


Graded response Hopfield network

■ The model of the basic neural element

Hopfield's 1982 paper was strongly criticized for having an unrealistic model of the neuron. In 1984, he published another influential paper with an improved neural model. The model was intended to capture the fact that neural firing rate can be considered a continuously valued response (recall that frequency of firing can vary from 0 to 500 Hz or so). The question was whether this more realistic model also showed well-behaved convergence properties. Earlier, in 1983, Cohen and Grossberg had published a paper showing conditions for network convergence.

Previously, we derived an expression for the rate of firing for the "leaky integrate-and-fire" model of the neuron. Hopfield adopted the basic elements of this model, together with the assumption of a non-linear sigmoidal output, represented below by an operational amplifier with non-linear transfer function $g()$. An operational amplifier (or "op amp") has a very high input impedance so its key property is that it essentially draws no current.

Below is the electrical circuit corresponding to the model of a single neuron. The unit's input is the sum of currents (input voltages weighted by conductances T_{ij} , corresponding to synaptic weights). I_i is a bias input current (which is often set to zero depending on the problem). There is a capacitance C_i and membrane resistance R_i —that characterizes the leakiness of the neural membrane.



The basic neural circuit

Now imagine that we connect up N of these model neurons to each other to form a completely connected network. Like the earlier discrete model, neurons are not connected to themselves, and the two conductances between any two neurons is the same. In other words, the weight matrix has a zero diagonal ($T_{ii}=0$), and is symmetric ($T_{ij}=T_{ji}$). We follow Hopfield, and let the output range continuously between -1 and 1 for the graded response network (rather than taking on discrete values of 0 or 1 as for the network in the previous lecture.)

The update rule is given by a set of differential equations over the network. The equations are determined by the three basic laws of electricity: Kirchoff's rule (sum of currents at a junction has to be zero, i.e. sum in has to equal sum out), Ohm's law ($I=V/R$), and that the current across a capacitor is proportional to the rate of change of voltage ($I=Cdu/dt$). Resistance is the reciprocal of conductance ($T=1/R$). As we did earlier with the integrate-and-fire model, we write an expression representing the requirement that the total current into the op amp be zero, or equivalently that the sum of the incoming currents at the point of the input to the op amp is equal to the sum of currents leaving:

$$\sum_j T_{ij} V_j + I_i = C_i \frac{du_i}{dt} + \frac{u_i}{R_i}$$

With a little rearrangement, we have:

$$C_i \frac{du_i}{dt} = \sum_j T_{ij} V_j - \frac{u_i}{R_i} + I_i \quad (1)$$

$$V_i = g(u_i) \quad (2)$$

The first equation is really just a slightly elaborated version of the "leaky integrate and fire" equation we studied in Lecture 3. We now note that the "current in" (I_i in lecture 3) is the sum of the currents from all the inputs.

■ Proof of convergence

Here is where Hopfield's main contribution lies. All parameters (C_i, T_{ij}, R_i, I_i, g) are fixed, and we want to know how the state vector V_i changes with time. We could imagine that the state vector could have almost any trajectory depending on initial conditions, and wander arbitrarily around state space--but it doesn't. In fact, just as for the discrete case, the continuous Hopfield network converges to stable attractors. Suppose that at some time t , we know $V_i(t)$ for all units ($i=1$ to N). Then Hopfield proved that the state vector will migrate to points in state space whose values are constant: $V_i \rightarrow V_i^s$. In other words, to state space points where $dV_i/dt=0$. This is a steady-state solution to the above equations. This is an important result because it says that the network can be used to store memories.

To prove this, Hopfield defined an energy function as:

$$E = -\frac{1}{2} \sum_{i \neq j} T_{ij} V_i V_j + \sum_i (1/R_i) \int_0^{V_i} g_i^{-1}(V) dV + \sum_i I_i V$$

The form of the sigmoidal non-linearity, $g()$, was taken to be an inverse tangent function (see below).

The expression for E looks complicated, but we want to examine how E changes with time, and with a little manipulation, we can obtain an expression for dE/dt which is easy to interpret.

