UiO : University of Oslo

Selected topics from cognitive science
(and their relevance for language technology)

Pierre Lison
University of Oslo, Dep. of Informatics

Språkteknologisk seminar @ LNS
May 2, 2011

Introduction

- Long-time personal interest in the relation between language technology and cognitively-oriented approaches to linguistics
- Present here a "patchwork" of interesting ideas, rather than a rigidly structured presentation
    - Review of literature on these topics, from the perspective of neuroscience, psychology, psycholinguistics, cognitive linguistics, and philosophy of mind
- Talk should be viewed as an open discussion - *Feel free* to interrupt me at anytime to provide comments, questions, remarks
- Note that I'm neither a brain scientist nor a cognitive psychologist, so forgive me if I'm not always as precise and accurate as I should
- Thanks to Alessandra Zarcone (Uni. Stuttgart) and Olga Kukina (Uni. Saarland) for useful comments on an earlier version of these slides

Outline

- Motivation
- The brain and its functions
- Three important ideas from cognitive science:
    - The mind is embodied
    - The mind is proactive
    - The mind is social and emotional
- Implications for (computational) linguistics
- Conclusion

Outline

- **Motivation**
- The brain and its functions
- Three important ideas from cognitive science:
  - The mind is embodied
  - The mind is proactive
  - The mind is social and emotional
- Implications for (computational) linguistics
- Conclusion

Cognitive science for LT

- Are insights from cognitive science in any way relevant for language technology?
  - My answer to this: ***yes***, definitely!
  - I will motivate my answer with three complementary arguments

Cognitive science for LT

- **Argument 1**: Language is not an isolated faculty in the human brain, but is *part and parcel* of human cognition
  - Learning how the mind works can help us understand what language is, and how its functions are instantiated in the neural tissue
  - Essential properties of natural languages are closely tied to peculiarities of the human brain (language is "shaped" by the brain)
    - Importance of *ontogenetic* and *phylogenetic* brain development
    - Increasing neuroscientific evidence that language functions rely to a large extent on the *recruitment* of evolutionarily pre-existing brain subsystems
  - Major shift of perspective compared to early cognitive science, which advocated the "*fundamental irrelevance of the specific hardware*" to understand cognition (cf. Fodor, Newel, Pylishyn)

[Anderson, M. (2003), Embodied cognition: A field guide, *Artificial intelligence*]
[Van Berkum, J. J. A. (2010). *Italian Journal of Linguistics*]

Cognitive science for LT

- **Argument 2**: The human mind seems to be incredibly good at processing language!
    - People are able to process language *robustly* and *accurately*
    - They understand and produce language very *rapidly*, *incrementally* and in *real-time* (often without conscious effort)
    - They automatically *adjust* to the context, online
    - And this despite strong *limitations* on processing resources (slow connections, limited memory, etc.)
- So, maybe we can get some insights from human language processing to improve our processing algorithms?

Cognitive science for LT

- **Argument 3**: Many language technology applications require some kind of interaction with human users
  - This is obviously the case for spoken dialogue systems
  - But is also true (although in a slightly different sense) for question answering systems, IR and IE applications, machine translation software, etc.
  - In order to interact *naturally* with humans using natural language, it is therefore useful/necessary to gain some knowledge about the fundamental cognitive factors driving human verbal interaction

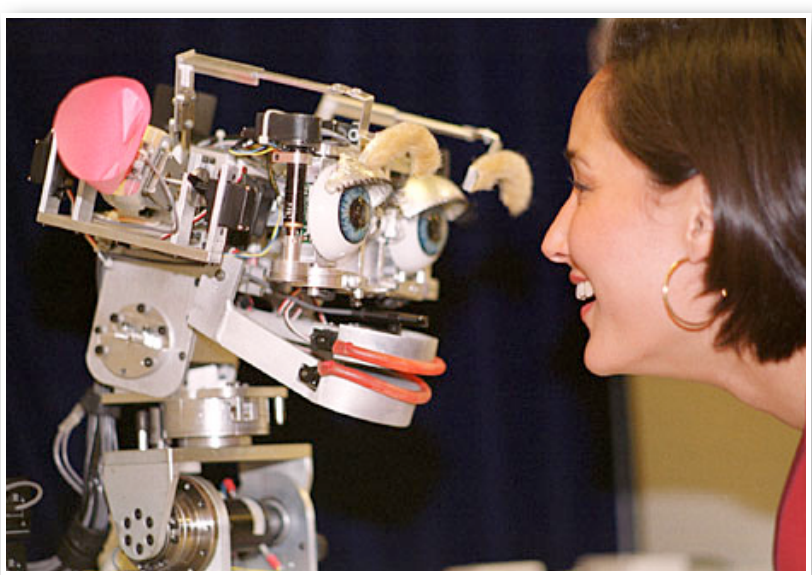

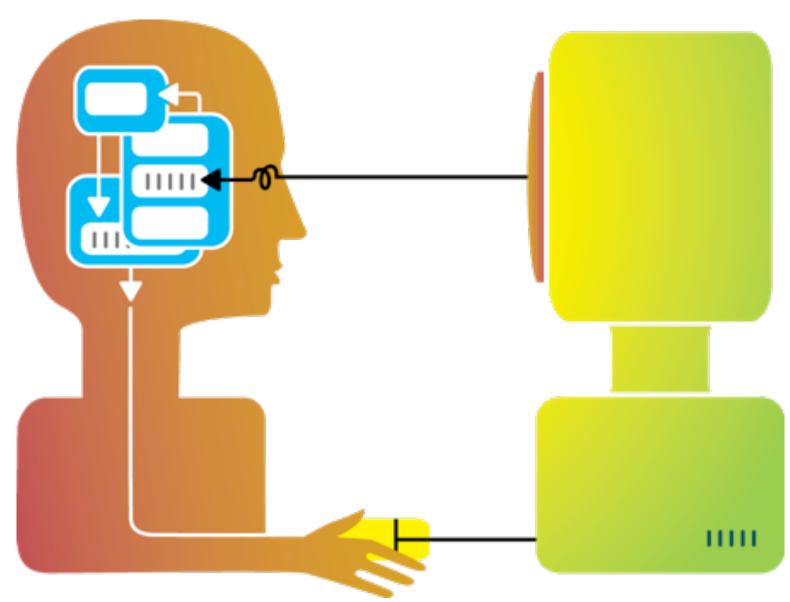

Cognitive science for LT

- Does that mean that, in order to perform their tasks efficiently, NLP algorithms must seek "cognitive plausibility" at all costs and try to mimic what is known about neural processes underlying human language processing?
- ***Of course not***! This is neither technically achievable (integrated circuits are *not* neural tissue and operate very differently from them) nor actually desirable
- But drawing inspiration from cognitive processes at the *functional* level is certainly useful
- This is what I will try to do in this talk

Outline

- Motivation
- **The brain and its functions**
- Three important ideas from cognitive science:
  - The mind is embodied
  - The mind is proactive
  - The mind is social and emotional
- Implications for (computational) linguistics
- Conclusion

The brain

- The brain is the control centre of the nervous system
- Divided in various brain regions:
    - Brain stem (“reptilian brain”): basic reflexes and vital functions
    - Cerebellum: movement and balance
    - Diencephalon (“midbrain”), which includes the thalamus and the hypothalamus
    - Basal ganglia, amygdala and hippocampus
    - And finally, the cerebral cortex, itself divided into two hemispheres and four lobes (frontal, parietal, occipital, and temporal)
- Contains about 80-90 billions neurons, and same number of glial cells

