Connectionist
Language Processing

Lecture 1: **Introduction**

Matthew W. Crocker
crocker@coli.uni-sb.de
Harm Brouwer
brouwer@coli.uni-sb.de

The Course

Can we construct computational models of language inspired by what we know about how computation takes place in the brain?

- Lectures: Tuesday 14-16
- Tutorials: Thursday 14-16
- Connectionist models of human language
  - McLeod et al. (1998). *Introduction of connectionist modelling of cognitive processes*. UK:OUP.
  - Plunkett and Elman (1997). Exercises in rethinking innateness: A Handbook for Connectionist Simulations. MIT Press. Chapters: 1-8, 11, 12.
  - Software: MESH (Mac/Linux), also on coli servers.

Recommended background

- General facility for math and statistics is helpful, but specifics will be covered in the lectures:
  - Basic math & statistics
  - Linear algerbra (vectors and matrixes)
  - Differentiation (Analysis) – not essential, but useful
- No programming required, but …
  - Simulations will be done using the MESH simulator
  - R may be be used for some graphical tasks
  - Ability to use the command line ;-)

Cognitive models of language

- Goal: model and understand human language processing and development
  - What mechanisms recover linguistic representations from the linguistic signal
  - How do these mechanism emerge/learn over time
  - Pathologies: model “errors” or “weaknesses” of human performance
- Language learning: Nature (innate) *versus* Nurture (experience)
- Cognitive Plausibility
  - Psychologically & *neurobiologically* plausible learning & processing
  - Plausible learning environments (what are children exposed to)

The “Traditional” Perspective

- Modern cognitive modelling has been heavily influenced by available theories of computation.
  - Computation: digital computers, logic based
  - Language: Chomskian theory (+criticism of statistical approaches)
  - AI & problem solving: Newell and Simon
  - Result: digital, symbolic, rule/logic-based accounts of cognition
  - Emphasis on explicit, symbolic rules, representations, processes
- Reconsider our models, based on what we know about the brain?

