Multilayer Perceptrons

A discussion of
The Algebraic Mind
Chapters 1+2

Jasmin Steinwender (jsteinwe@uos.de)
Sebastian Bitzer (sbitzer@uos.de)
Seminar Cognitive Architecture
University of Osnabrueck
April 16th 2003

The General Question

What are the processes and representations underlying mental activity?

Connectionism vs. Symbol manipulation

- Also referred to as **parallel**-distributed processing (PDP) or neural network models
- Hypothesis that cognition is a dynamic pattern of connections and activations in a 'neural net.'
- Model of the parallel processor and the relevance to the anatomy and function of neurons.
- Consists of simple neuron- like processing elements: units

- Biological plausible?

  brain consisting of neurons, evidence for hebbian learning in the brain

- „classical view“
- Production rules
- Hierarchical binary trees
- computer-like application of rules and manipulation of symbols
- Mind as symbol manipulator (Marcus)

- Biological plausible?

  Brain circuits as representation of generalization and rules

BUT…

Ambiguity of the term connectionism:

in the huge variety of connectionist models
some will also include symbol-manipulation

Two types of Connectionism

1. **implementational connectionism:**
   - a form of connectionism that would seek to understand how systems of neuron-like entities could implement symbols

2. **eliminative connectionism:**
   - which denies that the mind can be usefully understood in terms of symbol-manipulation

→ „ …eliminative connectionism cannot work(…): eliminativist models (unlike humans) provably cannot generalize abstractions to novel items that contain features that did not appear in the training set.”

Gary Marcus:
http://listserv.linguistlist.org/archives/info-childes/infochi/Connectionism/connectionist5.html and
http://listserv.linguistlist.org/archives/info-childes/infochi/Connectionism/connectionism11.html

Symbol manipulation
-3 separable Hypothesis-

- Will be explicitly explained in the whole book, now just mentioned

1. „The mind represents abstract relationships between **variables**“
2. „The mind has a system of **recursively structured representations**“
3. „ The mind distinguishes between mental representations of **individuals** and mental representation of **kinds**“

If the brain is a symbol-manipulator, then one of this hypotheses must hold.

Introduction to Multilayer Perceptrons

- simple perceptron
  - local vs. distributed
  - linearly separable
- hidden layers
- learning

