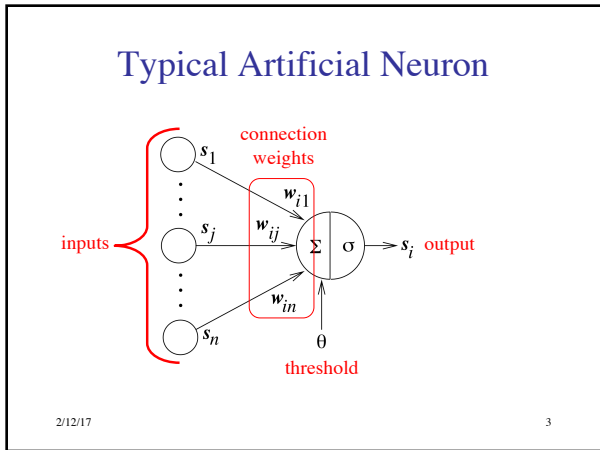


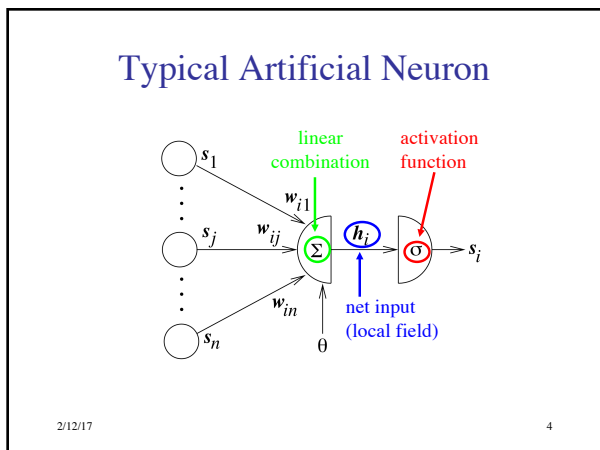
III. Recurrent Neural Networks

2/12/17 1

A. The Hopfield Network

2/12/17 2





Equations

Net input:
$$h_i = \left(\sum_{j=1}^n w_{ij} s_j \right) - \theta$$

$$\mathbf{h} = \mathbf{W}\mathbf{s} - \theta$$

New neural state:
$$s'_i = \sigma(h_i)$$

$$\mathbf{s}' = \sigma(\mathbf{h})$$

2/12/17 5

- ### Hopfield Network
- Symmetric weights: $w_{ij} = w_{ji}$
 - No self-action: $w_{ii} = 0$
 - Zero threshold: $\theta = 0$
 - Bipolar states: $s_i \in \{-1, +1\}$
 - Discontinuous bipolar activation function:
- $$\sigma(h) = \text{sgn}(h) = \begin{cases} -1, & h < 0 \\ +1, & h > 0 \end{cases}$$
- 2/12/17 6

What to do about $h = 0$?

- There are several options:
 - $\sigma(0) = +1$
 - $\sigma(0) = -1$
 - $\sigma(0) = -1$ or $+1$ with equal probability
 - $h_i = 0 \Rightarrow$ no state change ($s'_i = s_i$)
- Not much difference, but be consistent
- Last option is slightly preferable, since symmetric