The Brain

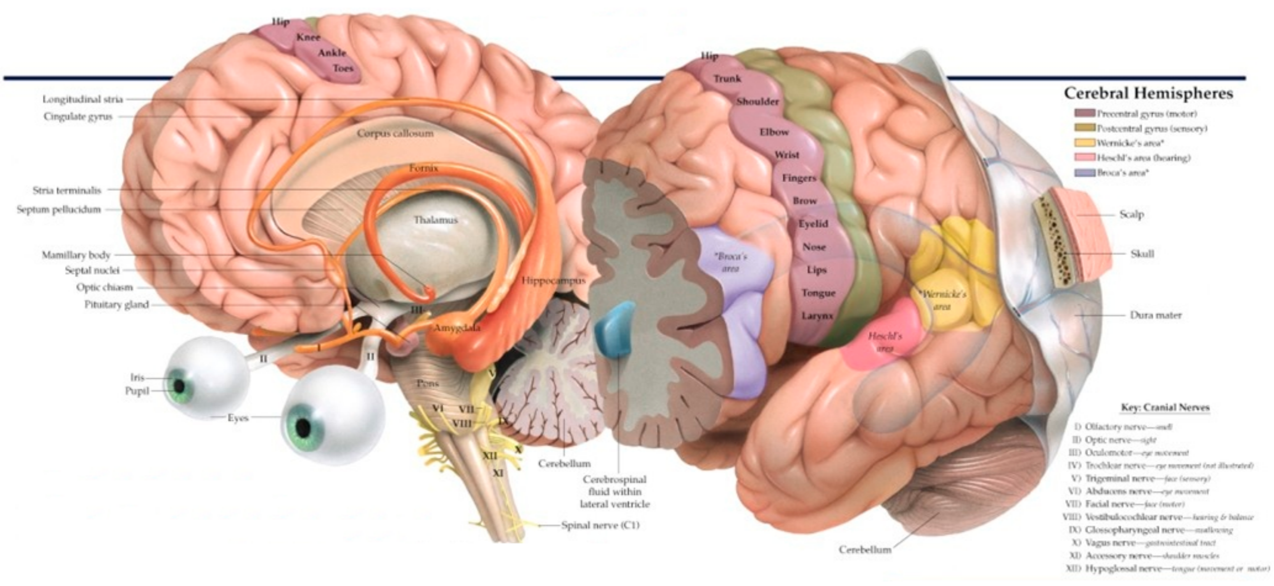

The Cortex

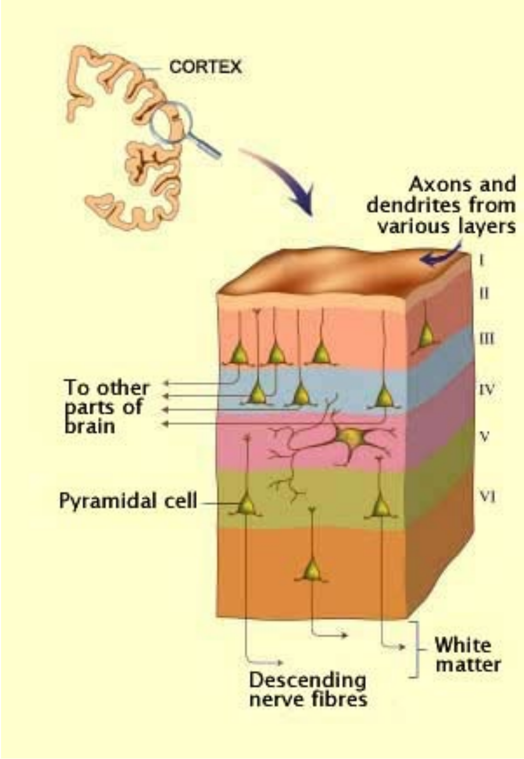

Neurons

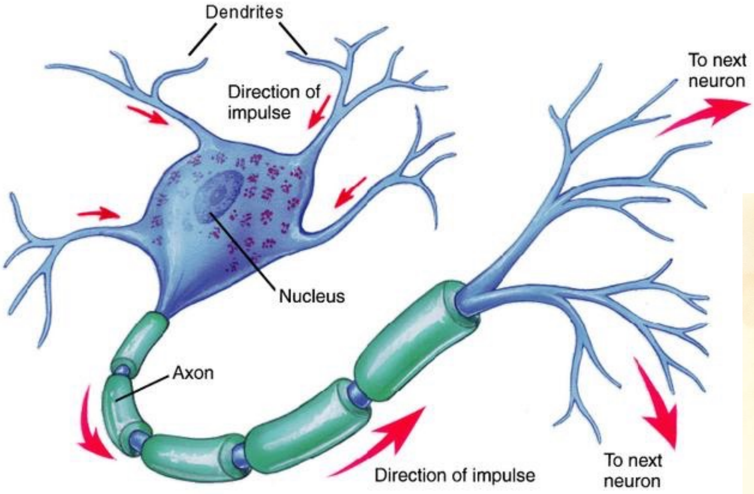

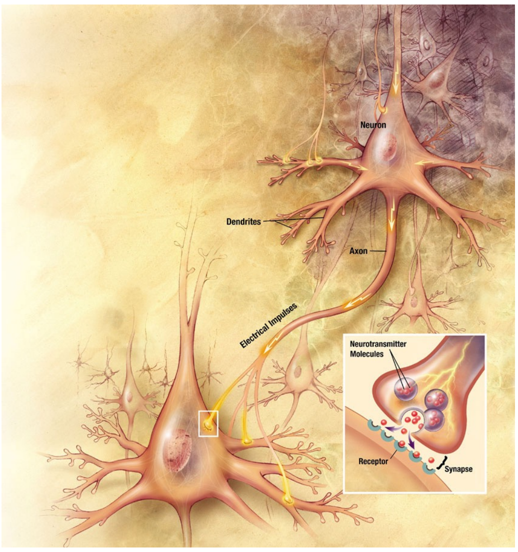

(Artificial) Neurons

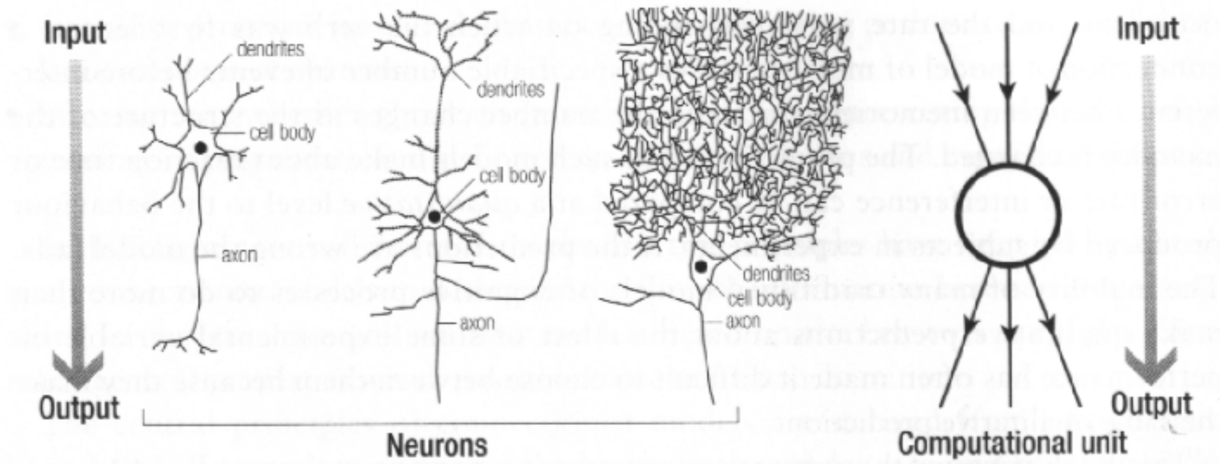

- Neurons receive signals (excitatory or inhibitory) from other neurons via synaptic connections to its dendrites.
- If the sum of these signals exceeds a certain threshold, then the neuron fires, sending a signal along its axon.

Brain versus Network