Let's get started. If we take the derivative of E with respect to time, then for symmetric T , we have (using the product rule, and the chain rule for differentiation of composite functions):

$$dE / dt = - \sum_i \frac{dV_i}{dt} \left(\sum_j T_{ij} V_j - u_i / R_i + I_i \right)$$

Where we've used equation 2 to replace $g^{-1}(V)$ by u (after taking the derivative of $\int_0^{V_i} g_i^{-1}(V) dV$ with respect to time)

Substituting the expression from the first equation (1) for the expression between the brackets, we obtain:

$$dE / dt = - \sum_i C_i (dV_i / dt) (du_i / dt)$$

And now replace du_i/dt by taking the derivative of the inverse g function (again using the chain rule):

$$= - \sum_i C_i \frac{dg_i^{-1}(V_i)}{dV_i} \left(\frac{dV_i}{dt} \right)^2 \quad (3)$$

Below, there is an exercise in which you can show that the derivative of the inverse g function is always positive.

And because capacitance and $(dV_i/dt)^2$ are also positive, and given the minus sign, the right hand side of the equation can never be positive--energy never increases. Further, we can see that stable points, i.e. where dE/dt is zero, correspond to attractors in state space. Mathematically, we have:

$$dE / dt \leq 0$$

$$dE / dt = 0 \rightarrow dV_i / dt = 0 \text{ for all } i.$$

E is a Lyapunov function for the system of differential equations that describe the neural system whose neurons have graded responses.

Use the product rule and the chain rule from calculus to fill in the missing steps

Simulation of a 2 neuron Hopfield network

■ Definitions

We will let the resistances and capacitances all be one, and the current input I_i be zero. Define the sigmoid function, $g[]$ and its inverse, $\text{inverseg}[]$:

```
In[10]:= a1 := (2/Pi); b1 := Pi 1.4 / 2;
g[x_] := N[a1 ArcTan[b1 x]];
inverseg[x_] := N[(1/b1) Tan[x/a1]];
```

Although it is straightforward to compute the inverse of $g()$ by hand, do it using the `Solve[]` function in *Mathematica*:


```
In[13]:= Solve[a ArcTan[b y]==x,y]
```

```
Out[13]= {{y -> \frac{\tan(\frac{x}{a})}{b}}}
```

As we saw in earlier lectures, dividing the argument of a squashing function such as $g[]$ by a small number makes the sigmoid more like a step or threshold function. This provides a bridge between the discrete (two-state) Hopfield and the continuous Hopfield networks:


```
In[14]:= Manipulate[Plot[g[1 / b2 x], {x, -\pi, \pi}, ImageSize -> Tiny, Axes -> False],
{{b2, 1}, .01, 15}]
```

```
Out[14]=
```



The derivative of the inverse of $g[]$ is **inverseg' [x]**. Here are plots of g , the inverse of g , and the derivative of the inverse of g :

```
GraphicsRow[{Plot[g[x], {x, -π, π}, ImageSize → Tiny],
  Plot[inverseg[x], {x, -0.9, 0.9}], Plot[inverseg'[x], {x, -1, 1}]}]
```



The third panel on the right shows that the derivative of the inverse is never zero. Because it is never zero, we saw from equation 3 that the rate of change of E must always be less than or equal to zero.

■ Initialization of starting values

The [initialization](#) section sets the starting output values of the two neurons.

$\mathbf{V} = \{0.2, -0.5\}$, and the internal values $\mathbf{u} = \text{inverseg}[\mathbf{V}]$, the step size, $\mathbf{dt} = 0.3$, and the 2x2 weight matrix, \mathbf{Tm} such that the synaptic weights between neurons are both 1. The synaptic weight between each neuron and itself is zero.

```
In[33]:= dt = 0.01;
Tm = {{0,1},{1,0}};
V = {0.2,-0.5};
u = inverseg[V];
result = {};
```

■ Main Program illustrating descent with discrete updates

Note that because the dynamics of the graded response Hopfield net is expressed in terms of differential equations, the updating is actually continuous (rather than asynchronous and discrete). In order to do a digital computer simulation, we will approximate the dynamics with discrete updates of the neurons' activities.

The following function computes the output (just up to the non-linearity) of the i th neuron for the network with a list of neural values \mathbf{uu} , and a weight matrix \mathbf{Tm} .

```
In[38]:= Hopfield[uu_,i_] := uu[[i]] +
  dt ((Tm.g[uu])[[i]] - uu[[i]]);
```

This follows from the above update equations with the capacitances (C_i) and resistances (R_i) set to 1.

$$C_i \frac{du_i}{dt} = \sum_j T_{ij} V_j - \frac{u_i}{R_i} + I_i$$

$$V_i = g(u_i)$$

Let's accumulate some results for a series of iterations. Then we will plot the pairs of activities of the two neurons over the iterations. The next line randomly samples the two neurons for asynchronous updating.