2/12/17 7

Positive Coupling

- Positive *sense* (sign)
- Large *strength*

2/12/17 8

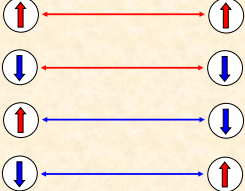
Negative Coupling

- Negative *sense* (sign)
- Large *strength*

2/12/17 9

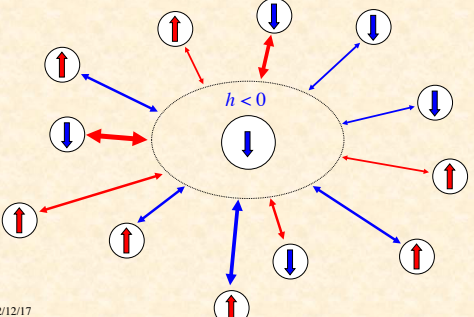
Weak Coupling

- Either *sense* (sign)
- Little *strength*



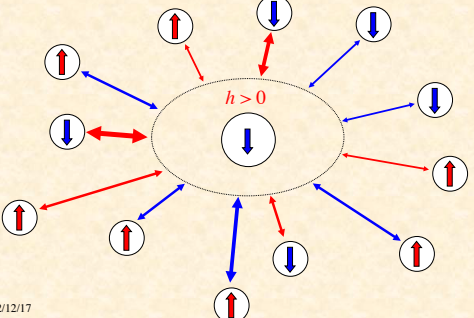
2/12/17 10

State = -1 & Local Field < 0

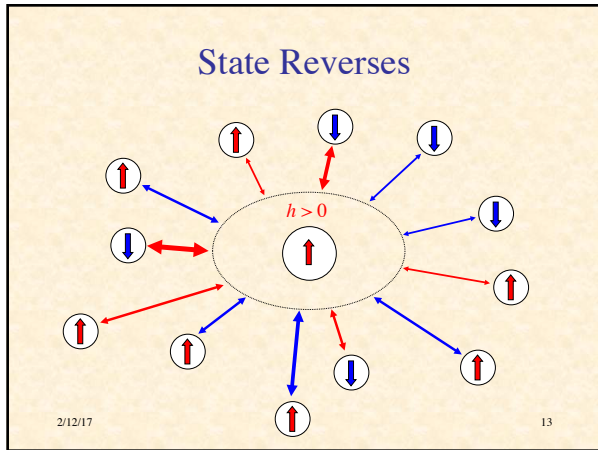


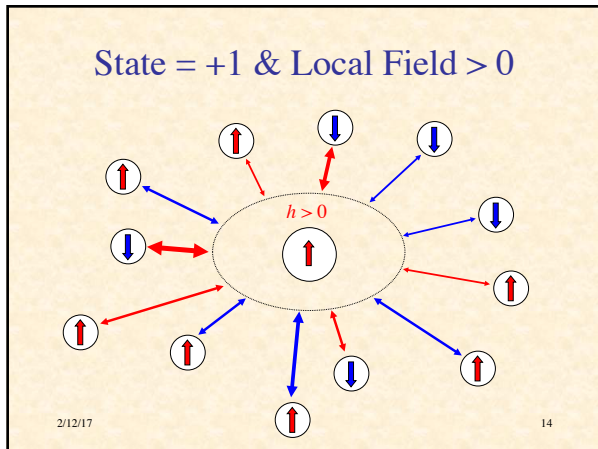
2/12/17 11

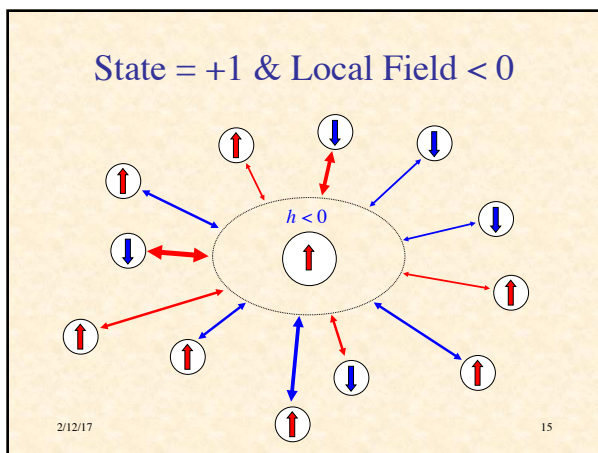
State = -1 & Local Field > 0

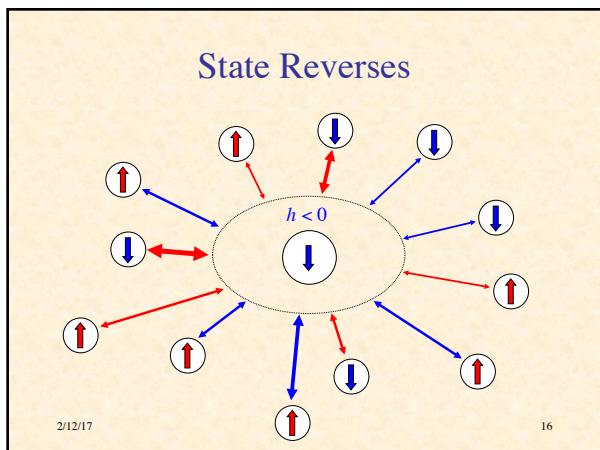


2/12/17 12









NetLogo Demonstration of Hopfield State Updating

[Run Hopfield-update.nlogo](#)

2/12/17 17

Hopfield Net as Soft Constraint Satisfaction System

- States of neurons as yes/no decisions
- Weights represent *soft constraints* between decisions
 - *hard* constraints *must* be respected
 - *soft* constraints have *degrees* of importance
- Decisions change to better respect constraints
- Is there an optimal set of decisions that best respects all constraints?

2/12/17 18

Demonstration of Hopfield Net Dynamics I
[Run Hopfield-dynamics.nlogo](#)

2/12/17 19

Convergence

- Does such a system converge to a stable state?
- Under what conditions does it converge?
- There is a sense in which each step relaxes the “tension” in the system
- But could a relaxation of one neuron lead to greater tension in other places?

