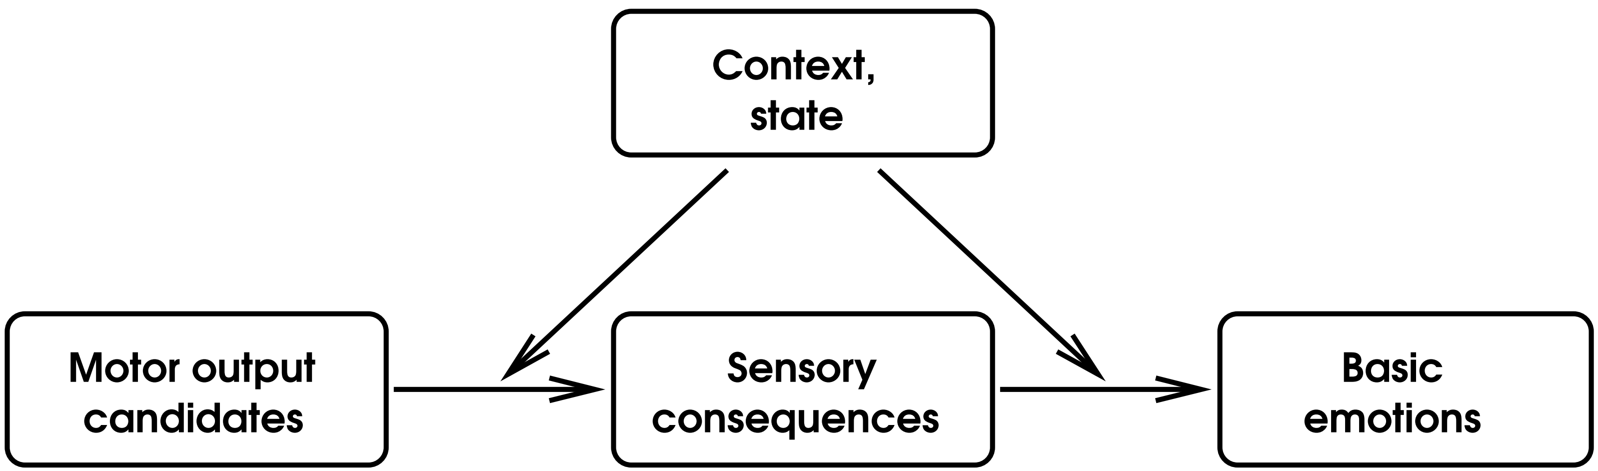

How to behave intelligently

- Solution: recognize the context based on lower-level activity and relations
- Problem: how? Did the problem become any easier?