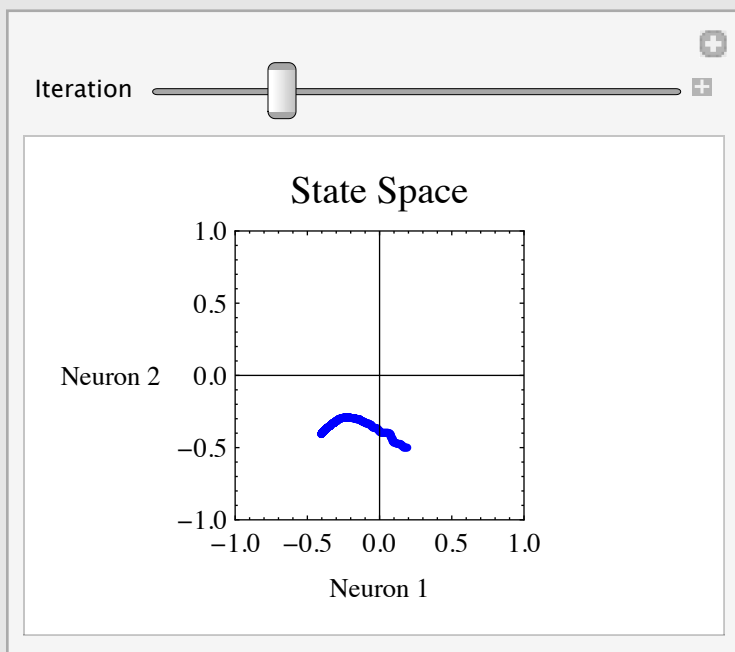
```
In[39]:= result = Table[{k = RandomInteger[{1, 2}], u[k] = Hopfield[u, k], u},
  {2800}];
```

```
In[40]:= result = Transpose[result][[3]];
```

We use ListPlot[] to visualize the trajectory of the state vector.

```
In[41]:= Manipulate[
  ListPlot[g[result[[1 ;; i]]], Joined -> False, AxesOrigin -> {0, 0},
  PlotRange -> {{-1, 1}, {-1, 1}}, FrameLabel -> {"Neuron 1", "Neuron 2"},
  Frame -> True, AspectRatio -> 1, RotateLabel -> False,
  PlotLabel -> "State Space", Ticks -> None, PlotStyle -> {RGBColor[0, 0, 1]},
  ImageSize -> Small], {{i, 1, "Iteration"}, 1, Dimensions[result][[1]], 1}]
```

Out[41]=



■ Energy landscape

Now let's make a contour plot of the [energy landscape](#). We will need a *Mathematica* function expression for the integral of the inverse of the g function ($\text{inverseg}'[x]$), call it `inig[]`. We use the `Integrate[]` function to find it. Then define `energy[x_,y_]` and use `ContourPlot[]` to map out the energy function.

```
Integrate[(1/b) Tan[x1/a], x1] /. x1 -> x
```

$$-\frac{a \log\left(\cos\left(\frac{x}{a}\right)\right)}{b}$$


Change the above output to input form:

```
a Log[Cos[x/a]]
- (-----)
      b
```

$$-\frac{a \log\left(\cos\left(\frac{x}{a}\right)\right)}{b}$$

And copy and paste the input form into our function definition for the integral of the inverse of g:

```
inig[x_] := -N[ $\frac{a1 \text{Log}\left[\text{Cos}\left[\frac{x}{a1}\right]\right]}{b1}$ ]; Plot[inig[x], {x, -1, 1}]
```



We write a function for the above expression for energy:

$$E = -\frac{1}{2} \sum_{i \neq j} T_{ij} V_i V_j + \sum_i (1/R_i) \int_0^{V_i} g_i^{-1}(V) dV + \sum_i I_i V_i$$

```
energy[Vv_] := -0.5 (Tm.Vv).Vv + Sum[inig[Vv][[i]], {i, Length[Vv]}];
```

And then define a contour plot of the energy over state space:

```

gcontour = ContourPlot[energy[{x,y}],{x,-1,1},
  {y,-1,1},AxesOrigin->{0,0},
  PlotRange->{-.1,.8}, Contours->32,ColorFunction -> "Pastel",
  PlotPoints->30,FrameLabel->{"Neuron 1","Neuron 2"},
  Frame->True, AspectRatio->1,RotateLabel->False,PlotLabel->"Energy over