The Simple Perceptron I

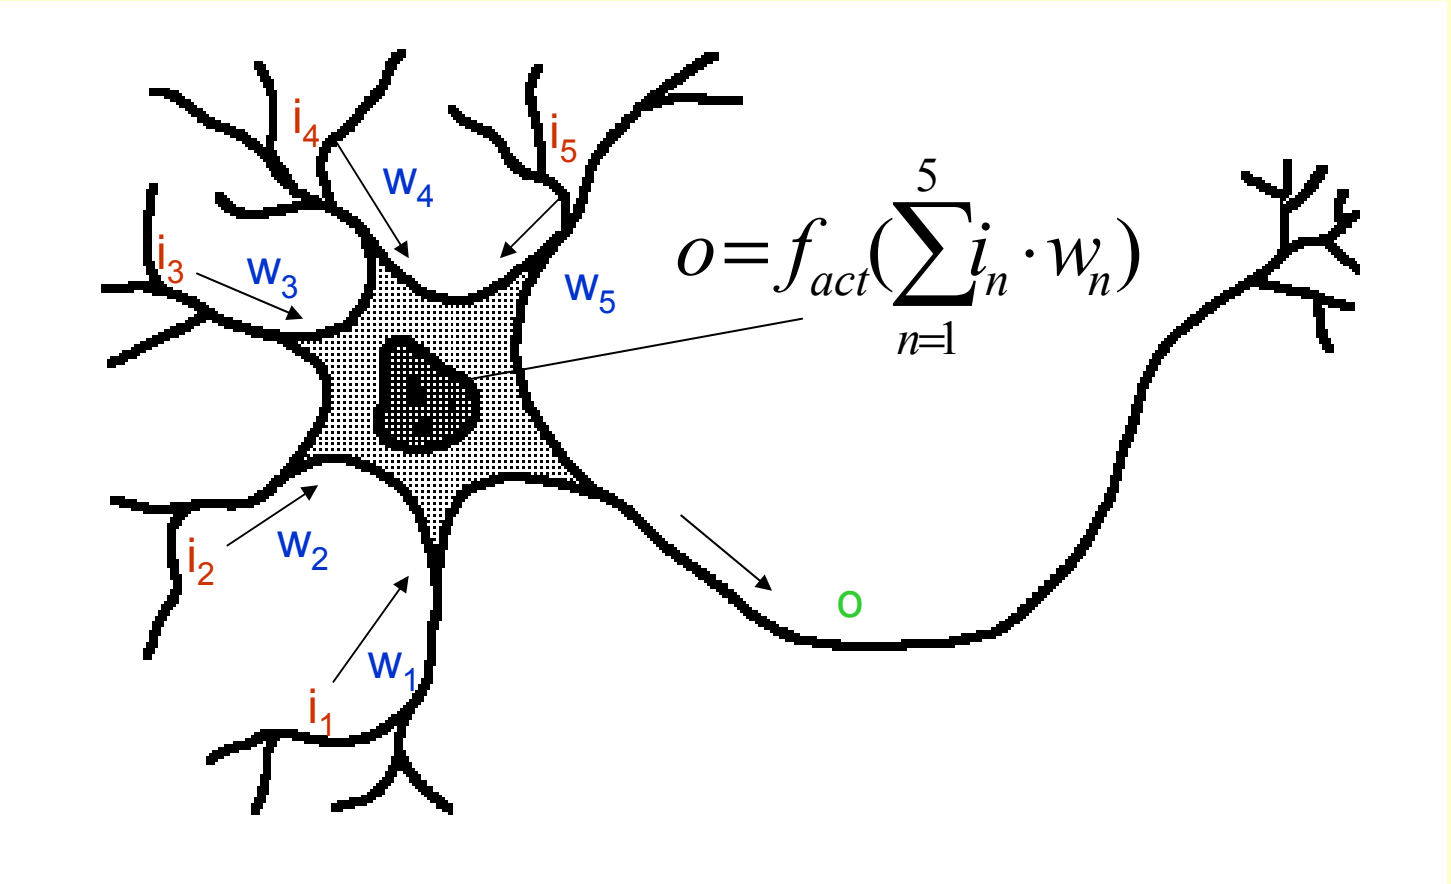

Activation functions

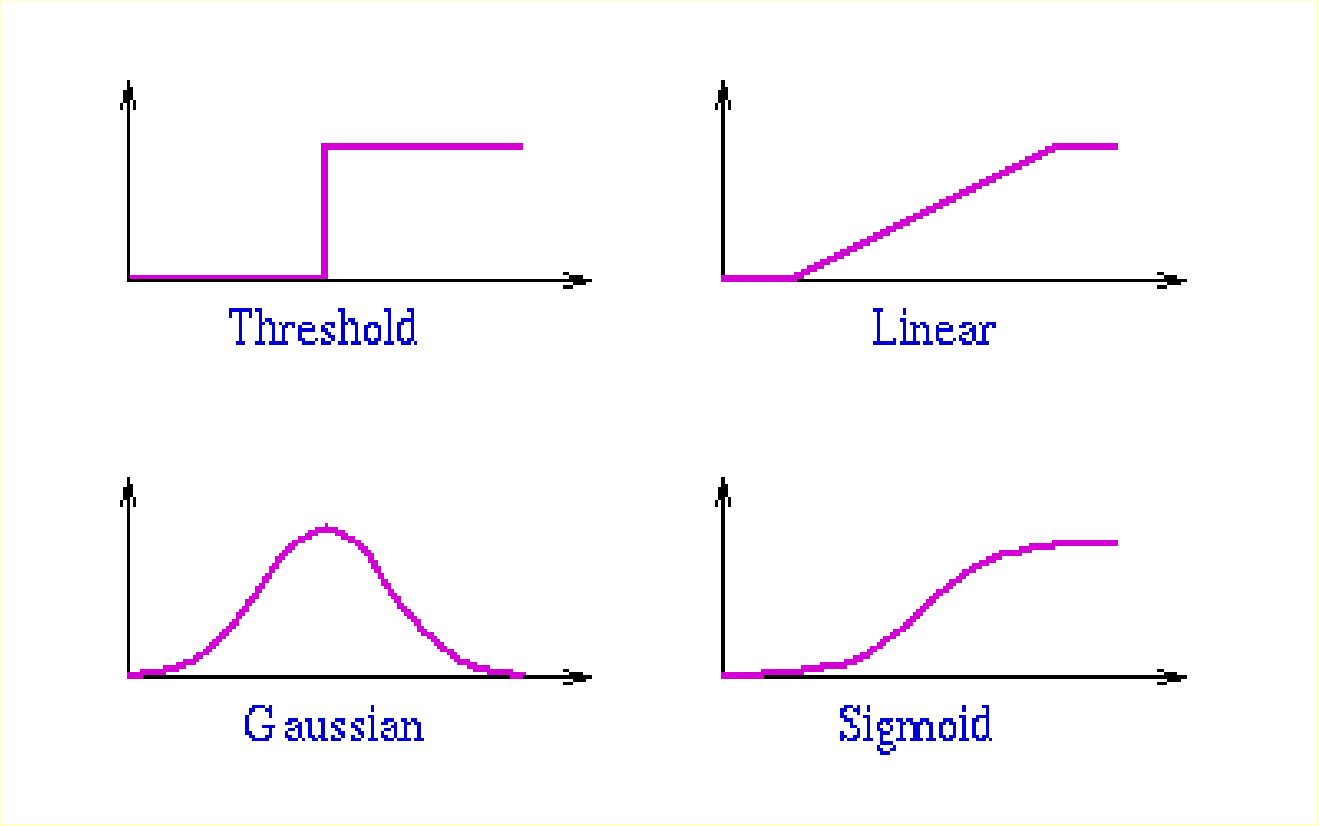

The Simple Perceptron II

a single-layer feed-forward mapping network

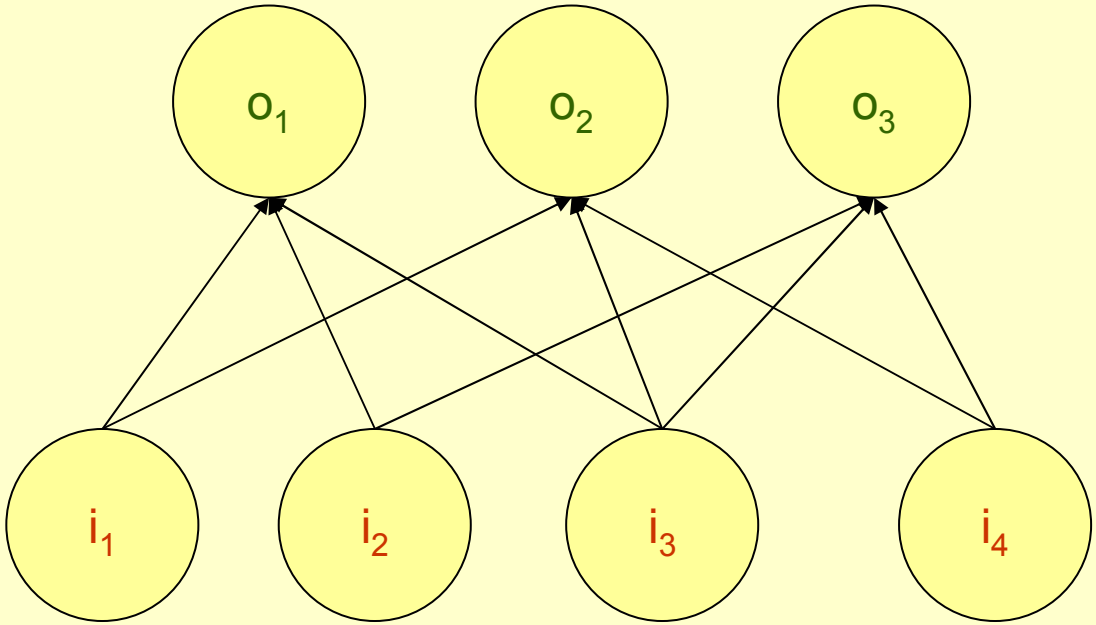

Local vs. distributed representations

representation of CAT

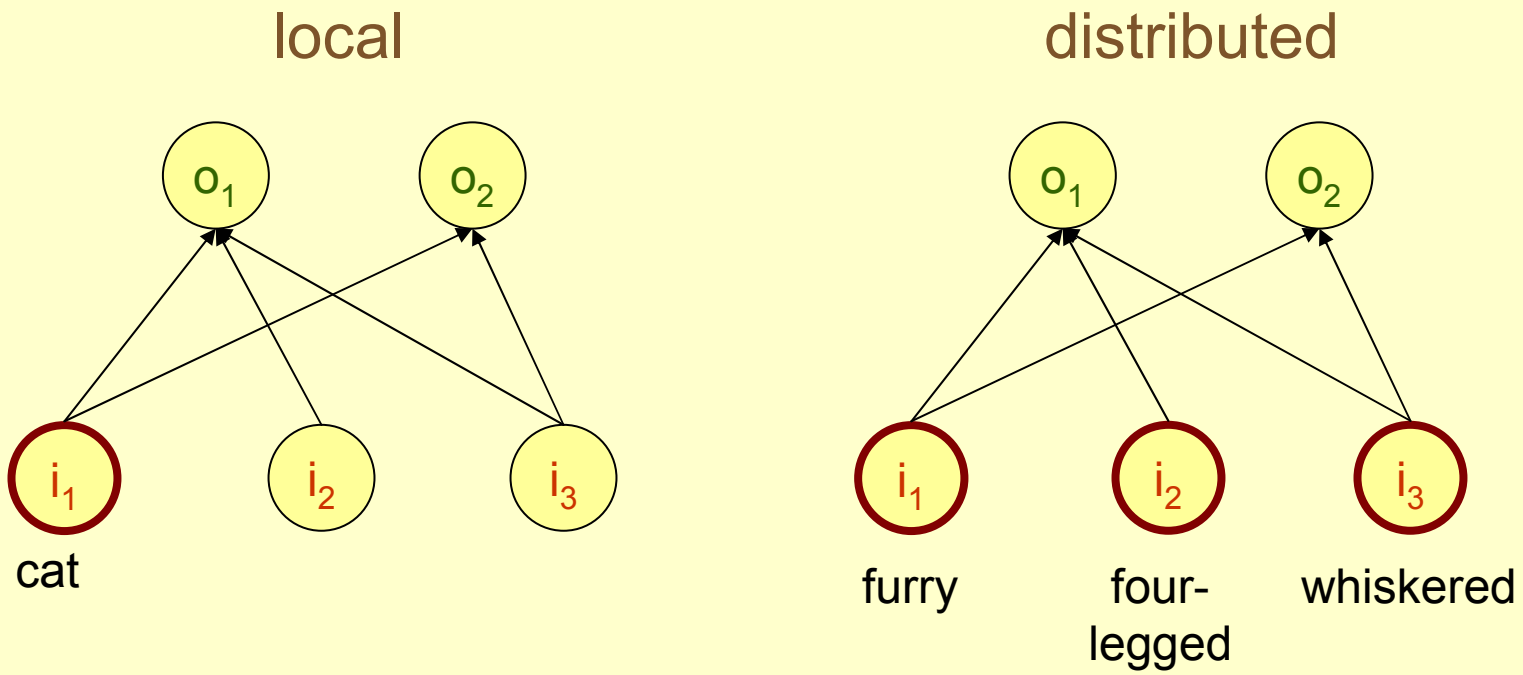

Linear (non-)separable functions I

A.

Boolean OR function

x_1	x_2	y
0	0	0
0	1	1
1	0	1
1	1	1

