Conexionism

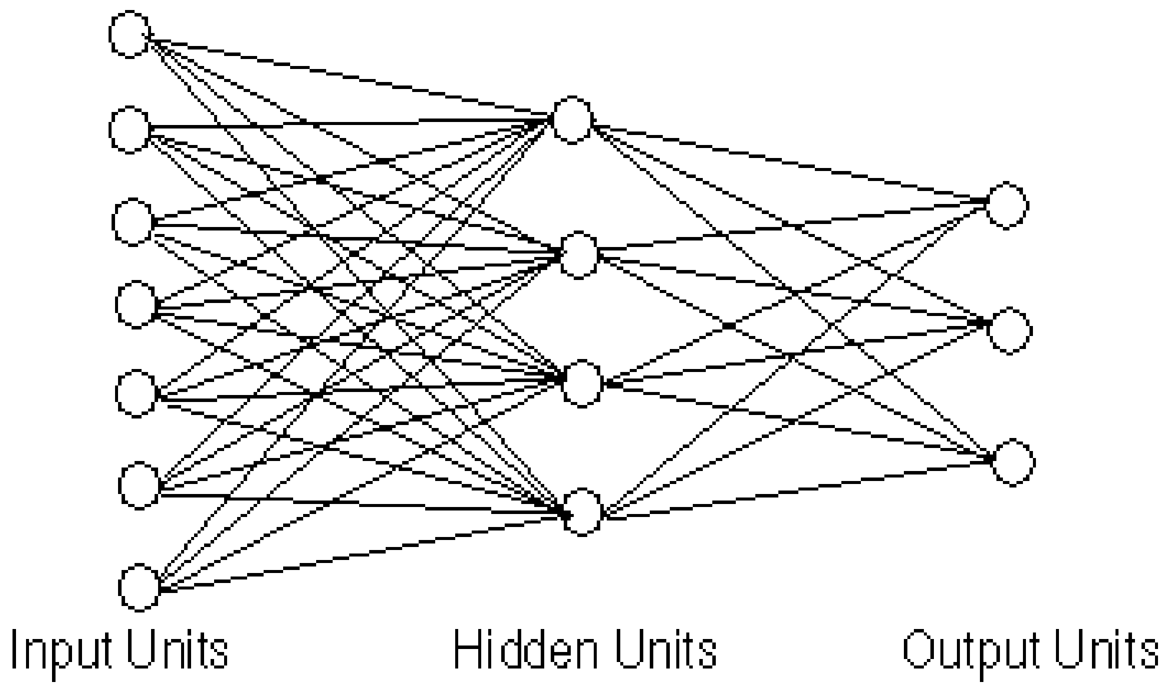
Input Units
Hidden Units
Output Units

- Retea larga neuroni: Computari + Rs
- Computarile = Proiectare vector input pe un vector output
- Diferenta paradigma simbolica si C - Natura reprezentarilor (Rs)
- Rs = Vectori (interpretati)
- Rs distribuite sau locale
- Reteaua lucreaza in paralel

- a_j = activarea nodului j care trimite output la nodul i
- Legatura dintre a_j si a_i este w_{ij}
- Inputul j la i este $a_i = w_{ij}\ a_j$

- Ex: Nodul j are output = 0.5; legatura j cu i = -2.0 →

  (0.5 x -2.0) = -1

- Un nod primeste inputuri alte noduri: *Net input*

  $$Net_i = \sum w_{ij}\ a_j$$

Input total primit de nodul i

McCulloch & Pitt ('43), Selfridge ('59)

Rosenblatt ('62) - **Perceptron Convergence Procedure,**

- Doua niveluri cu **regula Hebb**
- Minsky and Papert ('69):

- Functii logice SI, SAU = Linear separabile
- Exclusive SAU = Nonlinear separabil Elman (pp. 61-63)

- Solutia: Rs interne (prin unitati ascunse)

(Elman, p. 64-65)

- Reteaua invata ce legaturi are nevoie!

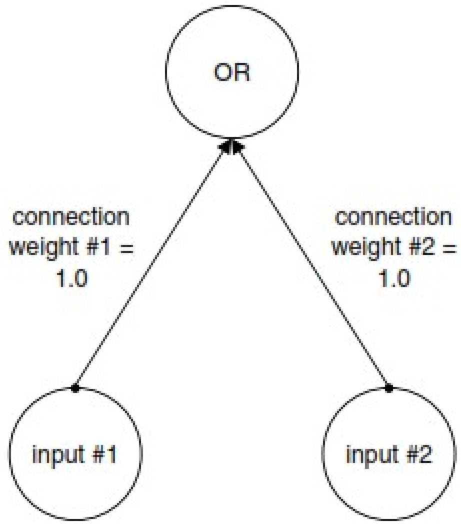

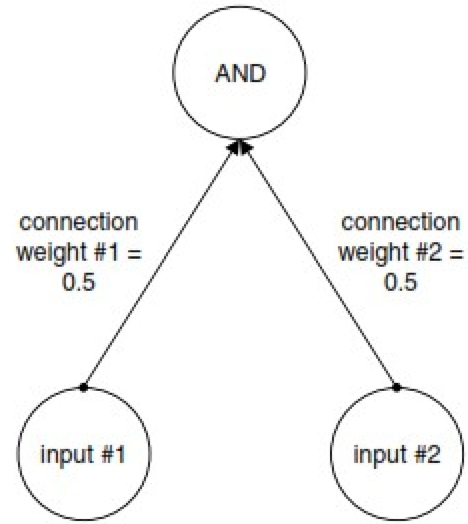

SAU

Input1	input2	output
0	0	0
1	0	1
0	1	1
1	1	1

SI

Input1	input2	output
0	0	0
1	0	0
0	1	0
1	1	1

FIGURE 2.4 two-layer network that cannot learn to solve the XOR problem

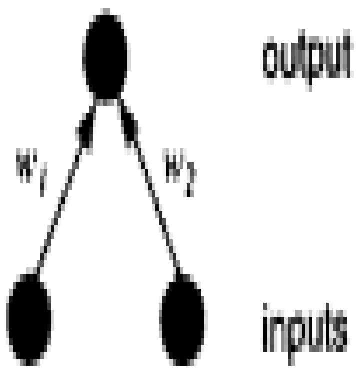

EXCLUSIVE SAU

input1	input2	output
0	0	0
1	1	0
0	1	1
1	0	1

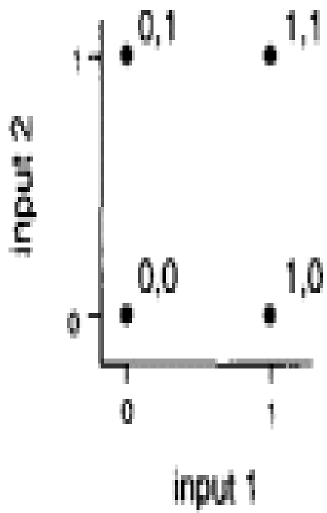

FIGURE 2.5 Spatial representation of the four input patterns used in XOR.