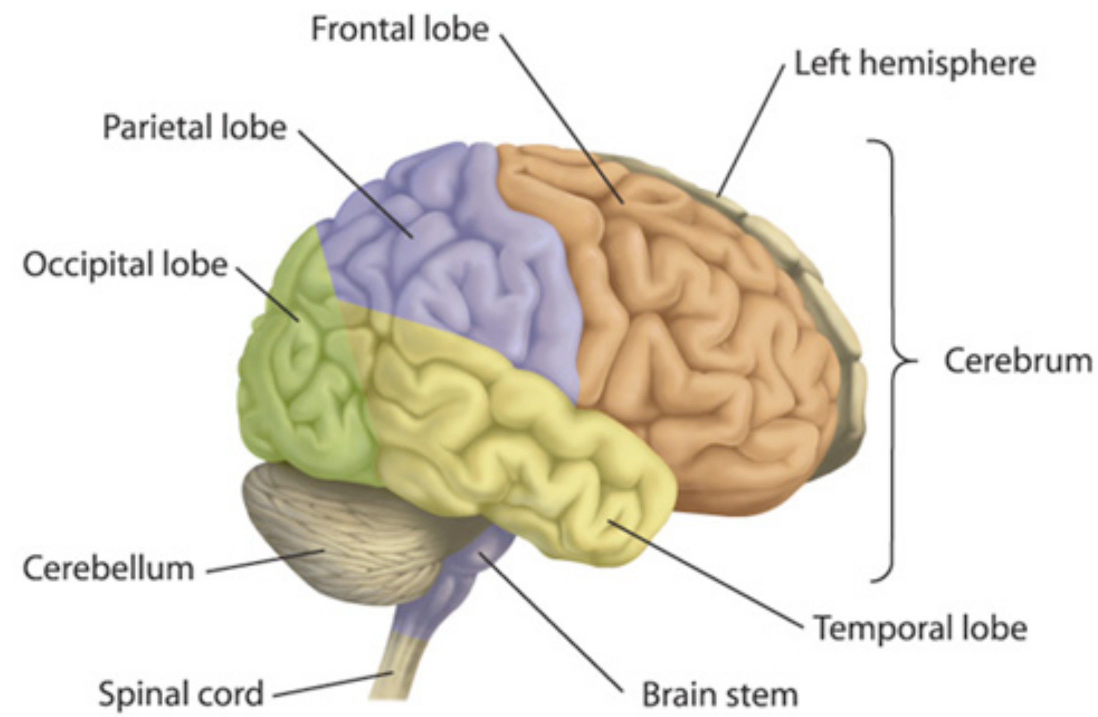

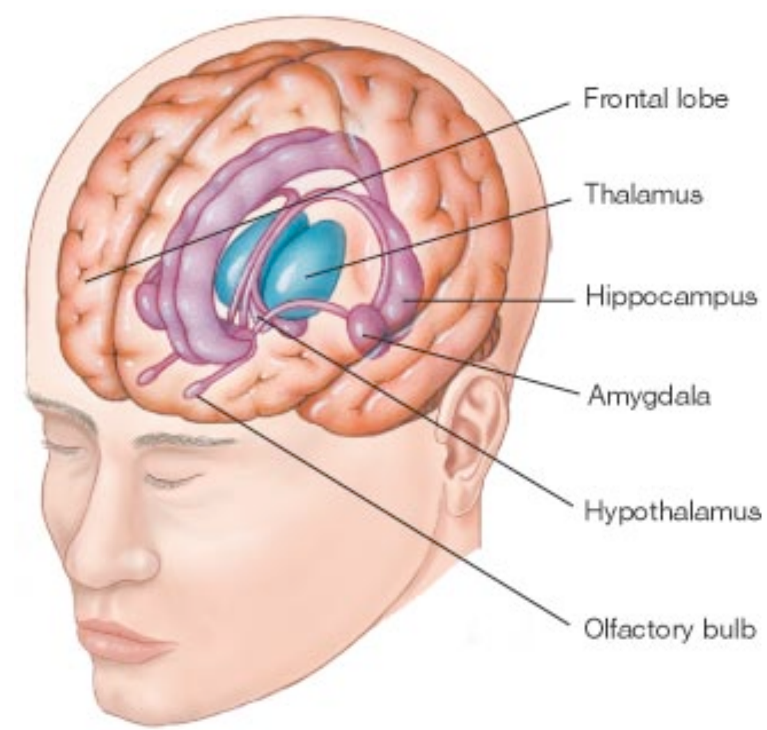

Functional areas in the cortex

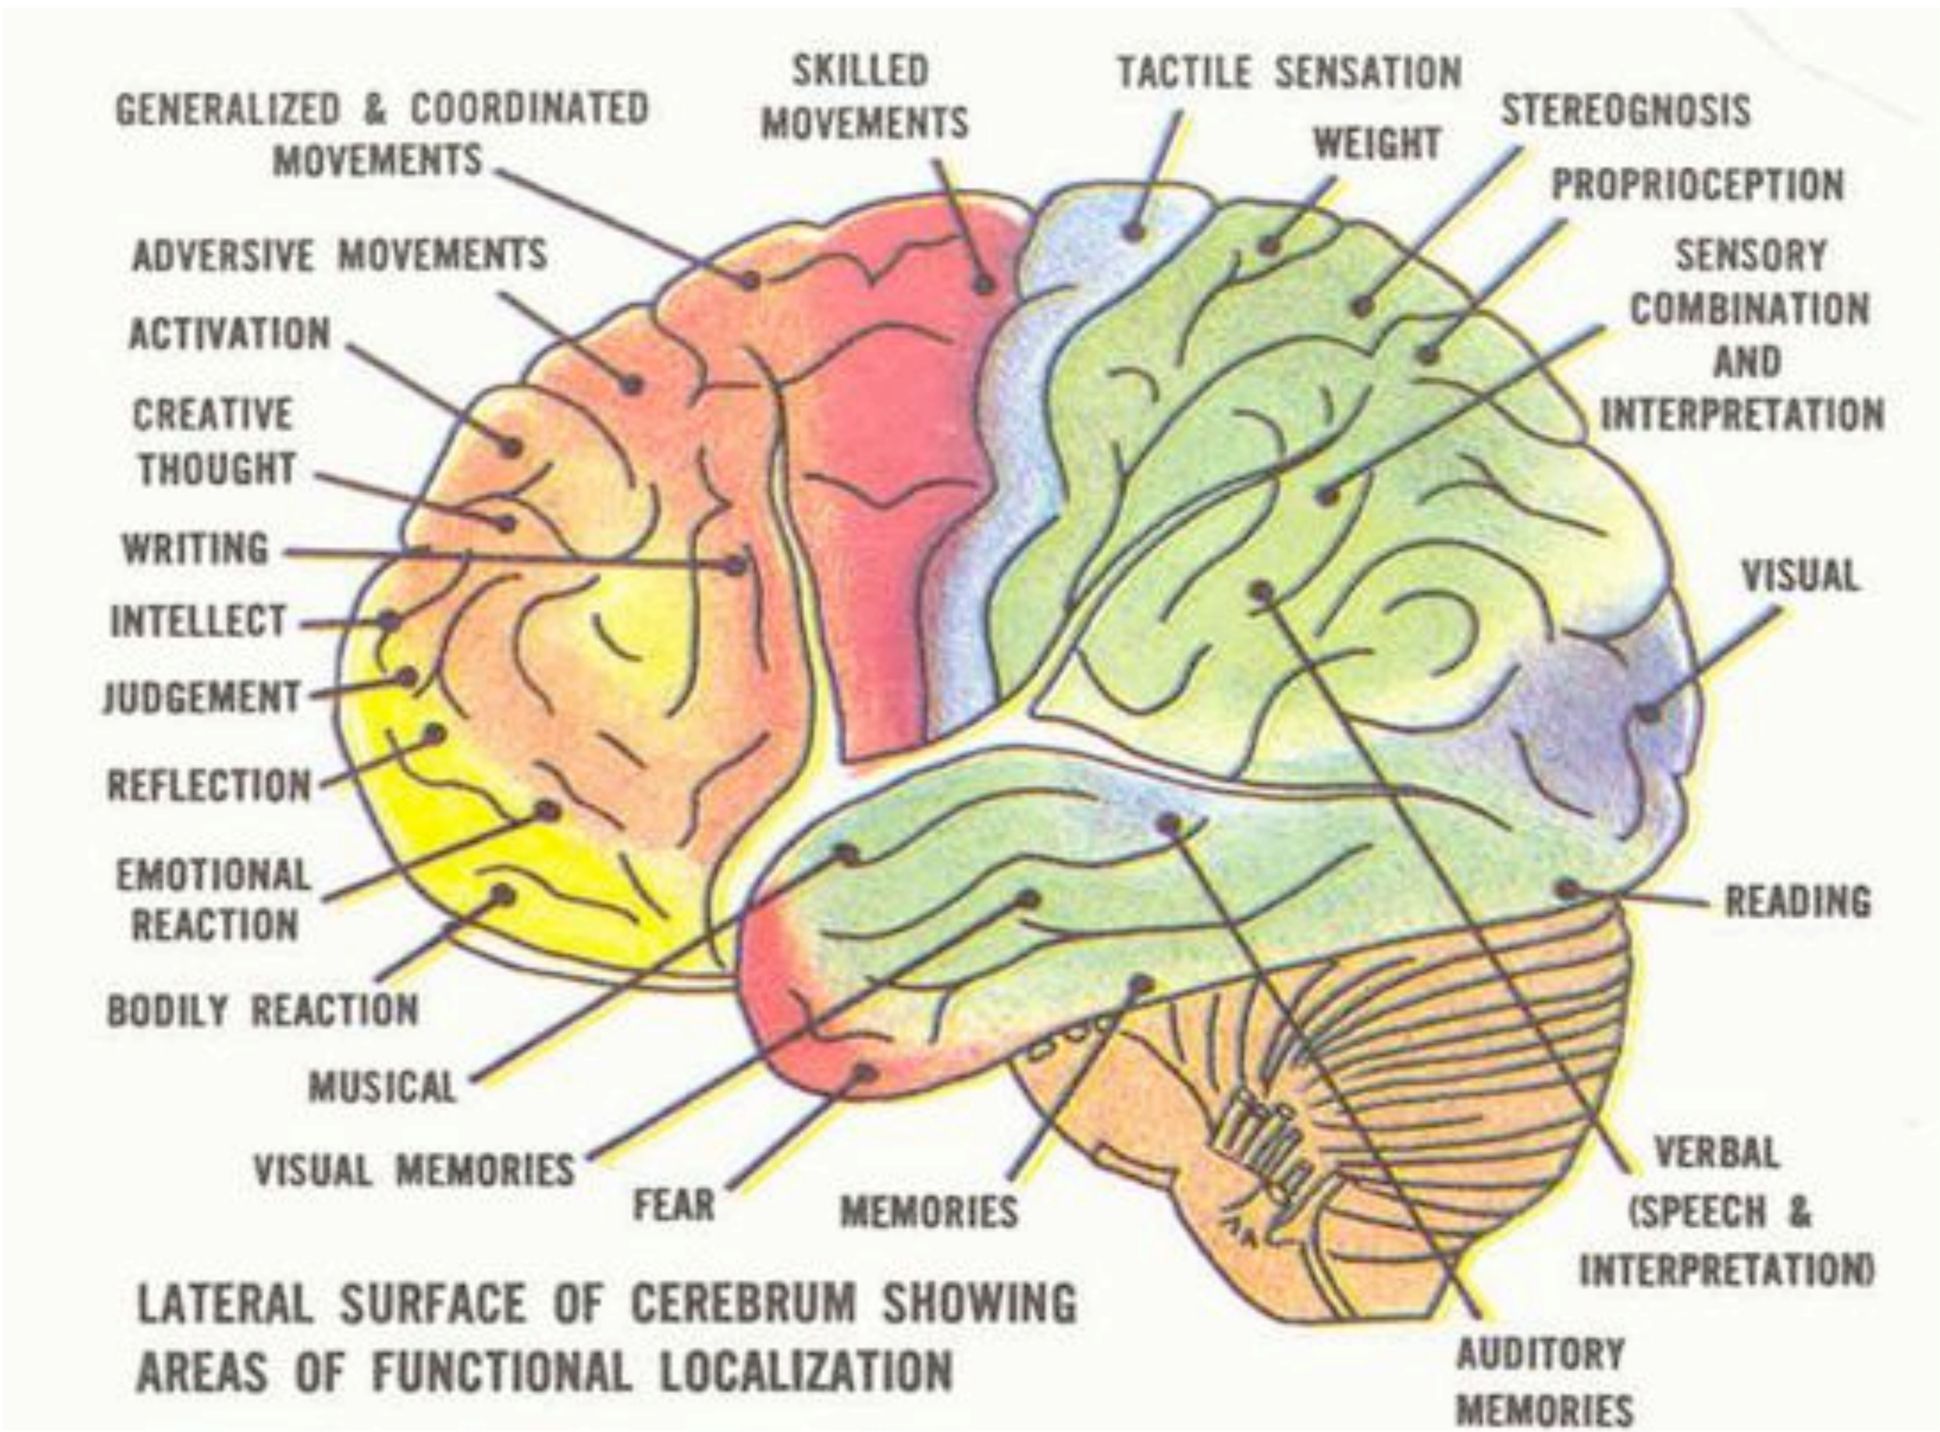

Functional areas in the cortex

(focus on graduate students)