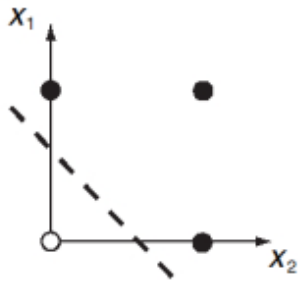

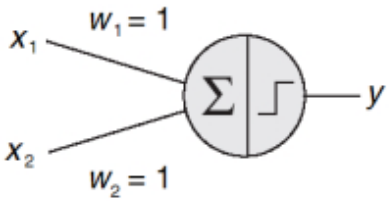

B.

Boolean XOR function

x_1	x_2	y
0	0	0
0	1	1
1	0	1
1	1	0

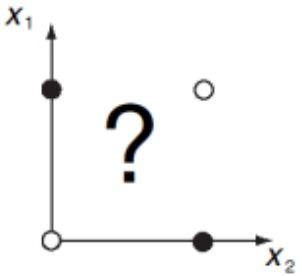

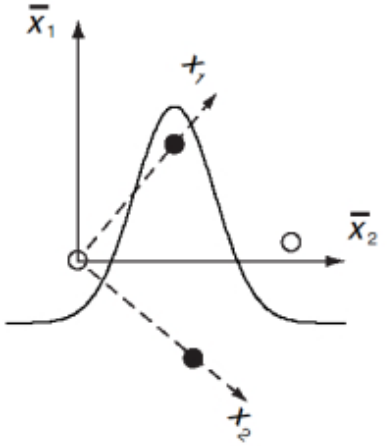

[Trappenberg]

Linear (non-)separable functions II

boolean functions

n	Number of linear separable functions	Number of linear non-separable functions
2	14	2
3	104	151
4	1,882	63654
5	94,572	$\sim 4.310^{9}$
6	15,028,134	$\sim 1.810^{19}$