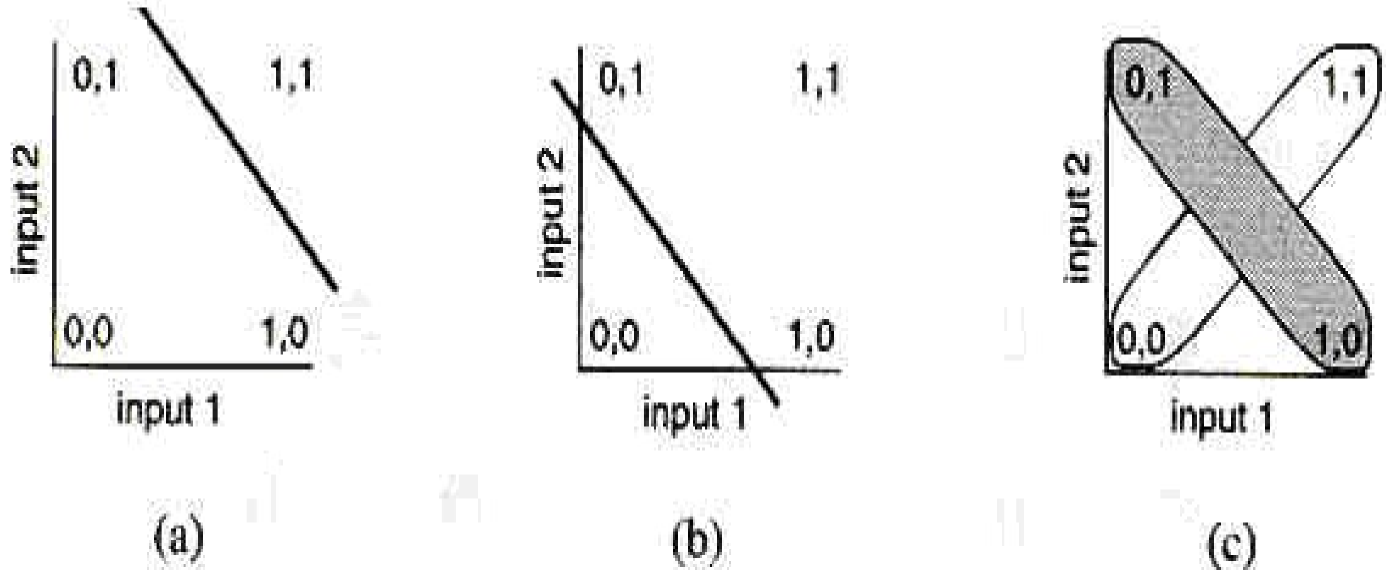

FIGURE 2.6 Geometric representation of the XOR problem. If the four input patterns are represented as vectors in a two-dimensional space, the problem is to find a set of weights which implements a linear decision boundary for the output unit. In (a), the boundary implements logical AND. In (b), it implements logical OR. There is no linear function which will simultaneously place 00 and 11 on one side of the boundary, and 01 and 10 on the other, as required for (c).

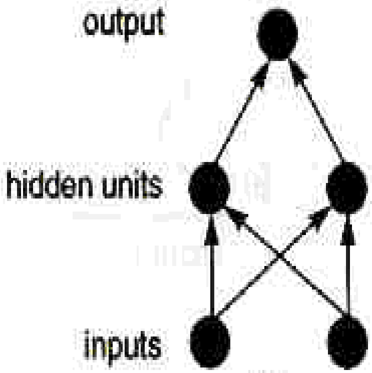
output
hidden units
inputs

(a)

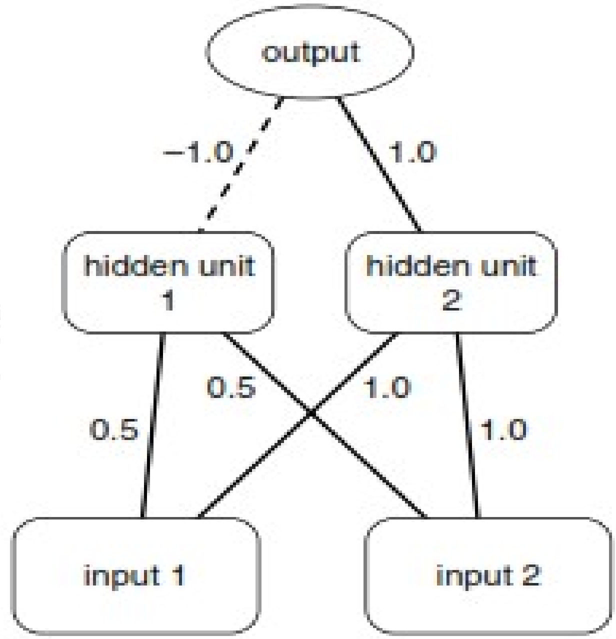
output
−1.0
1.0
hidden unit 1
hidden unit 2
0.5
0.5
1.0
1.0
input 1
input 2


0,1
1,1
input 2
0,0
1,0
input 1
0,1
1,1
0,0
hidden unit 2
1,0
hidden unit 1
1,0
0,1
1,1
0,0
output

(a) (b) (c)

- Diferete tipuri de retea

Elman p. 51

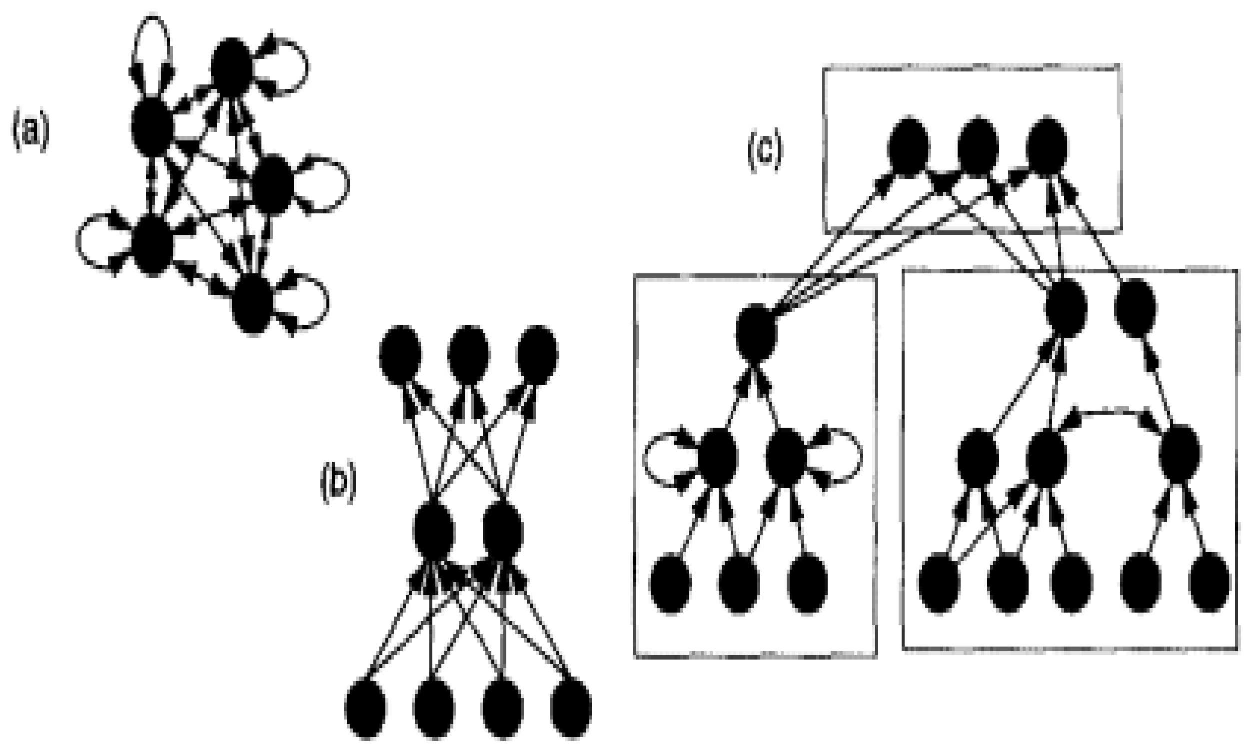

- **Paternul de activare a unei retele - Determinat de legaturile intre noduri**
- “Cunostiintele retelei sunt construite progresiv prin schimbarea legaturilor.” (Elman et al. ‘96)
- Fiecare unitate input primeste input extern
- Fiecare unitate ascunsa calculeaza valoarea proprie in functie de valorile de activare primite de la unitatile input. (Garson 2007)

- “Fiecare unitate ascunsa este sensitiva la regularitati complexe = **microtrasaturi**.” (Elman, p. 55)

Each layer of hidden units - providing a particular distributed encoding of input pattern (= an encoding in terms of pattern of microfeatures).” (Bechtel & Abrahamsen 2002 or Hinton, McClleland and Rumelhart 1986) →

Nivelul subsimbolic (Smolensky ‘88)

- Acelasi proces intre unitatile ascunse si cele output
- Activarea retelei = Legaturi pozitive (excitatorii) sau negative (inhibitorii)

- Valoarea de activare - fiecare unitate: “Functia aduna contributiile de la toate unitatile care trimit, unde contributia unei unitati este definita de legaturile intre unitatile de trimitere si cele de primite inmultit cu valoarea de activare a unitatii care trimite.”