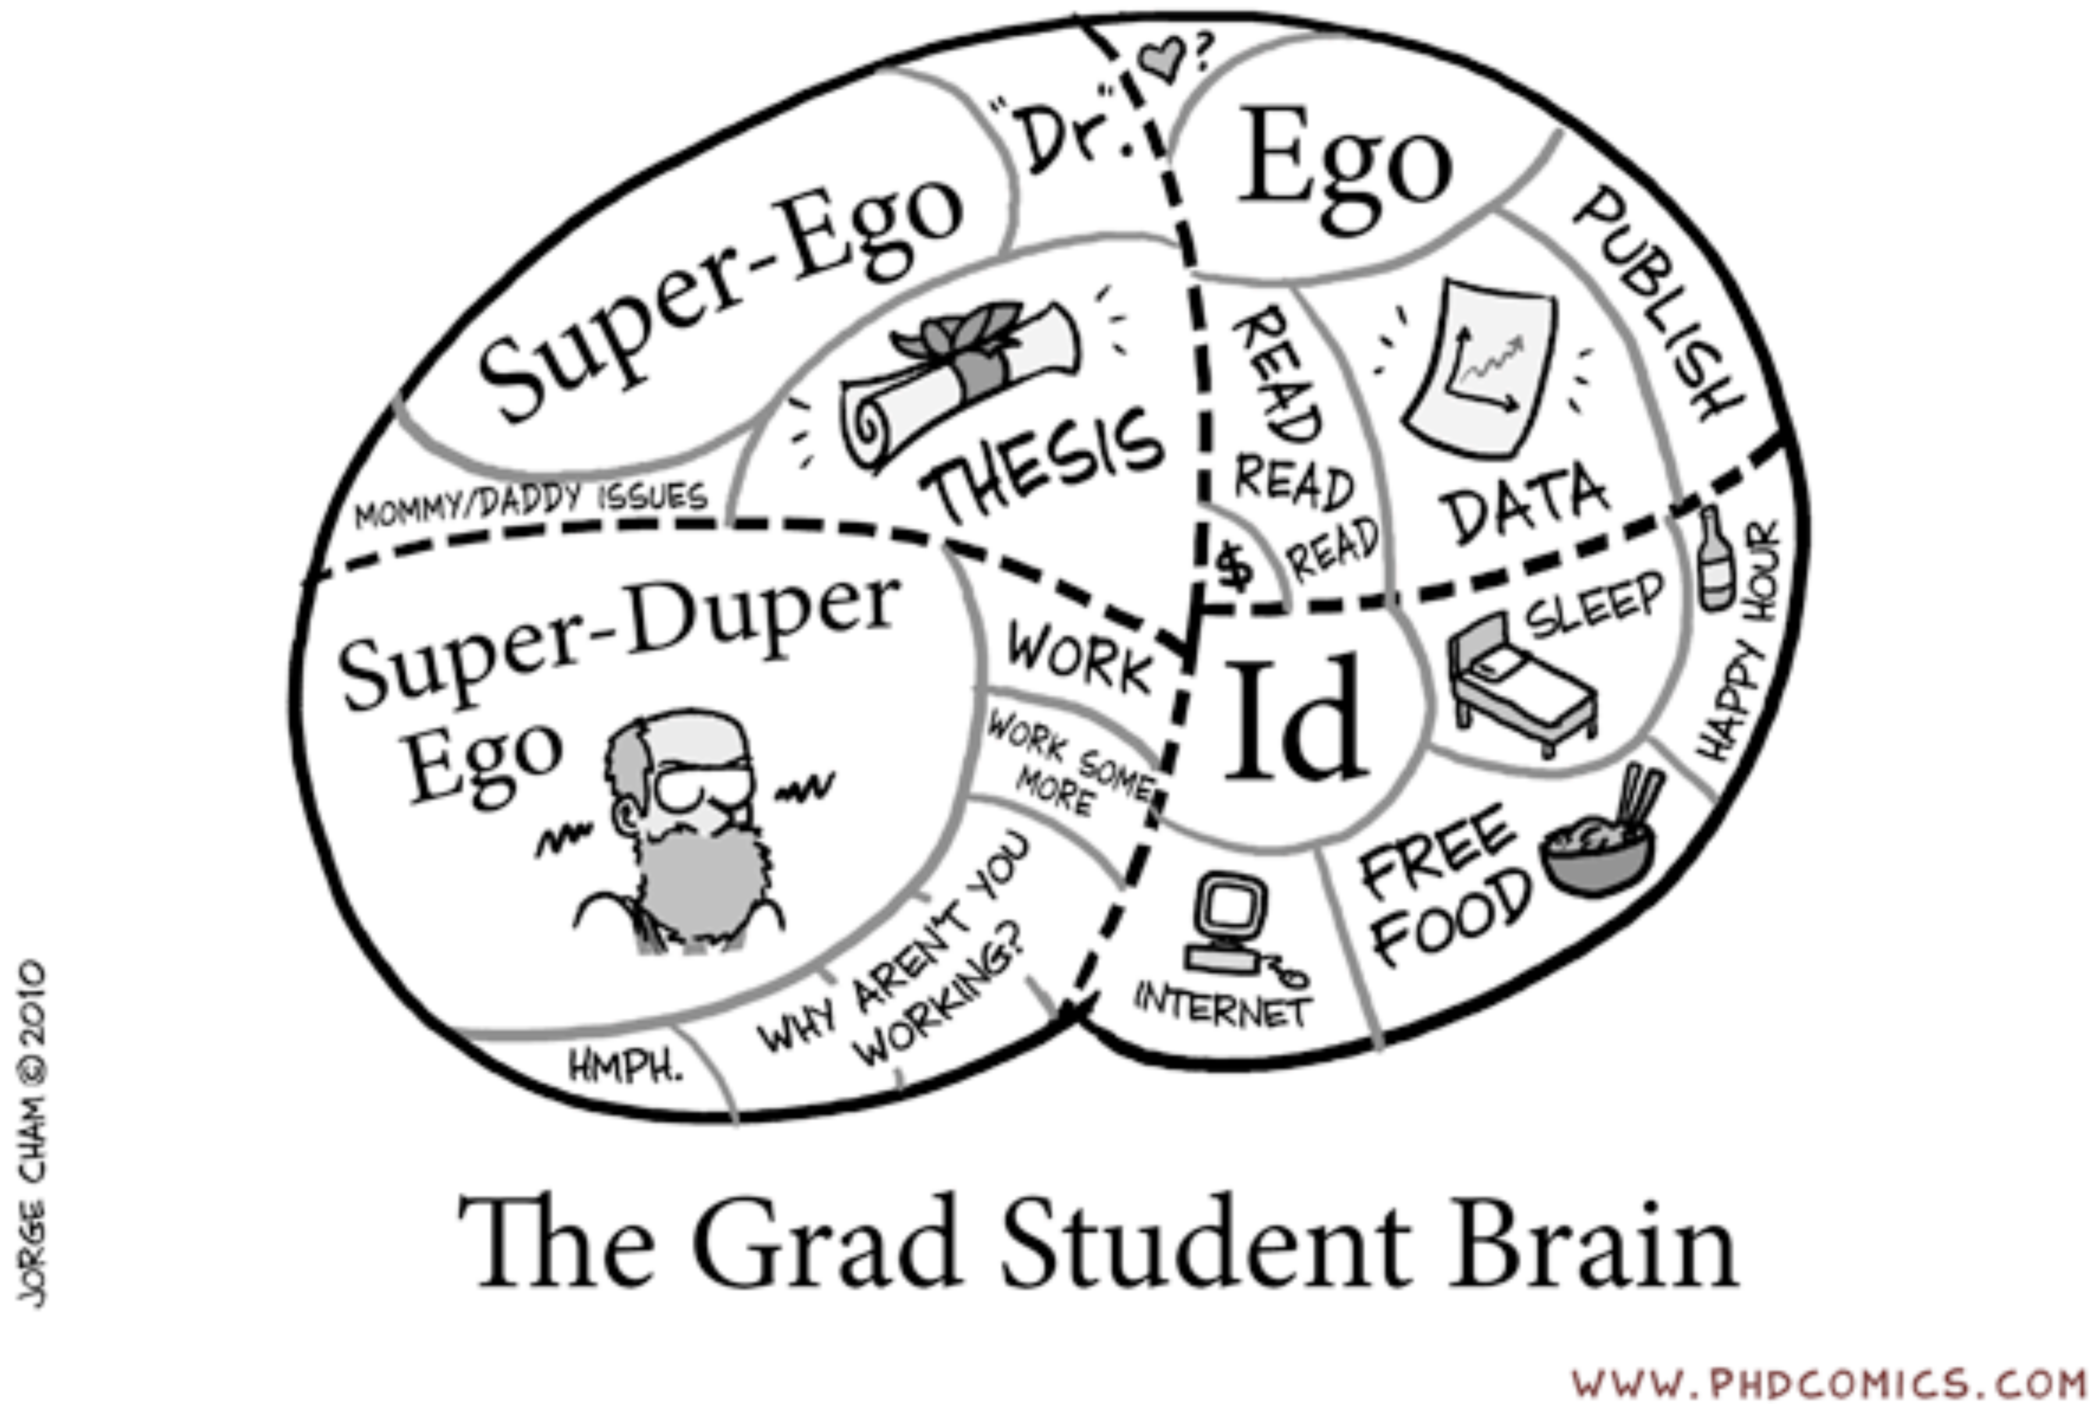

Neurons

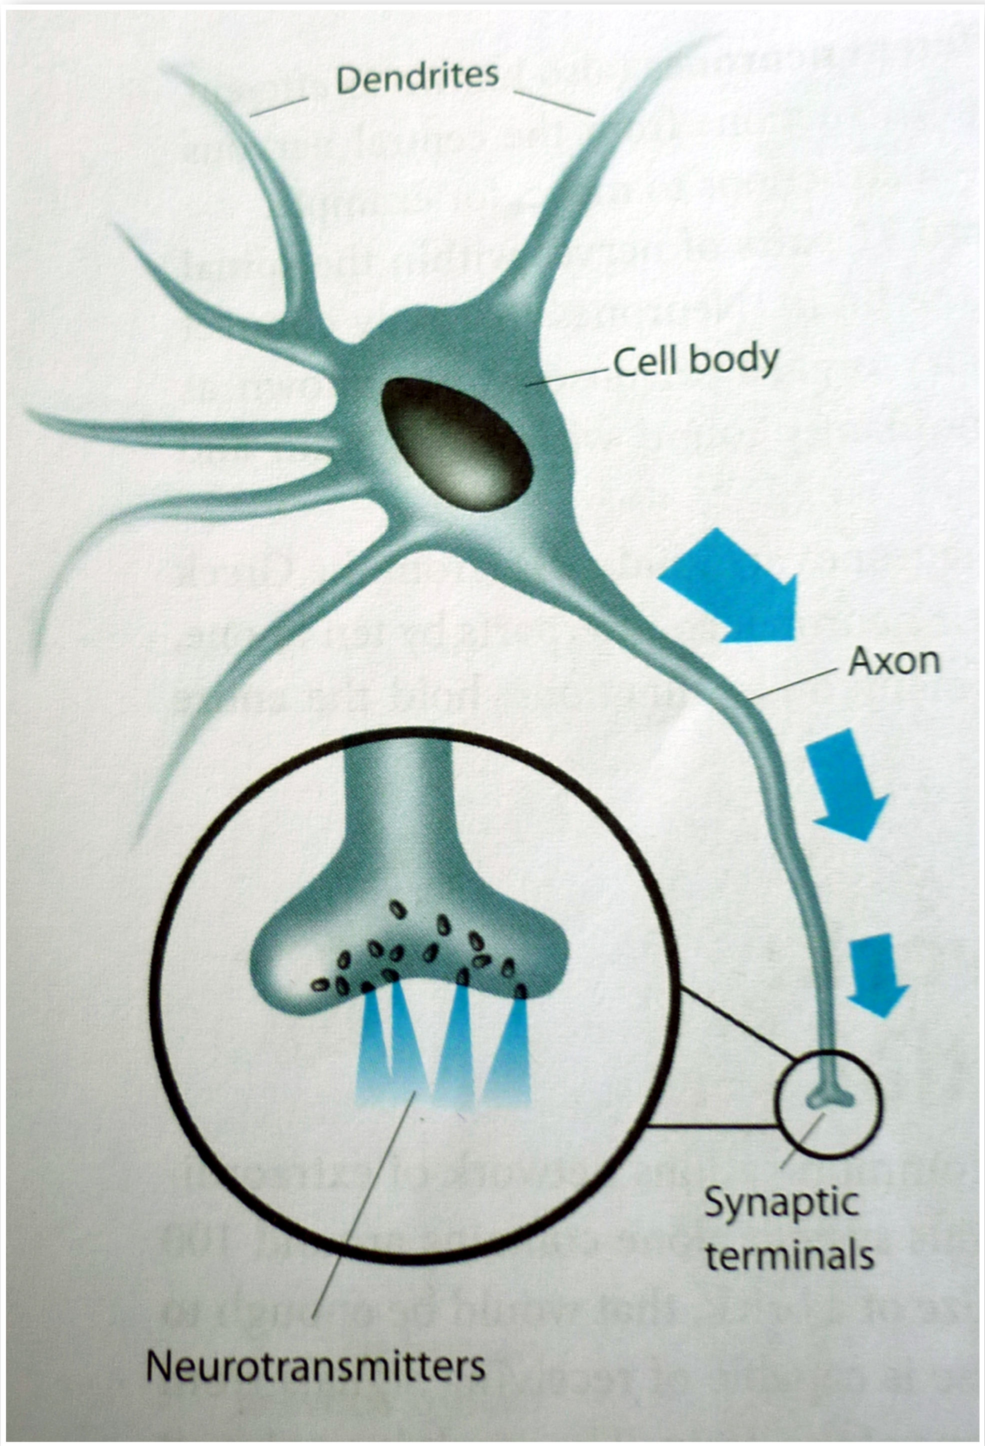

- What is a neuron?
  - Electrically excitable cell
  - made of a cell body, a set of dendrites (inputs) and an axon (output)
  - transmits information from one another by electro-chemical signalling
- The axon delivers the *electric charge* to the synaptic terminals
- Each neuron has on average 7,000 synaptic connections to other neurons
- The brain is highly *plastic*: the connections change continuously as a result of one's experience

[Gibb, B. J. (2007), *The Rough Guide to the Brain*.]