2/12/17 20

Quantifying “Tension”

- If $w_{ij} > 0$, then s_i and s_j want to have the same sign ($s_i s_j = +1$)
- If $w_{ij} < 0$, then s_i and s_j want to have opposite signs ($s_i s_j = -1$)
- If $w_{ij} = 0$, their signs are independent
- Strength of interaction varies with $|w_{ij}|$
- Define disharmony (“tension”) D_{ij} between neurons i and j :
 $D_{ij} = -s_i w_{ij} s_j$
 $D_{ij} < 0 \Rightarrow$ they are happy
 $D_{ij} > 0 \Rightarrow$ they are unhappy

2/12/17 21

Total Energy of System

The “energy” of the system is the total “tension” (disharmony) in it:

$$\begin{aligned}
 E\{\mathbf{s}\} &= \sum_{\langle ij \rangle} D_{ij} = - \sum_{\langle ij \rangle} s_i w_{ij} s_j \\
 &= -\frac{1}{2} \sum_i \sum_{j \neq i} s_i w_{ij} s_j \\
 &= -\frac{1}{2} \sum_i \sum_j s_i w_{ij} s_j \\
 &= -\frac{1}{2} \mathbf{s}^T \mathbf{W} \mathbf{s}
 \end{aligned}$$

2/12/17 22

Review of Some Vector Notation

$\mathbf{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix} = (x_1 \ \cdots \ x_n)^T$ (column vectors)

$\mathbf{x}^T \mathbf{y} = \sum_{i=1}^n x_i y_i = \mathbf{x} \cdot \mathbf{y}$ (inner product)

$\mathbf{x} \mathbf{y}^T = \begin{pmatrix} x_1 y_1 & \cdots & x_1 y_n \\ \vdots & \ddots & \vdots \\ x_m y_1 & \cdots & x_m y_n \end{pmatrix}$ (outer product)

$\mathbf{x}^T \mathbf{M} \mathbf{y} = \sum_{i=1}^m \sum_{j=1}^n x_i M_{ij} y_j$ (quadratic form)

2/12/17 23

Another View of Energy

The energy measures the disharmony of the neurons’ states with their local fields (i.e. of opposite sign):

$$\begin{aligned}
 E\{\mathbf{s}\} &= -\frac{1}{2} \sum_i \sum_j s_i w_{ij} s_j \\
 &= -\frac{1}{2} \sum_i s_i \sum_j w_{ij} s_j \\
 &= -\frac{1}{2} \sum_i s_i h_i \\
 &= -\frac{1}{2} \mathbf{s}^T \mathbf{h}
 \end{aligned}$$

2/12/17 24

Do State Changes Decrease Energy?

- Suppose that neuron k changes state
- Change of energy:

$$\begin{aligned} \Delta E &= E\{\mathbf{s}'\} - E\{\mathbf{s}\} \\ &= -\sum_{(ij)} s'_i w_{ij} s'_j + \sum_{(ij)} s_i w_{ij} s_j \\ &= -\sum_{j \neq k} s'_k w_{kj} s_j + \sum_{j \neq k} s_k w_{kj} s_j \\ &= -(s'_k - s_k) \sum_{j \neq k} w_{kj} s_j \\ &= -\Delta s_k h_k \\ &< 0 \end{aligned}$$

2/12/17 25

Energy Does Not Increase

- In each step in which a neuron is considered for update:
 $E\{\mathbf{s}(t+1)\} - E\{\mathbf{s}(t)\} \leq 0$
- Energy cannot increase
- Energy decreases if any neuron changes
- Must it stop?

2/12/17 26

Proof of Convergence in Finite Time

- There is a minimum possible energy:
 - The number of possible states $\mathbf{s} \in \{-1, +1\}^n$ is finite
 - Hence $E_{\min} = \min \{E(\mathbf{s}) \mid \mathbf{s} \in \{\pm 1\}^n\}$ exists
- Must reach in a finite number of steps because only finite number of states

2/12/17 27

Conclusion

- If we do asynchronous updating, the Hopfield net must reach a stable, minimum energy state in a finite number of updates
- This does not imply that it is a global minimum

2/12/17 28

Lyapunov Functions

- A way of showing the convergence of discrete- or continuous-time dynamical systems
- For discrete-time system:
 - need a Lyapunov function E ("energy" of the state)
 - E is bounded below ($E\{s\} > E_{\min}$)
 - $\Delta E < (\Delta E)_{\max} \leq 0$ (energy decreases a certain minimum amount each step)
 - then the system will converge in finite time
- Problem: finding a suitable Lyapunov function