- Outputul ≠ inputul (ca si neuronii)
- Ce “face” un nod (functia raspuns) = Valoarea de activare a nodului = Functii liniare, mai ales *nonliniare* (functia sigmoid, etc.) (Elman, p. 53)

- Nonliniar = “Valoare numerica a outputului nu este direct proportionala cu suma inputurilor.” (Clak)

- Rumelhart and McClleland (‘81): 3 niveluri - cuvant, litera, trasaturi ortografice

Ex: Nodul “trap” primeste input pozitiv de la nodurile literelor “t”. “r”, “a”, and “p”; celelalte noduri – inhibitate (Elman, p. 55)

↓

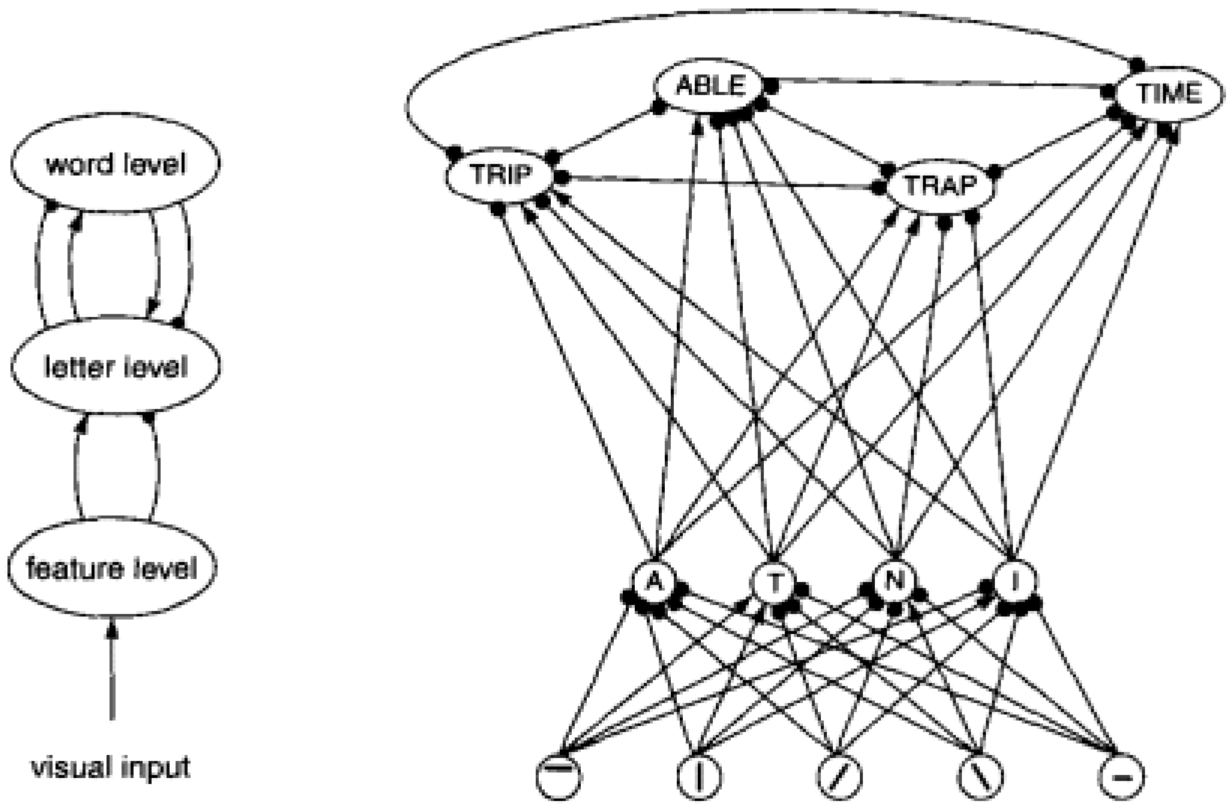
word level
letter level
feature level
visual input
ABLE
TIME
TRIP
TRAP
A
T
N
I

(a) (b)

- Principiu: ***Similaritatea*** → Generalizarea - Dincolo de cazurile intalnite!

- **Asociationism** – Regularitati statistice

- Reteaua clasifica pattern (1100), tendinta clasifice un nou pattern (1101) in mod **similar** (vezi Elman, p. 91)

FIGURE 2.16 Examples of distributed representations. Both (a) and (b) illustrate different patterns of activation of the same set of four units. Activation values for individual units are shown as numbers above each unit. Note that the second unit has the same activation value in both representations; in order to know which concept is represented, one has to look at the entire set of nodes.

[vs. “tirania similaritatii” (Mcleod, Rolls, Plunkett ‘96)]

- Telul: Legaturile pentru o sarcina
- Regula de invatare “**Hebb’s rule**” = invatarea prin corelari-pereche de activare

↓

$$\Delta w_{ij} = ήa_i a_j$$

(Δw_{ij} = schimbarea legaturii intre nod j si i, constanta ή = rata de invatare)

- Limitari: Invatarea doar prin **corelari pereche** (valabil pt retea 2 niveluri)
- Metoda de propagare intoarsa (**backpropagation rule**) (valabil pt retea cu hidden units) (Nu corespunde procesului uman de invatare!)
- Alte metode: self-supervised learning, unsupervised learning) (Elman)

Backpropagation rule

- Pasul I: Eroarea = Diferenta output *actual* si output *target*
- Pasul II: Ajustarea fiecarei legaturi → Descreste eroarea (Churchland p. 43)
- "Feedback-ul local - *supervisor* - crestere sau scadere usoara – legaturi → Imbunatatirea performantei retelei
- Repetata legatura cu legatura + nivel cu nivel → Panta descresterii erorii" (Clark)

- Algoritmul: “Propagam informatia erorii in spate, de la unitatile output la cele ascunse [si la cele input]” (Elman)

- “Abilitatea - invata - schimba in timp - NU functie a unei schimbari **explicite** in mechanism, ci consecinta **intrinseca** a invatarii insasi. Reteaua invata ca si copii.”
- **Invatarea** ca *gradient descendent* in spatiu legaturilor (Elman, p. 72)

↓

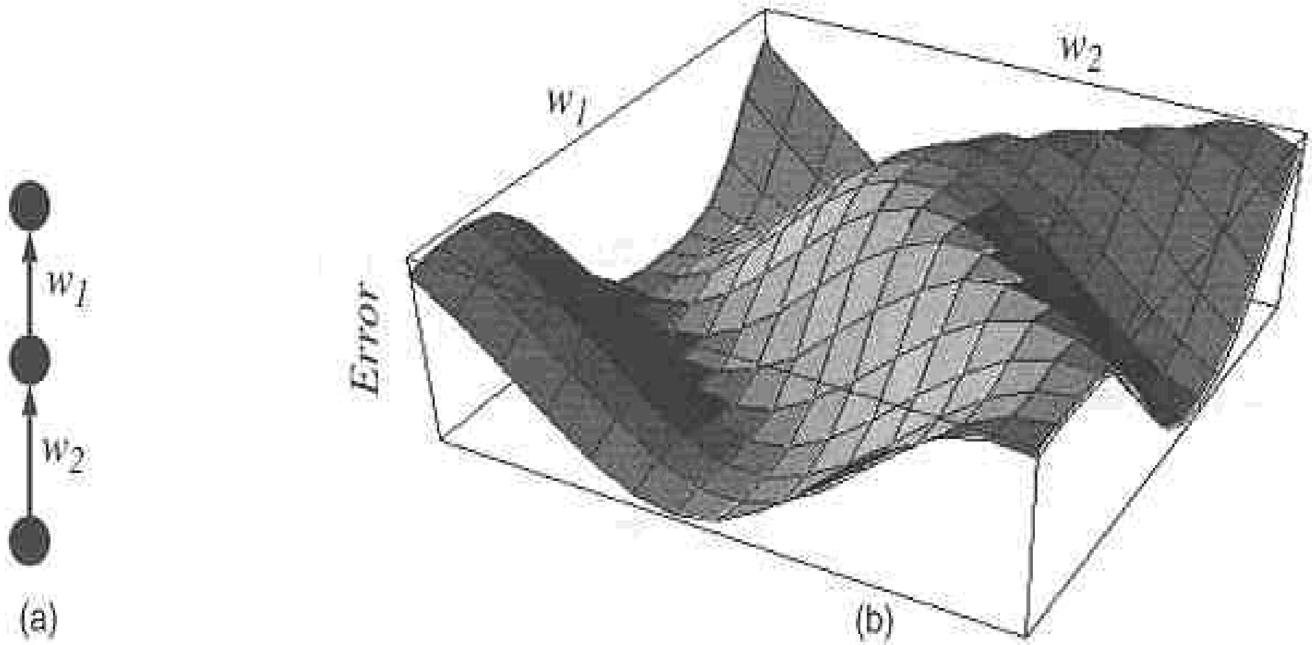

w_1
w_2
(a)
w_1
w_2
Error
(b)