Hidden Layers

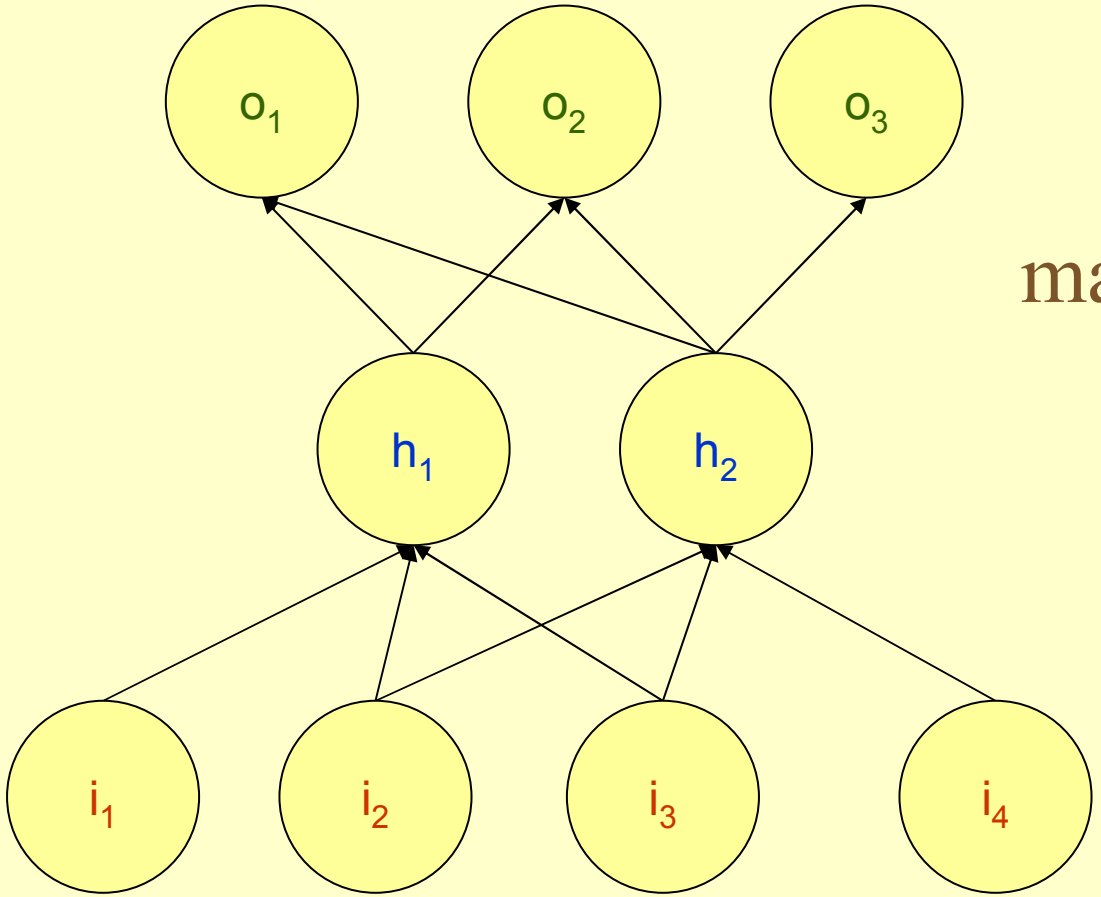

a two-layer feed-forward mapping network (a MLP)