2/12/17 29

Example Limit Cycle with Synchronous Updating

2/12/17 30

The Hopfield Energy Function is Even

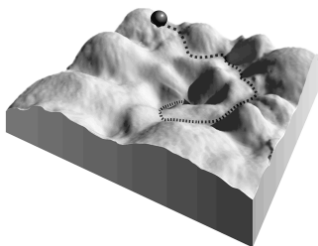
- A function f is **odd** if $f(-x) = -f(x)$, for all x
- A function f is **even** if $f(-x) = f(x)$, for all x
- Observe:

$$E\{-s\} = -\frac{1}{2}(-s)^T W(-s) = -\frac{1}{2}s^T Ws = E\{s\}$$

2/12/17

31

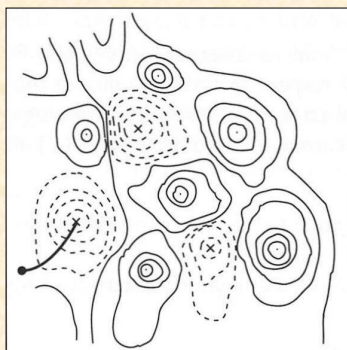
Conceptual Picture of Descent on Energy Surface



2/12/17

(fig. from Solé & Goodwin)

32

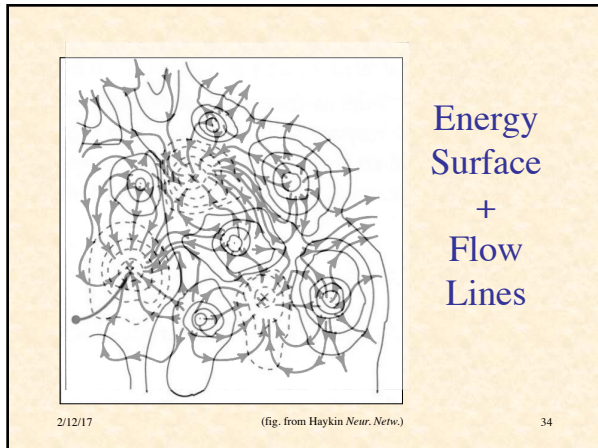


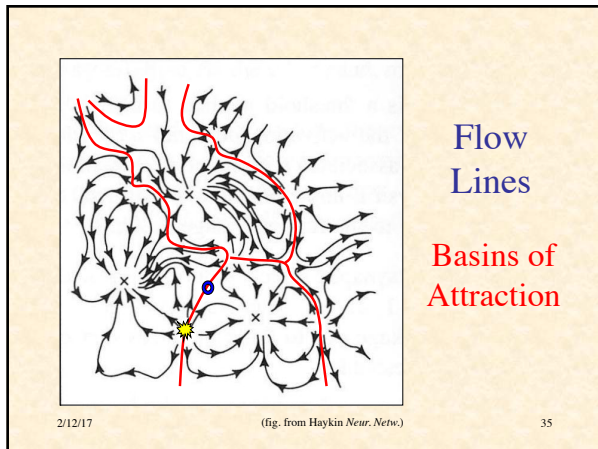
Energy Surface

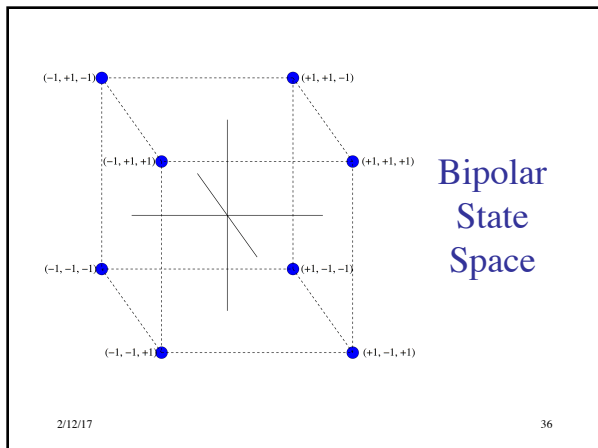
2/12/17

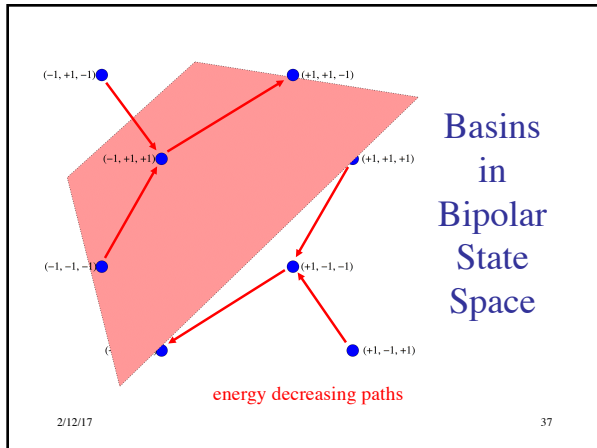
(fig. from Haykin *Neur. Netw.*)