Cotrell – Face recognition

- 64 X 64 input units (256 levels of activation/ “brigthness”); 80 hidden; 8 output
- Training set: 64 photos for 11 faces + 13 photos nonfaces
- 8 output units – 1: faces-nonfaces; 2: male-female; 3: name; restul 5: name

Each hidden unit – links to all input units: **holism**

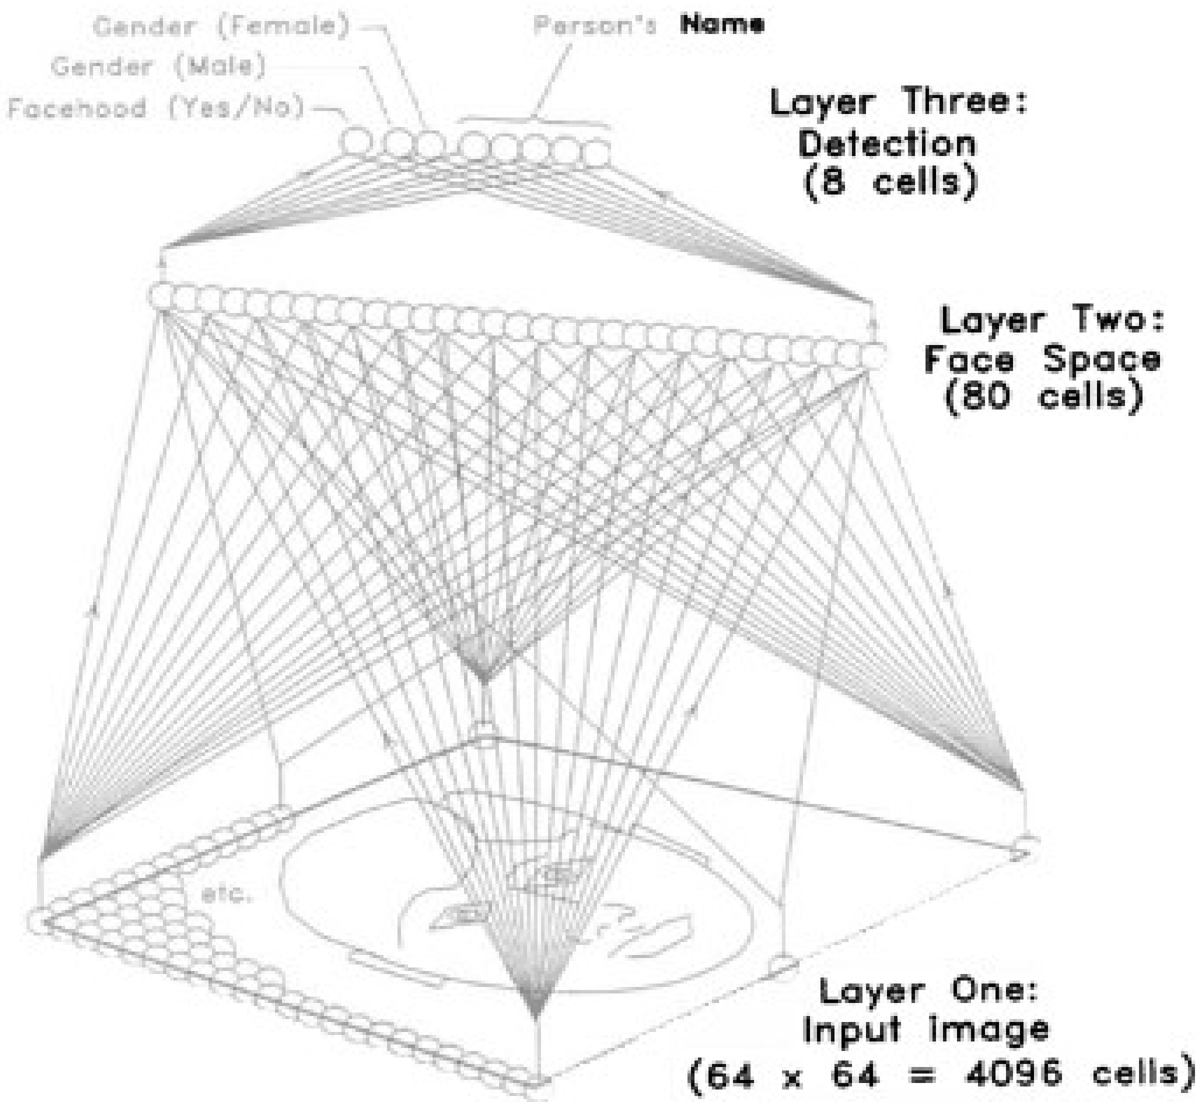

An artificial neural network for recognizing real faces.

Backpropagation

Eroarea dupa prima antrenare:

	(1,	0,	1,	.5,	1,	.5,	0,	0)	Desired output vector
−	(.23,	.8,	.39,	.2,	.03,	.19,	.66,	.96)	Actual output vector
=	(.77,	−.8,	.61,	.3,	.97,	.31,	−.66,	−.96)	Error vector
	(.59,	.64,	.37,	.09,	.94,	.09,	.44,	.94)	Squared-error vector

.5125 ← Mean of squared errors

- Se schimba valoarea fiecare legaturi (in plus sau in minus) si se calculeaza din nou eroarea → Se miscsoreaza eroarea.
- Se face la fel cu toate legaturile retelei si apoi la fiecare input-output

Performantele retelei:

- Recunoaste o foto din antrenament
- Recunoaste o persoana din antrenament in poza necunoscuta (= categorii (prototip) = clusturs) (98%)
- Discrimineaza femeie sau barbat din persoana/poza noua (81%)
- 1/5 lipsa din foto din antrenament (71%)

(Churchland ‘95, pp. 40- 49)

Alta retea: NETtalk (Sejnowski si Rosenberg 1986) 1000 cuvinte (78% corecte dupa antrenament)