Learning

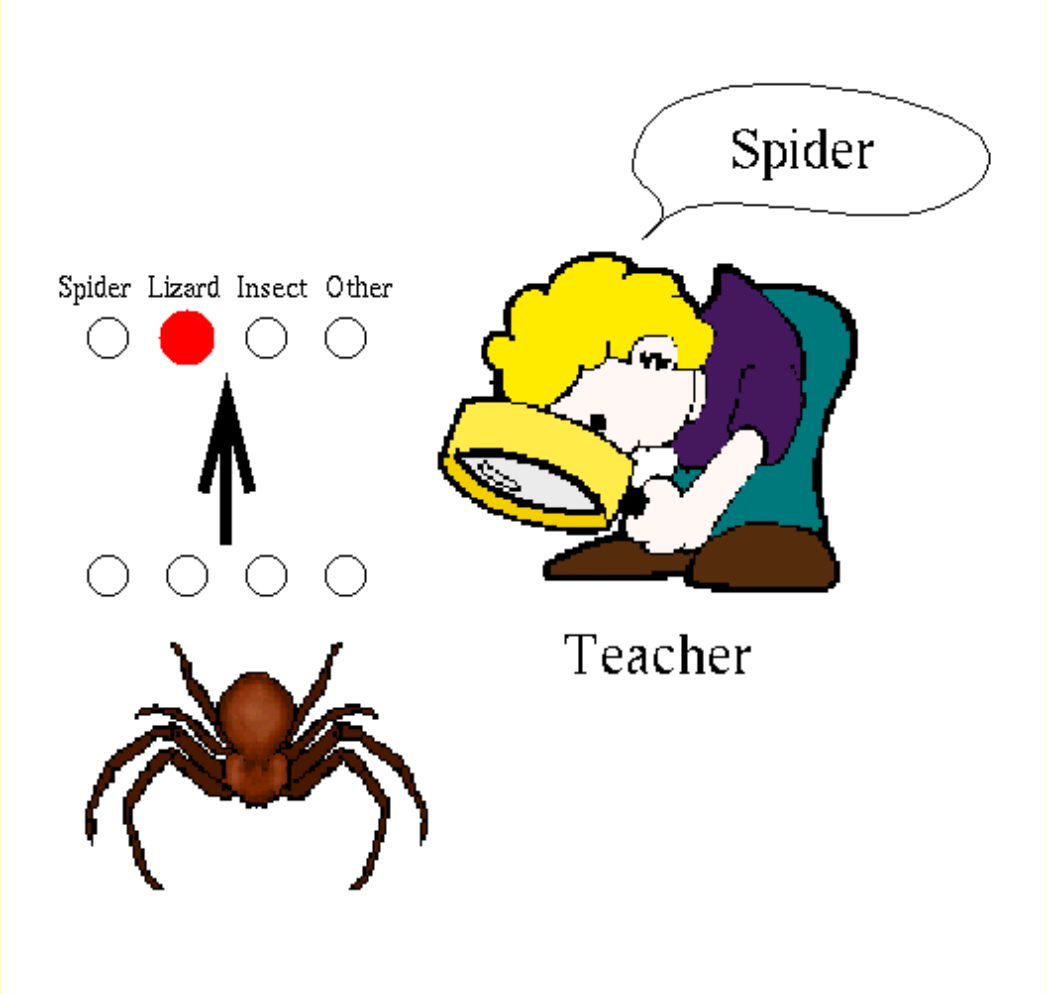

Backpropagation

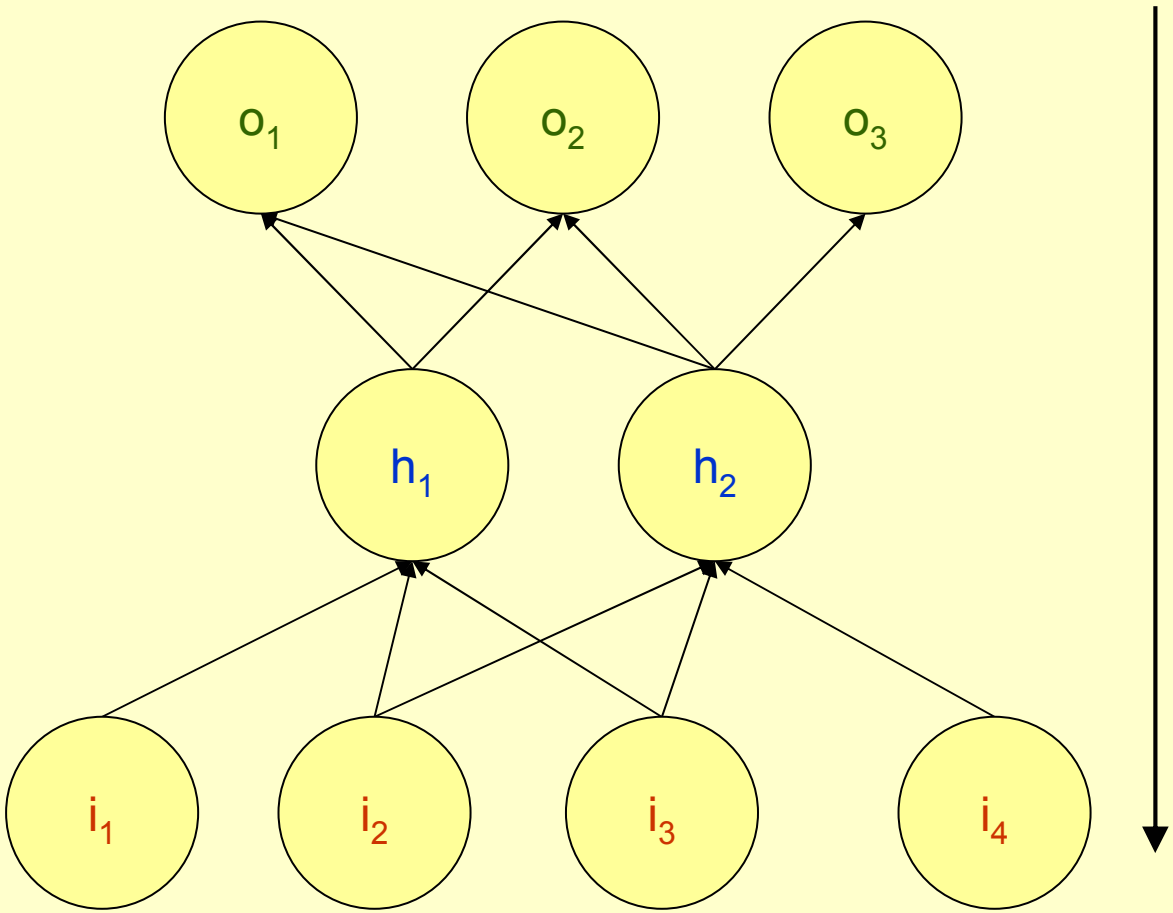

- compare actual output - right o., change weights
- based on comparison from above change weights in deeper layers, too

Multilayer Perceptron (MLP)

A type of feedforward neural network that is an extension of the perceptron in that it has at least one hidden layer of neurons. Layers are updated by starting at the inputs and ending with the outputs. Each neuron computes a weighted sum of the incoming signals, to yield a net input, and passes this value through its sigmoidal activation function to yield the neuron's activation value. Unlike the perceptron, an MLP can solve linearly inseparable problems.