Functional properties of the brain

- Different parts of the brain perform different *functions*
    - But no closed "bundles" - fully *integrated* system
    - The same function can be recruited/reused for different purposes
- Robust and fault-tolerant
- Distributed and parallel
- Slow
- Most cognitive operations are *unconscious*
    - i.e. they are performed beneath the level of cognitive awareness
    - That's why introspection alone is insufficient to understand the mind

Probing techniques for the human brain

Brain Scanning

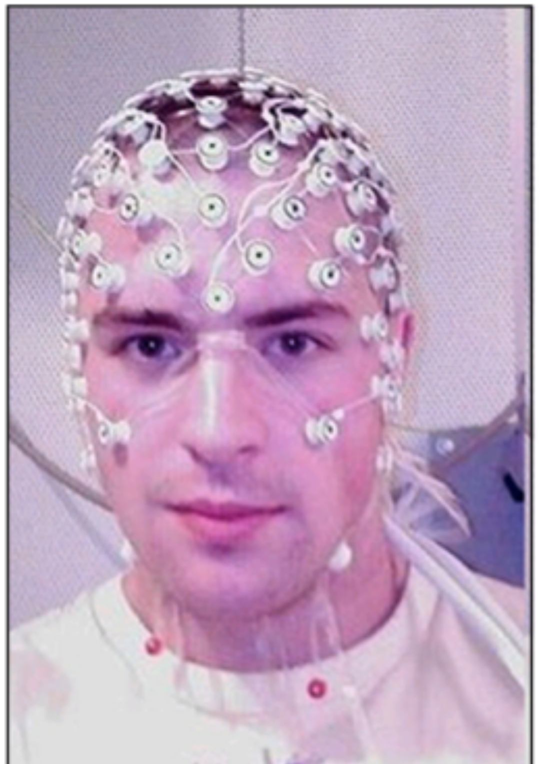

EEG (electro-encephalogram) measures electrical brain potentials from outside. Fast, but low spatial resolution

MEG (magneto-encephalogram) measures magnetic fields associated with ion flow. Fast, better resolution than EEG, expensive.

fMRI (functional Magnetic Resonance Imaging) measures nuclear magnetic resonance of protons and its relaxation times. Relates to local cerebral blood flow via BOLD (blood oxigen level dependence) effect. Slow, high spatial resolution, expensive.

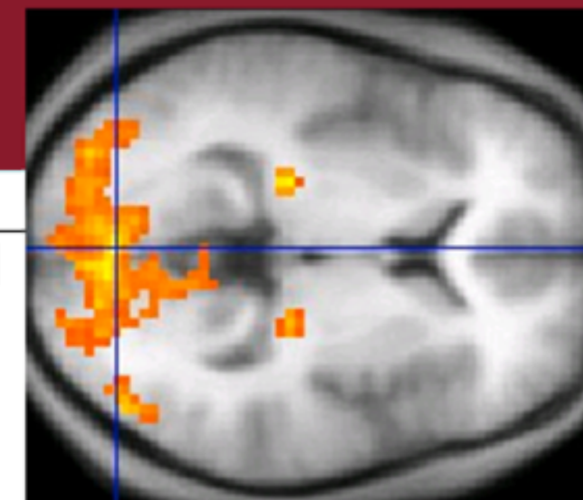

PET (Positron Emittance Tomography) measures β-decay of radioactive tracers and wash-out by local cerebral blood flow. Low spatial and temporal resolution. Expensve, requires radioactive tracers

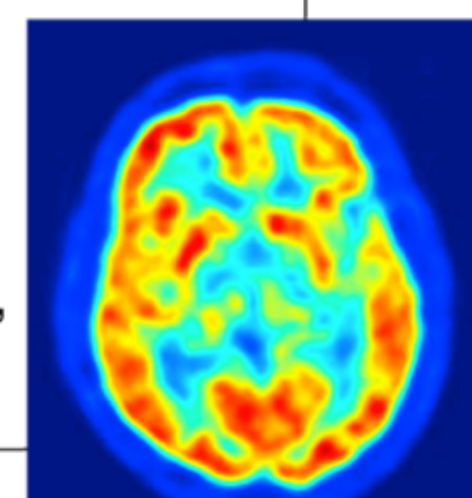

[slide borrowed from course given by Hanspeter Mallot, *Cognitive Neuroscience*, IK 2011]

Outline

- Motivation
- The brain and its functions
- **Three important ideas from cognitive science:**
  - The mind is embodied
  - The mind is proactive
  - The mind is social and emotional
- Implications for (computational) linguistics
- Conclusion

Outline

- Motivation
- The brain and its functions
- Three important ideas from cognitive science:
    - **The mind is embodied**
    - The mind is proactive
    - The mind is social and emotional
- Implications for (computational) linguistics
- Conclusion

Categorisation and embodiment

- A good place to start looking for embodiment is in **categorisation**
  - One of the most basic functions of biological organisms
- The nature of our bodies and brain ultimately *determines* what kind of categories we have and their structure
- Idea of embodied conceptual representations:
  - Concepts grounded in *modality-specific* networks
  - In other words: same neural networks are engaged in both perception/action and higher cognitive functions (categorisation, decision, language)
- Sensory-motor interaction with the environment is also known to play a key role in *child development*

[Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*]
[Piaget, J (1966), *La psychologie de l'enfant.*]

Categorisation: example of colour

- The categorisation of *colour* is an illustrative example of the role played by embodiment:
  - We all perceive colours as inherent properties of objects...
  - But there's no such thing in the physical world!
  - Our bodies and brain have evolved to *create* colour
  - Colour is an interactional property dependent on multiple factors: reflected wavelength, lightning conditions, colour cones and complex neural circuitry connected to these
  - Human perception of colour has a center-periphery structure, with certain colours being more "focal" than others
  - Contrast between different colours is a fact about our neural circuitry, not about the reflectance properties of surfaces!

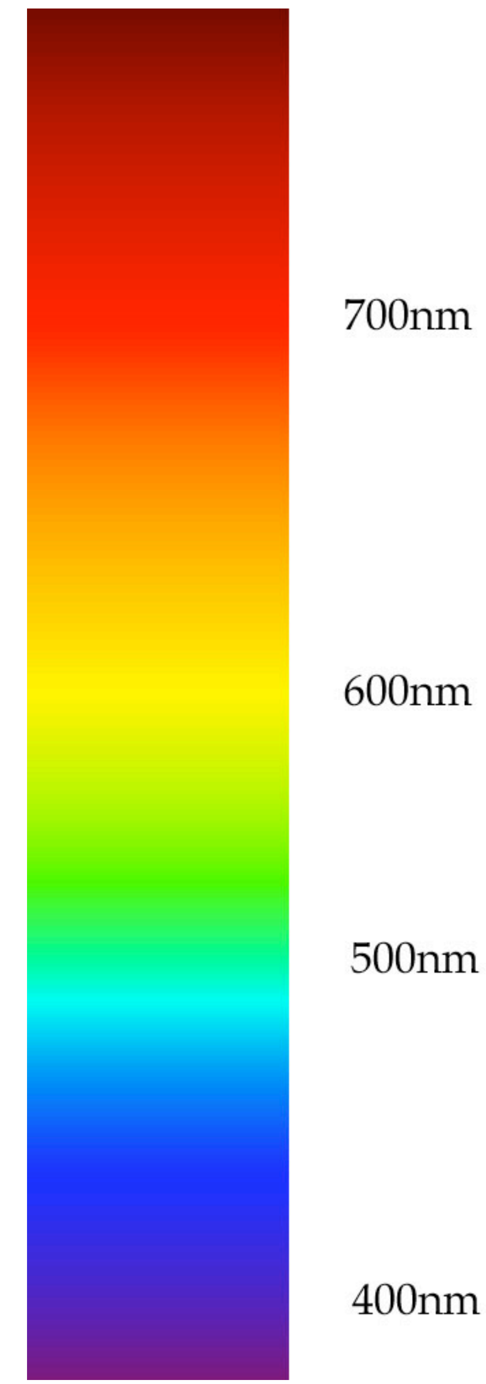

Human categorisation

- Two important findings from cognitive psychology:
  - Human categories are *graded* and conceptualised in terms of *prototypes*, and *not* using sufficient & necessary conditions
  - Some categories are "cognitively basic" and determine upper and lower hierarchies (generalisation and specialisation)
- Prototypes:
  - Category member seen as central or typical for the category
  - Different kinds of prototypes: typical cases, ideal cases, etc.
- Basic-level categories:
  - Example of "*chair*" (basic-level) compared to "*furniture*" (generalisation) and "*rocking chair*" (specialisation)
  - What distinguishes basic-level categories is their direct grounding in experience (mental image, sensory-motor affordances)

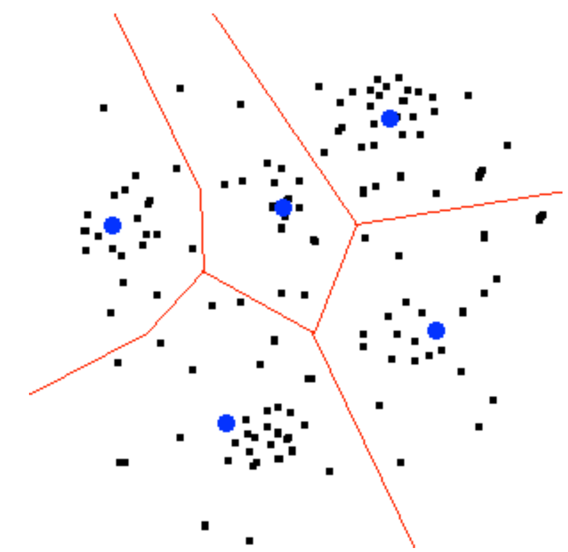

[Rosch, E.H. (1973): "Natural categories", *Cognitive Psychology*]

Conceptual metaphors

- So far, so good, but what about more abstract concepts? Are they also embodied?
    - Yes, and notably via a mechanism called *conceptual metaphor*
- Metaphors allow mental imagery from one domain (SOURCE) to be used for another domain (TARGET)
    - *Mapping* across conceptual domains
    - Source domain more concrete than target domain (typically related to subjective experience or abstract concepts)
    - Target domain inherits the *inferential structure* of the source
    - Pervasive mechanism in thought and language
- Small set of primary metaphors directly grounded in experience (spatial relations, senses, physical actions)

[Lakoff, G. and Johnson, M (1980), *Metaphors we live by*]

Complex metaphors

- Complex metaphors can be then formed by *blending* together primary metaphors
- Example of metaphor: LIFE is a JOURNEY
  - Made of two primary metaphors: PURPOSE are DESTINATIONS, and ACTION are MOTIONS
  - Mappings:

Journey	⟶	Life
Traveler	⟶	Person living a life
Destination	⟶	Life goals
Itinerary	⟶	sequence of decision, life plan

  - We can thus talk about: "a good start in life", "knowing where to go in life", "being at a crossroad", "going forward in life", "sharing the road with someone", "taking the right/wrong step", etc.