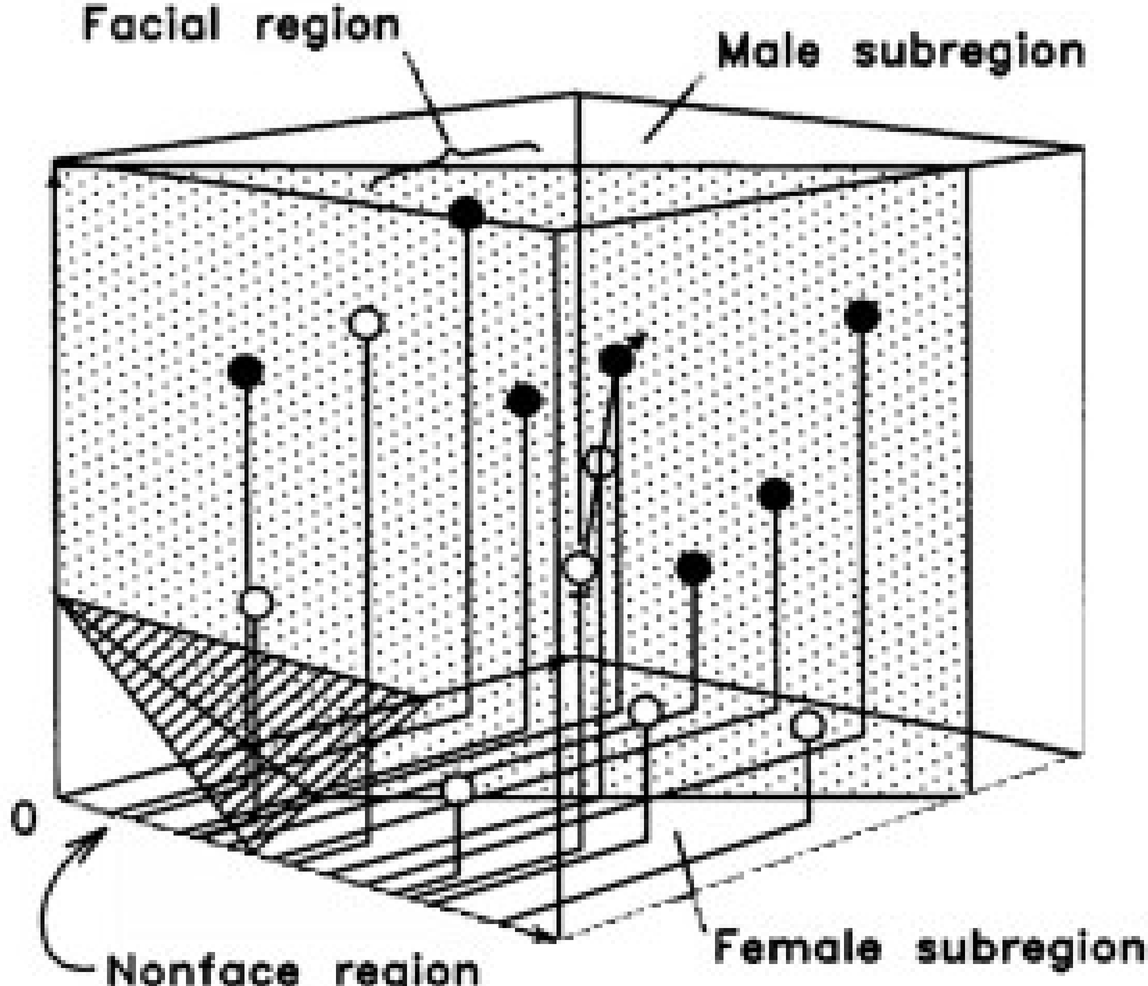

- Individual male face
- Individual female face
- Prototypical male face
- Prototypical female face
- Gender-ambiguous face

Layer 2 cu 3 unitati din 80 (hyper-space)

A. Memorie de superpozitie
("Superpositional storage")

- "2 Rs superpozitionate daca resursele ce reprezinta itemul 1 - **coextensiv** cu cele ale itemul 2"
- "Incodeaza informatia pt. item 2 prin **amendarea** setului de legaturi originale - **prezerva** - necesar *functional* (anumite patternuri input-output) sa reprezinte item 1 - **simultan** cu necesitatea functionala sa reprezinte item 2." (Clark)

Superpozitie - 2 trasaturi combinate:

(1) Folosirea Rs distribuite

(2) Folosirea regulii invatare - impune **metrica semantica** asupra Rs

- "Semantically related items - represented sintac-lly related (**partially overlapping**) patterns of activation." (Ex: pisica si caine vs. masina sau Elman, p. 91)
- Implicit–explicit sau potential-actual → "**Recreation"** of a R at t_1 (Horgan '97)

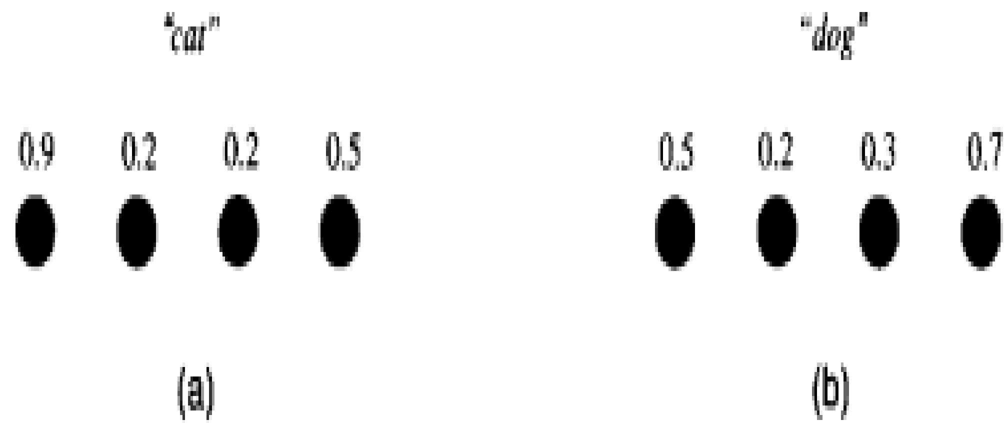

Examples of distributed representations. Both (a) and (b) illustrate different patterns of activation of the same set of four units. Activation values for individual units are shown as numbers above each unit. Note that the second unit has the same activation value in both representations; in order to know which concept is represented, one has to look at the entire set of nodes

- “Semantic similarity R-I contents - echoed

R-I vehicle.”→ **Prototype extraction** (category/concept) + Generalization

Prototype-knowledge (categorizare)

- Trasaturi comune ale inputurilor devin cele mai puternic asociate (= legaturi mutuale excitatorii) → Prin antrenament, reteaua “extrage” *tendinta centrala statistica* a trasaturilor comune = Prototipul (totusi reteaua nu “vede” acest prototip!!)

- Un exemplar nou plasat in functie de prototip (ex: cainele 3 picioare) → *Flexibilitate*
- In aceasi retea, *mai multe categorizari*!

- Reguli – antrenament/invatarea → "*Metrica semantica*" a Rs dobandite (Clark)

B. Sensitivitate de context intrinseca

(“Paradigma *subsimbolica*”, Smolensky)

Computationalism:

- *Transparenta semantica*
- LOT: Compositionalitate, sistematicitate, productivitate
- Rs simbolice: *Sintaxa* + *semantica combinatoriala*
- *C nu are sintaxa + semantica combinatoriala*
- Doar i**mplementare** sistem simbolic (Fodor & Pylyshyn)

vs. **Connectionism:**

- O *reprezentare* data de *activitatea patternului distribuit - contine subpatternuri pt. trasaturi* …
- In functie de context, reteaua reprezinta *instantieri* a unui astfel de item, care poate sa difere referitor la o trasatura sau mai multe.” (micro-trasaturi - subsimbolic!)

- Vecinii (clusterii) - reprezentati de Rs cu structuri similare (patternuri activate) - reflecta similaritatea semantica (metrica semantica (Clark)

- “Continutul elementelor in program subsimbolic nu recapituleaza direct conceptele” (Smolensky) si “**unitatile nu au continut semantic precum cuvintele in limbajului natural**” (Clark)

Diferentele de *activare* la fiecare unitate **oglindesc** detalii a diferitelor *functii mentale* in interactiune cu “contextul lumii reale”

- 1 stare conex. = 1 pattern activare (in spatiu activare) constituita din patt-nuri constituiente
- Pattern activare - NU e descompus in constituienti conceptuali (ca la simbolic) (Smolensky ‘88)

- Descompunerea conex. = Aproxim - Patternul complex (subp-uri constit.) nu poate fi definit, depinde de *context*

- Structura constituenta a subpatternurilor - puternic influentata de structura interna inclusa in system (Ex. Ceasca de cafea)

1) Simbolic: Compozitiona-te *concatenativa*

2) C: Compozitiona-te *functionala*

(van Gelder)

1) Concatenarea = “Unind constituienti succesivi fara alterarea lor” = Rs “tb. sa preserve tokens a unei expresii din constituienti si relatia lor secventiala”

2) Compozitionalitatea functionala = A avea R prin *recuperarea* ***partilor*** in anumite procese

Ex: O retea antrenata cu “caine”, alege una dintre starile interne *asociate* cu “caine” din antrenament →

Reteaua se bazeaza pe informatia contextuala: “caine” asociat cu

- “blana de lana protectoare” activeaza “poodle” feature-complex
- “maudled the burglar” activeaza “Rottweiler” feature-complex.

→ Reteaua construieste in functie de contextul curent in care este inclusa notiunea.