Gary William Flake,
The Computational Beauty of Nature,
MIT Press, 2000

Many other network structures

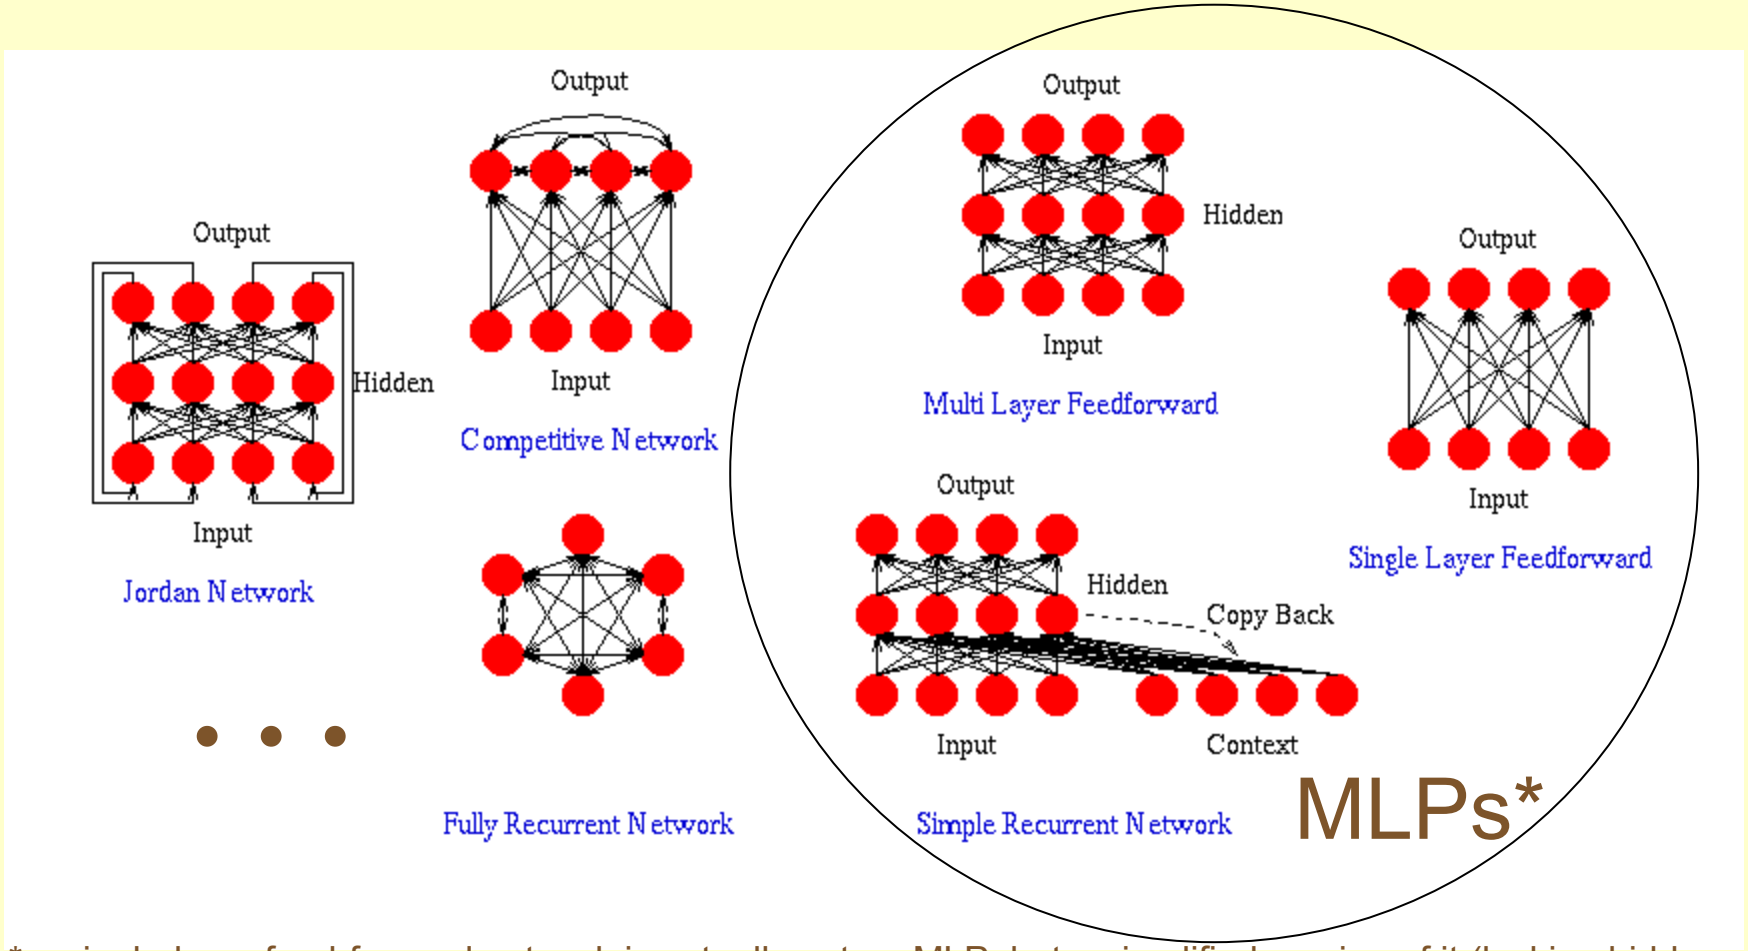

* a single-layer feed-forward network is actually not an MLP, but a simplified version of it (lacking hidden layers)