```

We can visualize the time evolution by color coding state vector points according to Hue[time], where H starts off red, and becomes "cooler" with time, ending up as violet (the familiar rainbow sequence: ROYGBIV).

```

gcolortraj =
Graphics [
  ( { Hue [  $\frac{\#1}{\text{Length}[g[\text{result}] ]}$  ],
    Point [ { g[result][[#1]][1], g[result][[#1]][2] } ] } & ) /@
  Range[1, Length[g[result]], AspectRatio -> Automatic];

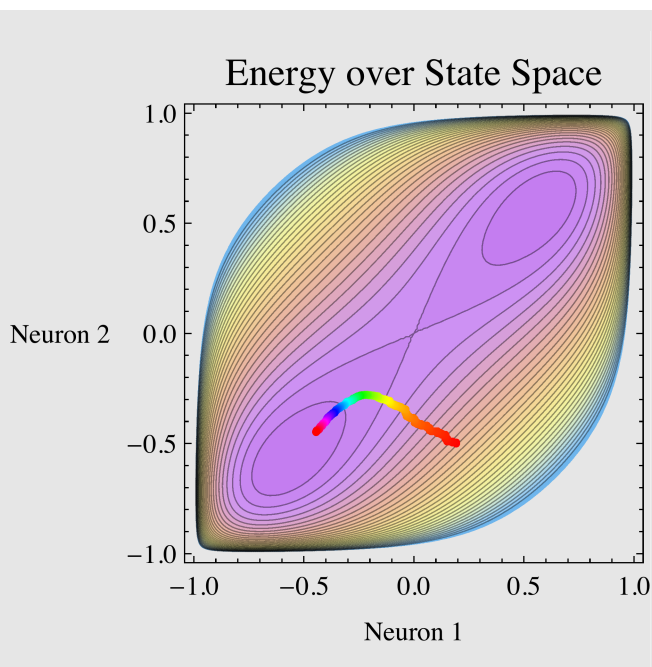
```

Now let's superimpose the trajectory of the state vector on the energy contour plot:

```

Show[gcontour, gcolortraj, ImageSize -> {240, 240}]

```



Has the network reached a stable state? If not, how many iterations does it need to converge?

Autoassociative memory example

Sculpting the energy for memory recall using a Hebbian learning rule. TIP example.


```
In[42]:= SetOptions[ArrayPlot, ImageSize -> Tiny, Mesh -> True, DataReversed -> True];
```

■ The stimuli

We will store the letters P, I, and T in a 25x25 element weight matrix.

```
In[43]:= ArrayPlot[
  Pmatrix = {{0, 1, 0, 0, 0}, {0, 1, 1, 1, 0}, {0, 1, 0, 1, 0},
    {0, 1, 0, 1, 0}, {0, 1, 1, 1, 0}}]
```

Out[43]=



```
In[44]:= ArrayPlot[
  Imatrix = {{0, 0, 1, 0, 0}, {0, 0, 1, 0, 0}, {0, 0, 1, 0, 0},
    {0, 0, 1, 0, 0}, {0, 0, 1, 0, 0}}];
```

```
In[45]:= ArrayPlot[
  Tmatrix = {{0, 0, 0, 1, 0}, {0, 0, 0, 1, 0}, {0, 0, 0, 1, 0},
    {0, 0, 1, 1, 1}, {0, 0, 0, 0, 0}}];
```

For maximum separation, we will put them near the corners of a hypercube.

```
In[46]:= width = Length[Pmatrix];
one = Table[1, {i,width}, {j,width}];
Pv = Flatten[2 Pmatrix - one];
Iv = Flatten[2 Imatrix - one];
Tv = Flatten[2 Tmatrix - one];
```

Note that the patterns are not normal, or orthogonal:

```
In[51]:= {Tv.Iv, Tv.Pv, Pv.Iv, Tv.Tv, Iv.Iv, Pv.Pv}
```

```
Out[51]= {7, 3, 1, 25, 25, 25}
```

■ Learning

Make sure that you've defined the sigmoidal non-linearity and its inverse above. The items will be stored in the connection weight matrix using the outer product form of the Hebbian learning rule:

```
In[52]:= weights =
Outer[Times, Tv, Tv] + Outer[Times, Iv, Iv] +
Outer[Times, Pv, Pv];
```