- The human brain contains approximately 10^{11} neurons
- Those neurons are densely interconnected:
  - 10^5 connections per neuron
  - Thus, 10^{15} - 10^{16} connections in total
- Connections can be both excitatory and inhibitory
- Learning involves modifying of synapses (connections)
- Connections can be both added and eliminated (pruning)

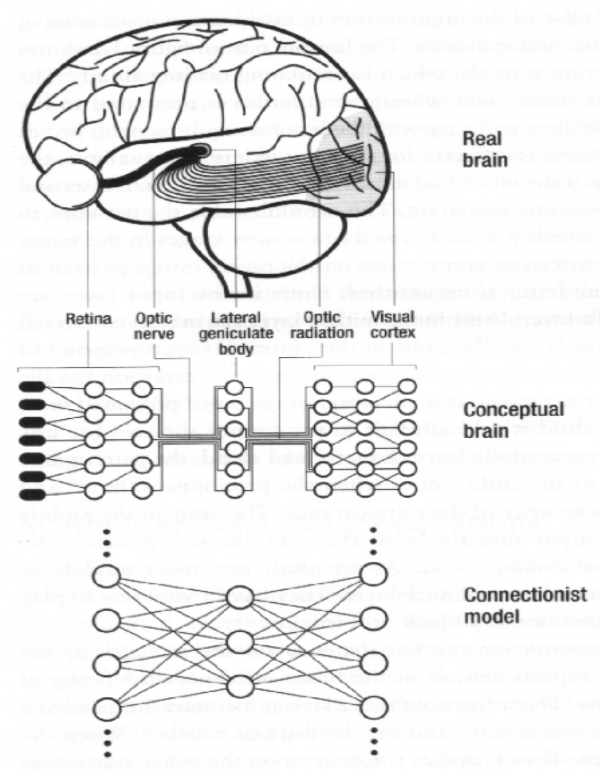

Connectionist Information Processing

Connectionist models of information processing can become complex, but the idea is based on simple neuronal processing in the brain:

- **Neurons integrate information**: All neuron types sum inputs and compute an output
- **Neurons pass information about the strength of their input**: Output encodes information about the degree of input: firing rate
- **Brain structure is layered**: Information passes through sequences of independent structures
- **Influence of one neuron upon another depends on connection strength**: A given neuron is connected to thousands of other neurons, but its influence on a particular node is determined by synaptic strength
- **Learning is accomplished through changing connection strengths**: There is evidence that this is so, but many connectionist learning rules are not biologically plausible.

Terms: connectionism, parallel distributed processing, neural networks, neurocomputing

The "Connectionist" Perspective

*"... implicit **knowledge of language may be stored among simple processing units organized into networks**. While the behaviour of such networks may be describable (at least approximately) as conforming to some system of rules, we suggest that an account of **the fine structure of the phenomena of language and language acquisition can be best formulated in models that make reference to the characteristics of the underlying networks**"*

(Rumelhart and McClelland, p. 196, 1987)

- Neurologically based (but not true models of the brain)
- Distributed, implicit representations
- Dense connectivity
- Communication of "real values" not "symbols"
- Representations and processing are the same
- Learning

Basic Structure of Nodes

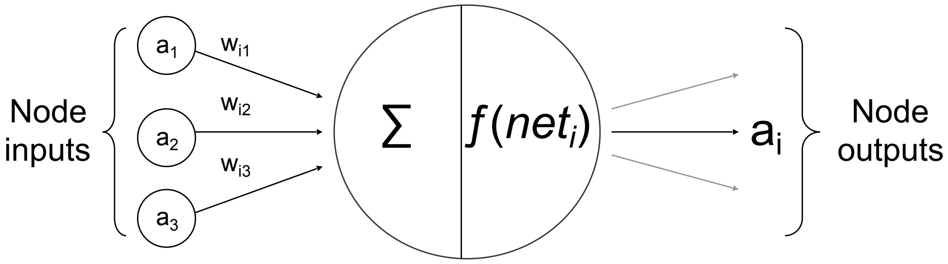

- A node can be characterised as follows:
  - Input connections representing the flow of activation from other nodes or some external source
  - Each input connection has its own weight, which determines how much influence that input has on the node
  - A node i has an output activation ai = f(neti) which is a function of the weighted sum of its input activations, net.
- The net input is determined as follows: $net_i = \sum_j w_{ij} a_j$

An example

- A one-layer feed-forward network:

$$net_i = \sum_j w_{ij} a_j$$

- So the net input for a_2 is:

$$\text{net input } a_2 = w_{20} \cdot a_0 + w_{21} \cdot a_1$$

- Consider a network with the following inputs and weights:
- The net input for node a_2 is:
  - 1 x .5 + 1 x .25 = 0.75

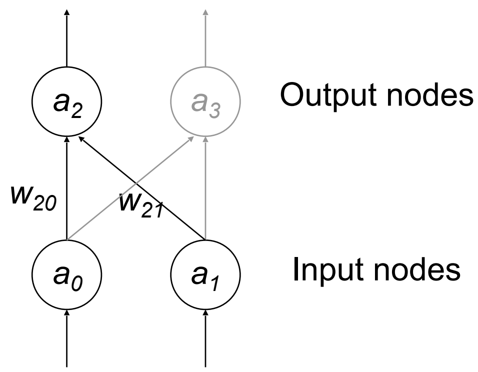

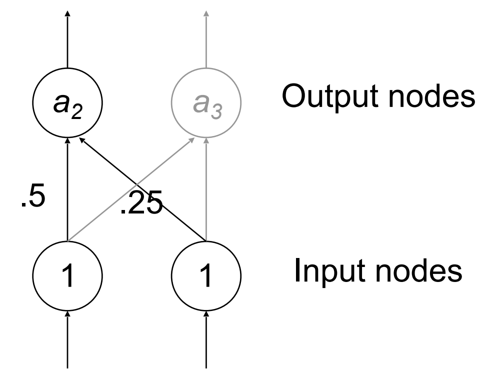

About weights

- Node *j* influences node *i* by passing information about its activity level.
- The degree of influence it has is determined by the weight connecting node *j* to node *i*.
  - A smaller weight corresponds to reduced influence of one node on another
  - A larger weight emphasises the influence of the node's activation
- Weights can be either positive or negative
  - Positive weights contribute activation to the net input
  - Negative weights lead to a reduction of the net input activation
  - Brain: excitatory versus inhibitory connections