1) “*Symbolic*: Context of a symbol is manifest **around** it and consist of other symbols

(1) Subsymbolic: Context - symbol - manifest **inside** it, and consist of subsymbols” (Smolensky, 91)

Ex: “Ceasca cu cafea”

- “Cup with coffee”- upright container, burnt odor, brown liquid contacting porcelain

Extragem “cup” ramane set de micro-trasaturi care include itemi “*contextually biased*” (brown contacting …)

- “can with coffee” extragem “coffee” ramane “granules contacting tin”

→ No context-independent core R of coffee!

- Structura compoz. = *Sens aproximativ*

- Nets - 2 levels:

1) Mechanisme algoritmice de procesare

2) Interpretarea semantica

→ 2 niveluri descriere (Smolensky 1991)

“1 nivel: Procesele mentale *reprezentate* de nivelul de **descriere numerice** a unit-lor, legat-lor, ecuatii de evolutie activarilor (NU interpretare semantica) (Smolensk*y)*

- “2 nivel: Activitati la *nivel larg* permit **interpretari,** dar patt-le astfel fixate nu sunt descrieri clare a procesarii”
- Si asta pt ca patter-le interpretate (grupul de vectori asociate cu “cafea”) “trimit la acele diferente microscopice care fac diferenta in procesare. (ex. “caine”)
- “Metrica semantica a sistemului impune o similaritate pt. *continut* [inteles] acolo unde exista o similaritate pt. v*ehicul* (= patternuri similare).” (Clark)

- Fodor and McLaughlin ‘90, McLaughlin ‘93

vs. Smolensky, Elman

Retelele recurente (SRN)

- Memoria termen scurt (timpul) (Elman ’90, ’91, ’93) (Elman, p. 75 sau 81): **Predictia**
- Elman ’90: Predictie cuvinte/litere
- Reteaua prezice nu un cuvant ci toate cuvintele care au aceasi functie gramaticala (subiect, predicat, etc.)

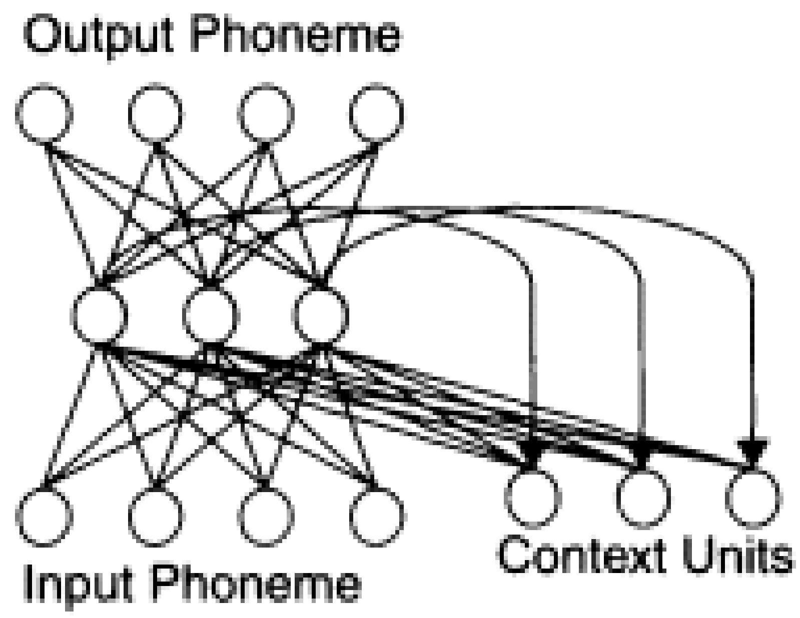

FIGURE 3.2 Simple recurrent network used by Elman (1990) to predict the next phoneme in a string of phonemes. The input to the network represents the current phoneme. The output represents the network's best guess as to the next phoneme in the sequence. The context units provide the network with a dynamic memory to encode sequence information.

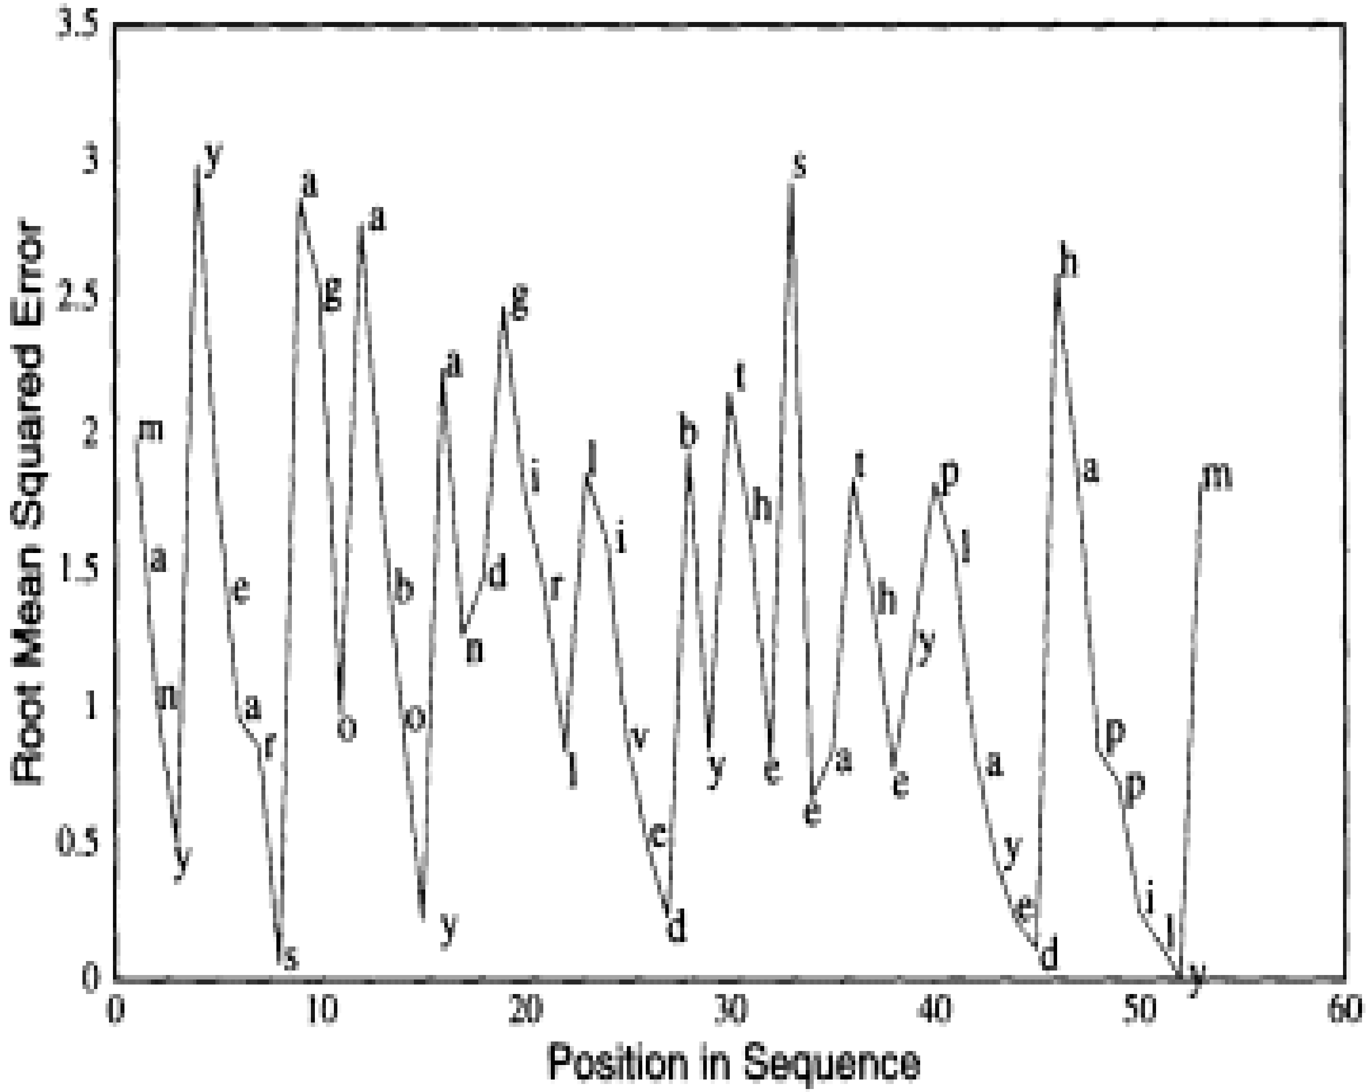

The root mean squared error for predicting individual letters in the string:

Manyyearsagoaboyandgirllivedbytheseatheyplayedhappily.