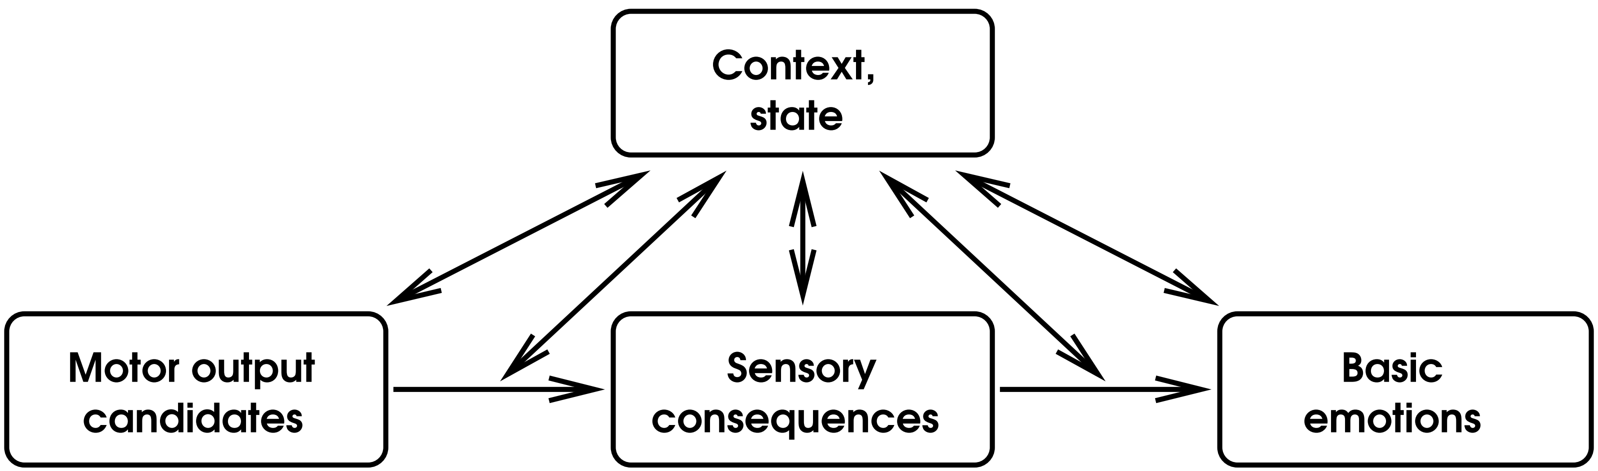

Let's study the solution adopted by the brain

Neuroscience

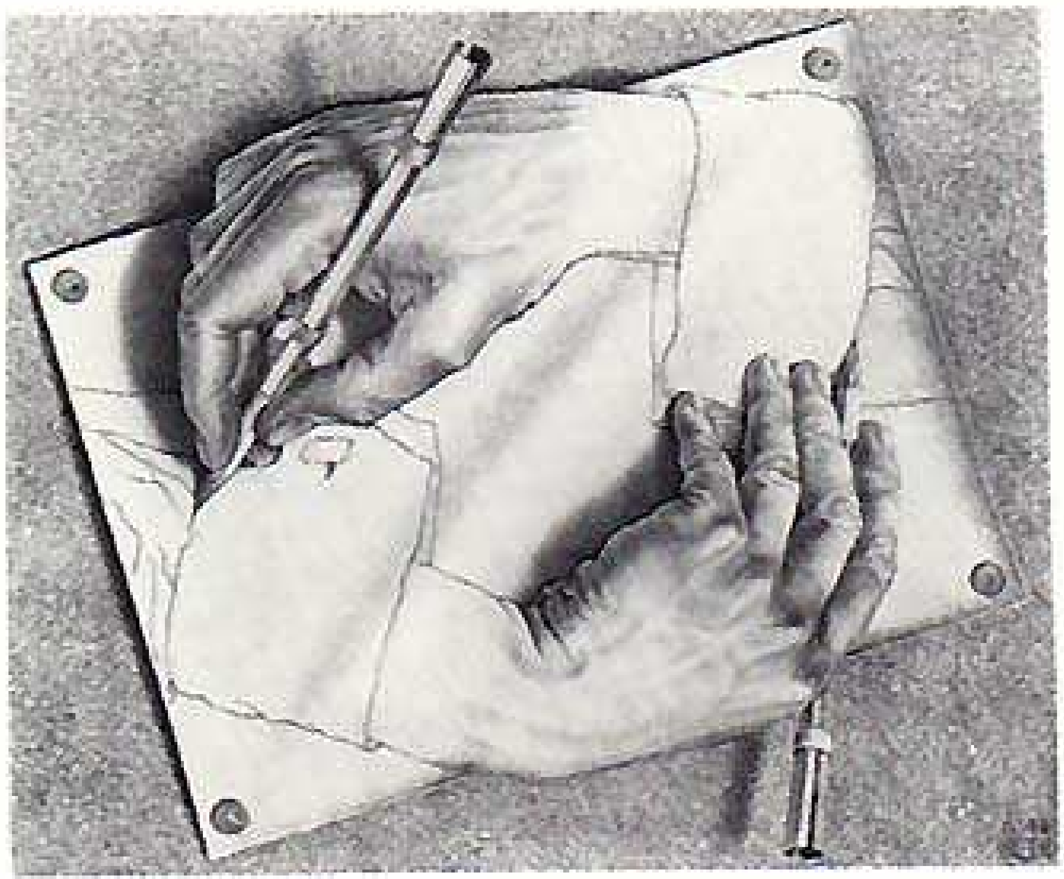

Technology

Levels of explanation

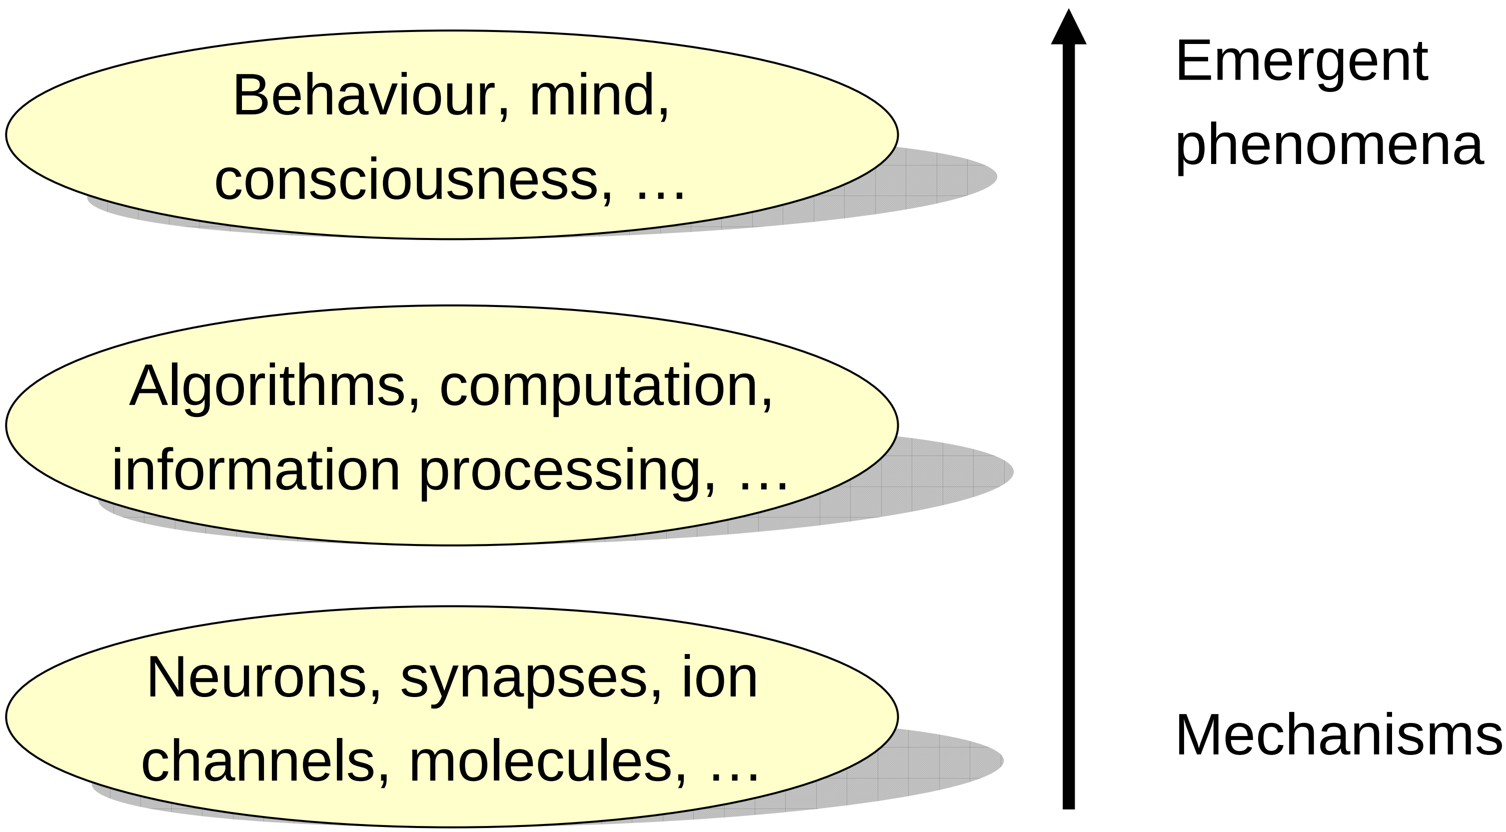

Levels of explanation

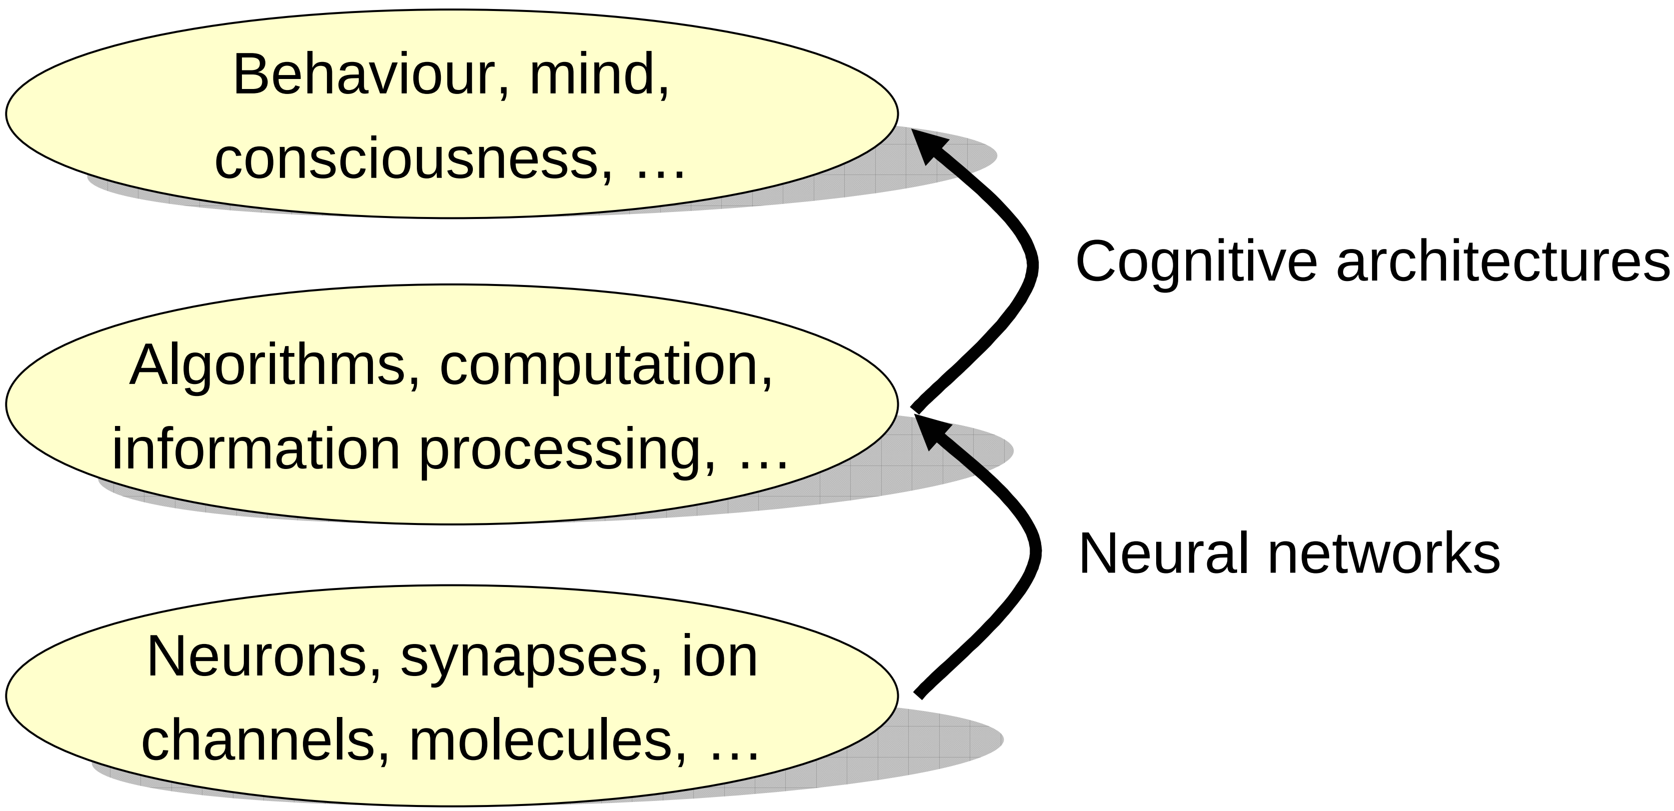

Vertebrate cognitive architecture

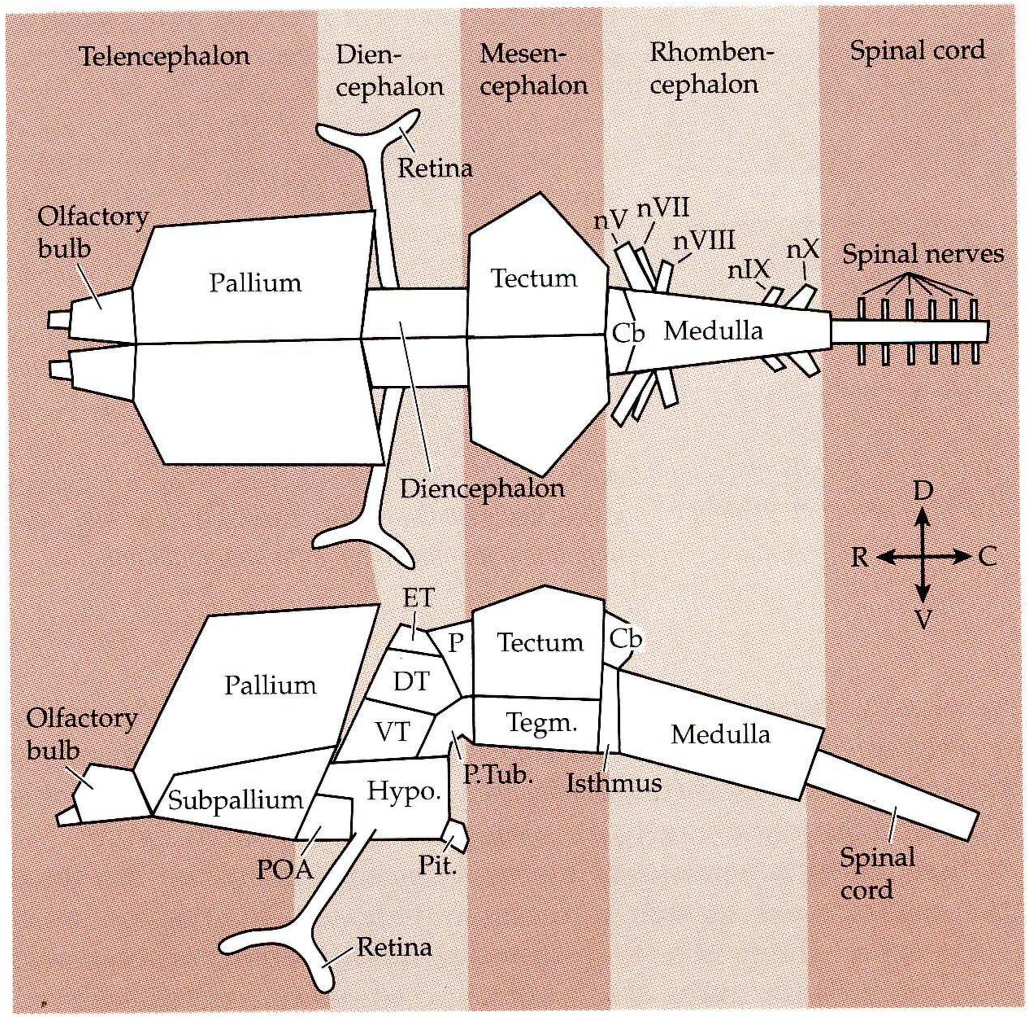

Computational neuroscience group

What: Figuring out how the brain works.

How: Building brains for robots = system-level modelling and implementation of a whole vertebrate/mammalian brain.

Intact cerebellum and sober

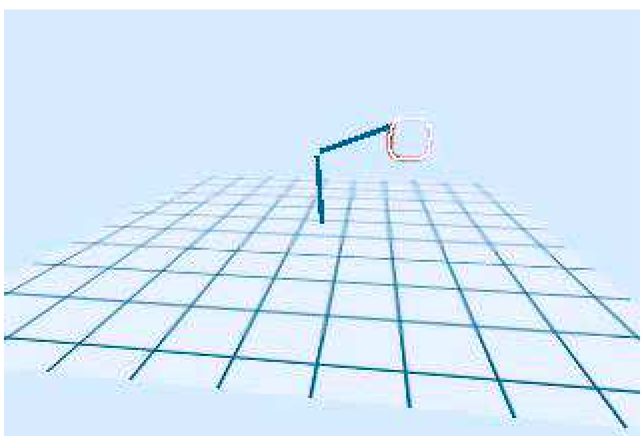

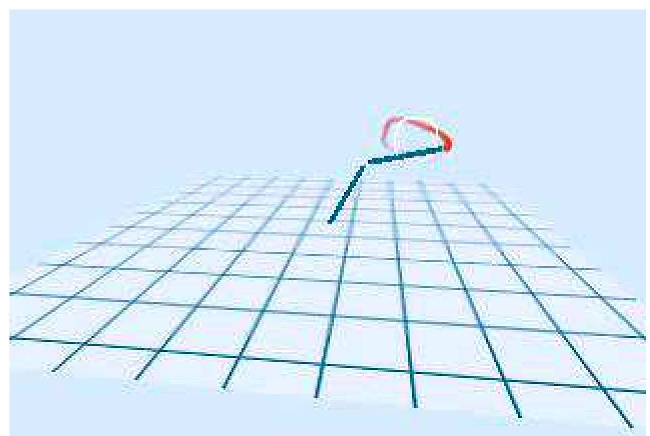