Conceptual metaphors

- Our most important concepts are structured through metaphors:
  - **Time**: as (external or internal) *motion*, or as *resource*
  - **Event structure**: *States* as *locations*, *changes* as *motion*, *causes* as *forces*, *means* as *paths*, *goals* as *destinations*
  - **Mental states**: *knowing* as *seeing*, *understanding* as *grasping*, *thinking* as *motion*
  - **Consciousness** and **morality** also have a complex set of metaphors

[Lakoff, G. and Johnson, M (1980), *Metaphors we live by*]

Conceptual metaphors (2)

- Four comments on conceptual metaphors:
  - Many concepts can be expressed through several alternative metaphors, which can be inconsistent with each other (example: time as motion or resource)
  - Interesting explanation for the widespread occurrence of certain linguistic patterns ("universals") in different cultures : humans have the same kind of bodies, they are therefore likely to structure the world in the same way
  - Metaphors relate to our very *conception* of reality! The use of metaphorical expressions in natural language is only a by-product
  - Still a lot of work to be done on the empirical side

[Lakoff, G. and Johnson, M (1999), *Philosophy in the Flesh.*]
[Gibbs, R. (2007) in *Methods in cognitive linguistics*]

Theoretical perspectives on embodiment

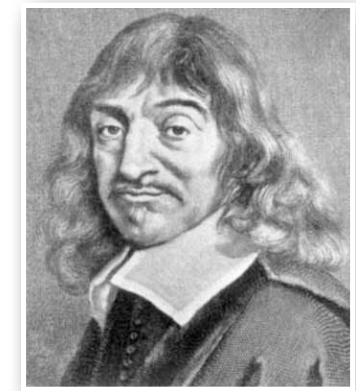

- Truth viewed as metaphorical construction, drawn from our experience of having a body, purposefully engaged in the world
- Representations as "sublimations" of bodily experience
- Contrasts with classical "cognitivist" ideas, which viewed the brain as some type of universal computer, solving all tasks with the same basic mechanisms (e.g., Turing machine).

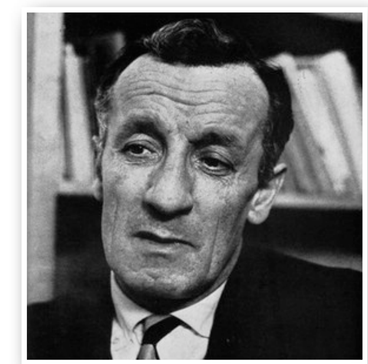

- Connections with philosophy of mind:
  - Radical shift away from mind/body dualism (Descartes)
  - *Merleau-Ponty*: our perception of reality is determined by our engagement with the world and its possibilities for bodily interaction
  - *Heidegger/Dreyfus*: human activity is not context-free manipulation of representations, but contextualised and purposeful experience of the body-environment system ("In-der-Welt-sein")

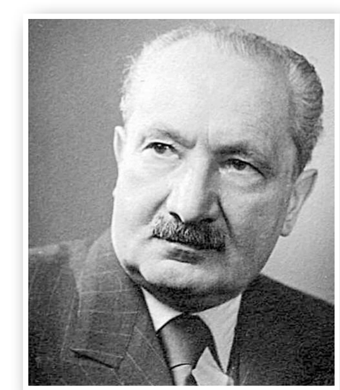

[Varela, F., Thompson, E. and Rosch E (1991), *The embodied mind: cognitive science and human experience*]

[Anderson, M. (2003), Embodied cognition: A field guide, *Artificial intelligence*]

Relation to AI and robotics

- Good old fashioned Artificial Intelligence (GOFAI) is centered on the notion of *representation*
  - A central representation (world model) describes all knowledge that is known to the agent
  - Large set of abstract rules operating on this representation
  - Sense-Model-Plan-Act (SMPA) framework
- Not scalable for real-world environments
- Modern approaches in AI and robotics focus on more contextualised and embodied approaches
  - Tight coupling between perception and action
  - "Use the world itself as its own model" [Brooks]
  - When necessary, use hybrid cognitive architectures integrating both *reactive* & *deliberative* control (sense-think-act models on top of a behaviour-based substrate)

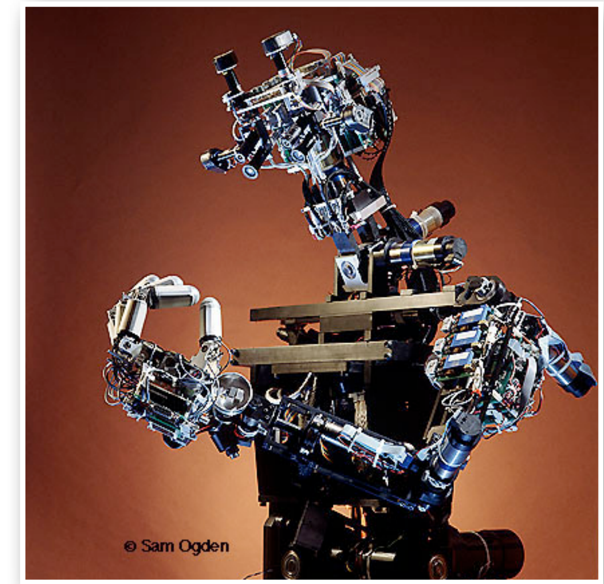

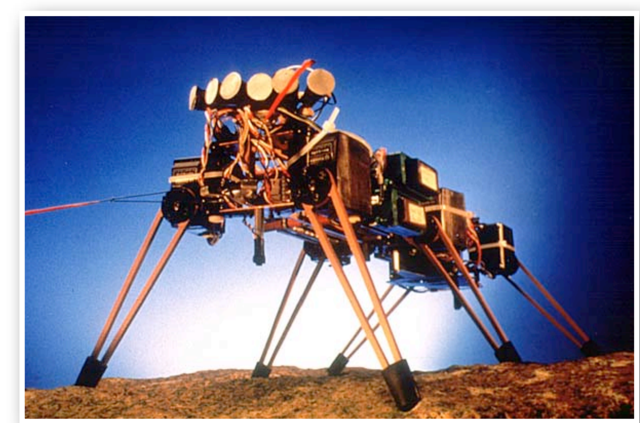

[Brooks, R. (1991), "Intelligence without representation", *Artificial Intelligence*]

Outline

- Motivation
- The brain and its functions
- Three important ideas from cognitive science:
  - The mind is embodied
  - **The mind is proactive**
  - The mind is social and emotional
- Implications for (computational) linguistics
- Conclusion

Predictive mechanisms

- The brain does not receive information passively, it continuously projects *hypotheses* and interprets things in a particular way (proactivity)
  - No strict boundary between perception and action (perception is always active, and action is always coupled with sensory feedback)
  - No closed perception-action loop: there are mechanisms which allow us to *simulate* actions without executing them
- Crucial role played by top-down **predictive** mechanisms in cognition (in motor behaviour, vision, social interaction)
- Shift from earlier cognitive theories which saw the brain as a bottom-up device without much top-down influence

[Berthoz, A. (1991), *Le sens du mouvement*]

Predictive mechanisms (2)

- Prediction is performed via *forward models* selected from context
- Predictive coding models are useful for several reasons:
  - Efficiency: raw signals too ambiguous/complex to deal with in a strictly bottom-up fashion
  - Allow us to perceive *stability* and *coherence* in our environment
- The purpose of *long-term memory* is precisely to support such predictions
  - Memory is not there to record the past, but to predict the future!

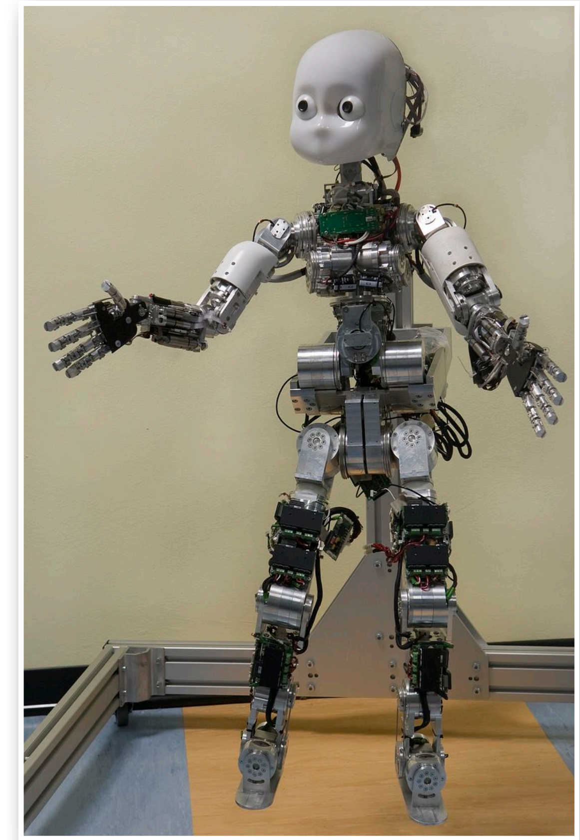

[Van Berkum, J. J. A. (2010) in *Italian Journal of Linguistics*]

Prediction in language processing

- Language comprehension is also strongly proactive
- People continuously predict what their interlocutor is going to say or talk about next, based on the current context
  - The prediction is performed incrementally and is gradually refined
- Psycholinguistics evidence for "anticipatory comprehension"
  - ERP experiments reveal that, during reading/listening, people routinely use their knowledge of the wider discourse context to *predict* upcoming words
  - N400 effects of semantic anomaly, which reflect the computational resources used in retrieving the coded meaning(s) stored in semantic long-term memory

[Van Berkum et al (2005), *Journal of experimental psychology*]