33





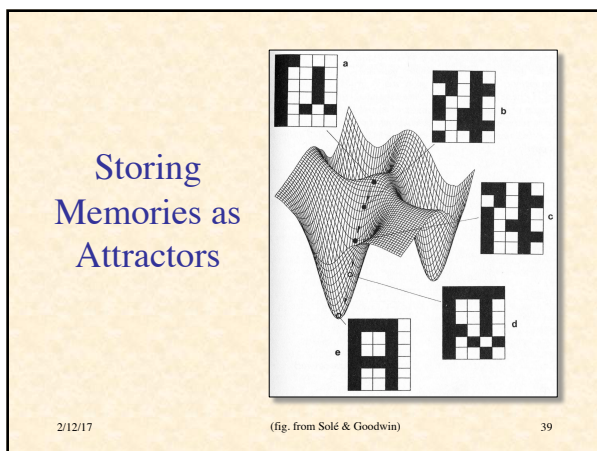





Demonstration of Hopfield Net Dynamics II

[Run initialized Hopfield.nlogo](#)

2/12/17 38




Example of Pattern Restoration




2/12/17 (fig. from Arbib 1995) 40

Example of Pattern Restoration




2/12/17 (fig. from Arbib 1995) 41

Example of Pattern Restoration




2/12/17 (fig. from Arbib 1995) 42

Example of
Pattern
Restoration




2/12/17 (fig. from Arbib 1995) 43

Example of
Pattern
Restoration



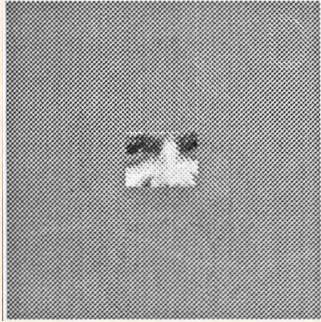
2/12/17 (fig. from Arbib 1995) 44

Example of
Pattern
Completion




2/12/17 (fig. from Arbib 1995) 45

Example of Pattern Completion




2/12/17 (fig. from Arbib 1995) 46

Example of Pattern Completion




2/12/17 (fig. from Arbib 1995) 47

Example of Pattern Completion




2/12/17 (fig. from Arbib 1995) 48

Example of
Pattern
Completion




2/12/17 (fig. from Arbib 1995) 49

Example of
Association




2/12/17 (fig. from Arbib 1995) 50

Example of
Association




2/12/17 (fig. from Arbib 1995) 51

Example of Association




2/12/17 (fig. from Arbib 1995) 52

Example of Association



2/12/17 (fig. from Arbib 1995) 53

Example of Association



2/12/17 (fig. from Arbib 1995) 54

Applications of Hopfield Memory

- Pattern restoration
- Pattern completion
- Pattern generalization
- Pattern association

2/12/17 55

Hopfield Net for Optimization and for Associative Memory

- For optimization:
 - we know the weights (couplings)
 - we want to know the minima (solutions)
- For associative memory:
 - we know the minima (retrieval states)
 - we want to know the weights

2/12/17 56

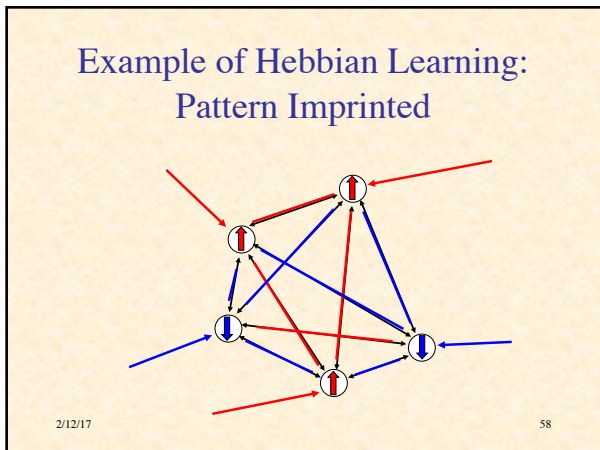
Hebb's Rule

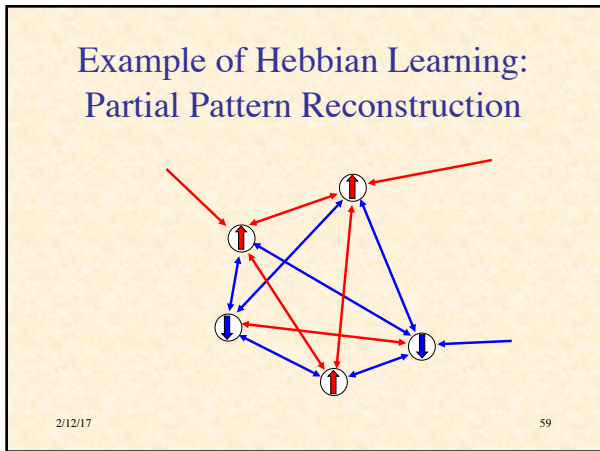
“When an axon of cell A is near enough to excite a cell B and repeatedly or persistently takes part in firing it, some growth or metabolic change takes place in one or both cells such that A's efficiency, as one of the cells firing B, is increased.”