Cerebellar lesion or drunk

Mammalian cerebral cortex

- A hierarchy of feature maps: increasing levels of abstraction

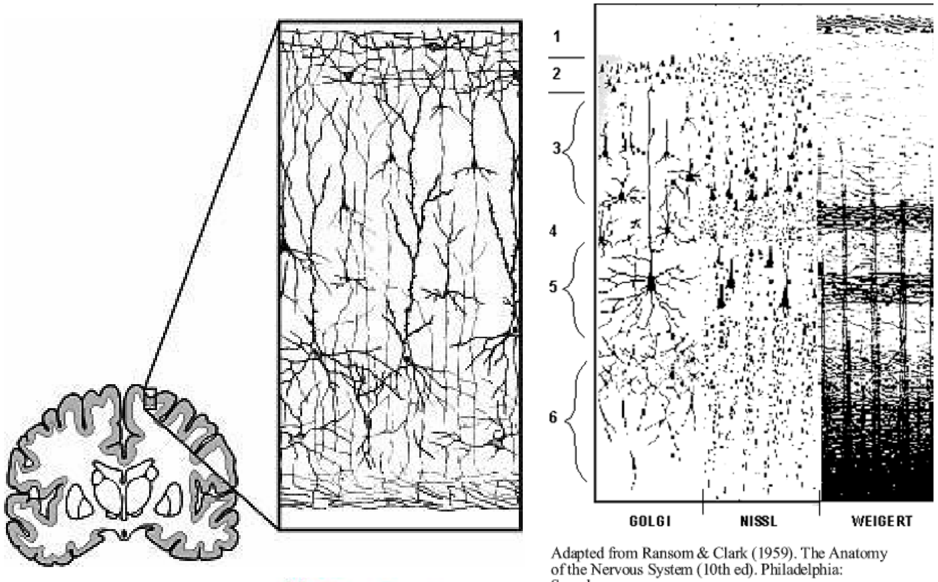

Adapted from Ransom & Clark (1959). The Anatomy of the Nervous System (10th ed). Philadelphia: Saunders.

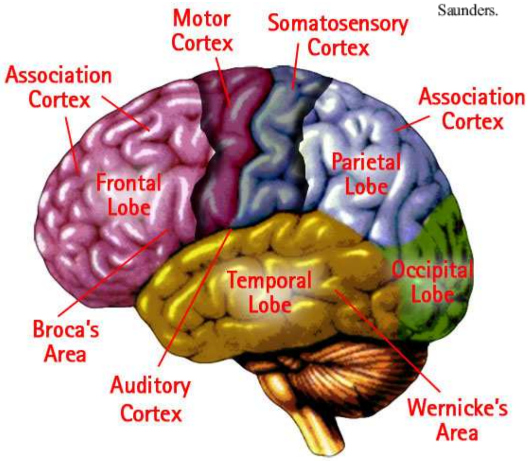

Cortical algorithm

- Stereotypical 6-layered structure, isocortex
- One algorithm, applied to different data?
- Functional unit: cortical column
- Different layers have different types of inputs and output

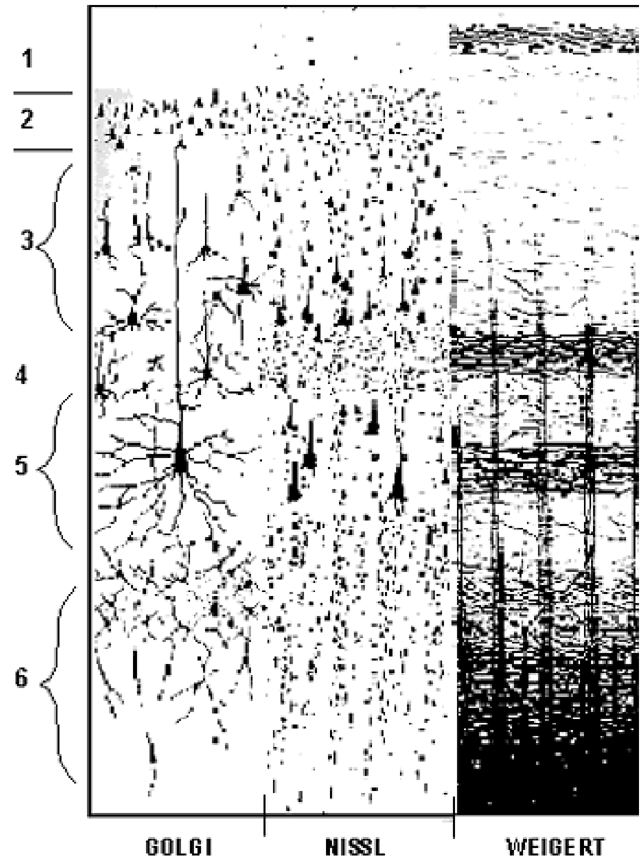

Adapted from Ransom & Clark (1959). The Anatomy of the Nervous System (10th ed). Philadelphia: Saunders.