The Family-Tree Model

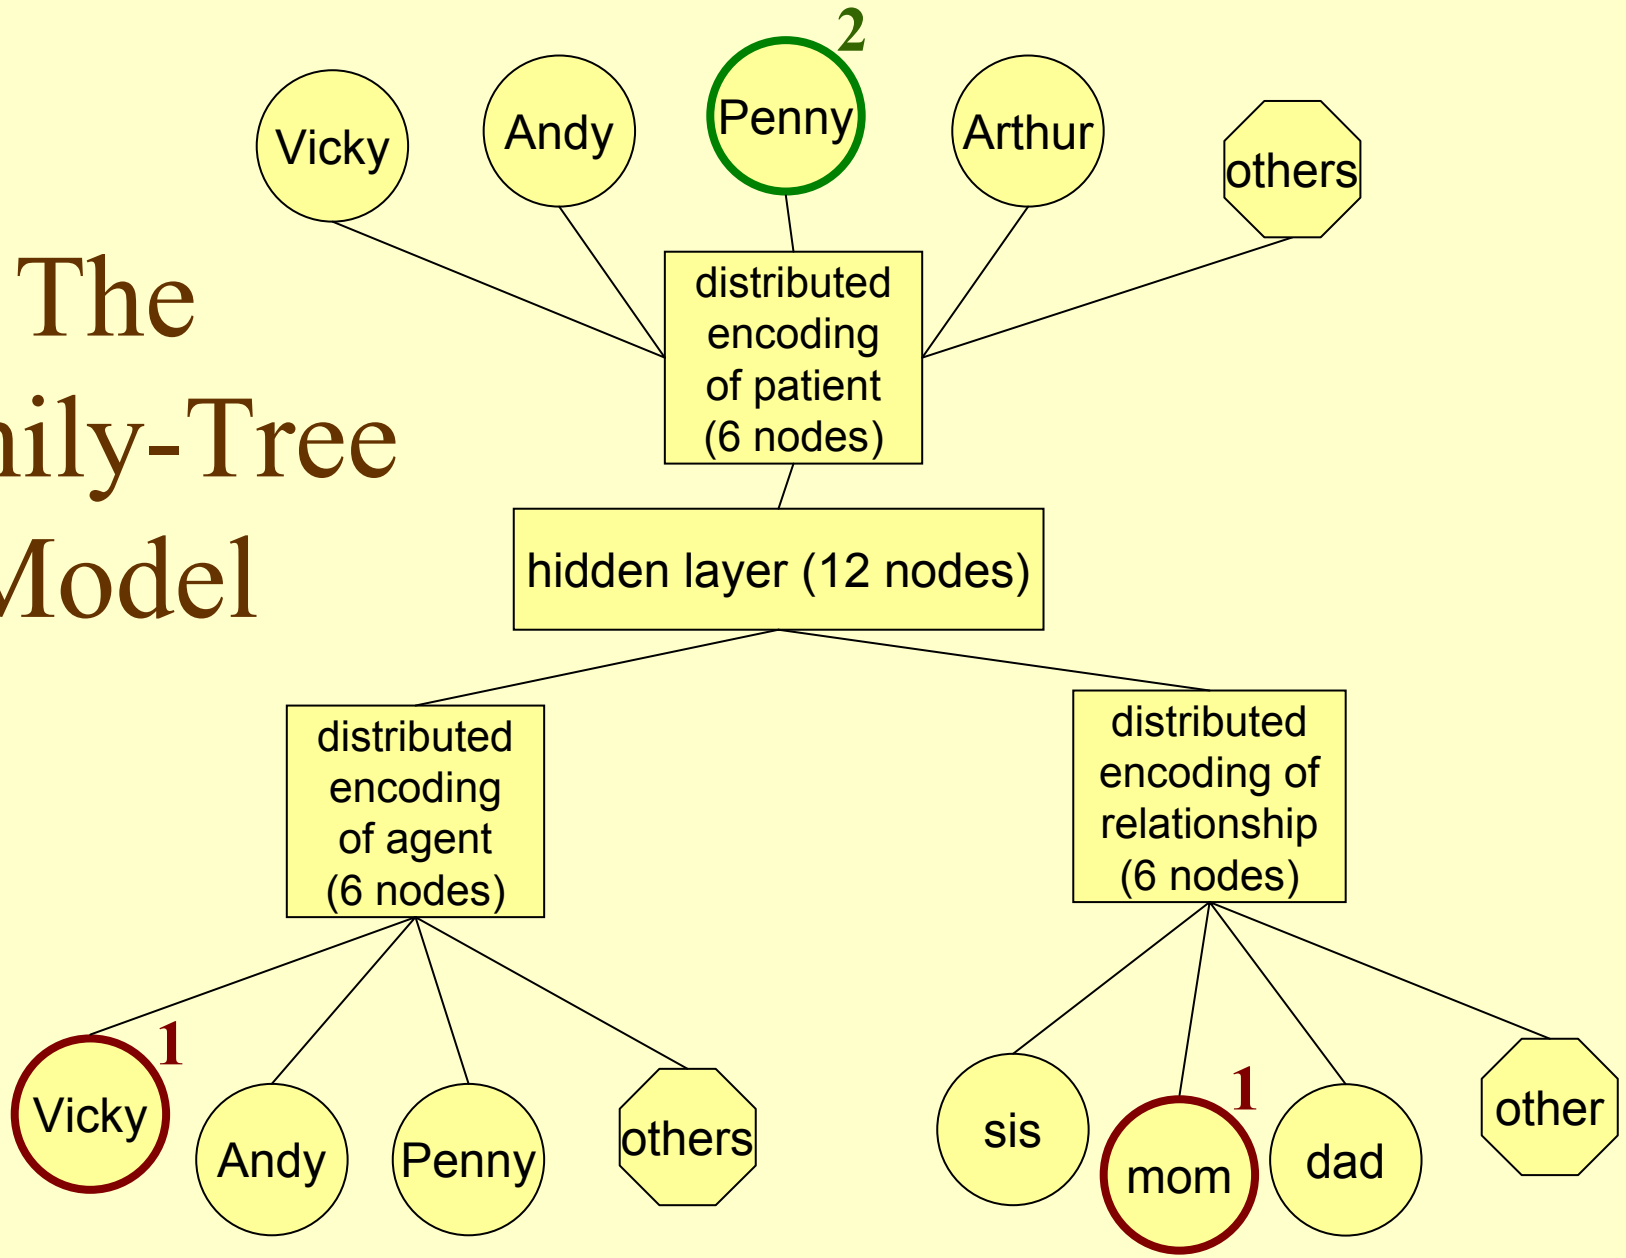

The sentence prediction model

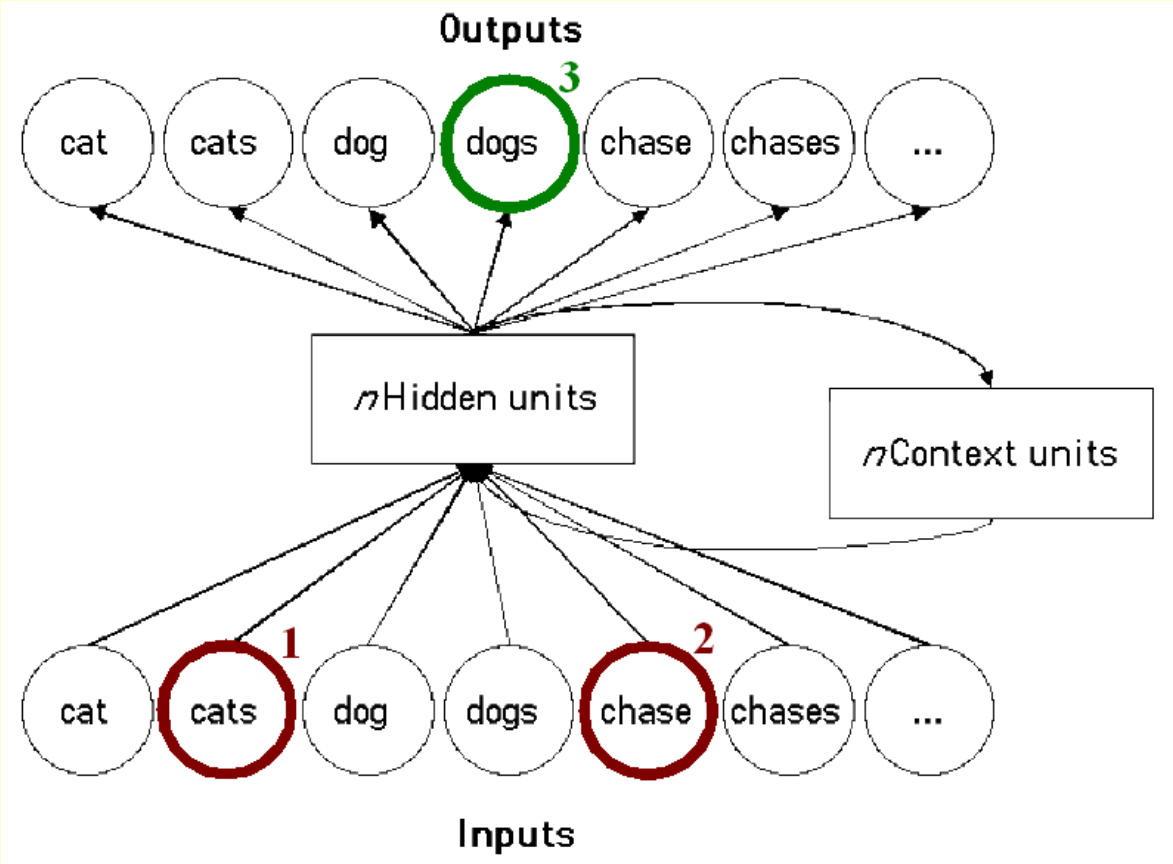

The appeal of MLPs (preliminary considerations)

1. Biological plausibility

- independent nodes
- change of connection weights resembles synaptic plasticity
- parallel processing

$\Rightarrow$ brain is a network and MLPs are too

Evaluation Of The Preliminaries

1. Biological plausibility

- Biological plausibility considerations make no distinction between eliminative and implementing connectionist models
- Multilayered perceptron as „more compatible than symbolic models“, BUT nodes and their connections only loosely model neurons and synapses
- Back-propagation MLP lacks brain-like structure and requires varying synapses (inhibitory and excitatory)
- Also symbol-manipulation models consist of multiple units and operate in parallel → brain-like structure
- Not yet clear what is biological plausible – biological knowledge changes over time

Remarks on Marcus

difficult to argue against his arguments:

- sometimes addresses comparison between eliminative and implementational connectionist models
- sometimes he compares connectionism and classical symbol-manipulation

Remarks on Marcus

1. Biological plausibility

(comparison MLPs – classical symbol-manipulation)

- MLPs are just an abstraction
- no need to model newest detailed biological knowledge
- even if not everything is biological plausible, still MLPs are more likely

Preliminary considerations II

2. Universal function approximators

- "multilayer networks can approximate any function arbitrarily well" [Trappenberg]
- "information is frequently mapped between different representations" [Trappenberg]
- mapping of one representation to another can be seen as a function

Evaluation Of The Preliminaries II

2. Universal function approximators

- MLP cannot capture all functions (f. e. partial recursive func. – models computational properties of human language)
- No guarantee of generalization ability from limited data like humans
- Unrealistic need of infinite resources for universal function approximation
- Symbol-manipulators could also approximate any function

Preliminary considerations III

3. Little innate structure

- children have relatively little innate structure

$\Rightarrow$ "simulate developmental phenomena in new and … exciting ways" [Elman et al., 1996]

e.g. model of balance beam problem [McClelland, 1989] fits data from children