- 10.000 propozitii simple (2-3 cuvinte) avand 15 tipuri diferite de combinari

- Reteaua invata sa grupeze, de exemplu, obiectele animate doar pt ca erau distribuite similar in “training corpus”. (B&A, p. 182)

SRN - Prezice cuvintele succesive:

- Inputul: 1 cuvant (sau litera)
- Outputul: Prezicerea cuvantului urmator
- Backprop: Adjustarea fiecarei legaturi la fiecare eroare, apoi urmatorul cuvant
- Proces repetat (mii propozitii)
- Unitatile ascunse: hyper-spatiu 150-dimens
- “Retea – invata sa reprezinte cuvinte care se comporta in moduri similare (= proprietati distribuite similare) cu vectori care sunt apropiati in spatiu intern reprezentational”

- Time in language-sentences, 2 features:

(1) Processed sequentially in time

(2) Exhibit long-distance dependencies - form of one word - depend on another that is located at an *indeterminate distance*." (Verbs agree subjects, a relative clause - between subject and verb) (Bechtel &A)

- Net - incorporate such relationships *without explicit Rs of linguistic structures*
- Hidden units – recycle information over multiple steps → *Abstract Rs of time*!

Spatiu = "Hierarchical **clustering** tree" a patt-lor activare unit. ascunse (cuvinte)

"Capturing hidden unit activation pattern corresponding each word, measuring distance betw. each pattern and every other pattern. These pattern distances = Euclidian distance betw. vectors in activation space" → VERBS, animates, NOUNS, inanimates (Elman, p. 96)

(nota la Elman, p. 96: sunt mai multi vectori diferiti pentru "dragon", "tree" arata o medie a fiecarui vector format in functie de context!)

↓

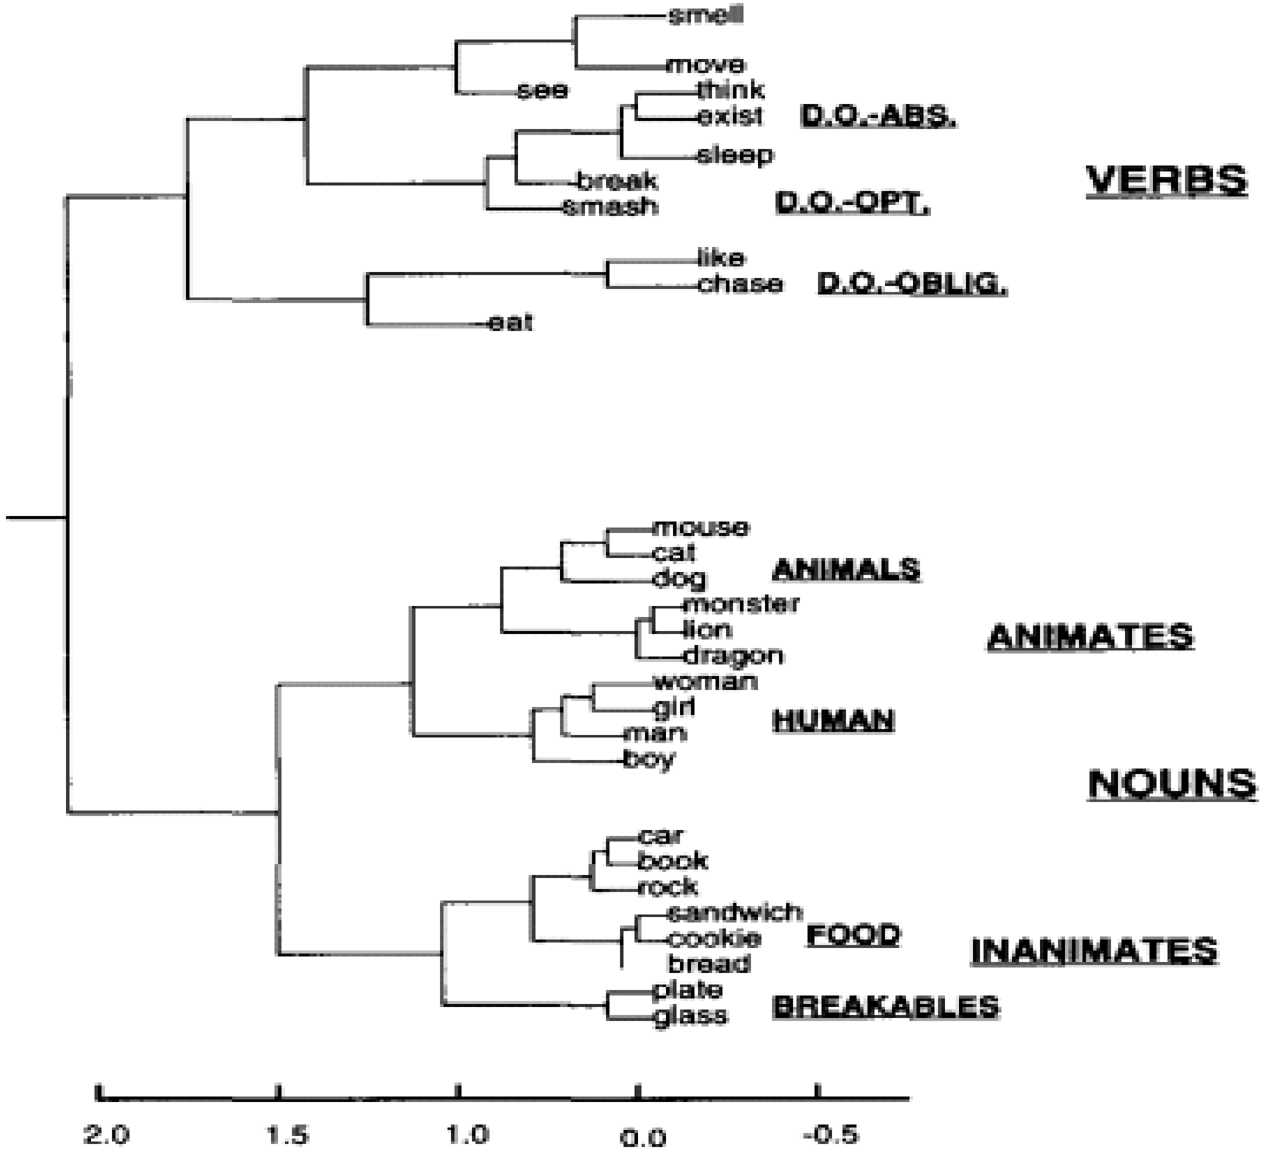

The network learns distributed representations for each word which reflects its similarity to other words. Words and groups of words which are similar are close in activation space, and close in position in the tree.