Activation functions

- The activation function determines the activation a_i for node *i* from the net input, net_i, to the node: f(net_i)
- Linear activation function:

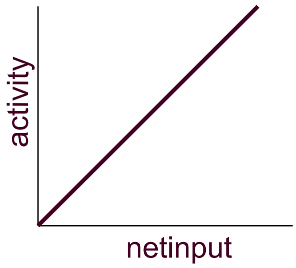

  - (McCulloch-Pitts neurode, perceptron)
  - Identity: the $a_i = net_i$

$$f(net_i) = net_i$$
$$f(0.75) = 0.75$$

- Threshold activation function:

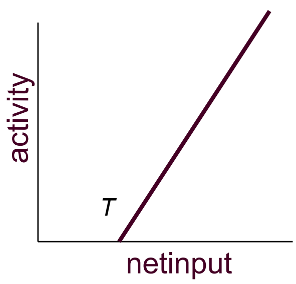

  - **IF** $net_i > T$ **THEN** $a_i := net_i - T$
  - **ELSE** $a_i := 0$

More Activation Functions

- Binary threshold activation function:
  - **IF** $net_i > T$ **THEN** $a_i := 1$
  - **ELSE** $a_i := 0$

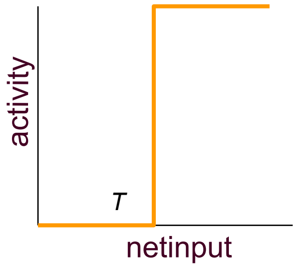

- Nonlinear activation function:
  - It is often more useful to use the "sigmoidal" logistic function:

$$a_i = f(net_i) = \frac{1}{1 + e^{-net_i}}$$

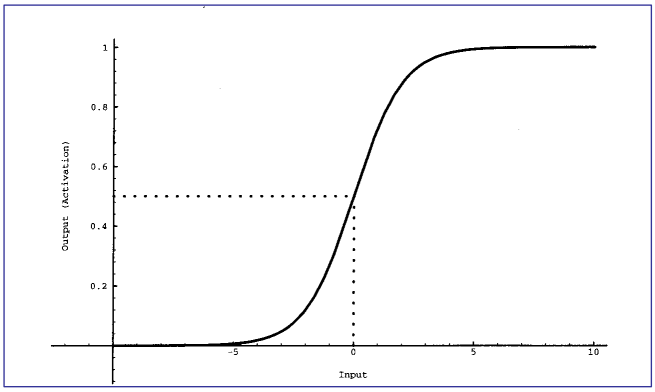

Logistic Function: bias and gain

Weighted sum of inputs

$$net_i = \sum_j w_{ij} a_j$$

Bias θ: threshold weight

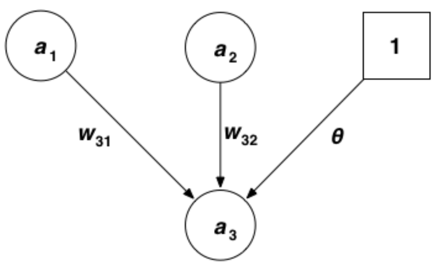

Gain γ: slope of logistic function

$$a_3 = \sigma(\gamma\, net_3 + \theta) = \frac{1}{1 + e^{-(\gamma\, net_3 + \theta)}}$$

