Hopfield Networks and Boltzmann Machines

15-496/782: Artificial Neural Networks
David S. Touretzky

Spring 2004

1

Properties of Hopfield Nets

- Special class of recurrent network.
- Fully connected; binary units (+1/–1 or 1/0.)
- The stable states are fixed point attractors.

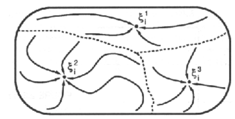

- Can act as a content-addressable memory.

2

Properties of Hopfield Nets (cont.)

- Analogous to spin glass systems (Ising models) in physics, like magnetic bubble memories.

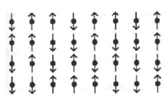
spin glass

  - Has an energy function.

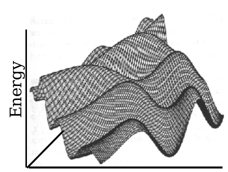

  - We can use physics to analyze a neural net!

3

Definition of a Hopfield Net

1. Binary threshold units:

$$S_i = \begin{cases} +1 & \text{if } net_i \geqslant 0 \\ -1 & \text{otherwise} \end{cases}$$

Can also use 0/1 states.

2. Symmetric weight matrix:

$$W_{ij} = W_{ji}$$

$$W_{ii} = 0$$

4

Definition of a Hopfield Net (cont.)

3. No **systematic** communication delays between units. In other words, updating must be asynchronous.

- Could update one at a time, in random order.
- Could update each unit at time t with probability p < 1.

'Update' means recompute S_i based on current net_i:

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

Storing One Pattern

When is a pattern ξ stable?

$$S_i = \mathrm{sgn}\left(\sum_j w_{ij}\xi_j\right) = \xi_i \quad \text{for all i}$$

Suppose $w_{ij} \propto \xi_i \xi_j$:

$$\begin{aligned} \xi_i &= \mathrm{sgn}\left(\sum_j (\xi_i \xi_j)\cdot\xi_j\right) \\ &= \mathrm{sgn}\left(\sum_j \xi_i \xi_j^2\right) \\ &= \mathrm{sgn}\left(\sum_j \xi_i\right) \quad \text{since } \xi_j^2 = 1 \\ &= \mathrm{sgn}(N\xi_i) \quad \text{where N = pattern size} \end{aligned}$$

So set $w_{ij} = \frac{1}{N}\xi_i \xi_j$

Reversal States Are Also Stable

If ξ is a stable state, then so is $-\xi$.

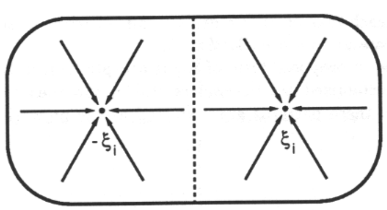

Storing Multiple Patterns

$$w_{ij} = \frac{1}{N}\sum_{\mu=1}^{P} \xi_i^\mu \xi_j^\mu$$

Is ξ^ν stable?

$$\begin{aligned} \xi_i^\nu &= \mathrm{sgn}\left(\sum_j w_{ij}\xi_j^\nu\right) \\ &= \mathrm{sgn}\left(\frac{1}{N}\sum_j \sum_\mu \underbrace{\xi_i^\mu \xi_j^\mu \xi_j^\nu}_{\text{when } \mu=\nu \text{ this is just } \xi_i^\nu}\right) \\ &= \mathrm{sgn}\left(\underbrace{\xi_i^\nu}_{\text{original pattern}} + \underbrace{\frac{1}{N}\sum_j \sum_{\mu\neq\nu} \xi_i^\mu \xi_j^\mu \xi_j^\nu}_{\text{noise or crosstalk term}}\right) \end{aligned}$$

ξ^ν is stable if |noise| < 1.

Stability

- Will units keep flipping state forever?
  - No: there are **stable states.**
- Are we guaranteed to reach a stable state from any starting point?
  - Yes, within a **finite number of flips**.
- Prove it!

Lyapunov Function

A **Lyapunov function** assigns a numerical value to each possible state of the system.

Also called an **energy function**.

To prove stability, show that each state transition reduces the value of the Lyapunov function.

Result: stable states must exist.

- Minimum energy states are stable.
- But local minima may also exist.