ERP experiments on word prediction

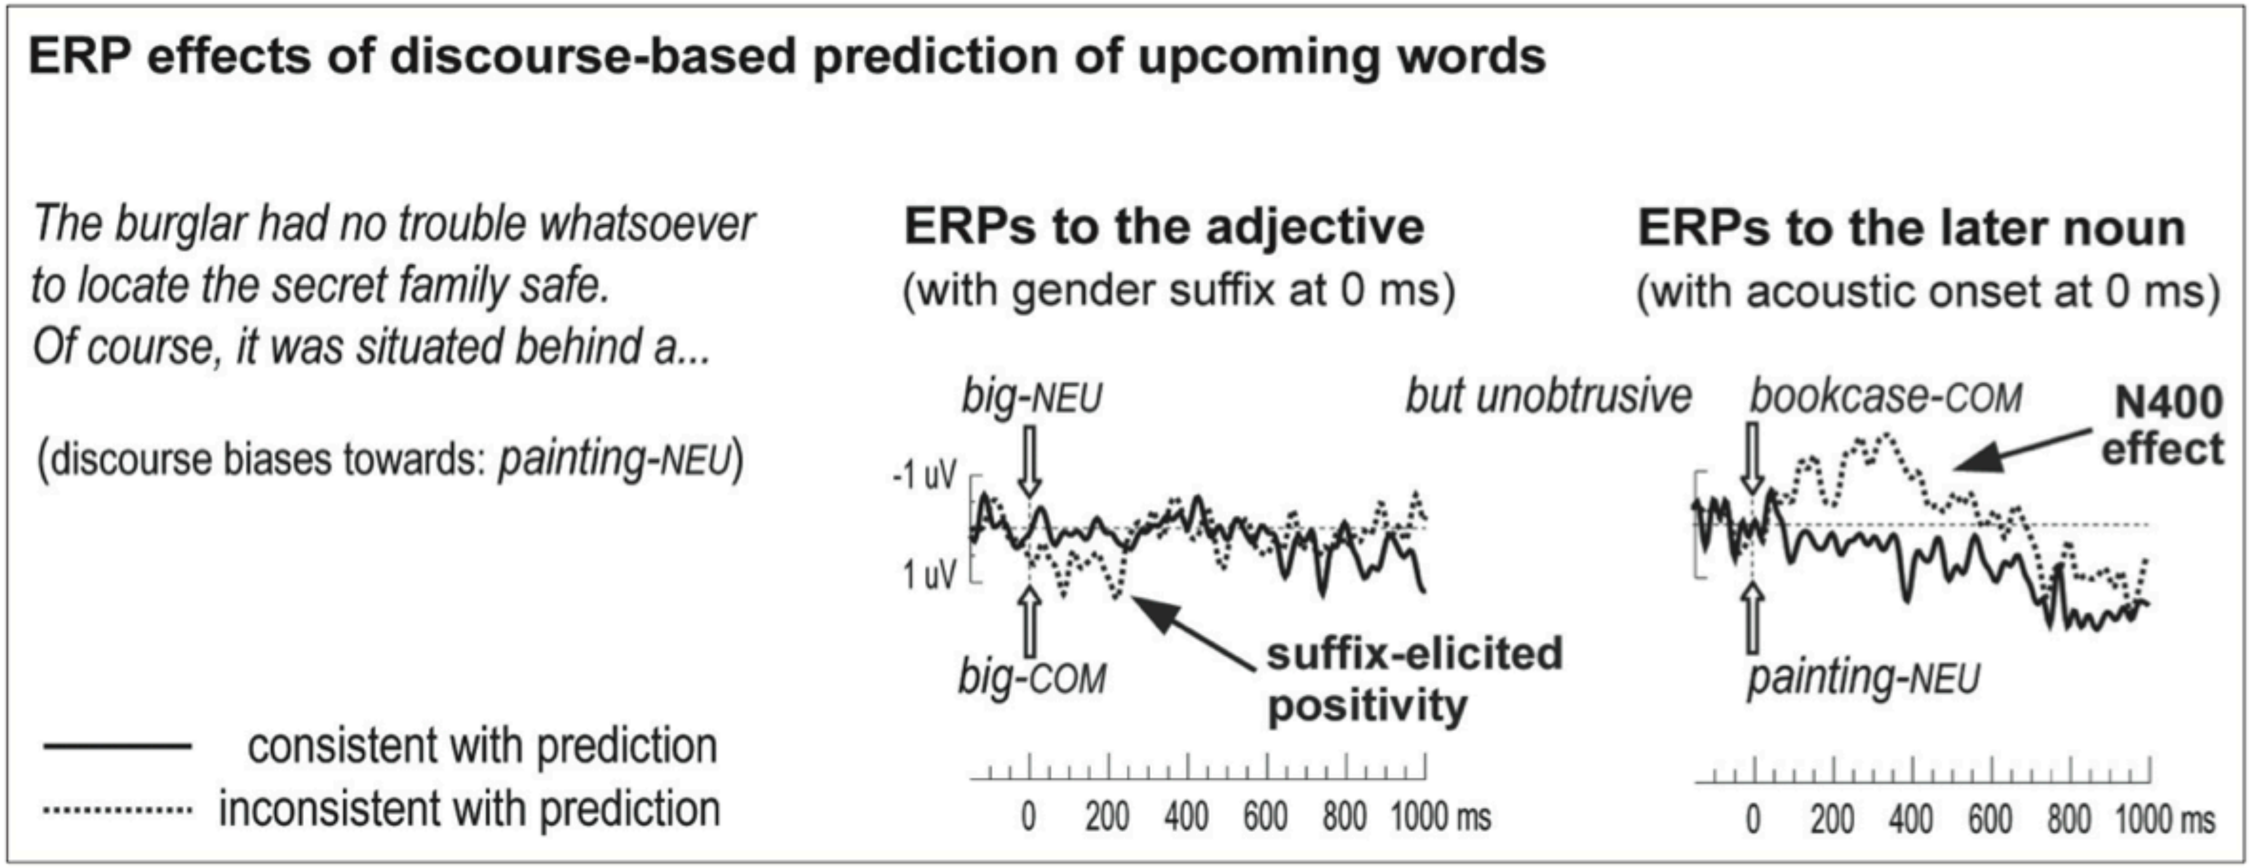

Figure 1. Left: The ERP effect to spoken adjectives whose morphosyntactic gender suffix did (solid line) or did not (dotted line) match discourse-based expectations for specific upcoming nouns (e.g., the neuter Dutch equivalent of *painting*, preceded by a prenominal adjective with common gender suffix). Right: The N400 effect elicited by the actual spoken nouns presented later in the sentence, with a coherent but less expected noun (e.g., *bookcase*, dotted line) eliciting a much larger N400 than the discourse-predictable noun (e.g., *painting*, solid line). Acoustic onset of the critical suffix (left) or later noun (right) is at 0 ms.

[Van Berkum et al (2005), *Journal of experimental psychology*]

Prediction in language processing

- Predictive mechanisms have important implications for how we view and handle linguistic interaction and pragmatics (in turn-taking, coordination patterns, pragmatic effects, etc.)
- It is important to note that the prediction-based exploitation of the broader context occurs *very rapidly* in the interpretation process, and guides all processing
- Context is a rich combination of various linguistic constraints (syntax, semantics) and general background knowledge
- Predictive forward models are crucial for efficiency:
    - Rich context is not a barrier, but the *key* to speed!

Kutas, M. and Federmeier, K. D. (2000) in *Trends in Cognitive Sciences*]
[McRae, K., & Matsuki, K. (2009) in *Language & Linguistics Compass*]
[Van Berkum, J. J. A. (2010) in *Italian Journal of Linguistics*]

Understanding as mental simulation

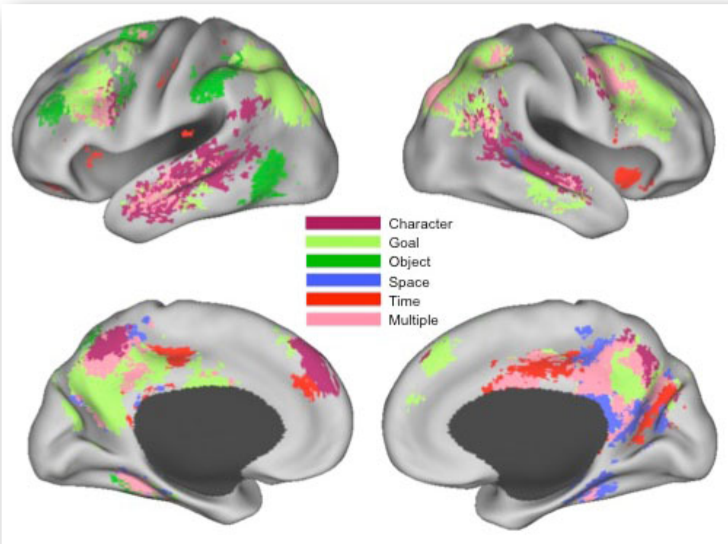

- Imagining and doing use a shared neural substrate
- Evidence from brain imaging research that imagining (and thus understanding) use many of the same neurons as actually acting or perceiving
- Idea of understanding as *mental simulation* of action or perception
  - Example: Hearing the sentence "he kicked the ball" activates the foot area of the primary motor cortex
- *Simulation-based view of meaning*: understanding a piece of language is hypothesized to entail performing mental perceptual and motor simulations of its content
  - Contrast with classical view of semantics, where language is essentially characterized as the manipulation of amodal ("disembodied") abstract symbols

[Barsalou, L.W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*]
[Bergen, B. and Wheeler, K. (2010). Grammatical aspect and mental simulation. *Brain & Language*]

Mirror neurons

- **Mirror neurons** fire both when an animal executes a goal-oriented *action*, and when the animal *observes* the same action performed by another
  - Encodes an internal *representation* of the action
  - Another example of how perception & action are intertwined
- First discovered in macaque monkeys, later evidenced in the human brain (from behavioural and brain imaging research)
- Possible cognitive functions:
  - Understanding *actions* and *intentions*
  - *Empathy*: certain brain regions are active both when someone experience an emotion, or see another experience an emotion
  - *Imitation*: observing and replicating another person's behaviour
  - *Theory of Mind*: inferring another person's mental state
- Role in language?