Part 1: Modelling correlation structures

Model for correlation structure

- Recurring template:

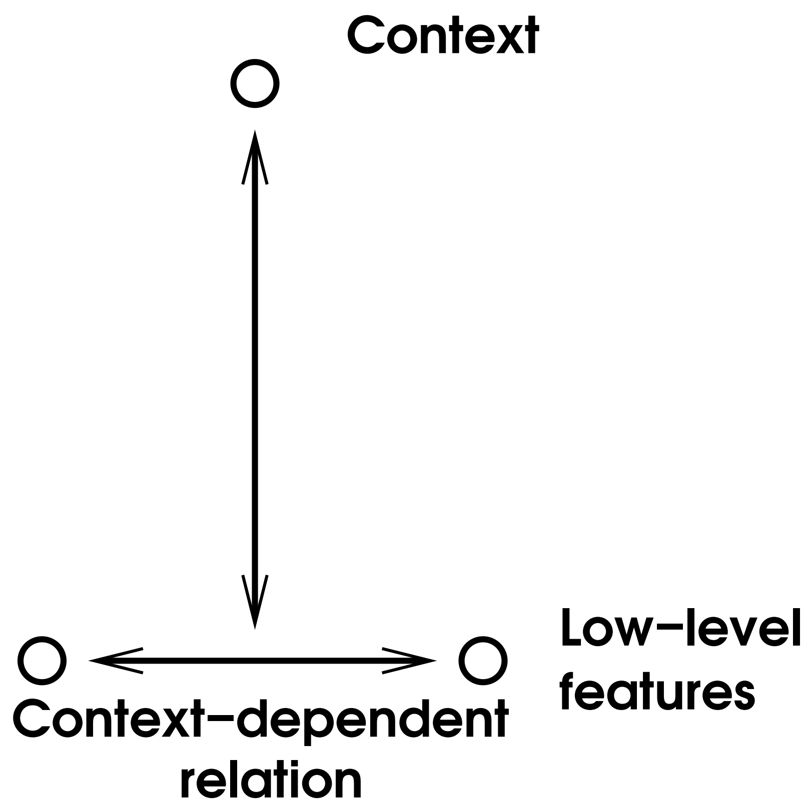

Model for correlation structure

- Key problem: how to learn this efficiently?

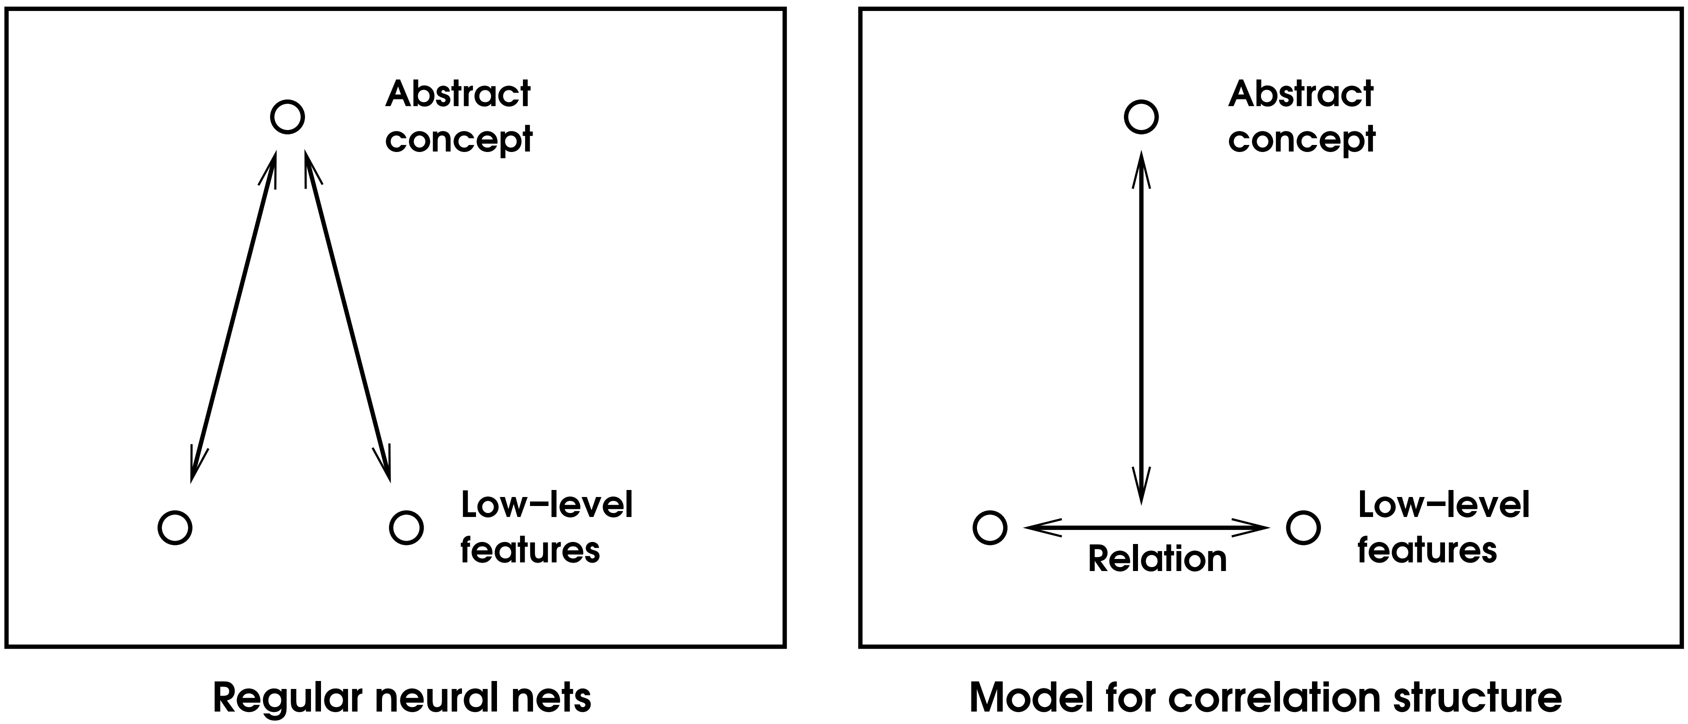

Brain-inspired learning mechanism

- The cerebral cortex has come up with a learning algorithm that avoids the combinatorial explosion
- Seems to work fine in simulations
- Details later… Sorry!

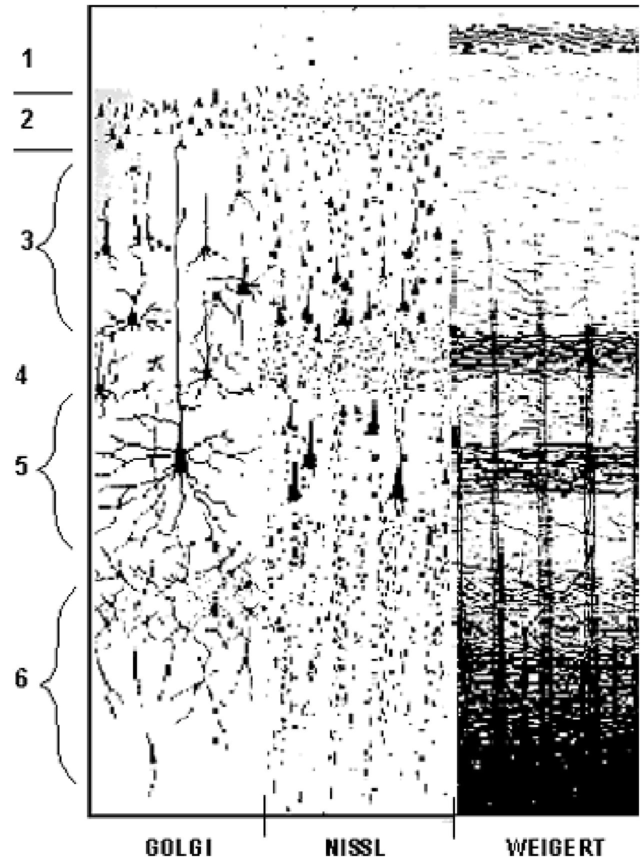

Adapted from Ransom & Clark (1959). The Anatomy of the Nervous System (10th ed). Philadelphia: Saunders.