Energy Landscape

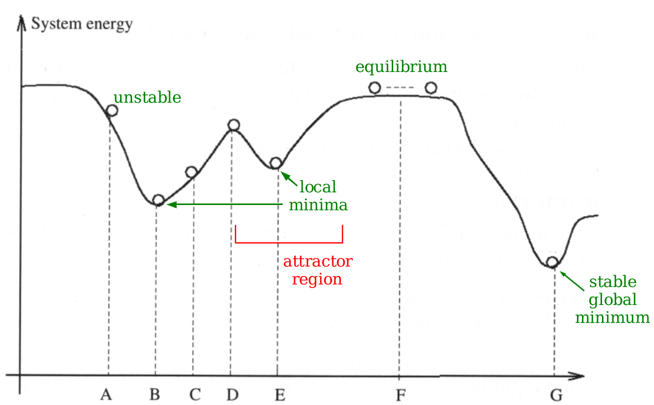

Define an Energy Measure

$$E = -\frac{1}{2}\sum_{i,j} S_i S_j w_{ij}$$

Update step: should unit S_i be set to $+1$ or to -1?

$$\left.\begin{aligned} E(S_i{=}{+1}) &= -\frac{1}{2}\overbrace{\boxed{\sum_j S_j w_{ij}}}^{net_i} \\ E(S_i{=}{-1}) &= -\frac{1}{2}\sum_j -S_j w_{ij} \end{aligned}\right\} + \left(-\frac{1}{2}\sum_{j,k\neq i} S_j S_k w_{jk}\right)$$

If $net_i > 0$, then $E(S_i{=}{+1}) < E(S_i{=}{-1})$.
And... If $net_i \geq 0$, state update rule sets S_i to $+1$.
So with every update, the E goes down or stays the same.
Only 2^N possible states, so a stable state must be reached.