Note that in order to satisfy the requirements for the convergence theorem, we should enforce the diagonal elements to be zero. (Is this necessary for the network to converge?)

```
In[53]:= For[i = 1, i <= Length[weights], i++, weights[[i, i]] = 0];
```

Hopfield graded response neurons applied to reconstructing the letters T,I,P. The non-linear network makes decisions

In this section, we will compare the Hopfield network's response to the linear associator. First, we will show that the Hopfield net has sensible hallucinations--a random input can produce interpretable stable states. Further, remember a major problem with the linear associator is that it doesn't make proper discrete decisions when the input is a superposition of learned patterns. The Hopfield net makes a decision.

■ Sensible hallucinations to a noisy input


We will set up a version of the graded response Hopfield net with synchronous updating.

```
In[54]:= dt = 0.03;
Hopfield[uu_] := uu + dt (weights.g[uu] - uu);
```

Let's first perform a kind of "Rorschach blob test" on our network (but without the usual symmetry to the input pattern). We will give as input uncorrelated uniformly distributed noise and find out what the network sees:


```
In[56]:= forgettingT = Table[2 RandomReal[] - 1, {i, 1, Length[Tv]}];
ArrayPlot[Partition[forgettingT, width]]
```

Out[56]=



```
In[57]:= rememberingT = Nest[Hopfield, forgettingT, 30];
ArrayPlot[Partition[g[rememberingT], width], PlotRange -> {-1, 1}]
```

Out[57]=




In this case, the random input produces sensible output--the network saw a specific letter.

Do you think you will always get the same letter as network's hallucination?

```
In[58]:= DocumentNotebook[
Dynamic[ArrayPlot[Partition[g[rememberingT], width],
PlotRange -> {-1, 1}]]]
```

Out[58]=



```
In[59]:= Button["Re-initialize",
  forgettingT = Table[2 RandomReal[] - 1, {i, 1, Length[Tv]}];
  rememberingT = Nest[Hopfield, forgettingT, 30]
```

Out[59]=


Click the button for random initial conditions

■ Comparison with linear associator

What is the linear associator output for this input? We will follow the linear output with the squashing function, to push the results towards the hypercube corners:

```
In[60]:= ArrayPlot[Partition[g[weights.forgettingT], width], PlotRange -> {-1, 1}]
```

Out[60]=




The noisy input doesn't always give a meaningful output. You can try other noise samples. Because of superposition it will typically not give the meaningful discrete memories that the Hopfield net does. However, sometimes the linear matrix memory will produce meaningful outputs and sometimes the Hopfield net will produce nonsense depending on how the energy landscape has been sculpted. If the "pits" are arranged badly, one can introduce valleys in the energy landscape that will produce spurious results.

■ Response to superimposed inputs

Let us look at the problem of superposition by providing an input which is a linear combination of the learned patterns. Let's take a weighted sum, and then start the state vector fairly close to the origin of the hypercube:

```
In[61]:= forgettingT = 0.1 (0.2 Tv - 0.15 Iv - 0.3 Pv);
ArrayPlot[Partition[forgettingT, width]]
```

Out[61]=




Now let's see what the Hopfield algorithm produces. We will start with a smaller step size. It sometimes takes a little care to make sure the descent begins with small enough steps. If the steps are too big, the network can converge to the same answers one sees with the linear associator followed by the non-linear sigmoid.

```
In[62]:= dt = 0.01;
Hopfield[uu_] := uu + dt (weights.g[uu] - uu);
rememberingT = Nest[Hopfield, forgettingT, 30];
```

```
In[65]:= ArrayPlot[Partition[g[rememberingT], width], PlotRange -> {-1, 1}]
```

Out[65]=




■ Comparison with linear associator

The linear associator followed by the squashing function gives:

```
In[66]:= ArrayPlot[Partition[g[weights.forgettingT], width], PlotRange -> {-1, 1}]
```

```
Out[66]=
```



For the two-state Hopfield network with Hebbian storage rule, the stored vectors are stable states of the system. For the graded response network, the stable states are near the stored vectors. If one tries to store the items near the corners of the hypercube they are well separated, and the recall rule tends to drive the state vector into these corners, so the recalled item is close to what was stored.