- Context-sensitivity: **contextul e parte a R unui cuvant**. Paternul activat reflecta si cuvantul de dinainte! Fiecare aparitie a acelasi cuvant depinde de context. (Ex. John1, John2, etc.)
- “Tokens of same type are all spatially proximal, and closer to each other than to tokens of any other type.” Dar nu avem reprezentare a unui cuvant in isolatie.
- “dragon”-as-subject and -as-object → Relative positions of -two sub-classes – identical across nouns →

“The internal R-on of a word reflects not only its core identity, but provides a grammatically interpretable context-sensitive shading” (Elman, p. 97)

- No separate stage of lexical retrieval
- No Rs of words in isolation
- Internal states following input of a word - reflect input taken together with prior state

- Rs = Not propositional; their information content changes over time depending on current task (Elman '96)

- SRN: **No semantics, net learns to group encoding of animate objects together only because they were distributed similarly in training corpus!** (Bechtel & Abrahamsen '02)

- Rumelhart, McClelland (‘86) – *Prediction of past tense English verbs*”: Regular vs. irregular verbs
- 2 layers → Problema: Past tense = Non-linear

1) **Simbolic**: 2 mecanisme (vb-s regulate/ neregulate

2) **C:** 1 mechanism = Retea de conexiuni pt. vb-s regulate/neregulate

- Retele - 1 nivel rezolva probleme liniar separabile
- Past tense = Problema nonlineara

- Diferentieri la nivel de comportament nu necesita diferentieri la nivel de mecanism.
- Verbele regulate si iregulate se comporta diferit chiar daca sunt reprezentate si proceseaza similar in acelasi mechanism. (Elman, p. 139)
- Plunkett-Marchman (1991): “U-shape form” in dezvoltare

- No explicit rule:

"… the mechanisms that process language and make judgments of grammaticality are constructed in such a way that their performance is characterizable by rules, but that the rules themselves are not written in explicit forma anywhere in the mechanism." (Rumelhart and McClelland 1986)

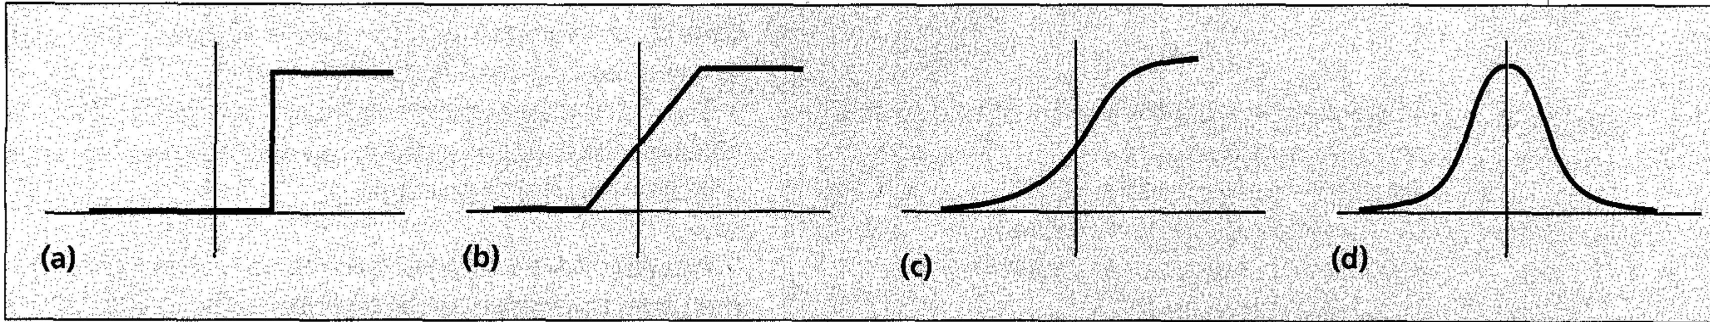

Figure 3. Different types of activation functions: (a) threshold, (b) piecewise linear, (c) sigmoid, and (d) Gaussian.

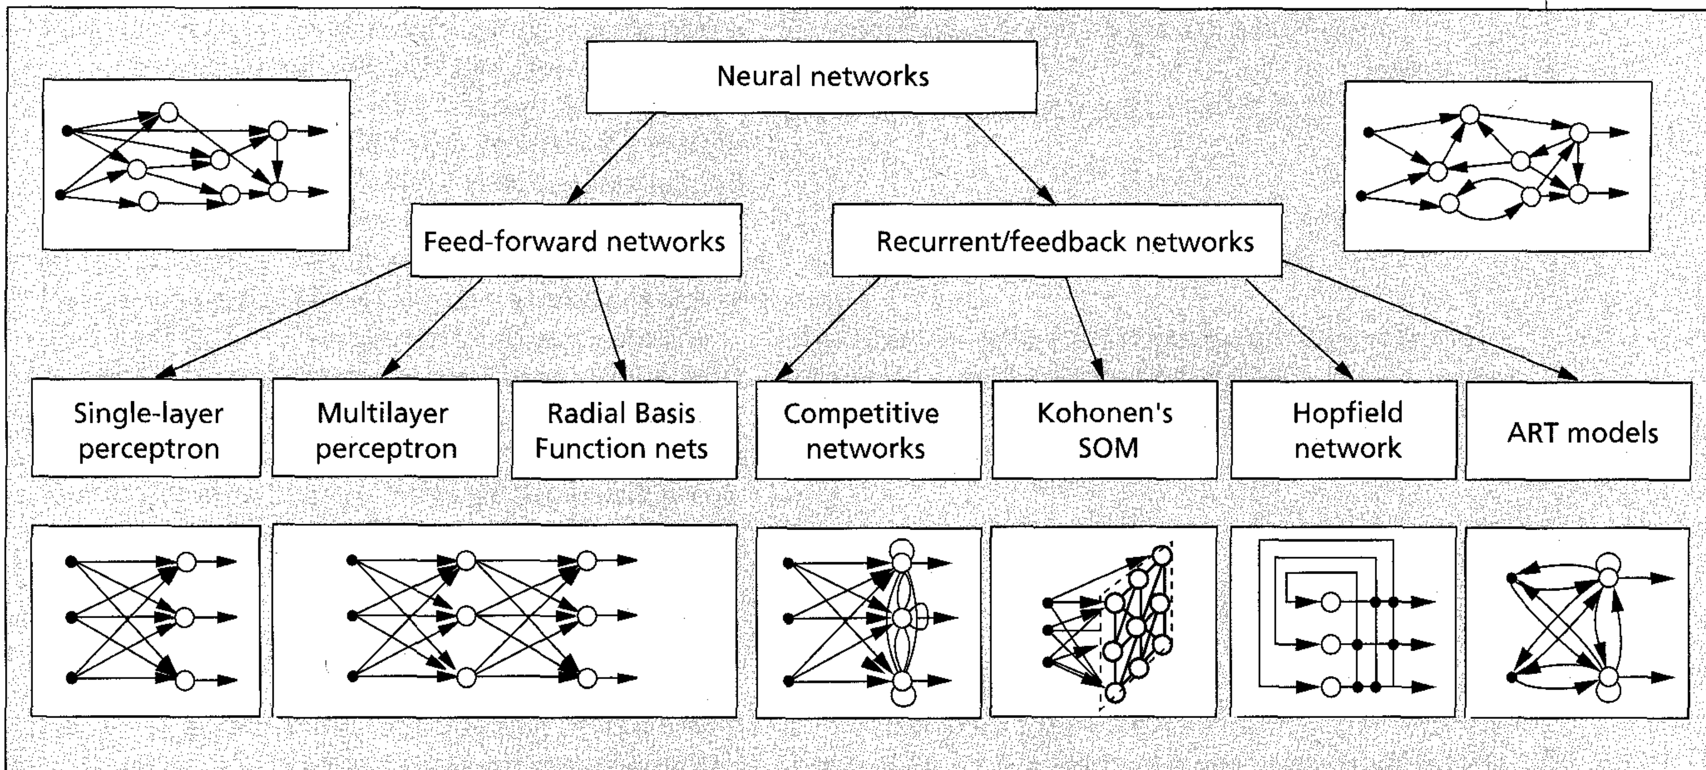

Figure 4. A taxonomy of feed-forward and recurrent/feedback network architectures.

Artificial Neural Networks: A tutorial by A. K. Jain, J. Mao, K. Mohiuddin, 1996

Dezbateri

- Pinker & Prince (‘88) – Nets: Good making associations + matching patterns

- Limitations in mastering general rules such as formation of regular past tense. (Garson or Bechtel & Abrahamsen)

Sistematicitatea (S) (Fodor & P)

- Fodor’s LOT → Rs simbolice au sintaxa si semantica combinatoriala
- Regulile care guverneaza procesele sunt sintactice: se aplica la forma si nu la continut!
- Semantica reflecta sintaxa: semantica intregului depinde de a partilor
- Fodor si Pylyshyn: C lipseste aceasta sintaxa si semantica combinatoriala
- Nefiind combinatorial sintactic nu sunt nici semantic

→ Retelele nu au productivitate, sistematicitate

Contrareplici

- Smolensky (90); Shastri si Ajjanagadde (93) (cu binding); Dyer (91) cu ierarhie de niveluri de cautare (Bechtel si Abrahamsen 2002)
- Compozitionalitate concatenativa si functionala; Pollack's RAAM nets, Elman recurrent nets