Settling Process

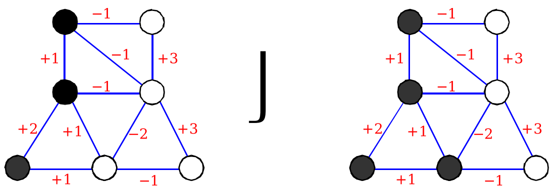

Energy E = −13 Energy E = −14

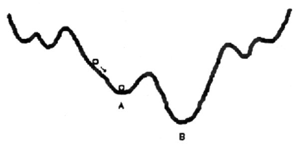

13

All Stable States

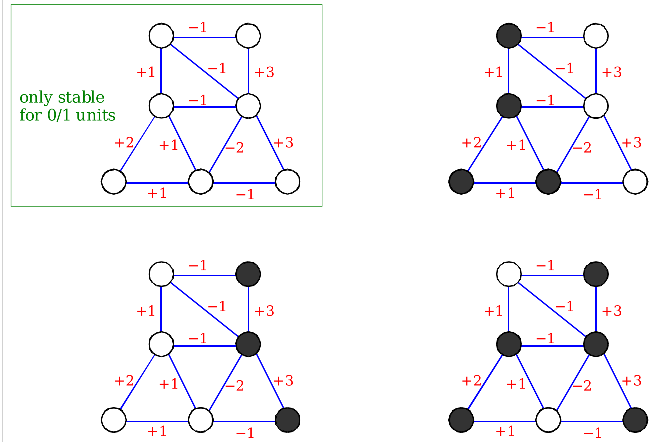

14

Hopfield with 0/1 Units

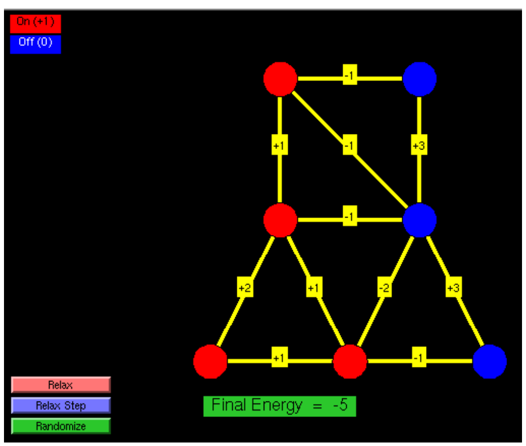

15

Associative Retrieval: Learned Patterns

5x5 = 25 units

4 patterns

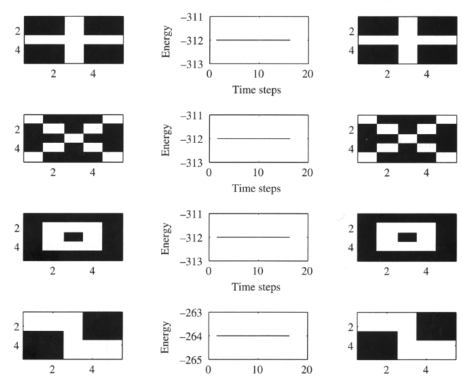

16

Associatve Retrieval: Noisy Cues

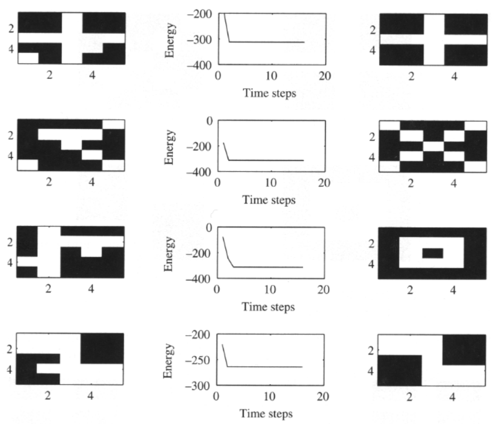

Image Retrieval From Partial Cues

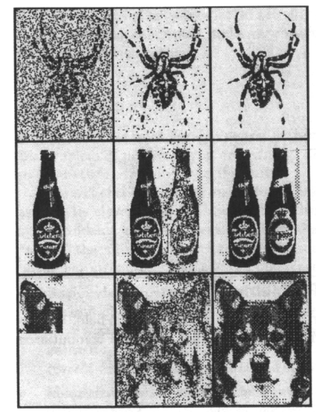

130 x 180 binary pixels =
23,400 bit patterns

sparsely connected network

7 stored patterns

Why No Self-Links?

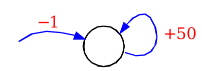

$w_{ii} > 0$ causes spurious stable states

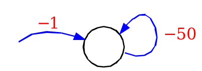

$w_{ii} < 0$ oscillates; no Lyapunov fn.

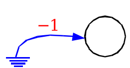

bias term is okay

Setting the Weights: A Heuristic

$$w_{ij} = \sum_{\mu} \xi_i^{\mu} \xi_j^{\mu} \quad \text{for } i \neq j$$

Note: this is just an outer product Hebbian learning rule.

$w_{ii} = 0$ simplifies analysis; gives better performance
$w_{ii} > 0$ allowed, but may cause spurious stable states
$w_{ii} < 0$ no Lyapunov function; can cause oscillations

Stored Patterns Are Energy Minima

Consider the case of one stored memory ξ.
Show that $S_i = \xi_i$ (for all i) is an energy minimum.

$$w_{ij} = \xi_i \xi_j \quad \text{for } i \neq j$$

$$\begin{aligned} E &= -\frac{1}{2}\sum_{i,j} S_i S_j w_{ij} \\ &= -\frac{1}{2}\sum_{i\neq j} S_i S_j (\xi_i \xi_j) \end{aligned}$$

When $S_i = \xi_i$ and $S_j = \xi_j$, all terms are positive, so E is minimal. Any state change would increase E.

Memory Capacity

How many patterns can we store in a net of N units?

- Each pattern is a vector of length N.
- Assume vectors are random (uncorrelated).

Hopfield: capacity C is ~ 0.15 N.

Tighter bound:

$$\frac{N}{4 \ln N} < C < \frac{N}{2 \ln N}$$

100 neurons can reliably store about 8 patterns.

Types of Stable States

1. Retrieval states: ξ^μ

2. Reversed states: $-\xi^\mu$

3. Mixture states: any linear combination of an odd number of patterns.

$$\xi^{mix} = \mathrm{sgn}\left(\pm\xi^1 \pm \xi^2 \pm \xi^3\right)$$

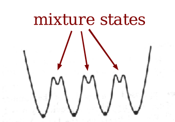

4. 'Spinglass' states: local minima not derivable from finite mixtures of patterns ξ.

Types 3 & 4 are spurious states. Spinglass states occur when too many patterns are stored.

An Aside: Optimization by Simulated Annealing

Simulated annealing is a stochastic search technique introduced by Kirkpatrick, Gelatt, & Vecchi in 1983.

Define some cost function C we want to minimize.

Try to make moves that lower C.

But accept moves that raise C with some probability that depends on a “temperature” parameter T.

Start out at high T; “anneal” by slowly lowering T.

Can escape from local minima!

Chip Layout by Simulated Annealing

Illustrations from Sechen (1988), inspired by Kirkpatrick, Gelatt, & Vechhi's work:

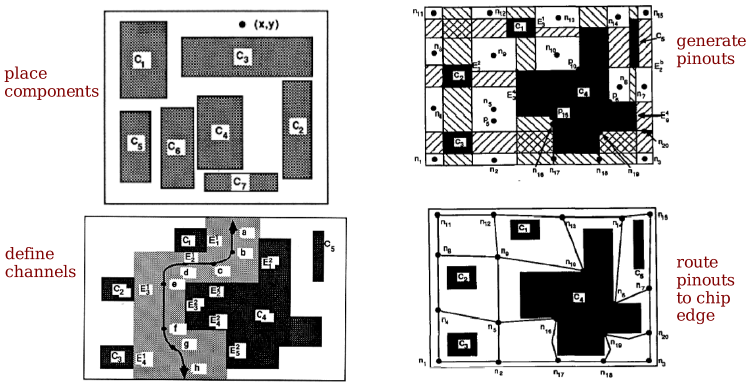

Back to Neural Networks

$$\text{Energy gap } \Delta E_i = E(S_i{=}{+}1) - E(S_i{=}{-}1)$$

$$= -\sum_j S_j w_{ij} = -net_j$$

$$= \text{change in E when } S_i \text{ turns on.}$$

Hopfield: $S_i \leftarrow \text{sgn}(net_i)$ always decreases E.