Illustrate how the Hopfield net can be used for error correction in which the input is a noisy version of one of the letters

■ Using the formula for energy

The energy function can be expressed in terms of the integral of the inverse of g , and the product of the weight matrix times the state vector, times the state vector again:

```
In[67]:= inig[x_] := -N[(a1*Log[Cos[x/a1]])/b1];
energy[Vv_] := -0.5 ((weights.Vv).Vv) +
               Sum[inig[Vv][[i]], {i,Length[Vv]}];
```

```
In[69]:= energy[.99*Iv]
energy[g[forgettingT]]
energy[g[rememberingT]]
```

```
Out[69]= -263.969
```

```
Out[70]= -0.528994
```

```
Out[71]= -219.84
```

One of the virtues of knowing the energy or Lyapunov function for a dynamical system, is being able to check how well you are doing. We might expect that if we ran the algorithm a little bit longer, we could move the state vector to an even lower energy state if what we found was not the minimum.

```
In[72]:= dt = 0.01;  
Hopfield[uu_] := uu + dt (weights.g[uu] - uu);  
rememberingT = Nest[Hopfield, forgettingT, 600];
```

```
In[75]:= energy[g[rememberingT]]
```

```
Out[75]= -244.693
```

Conclusions

Although Hopfield networks have had a considerable impact on the field, they have had limited practical application. There are more efficient ways to store memories, to do error correction, and to do constraint satisfaction or optimization. As models of neural systems, the theoretical simplifications severely limit their generality. Although there is current theoretical work on Hopfield networks, it has involved topics such as memory capacity and efficient learning, and relaxations of the strict theoretical presumptions of the original model. The field has moved on to more general probabilistic theories of networks that include older models as special cases. In the next few lectures we will see some of the history and movement in this direction.

References

Bartlett, M. S., & Sejnowski, T. J. (1998). Learning viewpoint-invariant face representations from visual experience in an attractor network. *Network*, 9(3), 399-417.

Chengxiang, Z., Dasgupta, C., & Singh, M. P. (2000). Retrieval properties of a Hopfield model with random asymmetric interactions. *Neural Comput*, 12(4), 865-880.

Cohen, M. A., & Grossberg, S. (1983). Absolute Stability of Global Pattern Formation and Parallel Memory Storage by Competitive Neural Networks. *IEEE Transactions SMC-13*, 815-826.

Hertz, J., Krogh, A., & Palmer, R. G. (1991). *Introduction to the theory of neural computation* (Santa Fe Institute Studies in the Sciences of Complexity ed. Vol. Lecture Notes Volume 1). Reading, MA: Addison-Wesley Publishing Company.

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A*, 79(8), 2554-2558.

Computational properties of use of biological organisms or to the construction of computers can emerge as collective properties of systems having a large number of simple equivalent components (or neurons). The physical meaning of content-addressable memory is described by an appropriate phase space flow of the state of a system. A model of such a system is given, based on aspects of neurobiology but readily adapted to integrated circuits. The collective properties of this model produce a content-addressable memory which correctly yields an entire memory from any subpart of sufficient size. The algorithm for the time evolution of the state of the system is based on asynchronous parallel processing. Additional emergent collective properties include some capacity for generalization, familiarity recognition, categorization, error correction, and time sequence retention. The collective properties are only weakly sensitive to details of the modeling or the failure of individual devices.

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proc. Natl. Acad. Sci. USA*, 81, 3088-3092.

A model for a large network of "neurons" with a graded response (or sigmoid input-output relation) is studied. This deterministic system has collective properties in very close correspondence with the earlier stochastic model based on McCulloch - Pitts neurons. The content-addressable memory and other emergent collective properties of the original model also are present in the graded response model. The idea that such collective properties are used in biological systems is given added credence by the continued presence of such properties for more nearly biological "neurons." Collective analog electrical circuits of the kind described will certainly function. The collective states of the two models have a simple correspondence. The original model will continue to be useful for simulations, because its connection to graded response systems is established. Equations that include the effect of action potentials in the graded response system are also developed.

Hopfield, J. J. (1994). Neurons, Dynamics and Computation. *Physics Today*, February.

Hentschel, H. G. E., & Barlow, H. B. (1991). Minimum entropy coding with Hopfield networks. *Network*, 2, 135-148.

Mead, C. (1989). *Analog VLSI and Neural Systems*. Reading, Massachusetts: Addison-Wesley.

Saul, L. K., & Jordan, M. I. (2000). Attractor dynamics in feedforward neural networks. *Neural Comput*, 12(6), 1313-1335.