$$\sigma'(x) = \gamma\, \sigma(x)\, [1 - \sigma(x)]$$

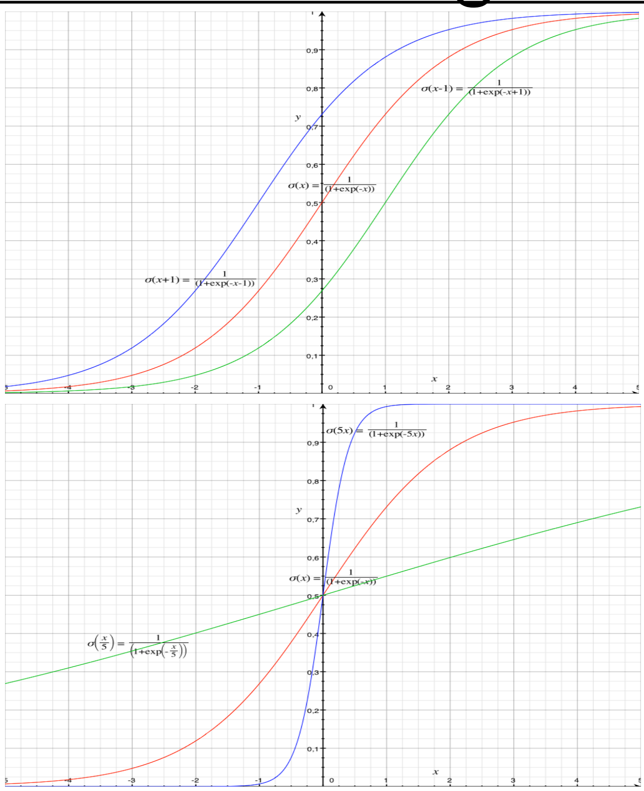

About activation functions

- Defines the relationship between the net input to a node, and its activation level/output.
- Neurons in the brain have thresholds, only fire with sufficient net input.
- Nonlinearity can be useful to reduce the effects of spurious inputs, noise.
  - e.g. where a small change in input can result in large change in output
- Most common in connectionist modelling: sigmoid/logistic
  - Activation ranges between 0 and 1
  - Rate of activation change is highest for net inputs around 0
  - Models neurons: thresholding, a maximum activity, smooth transition between states.
- The sigmoid function also has nice mathematical properties

Summary of network architecture

- The activation of a unit *i* is represented by the symbol a_i.
- The extent to which unit *j* influences unit *i* is determined by the weight w_{ij}
- The input from unit j to unit i is the product: $a_i * w_{ij}$
- For a node i in the network:

$$net_i = \sum_j w_{ij} a_j$$

The output activation of node i is determined by the activation function, e.g. the logistic:

$$a_i = f(net_i) = \frac{1}{1 + e^{-net_i}}$$

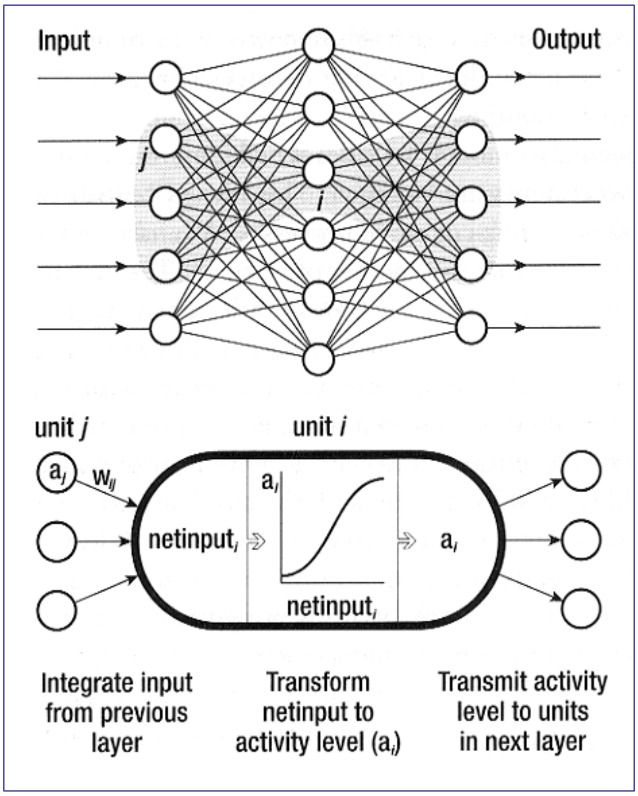

Central issues

- Language, like other cognitive and perceptual faculties, is "implemented" in the neural-tissues of the brain.
- What is the right computational level at which to develop our theories?
  - Is connectionist versus symbolic simply a matter of abstraction?
  - Can connectionism fully replace symbolic accounts?
  - Should they be viewed as complementary?
- Is there a clear boundary between connectionist and symbolic computation in the brain, or does symbol/rule-like behaviour emerge gradually?
- What kinds/levels of cognitive functions require connectionist explanation, and what are best suited to symbolic accounts?

Connectionist language processing

- Simple connectionist models and their properties: The perceptron
  - Learning in single layer networks
- Multi-layer perceptrons: feed-forward networks and internal representations
  - Learning in multi-layer networks (AKA *"deep learning"*)
- The encoding problem: Localist and distributed representations
- Generalisation, association, and translational invariance
- The role of representations for inputs and outputs

Dual Route vs Network Models

The standard model of reading posits two independent routes leading to pronunciation of a word, because …

- People can pronounce words they have never seen: SLINT or MAVE
    - *One mechanism uses general rules*
- People can pronounce words which break the rules: PINT or HAVE
    - *Another stores pronunciation information with specific words*

Can a "single-route" network account for this behaviour?

Also, similar for modeling past-tense formation

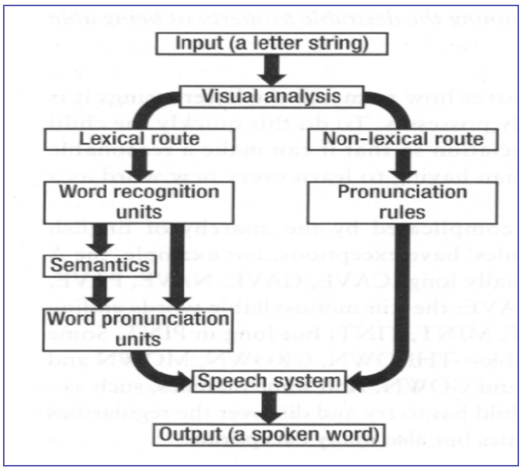

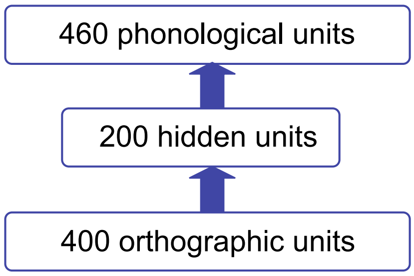

Simple Recurrent Networks

- Simple recurrent networks can learn sequences given as input
- We can tell they've learned by training them to predict the next item, i.e.
    - The next letter or sound of a word
    - The next word of a sentence
- Analyse what they've learned
- We can also train them to map sentences to meaning representations