A building block for hierarchical models

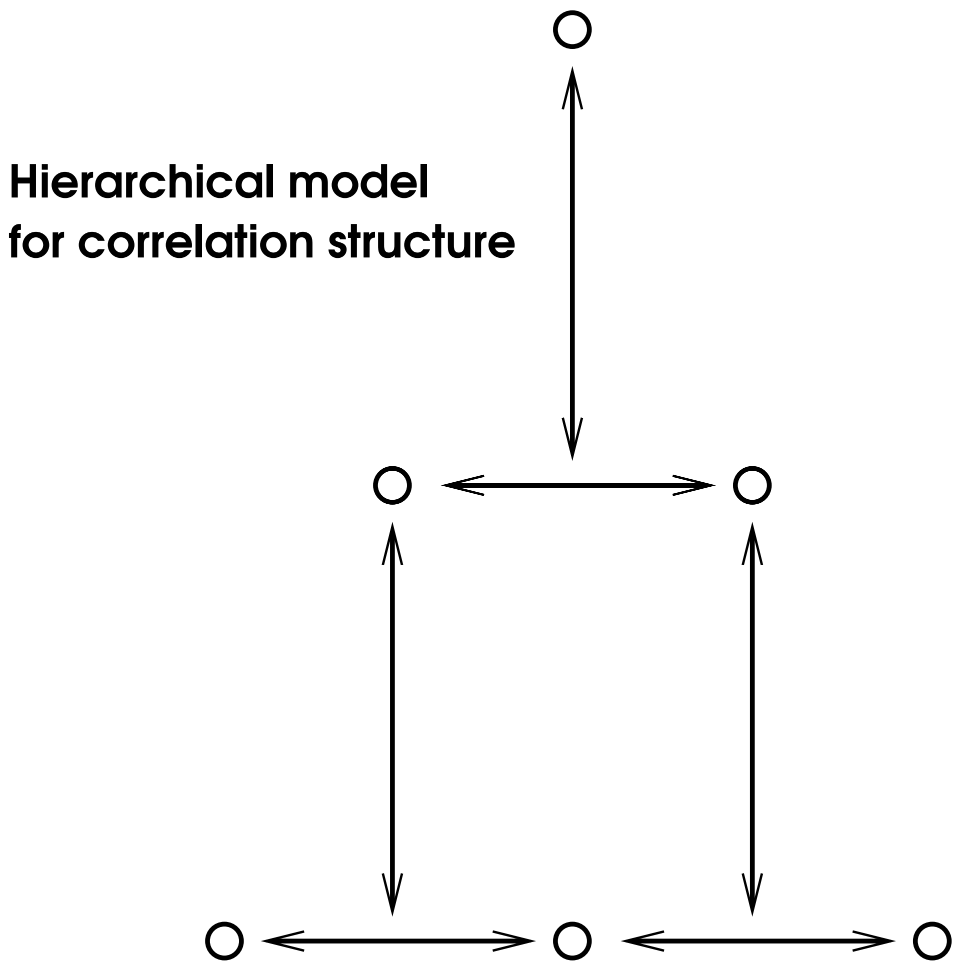

“Normal” connections can be included

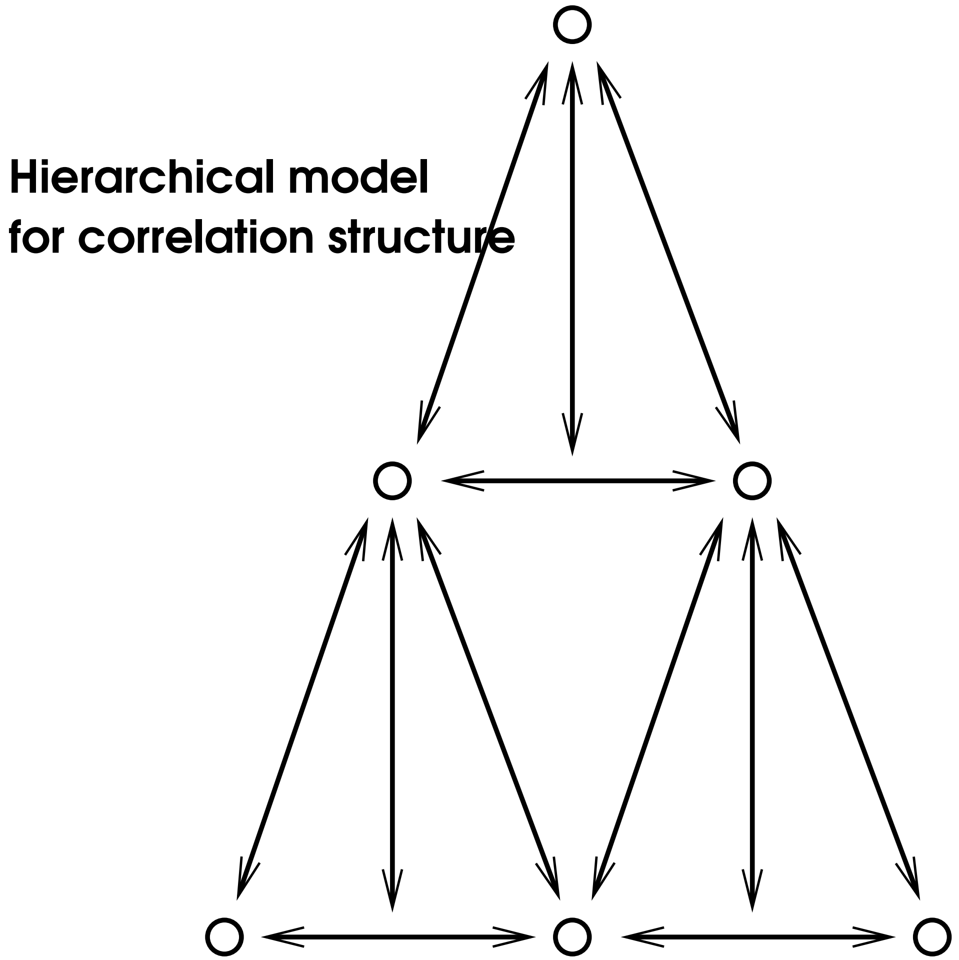

Revisiting the original problem

- Do we now have a better idea of how the context might be learned and used?