- domain-specific representations from domain-general architectures

Evaluation Of The Preliminaries III

3. Little innate structure

- There also exist symbol-manipulating models with little innate structure
- Possibility to prespecify the connection weights of MLP

Preliminary considerations IV

4. Graceful degradation

- tolerate noise during processing and in input
- tolerate damage (loss of nodes)

Evaluation Of The Preliminaries IV

4. Learning and graceful degradation

- No unique ability of all MLP
- Symbol-manipulation models which can also handle degradation
- No yet empirical data that humans recover from degraded input

Preliminary considerations V

5. Parsimony

- one just has to give the architecture and examples
- more generally applicable mechanisms (e.g. inflecting verbs)

Evaluation Of The Preliminaries V

5. Parsimony

- MLP connections interpreted as free parameters → less parsimonious
- Complexity may be more biological plausible than parsimony
- Parsimony as criterion only if both models cover the data adequately

What truly distinguishes MLP from Symbol -manipulation

- Is not clear, because…

…both can be context independent

…both can be counted as having symbols

…both can be localist or distributed

We are left with the question:

Is the mind a system that represents

- abstract relationships between variables OR*
- operations over variables OR*
- structured representations
- and distinguishes between mental representations of individuals and of kinds

We will find out later in the book…

*inclusive

Discussion

"… I agree with Stemberger that connectionism can make a valuable contribution to cognitive science. The only place that we differ is that, first, he thinks that the contribution will be made by providing a way of *eliminating* symbols, whereas I think that connectionism will make its greatest contribution by accepting the importance of symbols, seeking ways of supplementing symbolic theories and seeking ways of explaining how symbols could be implemented in the brain. Second, Stemberger feels that symbols may play no role in cognition; I think that they do."

Gary Marcus:
http://listserv.linguistlist.org/archives/info-childes/infochi/Connectionism/connectionist8.html

References

- Marcus, Gary F.: *The Algebraic Mind*, MIT Press, 2001
- Trappenberg, Thomas P.: *Fundamentals of Computational Neuroscience,* OUP, 2002
- Dennis, Simon & McAuley, Devin: *Introduction to Neural Networks*, http://www2.psy.uq.edu.au/~brainwav/Manual/WhatIs.html