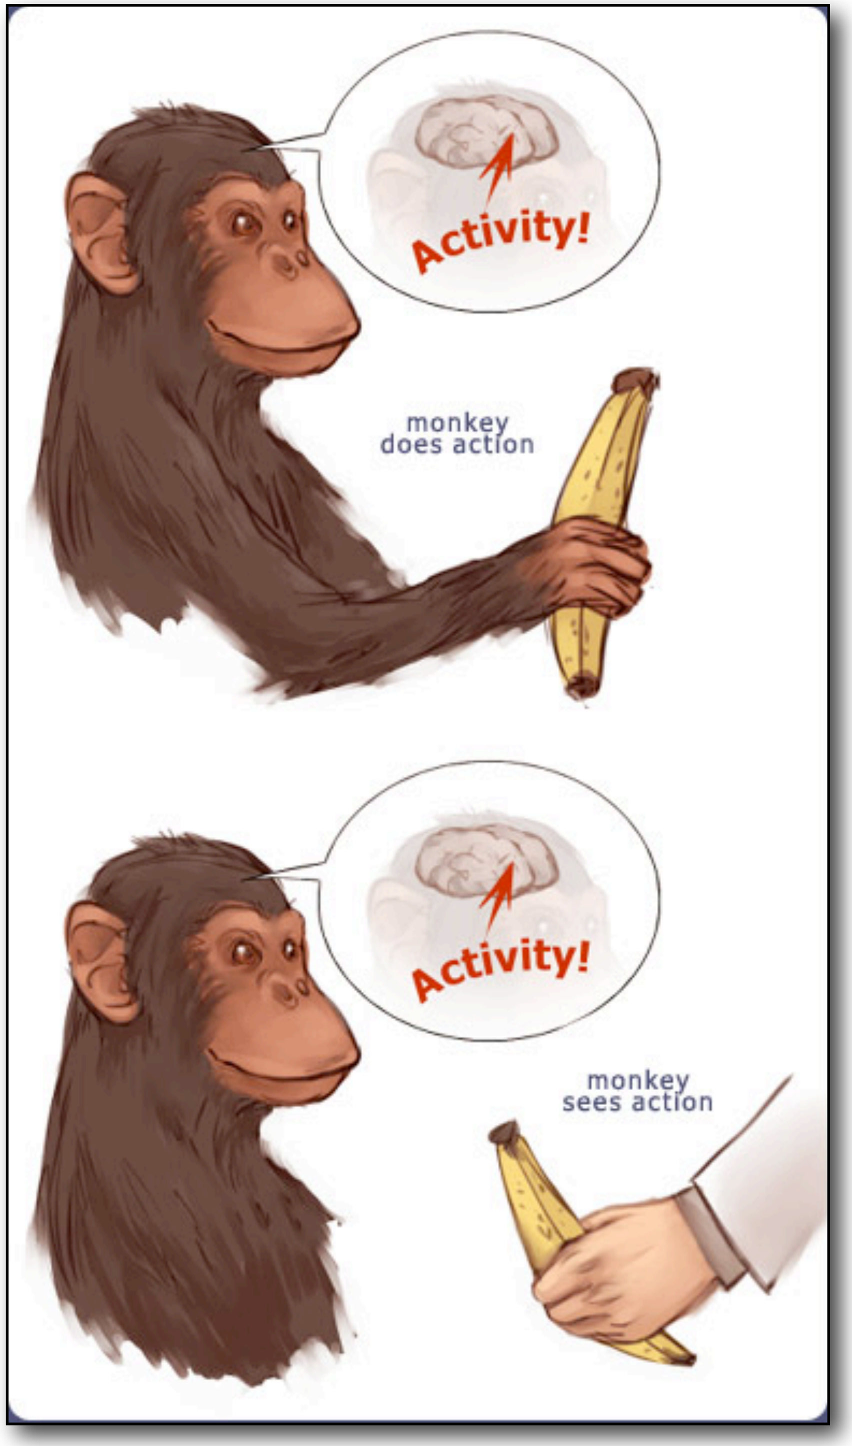

[Rizzolatti, G, Sinigaglia, C. (2008), *Mirrors in the Brain. How We Share our Actions and Emotions*]

Outline

- Motivation
- The brain and its functions
- Three important ideas from cognitive science:
  - The mind is embodied
  - The mind is proactive
  - **The mind is social and emotional**
- Implications for (computational) linguistics
- Conclusion

The social brain

- It is hard to look at the human brain without talking about the crucial role played by *social* and *emotional* processes
- The *social brain hypothesis* suggests that the primary factor for the significant growth of human brains compared to their nearest primate relatives is **social**:
  - need to navigate in the "social currents" of increasingly large groups
  - participate in *collaborative activities* with shared goals and intentions (cooperative hunting, grooming)
  - which in turn require more elaborate systems of communication and cultural practices (cf. theories on evolution of language)

[Bickerton, D. (2009). *Adam's Tongue*.]

[Dunbar, R (1998), The social brain hypothesis. *Evolutionary Anthropology*]

The social brain (2)

- Humans naturally view each other as goal-directed, intentional agents (cf. *intentional stance*) with whom they can share emotions, experience and participate in common activities
    - Construction of predictive models of others in terms of beliefs, desires and intentions (*theory of mind*)
    - Notions of *shared intentionality* and *joint attention,* giving rise to ever more complex *cultural* artefacts
- Intelligence lies less in the individual brain, and more in the dynamic *interaction* of brains with the wider world
    - Role of external scaffolds
    - *Imitation* as essential form of cultural learning and transmission

[Dennett, D (1996), *The intentional stance.*]

[Tomasello, M (1999), *The cultural origins of human cognition.*]

Dialogue as joint activity

- Dialogue seen as a joint activity percolating at different levels
  - where interlocutors continuously align their mental representations
  - both linguistic and non-linguistic processes, very tight coupling
  - Imitation and entrainment occupy central stage in verbal interaction
- Common coding across production and comprehension
  - comprehension involves *imitating* what has been heard using the production system, and using those representations to make *predictions*
  - Listeners "replay" what they hear, and work out what they would say next
  - Production would thus play the role of forward model
- Considerable empirical evidence for such phenomena

[Garrod, S., & Pickering, M. J. (2009). Joint action, interactive alignment, and dialog. *Topics in Cognitive Science*]

The affect system

- The brain is essentially a *control system* evolved to make us successful in negotiating a complex, dynamic physical and social environment
  - In order to do so, it must be able to *quickly* distinguish between what's good and what's bad for the individual
  - *Valence* or *emotions* are essential to survival
  - Special neural circuitry has evolved for this purpose, often called the emotional brain or the affect system (which notably includes the limbic system, plus other parts)
- The affect system usually reacts very quickly (and unconsciously) to external stimuli

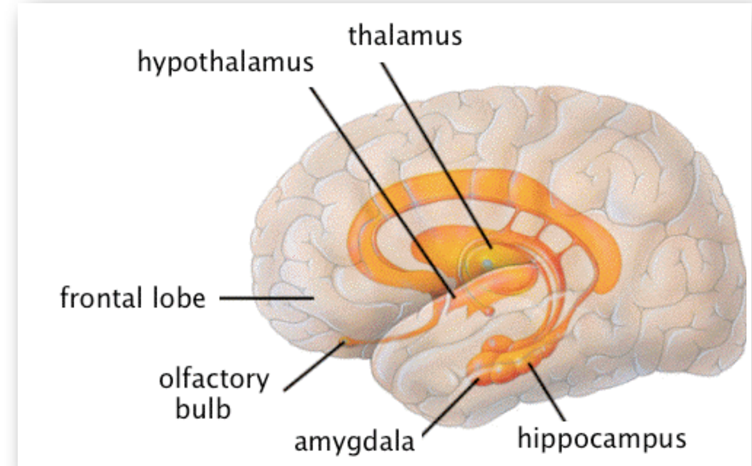

[Van Berkum, J. J. A. (2010) in *Italian Journal of Linguistics*]

Role of the affect system

- The affect system *modulates* many other cognitive processes (perception, attention, reasoning)
- The emerging idea in neuroscience is that the affect system is there to *value* information relative to our (conscious or unconscious) goals
- Far from being an impediment to knowledge and reasoning, emotions are actually the *driving factor* behind them
  - Somatic markers hypothesis (SMH): decision-making is impossible without emotions as a guide or bias
  - Reason always emotionally engaged
- Social and emotional cognition are strongly intertwined (empathy, social awareness)

[Damasio, A (1994), *Descartes' Error: Emotion, Reason, and the Human Brain*]

Valence as part of meaning

- Semantic content can also be grounded in experience related to valence/affects
- Recent ERP experiments on the processing of value-dependent meaning
  - Very interesting results on the influence of emotional responses in language processing
  - Not only is the emotional evaluation very fast, but it actually seems to *modulate* some aspect of semantic analysis itself
- Clear neurological evidence of strong interconnections between the affect system and language processing

[Van Berkum, J. J. A. (2010) in *Italian Journal of Linguistics*]

Van Berkum 2009 ERP experiment

Figure 3. ERPs to value-consistent (solid line) and value-inconsistent (dotted line) critical words in opinion poll statements, for members of a relatively strict Dutch Christian party (left), as well as a non-religious control group with opposing moral value systems. Morally objectionable words are rapidly perceived as emotionally aversive (LPP effect) and affect the ongoing semantic analysis (N400 effect); the two effects partially overlap. Written word onset is at 0 ms.