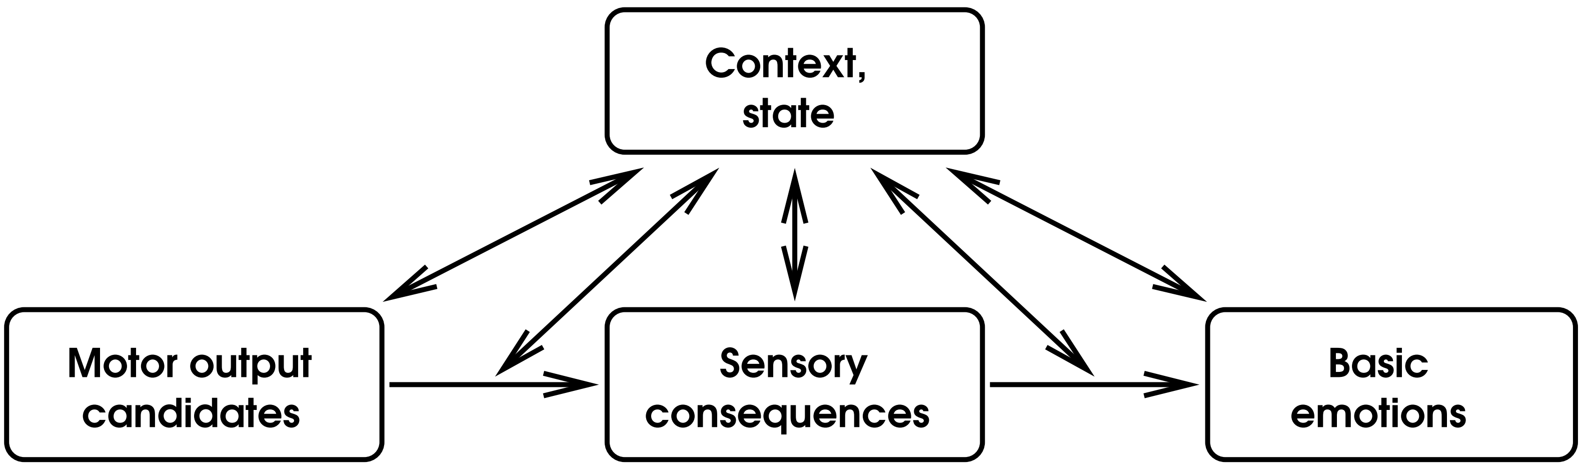

Example 1: from movement to action

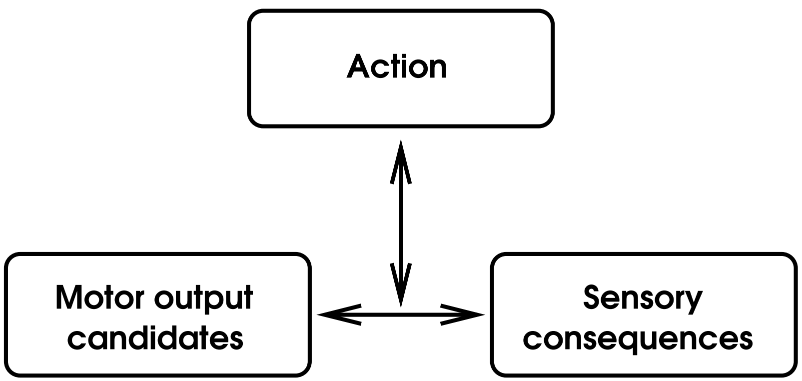

Example 2: shape from texture

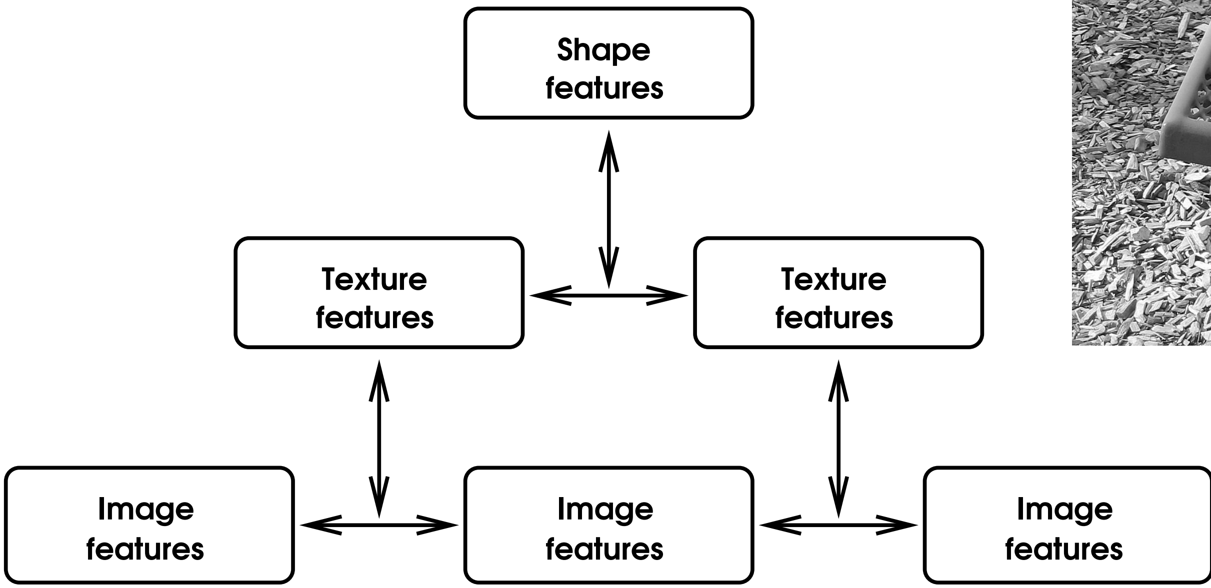

Example 3: feelings vs. emotions

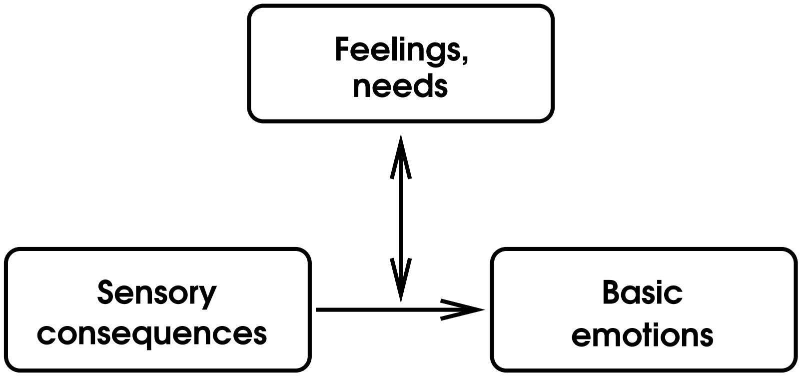

Part 2: Selection

Bayesian theory says:

- Decisions are based on
  1. Beliefs (measured by probability)
  2. Utilities
- The recipe:
  1. Evaluate the probabilities of all possible states of the world (probabilistic inference)
  2. Evaluate the probabilities of all outcomes for each and every potential action (probabilistic inference)
  3. Choose the action which maximises the expected utility
- This is optimal if there are no restrictions on the available computational resources

Key problem:
How to select useful information?

- But… computational resources *are* restricted →
- It is impossible to consider all the states and actions →
- It is necessary to select information in order to make decisions

- Selection is a type of decision, in other words:
- In order to decide we need to decide… Infinite regress!

Why does the cortex need such a large number of feedback connections?

- Primary input usually from bottom-up (from the senses)
- Feedback connections are far more numerous (order of 10 x)
- Where are all the "modulatory" connections needed for?

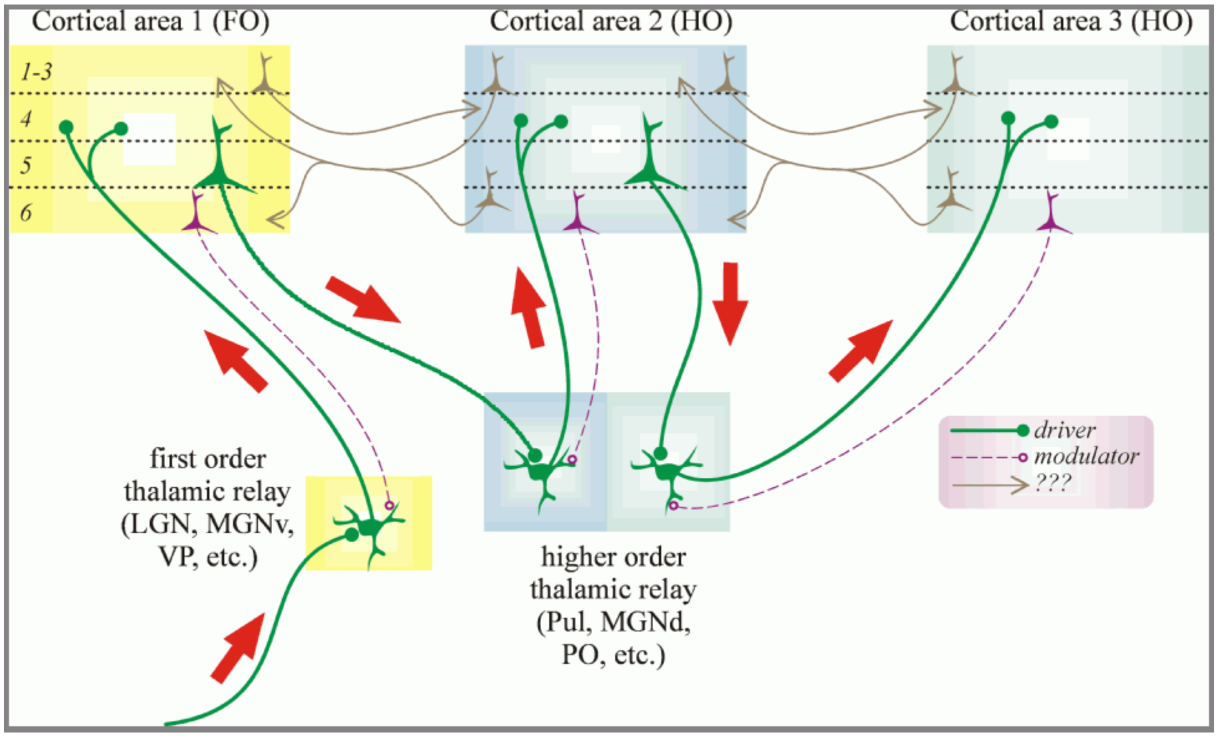

Hierarchy of areas

- The cerebral cortex is connected as a hierarchy of areas
- The representations get more abstract on higher levels

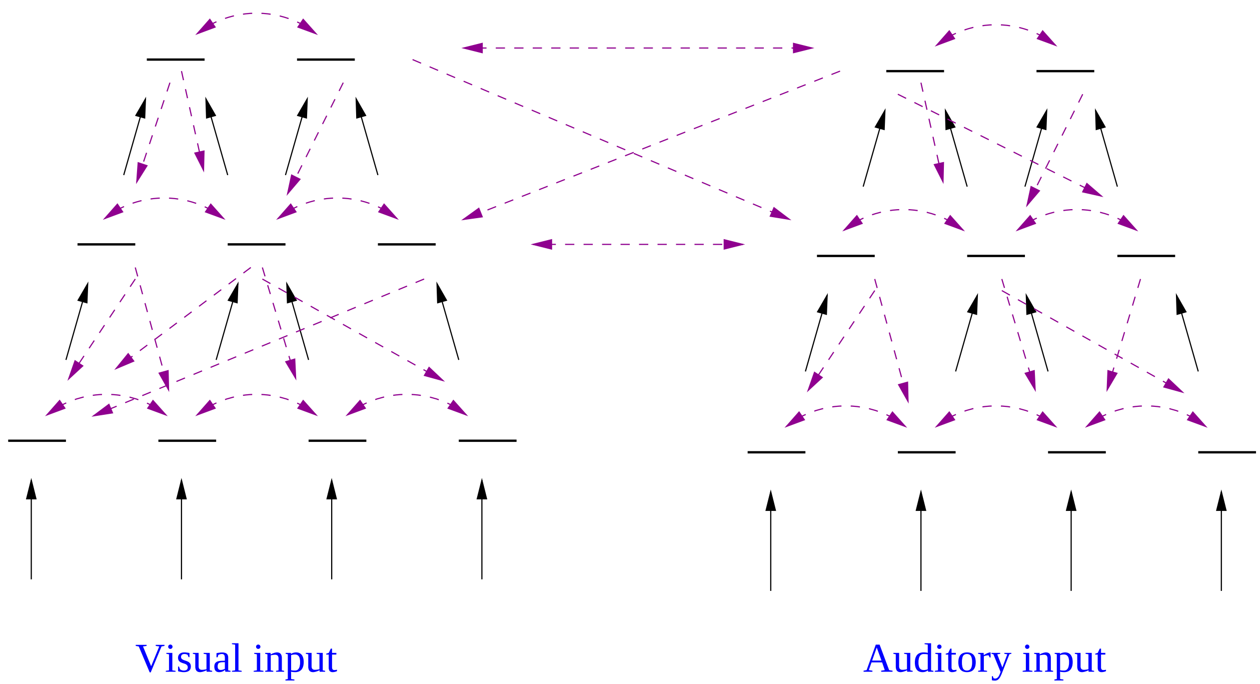

Brain's solution: distributed selection

- Each cortical area selects information to be represented
- Biased-competition model of attention: attention emerges from local selection and global communication