—Donald Hebb (*The Organization of Behavior*, 1949, p. 62)

“Neurons that fire together, wire together”

2/12/17 57





Mathematical Model of Hebbian Learning for One Pattern

Let $W_{ij} = \begin{cases} x_i x_j, & \text{if } i \neq j \\ 0, & \text{if } i = j \end{cases}$

Since $x_i x_i = x_i^2 = 1$, $\mathbf{W} = \mathbf{xx}^T - \mathbf{I}$

For simplicity, we will include self-coupling:

$$\mathbf{W} = \mathbf{xx}^T$$

2/12/17 60

A Single Imprinted Pattern is a Stable State

- Suppose $\mathbf{W} = \mathbf{x}\mathbf{x}^T$
- Then $\mathbf{h} = \mathbf{W}\mathbf{x} = \mathbf{x}\mathbf{x}^T\mathbf{x} = n\mathbf{x}$
since $\mathbf{x}^T\mathbf{x} = \sum_{i=1}^n x_i^2 = \sum_{i=1}^n (\pm 1)^2 = n$
- Hence, if initial state is $\mathbf{s} = \mathbf{x}$, then new state is $\mathbf{s}' = \text{sgn}(n\mathbf{x}) = \mathbf{x}$
- For this reason, scale \mathbf{W} by $1/n$
- May be other stable states (e.g., $-\mathbf{x}$)

2/12/17

61

Questions

- How big is the basin of attraction of the imprinted pattern?
- How many patterns can be imprinted?
- Are there unneeded *spurious* stable states?
- These issues will be addressed in the context of multiple imprinted patterns

2/12/17

62

Imprinting Multiple Patterns

- Let $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^p$ be patterns to be imprinted
- Define the sum-of-outer-products matrix:

$$W_{ij} = \frac{1}{n} \sum_{k=1}^p x_i^k x_j^k$$

$$\mathbf{W} = \frac{1}{n} \sum_{k=1}^p \mathbf{x}^k (\mathbf{x}^k)^T$$

2/12/17

63

Definition of Covariance

Consider samples $(x^1, y^1), (x^2, y^2), \dots, (x^N, y^N)$
 Let $\bar{x} = \langle x^k \rangle$ and $\bar{y} = \langle y^k \rangle$
 Covariance of x and y values:

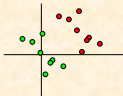
$$C_{xy} = \langle (x^k - \bar{x})(y^k - \bar{y}) \rangle$$

$$= \langle x^k y^k - \bar{x} y^k - x^k \bar{y} + \bar{x} \bar{y} \rangle$$

$$= \langle x^k y^k \rangle - \bar{x} \langle y^k \rangle - \langle x^k \rangle \bar{y} + \bar{x} \bar{y}$$

$$= \langle x^k y^k \rangle - \bar{x} \bar{y} - \bar{x} \bar{y} + \bar{x} \bar{y}$$

$$C_{xy} = \langle x^k y^k \rangle - \bar{x} \bar{y}$$



2/12/17 64

Weights & the Covariance Matrix

Sample pattern vectors: $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^p$
 Covariance of i^{th} and j^{th} components:

$$C_{ij} = \langle x_i^k x_j^k \rangle - \bar{x}_i \bar{x}_j$$

If $\forall i: \bar{x}_i = 0$ (± 1 equally likely in all positions):

$$C_{ij} = \langle x_i^k x_j^k \rangle = \frac{1}{p} \sum_{k=1}^p x_i^k x_j^k$$

$\therefore n\mathbf{W} = p\mathbf{C}$

2/12/17 65

Characteristics of Hopfield Memory

- Distributed (“holographic”)
 - every pattern is stored in every location (weight)
- Robust
 - correct retrieval in spite of noise or error in patterns
 - correct operation in spite of considerable weight damage or noise

2/12/17 66

Demonstration of Hopfield Net

[Run Malasri Hopfield Demo](#)

2/12/17 67

Stability of Imprinted Memories

- Suppose the state is one of the imprinted patterns \mathbf{x}^m
- Then: $\mathbf{h} = \mathbf{W}\mathbf{x}^m = \left[\frac{1}{n} \sum_k \mathbf{x}^k (\mathbf{x}^k)^T \right] \mathbf{x}^m$

$$= \frac{1}{n} \sum_k \mathbf{x}^k (\mathbf{x}^k)^T \mathbf{x}^m$$

$$= \frac{1}{n} \mathbf{x}^m (\mathbf{x}^m)^T \mathbf{x}^m + \frac{1}{n} \sum_{k \neq m} \mathbf{x}^k (\mathbf{x}^k)^T \mathbf{x}^m$$

$$= \mathbf{x}^m + \frac{1}{n} \sum_{k \neq m} (\mathbf{x}^k \cdot \mathbf{x}^m) \mathbf{x}^k$$