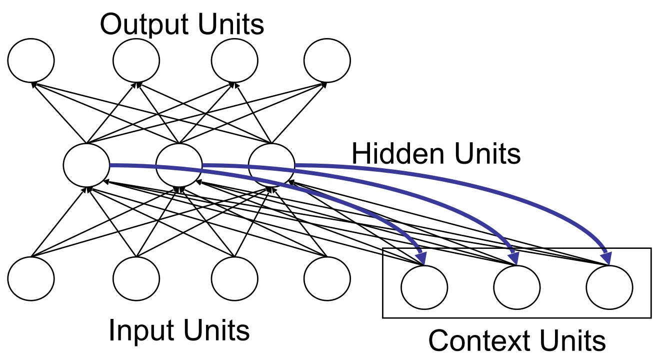

Neurocomputational model of comprehension

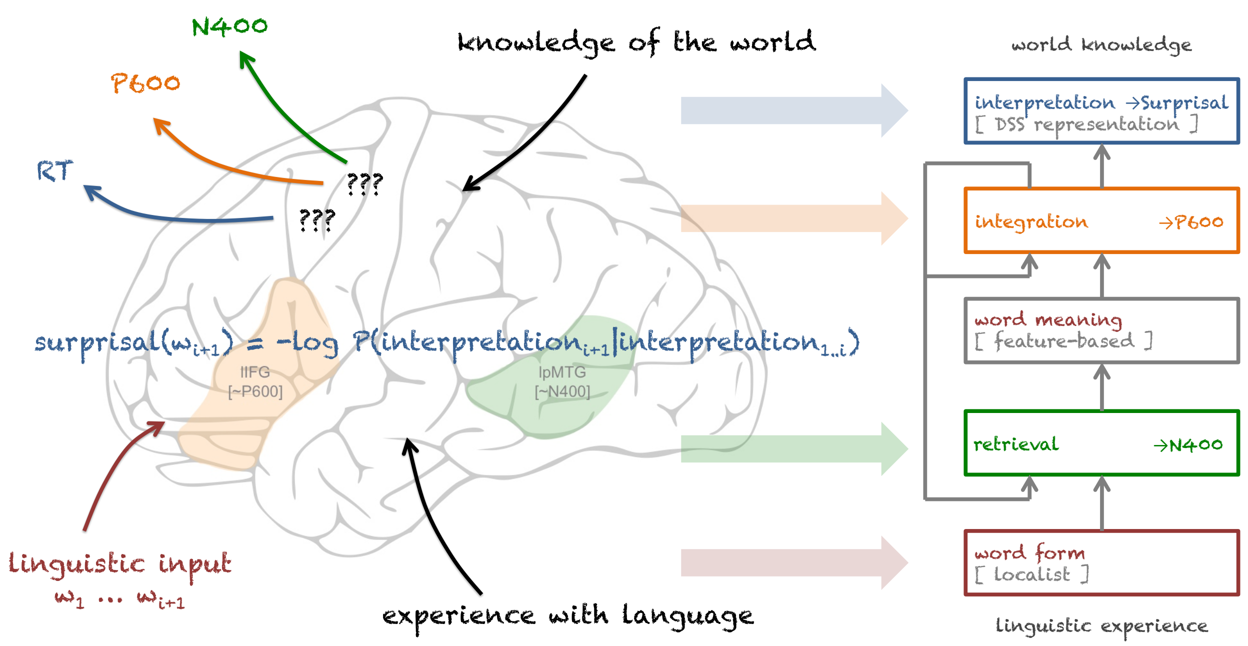

Properties of Connectionist Networks

- Learning: there is usually no predetermined (innate) knowledge of language, but ...
    - Input/output representation are often specified
    - The architecture of the network may be "suited" to a particular task
    - The learning mechanism and parameters provide degrees of freedom
    - Learning is takes place in direct response to experience
- Generalisation
    - Networks are able to learn generalisations not just by rote
    - More efficient representation of information
    - Novel inputs can be processed

Properties continued

- Representation
  - Learned automatically, and typically distributed
  - Single mechanism to explain both general rules and also exceptions
- Graded:
  - Can often give a useful output to new, partial, noisy input (pattern completion)
  - Damage is distributed, and some performance is still possible:
    - Modelling of brain damage and neurological disorders in possible
- Frequency effects
  - Model response time behaviours where high frequency inputs are recognised faster than low frequency ones

Course details

- Weekly lectures (Tues 2-4pm) and tutorials (Thurs 2-4pm) – USUALLY!
  - Participation in, and completion of, tutorials is required!
- Assessment: Final Exam (100%), Date: Tues, July 16, 2019
  - All tutorial assignments must be successfully completed to sit the exam
- Course materials (overheads and most readings) will be made available on the course homepage (linked from general course page)
- Contact: crocker@coli.uni-sb.de, brouwer@coli.uni-saarland.de