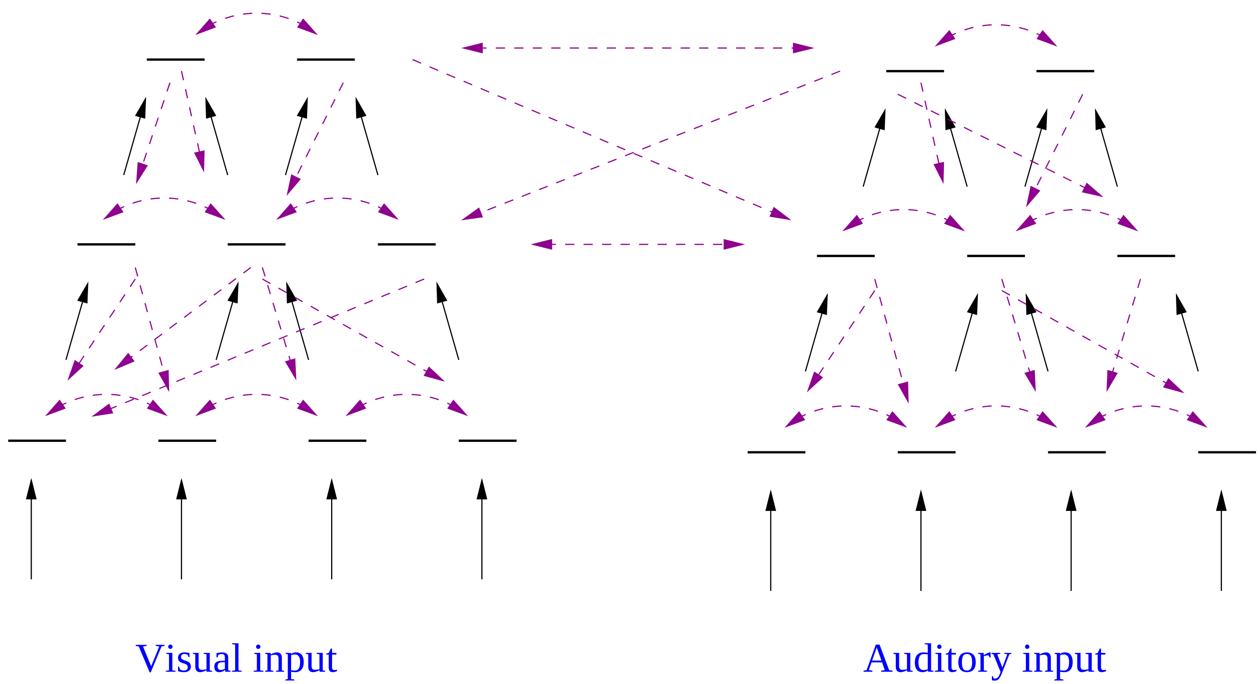

Attention and learning: selection on different timescales

- Our work: biased competition + competitive learning
- Within the Bayesian framework, the only difference between perceptual inference and learning is the timescale
- Attention and learning in the cortex are intimately coupled
- Both are a form of selection, only timescales differ

The value of information

- Motor areas may be able to rely (at least partly) on global reward signals (reinforcement learning)
- Sensory areas or a large brain: credit assignment problem
- More specific but locally available information: predictive power or “are the others listening?”