What if we were to allow E to increase occasionally?

The Boltzmann Machine

Hinton and Sejnowski combined two great ideas:

Spin glass neural net models (Hopfield)

Simulated annealing search (Kirkpatrick et al.)

The result is the **Boltzmann Machine**: a stochastic Hopfield net that avoids local minima.

$$P[S_i{=}{+}1] = \frac{1}{1+e^{\Delta E_i/T}} = \frac{1}{1+e^{-net_i/T}}$$

where T is the temperature.

Stochastic Units

$$P[S_i{=}{+}1] = \frac{1}{1+e^{-net_i/T}}$$

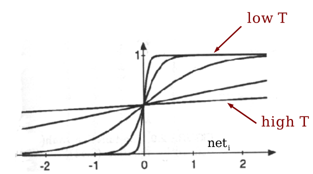

If $net_i = 0$, unit fluctuates randomly.

For large $|net_i|$, unit is mostly on (or mostly off).

We can use this randomness to jump out of local minima!

Boltzmann Machine Stochastic Search

Start at high temperature.

$P[S_i=1]$ is close to 0.5. Units fluctuate a lot.

Gradually cool to lower temperatures.

Units fluctuate less as P moves closer to 1 or 0.

At zero temperature, we have a Hopfield net.

Annealing schedule:

$$T_{i+1} \leftarrow 0.9\,T_i$$

29

Boltzmann Energy States

Given states $\mathbf{x_a}, \mathbf{x_b}$ with energies $E(\mathbf{x_a}), E(\mathbf{x_b})$, the ratio of their probabilities is given by the Boltzmann distribution:

$$\frac{P(\mathbf{x_a})}{P(\mathbf{x_b})} = \exp\left(-\left[E(\mathbf{x_a})-E(\mathbf{x_b})\right]/T\right)$$

States with equal energy are equally probable.

At low T, low energy states are most probable.
Slow annealing is 'guaranteed' to find the global minimum.

30

Variations on Hopfield/Boltzmann

$$\text{Hopfield:}\quad S_i \leftarrow \begin{cases} +1 & \text{if } net_i > 0 \\ \text{unchanged} & \text{if } net_i = 0 \\ -1 & \text{if } net_i < 0 \end{cases}$$

Can also choose randomly if $net_i = 0$

$$\text{Boltzmann:}\quad P(\text{flip}) = \begin{cases} 1 & \text{if } \Delta E(\text{flip}) < 0 \\ f(net_i) & \text{if } \Delta E(\text{flip}) > 0 \end{cases}$$

Settles to local minima more rapidly: always flips state if a flip would move downhill in energy.

31

Boltzmann Machines and Hidden Units

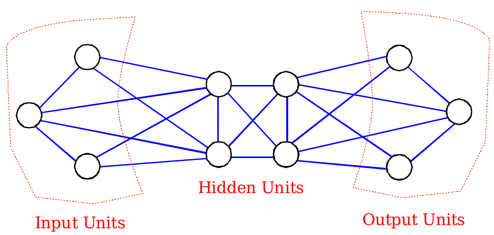

Clamp the input units to the desired state.

Anneal. Read the answer on the output units.

Hidden units add extra representational power.

32

Boltzmann Machines Are Universal

A Boltzmann machine with enough hidden units can compute any computable function.

But annealing may have to be very slow.

Mean field approximation to Boltzmann machine:

Replace S_i by $<S_i>$ as in continuous Hopfield net.

Faster than regular Boltzmann since we don't have to wait a long time to reach **equilibrium state**.

But not as good as avoiding local minima.

Boltzmann Learning Algorithm

1. Clamp, anneal, measure $\langle S_i S_j \rangle^+$ — 'wake' state
2. Unclamp, anneal, measure $\langle S_i S_j \rangle^-$ — 'sleep' state
3. $\Delta w_{ij} = \eta \left[\langle S_i S_j \rangle^+ - \langle S_i S_j \rangle^- \right]$ — weight update

Hebbian learning in wake state; antihebbian in sleep state.

Very, very slow, because each learning step requires many annealings to estimate $\langle S_i S_j \rangle$.