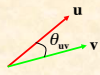
2/12/17 68

Interpretation of Inner Products

- $\mathbf{x}^k \cdot \mathbf{x}^m = n$ if they are identical
 - highly correlated
- $\mathbf{x}^k \cdot \mathbf{x}^m = -n$ if they are complementary
 - highly correlated (reversed)
- $\mathbf{x}^k \cdot \mathbf{x}^m = 0$ if they are orthogonal
 - largely uncorrelated
- $\mathbf{x}^k \cdot \mathbf{x}^m$ measures the *crossstalk* between patterns k and m

2/12/17 69

Cosines and Inner products

$\mathbf{u} \cdot \mathbf{v} = \|\mathbf{u}\| \|\mathbf{v}\| \cos \theta_{\mathbf{uv}}$


If \mathbf{u} is bipolar, then $\|\mathbf{u}\|^2 = \mathbf{u} \cdot \mathbf{u} = n$

Hence, $\mathbf{u} \cdot \mathbf{v} = \sqrt{n} \sqrt{n} \cos \theta_{\mathbf{uv}} = n \cos \theta_{\mathbf{uv}}$

Hence $\mathbf{h} = \mathbf{x}^m + \sum_{k \neq m} \mathbf{x}^k \cos \theta_{km}$

2/12/17 70

Conditions for Stability

Stability of entire pattern :

$$\mathbf{x}^m = \text{sgn} \left(\mathbf{x}^m + \sum_{k \neq m} \mathbf{x}^k \cos \theta_{km} \right)$$

Stability of a single bit :

$$x_i^m = \text{sgn} \left(x_i^m + \sum_{k \neq m} x_i^k \cos \theta_{km} \right)$$

2/12/17 71

Sufficient Conditions for Instability (Case 1)

Suppose $x_i^m = -1$. Then unstable if :

$$(-1) + \sum_{k \neq m} x_i^k \cos \theta_{km} > 0$$

$$\sum_{k \neq m} x_i^k \cos \theta_{km} > 1$$

2/12/17 72

Sufficient Conditions for Instability (Case 2)

Suppose $x_i^m = +1$. Then unstable if :

$$(+1) + \sum_{k \neq m} x_i^k \cos \theta_{km} < 0$$

$$\sum_{k \neq m} x_i^k \cos \theta_{km} < -1$$

2/12/17 73

Sufficient Conditions for Stability

$$\left| \sum_{k \neq m} x_i^k \cos \theta_{km} \right| \leq 1$$

The crosstalk with the sought pattern must be sufficiently small

2/12/17 74

Capacity of Hopfield Memory

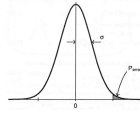
- Depends on the patterns imprinted
- If orthogonal, $p_{\max} = n$
 - weight matrix is identity matrix
 - hence every state is stable \Rightarrow trivial basins
- So $p_{\max} < n$
- Let **load parameter** $\alpha = p / n$

2/12/17 75

equations

Single Bit Stability Analysis

- For simplicity, suppose \mathbf{x}^k are random
- Then $\mathbf{x}^k \cdot \mathbf{x}^m$ are sums of n random ± 1
 - binomial distribution \approx Gaussian
 - in range $-n, \dots, +n$
 - with mean $\mu = 0$
 - and variance $\sigma^2 = n$
- Probability sum $> t$:



$$\frac{1}{2} \left[1 - \operatorname{erf} \left(\frac{t}{\sqrt{2n}} \right) \right]$$

[See "Review of Gaussian (Normal) Distributions" on course website]

2/12/17

76

Approximation of Probability

Let crosstalk $C_i^m = \frac{1}{n} \sum_{k \neq m} x_i^k (\mathbf{x}^k \cdot \mathbf{x}^m)$

We want $\Pr\{C_i^m > 1\} = \Pr\{nC_i^m > n\}$

Note: $nC_i^m = \sum_{k=1}^p \sum_{j=1}^n x_i^k x_j^k x_j^m$

A sum of $n(p-1) \approx np$ random ± 1 s

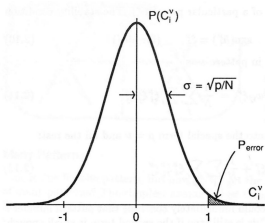
Variance $\sigma^2 = np$

2/12/17

77

Probability of Bit Instability

$$\begin{aligned} \Pr\{nC_i^m > n\} &= \frac{1}{2} \left[1 - \operatorname{erf} \left(\frac{n}{\sqrt{2np}} \right) \right] \\ &= \frac{1}{2} \left[1 - \operatorname{erf} \left(\sqrt{n/2p} \right) \right] \\ &= \frac{1}{2} \left[1 - \operatorname{erf} \left(\sqrt{1/2\alpha} \right) \right] \end{aligned}$$