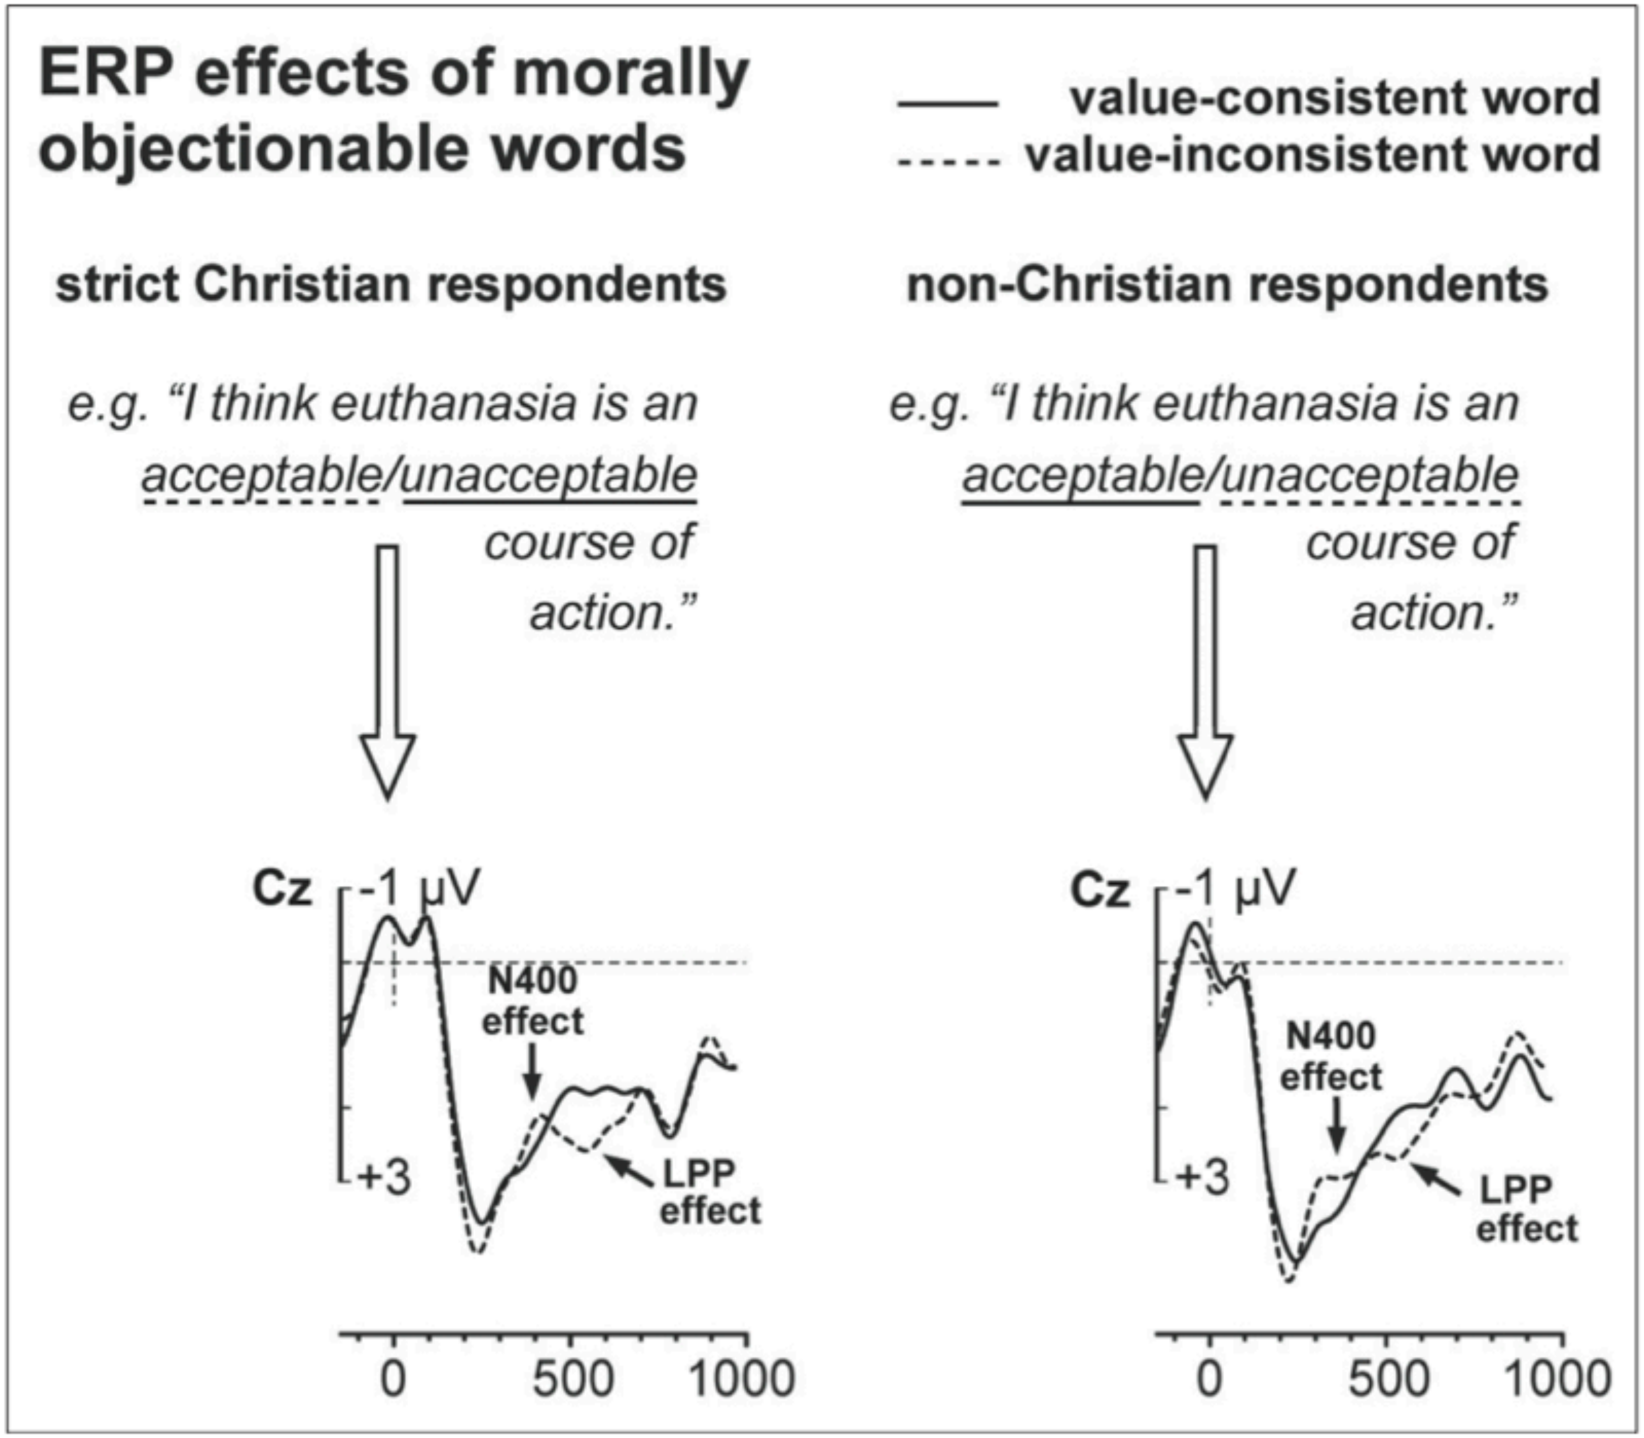

[Van Berkum, J. J. A. et al (2009). *Psychological Science*]

Outline

- Motivation
- The brain and its functions
- Three important ideas from cognitive science:
  - The mind is embodied
  - The mind is proactive
  - The mind is social and emotional
- **Implications for (computational) linguistics**
- Conclusion

Implications for NLP and LT applications

Photo by:Kiyoshi Takahase Segundo

- We can draw some inspiration from what we've just seen about about human cognition
- I'll group them in three types:
  - Implications about cognitively-informed software architectures
  - Implications about meaning representations
  - Implications about the design of interactive systems

Implications for software architectures

- We have seen that human cognition (and human language processing in particular) exhibits the following functional properties:
  - It combines a reactive substrate complemented with more deliberative control
  - It is both bottom-up and top-down (predictive mechanisms)
  - Representations are not a goal per se, they are only built as needed

Implications for software architectures

- We have seen that human cognition (and human language processing in particular) exhibits the following functional properties:
  - It combines a reactive substrate complemented with more deliberative control
  - It is both bottom-up and top-down (predictive mechanisms)
  - Representations are not a goal per se, they are only built as needed

→ **Insights** for natural language processing algorithms:

- NLP algorithms should be able to handle linguistic inputs with *varying levels of granularity*, from very shallow ("reactive") to relatively deep analysis. Balance between *discriminative* (model-free) and *generative* (model-based) approaches
- Good NLP algorithms should only construct representations *if necessary*
- Exploitation of contextual knowledge to *guide processing* at all levels of analysis
- Instead of generating all alternatives, it is more efficient to *focus* on the most likely analyses, and *backtrack* later if needed

Implications for semantics

- Cognitive semanticists often emphasise that language itself does not encode or "represent" meaning. Rather, linguistic units serve as a *prompt* for the dynamic *construction* of a particular meaning in a given context
  - Example from G. Fauconnier's work on conceptual blending:
    - (a) The child is safe
    - (b) The beach is safe
    - (c) The shovel is safe
  - Moreover, this dynamic construction of meaning often relies on *metaphors*

[Fauconnier, G. and Turner, M (2003). *The way we think.*]

Implications for semantics

- Cognitive semanticists often emphasise that language itself does not encode or "represent" meaning. Rather, linguistic units serve as a *prompt* for the dynamic *construction* of a particular meaning in a given context
  - Example from G. Fauconnier's work on conceptual blending:
    - (a) The child is safe
    - (b) The beach is safe
    - (c) The shovel is safe
  - Moreover, this dynamic construction of meaning often relies on *metaphors*

→ **Insights** for computational semantics:

- Abandon the GOFAI idea of representing all common-sense knowledge in terms of abstract symbols and rules (see e.g. the Cyc project)
- Rather, *tailor* and *ground* semantic representations in the specific context
- Use metaphorical constructions to perform *inference*

[Fauconnier, G. and Turner, M (2003). *The way we think.*]

Implications for interactive systems

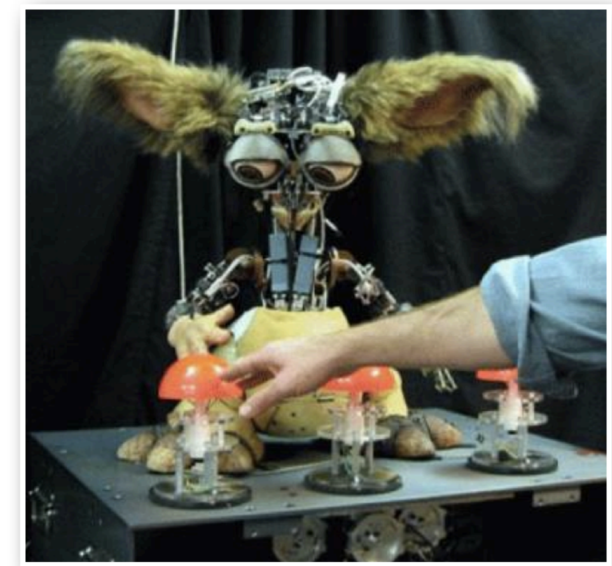

- We have seen that dialogue is a form of *collaborative activity*, where participants continuously align their representations
- We have also seen the key cognitive role played by *emotions*

Implications for interactive systems

- We have seen that dialogue is a form of *collaborative activity*, where participants continuously align their representations
- We have also seen the key cognitive role played by *emotions*

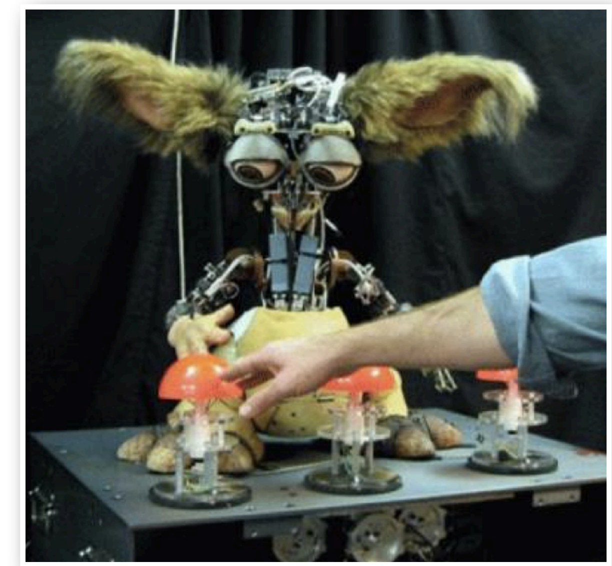

→ **Insights** for interactive systems

  - *Align* with your interlocutor at all linguistic levels
  - Process your inputs *incrementally*, allow for reactive *feedback*
  - Build a representation of the conversation's *common ground*
  - Try to achieve maximum *transparency* regarding your system capacities and current state of understanding
  - Attend to (and acknowledge) *emotional cues* in the inputs: prosody, emotion words, affective display
  - ... and *react* to them appropriately in generation (via e.g. emotional speech synthesis)

Outline

- Motivation
- The brain and its functions
- Three important ideas from cognitive science:
  - The mind is embodied
  - The mind is proactive
  - The mind is social and emotional
- Implications for (computational) linguistics
- **Conclusion**

Conclusion

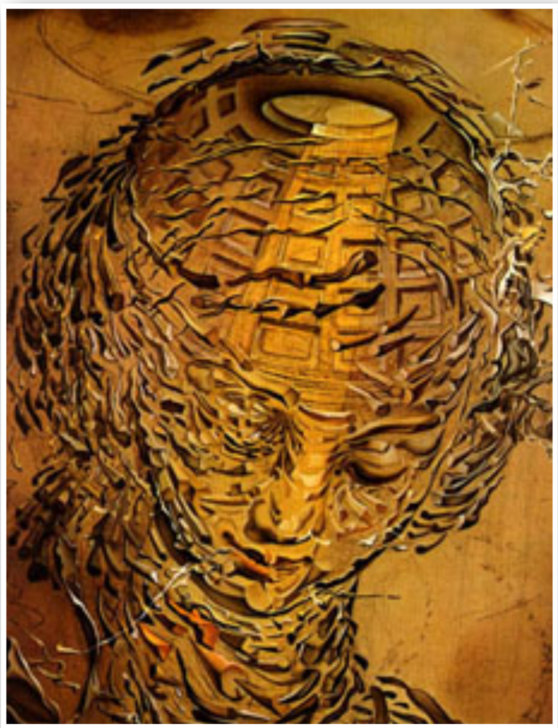

- We have reviewed a variety of functional properties about human cognition
  - the human mind is inherently embodied; linguistic concepts are metaphorically structured
  - the mind is fundamentally proactive - it continuously build predictions and performs mental simulations
  - Social and emotional processes play a major role and are driving factors for generic reasoning
- Relevance for computational linguistics?
  - Cognitive architectures of language processing
  - Grounded models of meaning
  - Interactive systems designed for human cooperation