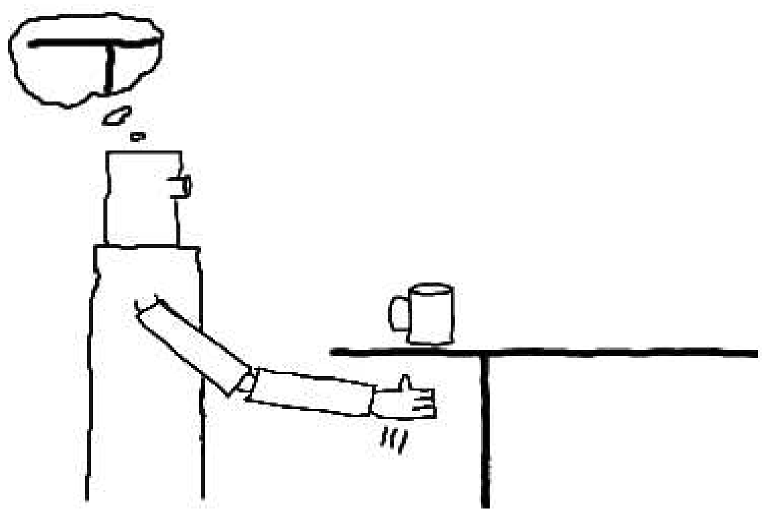

A model of a cortical area

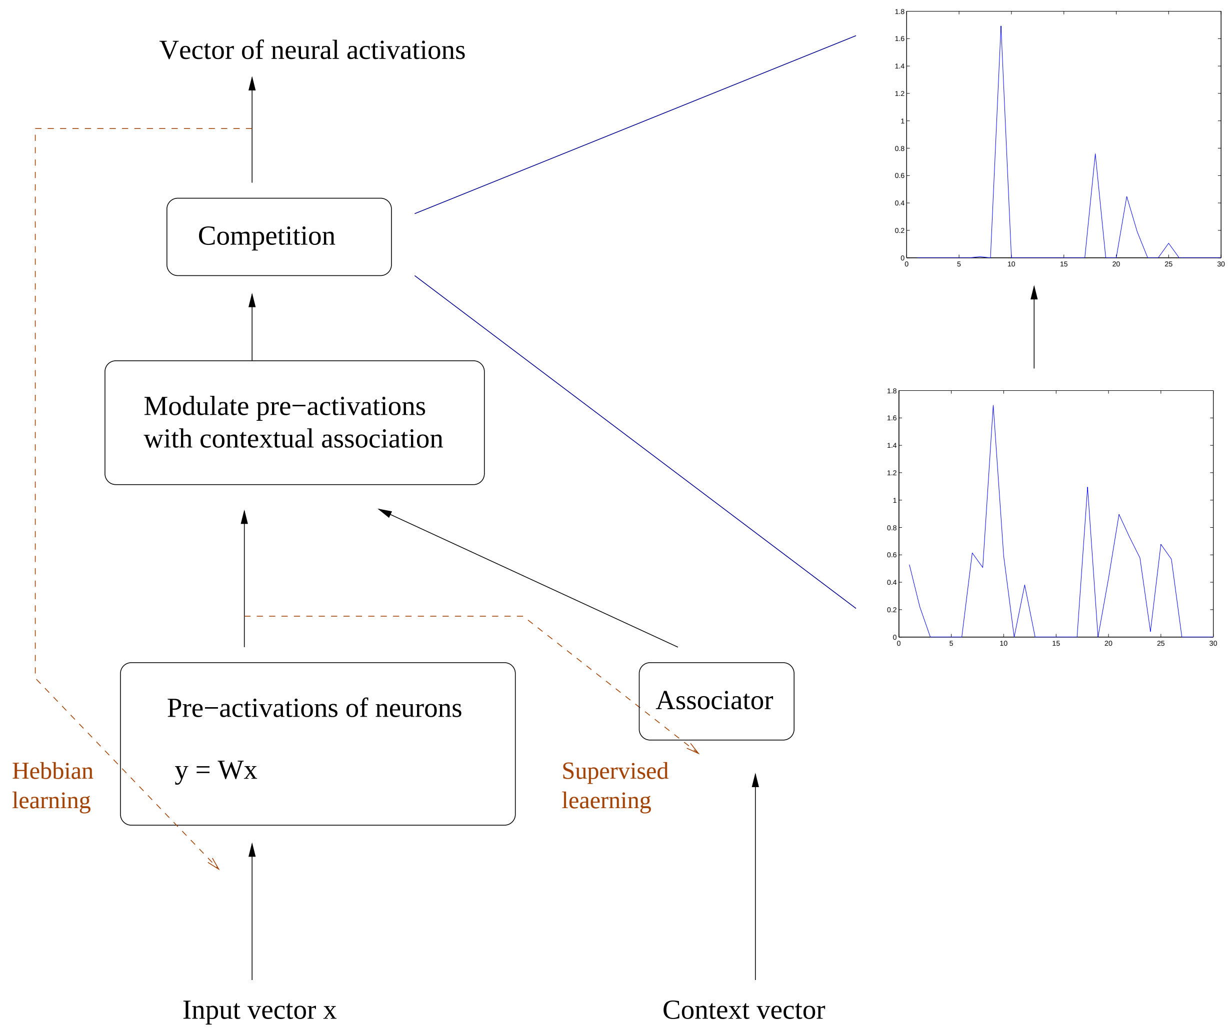

Results 1: data

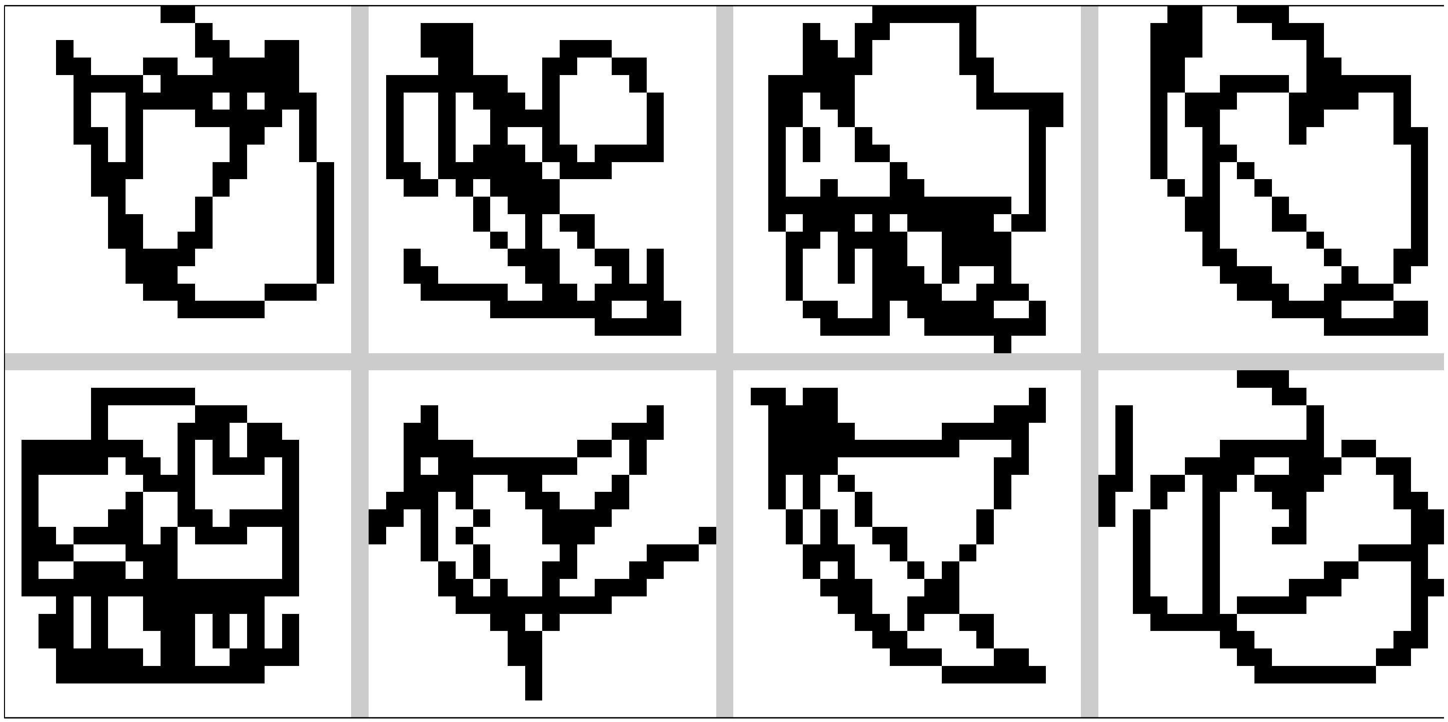

Results 2: abstract categories

- Four samples of test data, each of which have activated the same coalition of neurons at the highest level
- This invariant recognition of abstract categories was achieved without any supervision even if the objects never appeared in isolation

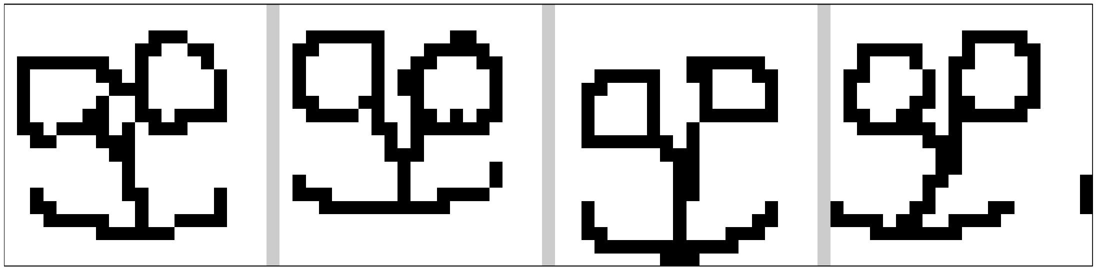

Results 3: switching attention

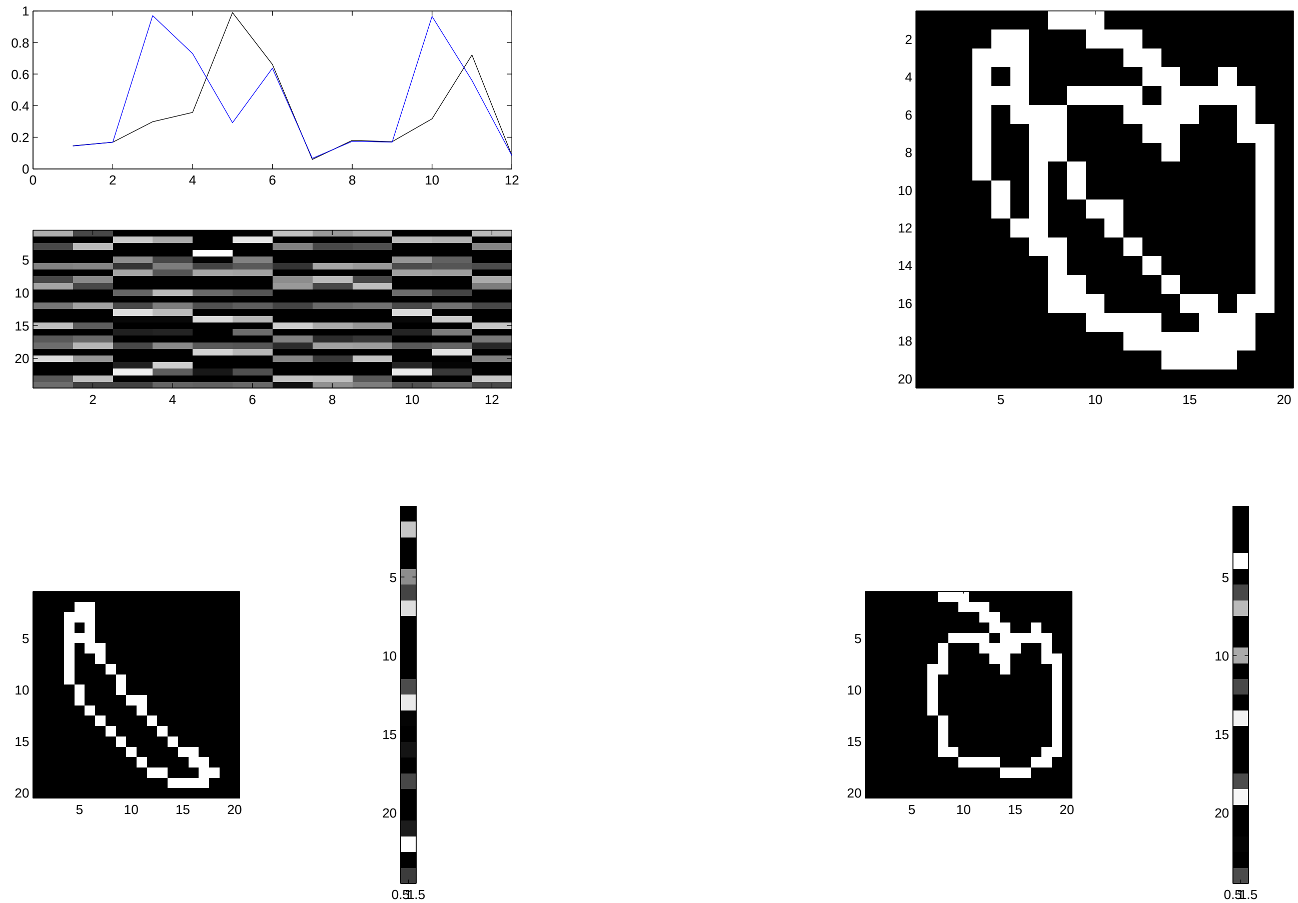

Conclusion

- We are beginning to understand the information processing on the cortex (perception, attention, learning, imagination, decision making, …)
- Learning from the brain:
  - Overall structure of the cognitive architecture
- Learning from the cortex:
  - Learn and use abstract concepts
  - Select relevant information
- Useful solutions that work right now + Potential for artificial general intelligence

Thank you!

www.zendroid.com