2/12/17 (fig. from Hertz & al. *Intr. Theory Neur. Comp.*)

78

Tabulated Probability of Single-Bit Instability

P_{error}	α
0.1%	0.105
0.36%	0.138
1%	0.185
5%	0.37
10%	0.61

2/12/17 (table from Hertz & al. *Intr. Theory Neur. Comp.*) 79

Orthogonality of Random Bipolar Vectors of High Dimension

- 99.99% probability of being within 4σ of mean $|\mathbf{u} \cdot \mathbf{v}| < 4\sigma$
- It is 99.99% probable that random n -dimensional vectors will be within $\epsilon = 4/\sqrt{n}$ orthogonal iff $\|\mathbf{u}\| \|\mathbf{v}\| |\cos\theta| < 4\sqrt{n}$
- $\epsilon = 4\%$ for $n = 10,000$ iff $n|\cos\theta| < 4\sqrt{n}$
- Probability of being less orthogonal than ϵ decreases exponentially with n iff $|\cos\theta| < 4/\sqrt{n} = \epsilon$
- The brain gets approximate orthogonality by assigning random high-dimensional vectors $\Pr\{|\cos\theta| > \epsilon\} = \text{erfc}\left(\frac{\epsilon\sqrt{n}}{\sqrt{2}}\right)$

$\approx \frac{1}{6}\exp(-\epsilon^2 n / 2) + \frac{1}{2}\exp(-2\epsilon^2 n / 3)$

2/12/17 80

Spurious Attractors

- **Mixture states:**
 - sums or differences of odd numbers of retrieval states
 - number increases combinatorially with p
 - shallower, smaller [basins](#)
 - basins of mixtures swamp basins of retrieval states \Rightarrow overload
 - useful as combinatorial generalizations?
 - self-coupling generates spurious attractors
- **Spin-glass states:**
 - not correlated with any finite number of imprinted patterns
 - occur beyond overload because weights effectively random

2/12/17 81

Basins of Mixture States

$x_i^{\text{mix}} = \text{sgn}(x_i^{k_1} + x_i^{k_2} + x_i^{k_3})$

2/12/17 82 ↷

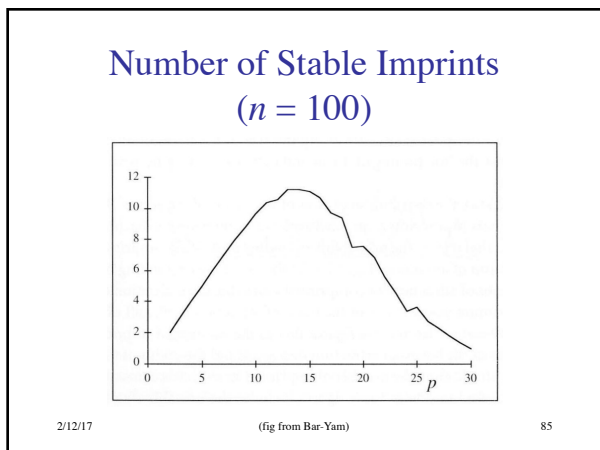
Run Hopfield-Capacity Test

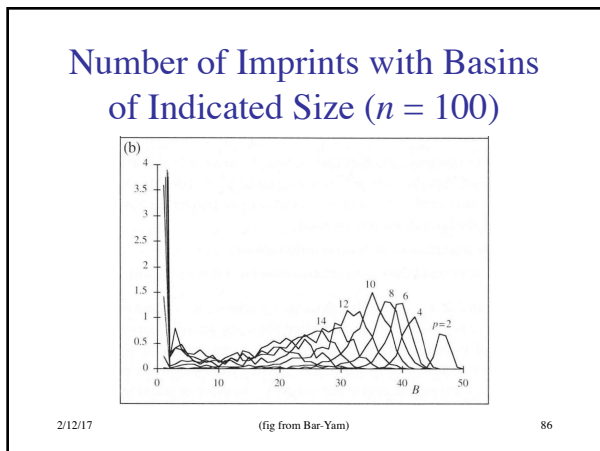
2/12/17 83

Fraction of Unstable Imprints ($n = 100$)

p	Fraction of Unstable Imprints
0	0.00
5	0.00
10	0.10
15	0.30
20	0.60
25	0.85
30	0.95
35	0.98
40	0.99
45	1.00
50	1.00

2/12/17 (fig from Bar-Yam) 84





Summary of Capacity Results

- Absolute limit: $p_{\max} < \alpha_c n = 0.138 n$
- If a small number of errors in each pattern permitted: $p_{\max} \propto n$
- If all or most patterns must be recalled perfectly: $p_{\max} \propto n / \log n$
- Recall: all this analysis is based on *random* patterns
- Unrealistic, but sometimes can be arranged

2/12/17III B87
