

# Introduction to Neural Networks

U. Minn. Psy 5038

Psy 5038, Spring, 1998

## Self-organizing Adaptive Maps

---

### Initialization

```
Off[SetDelayed::write]
Off[General::spell1]
(*<<Statistics`ContinuousDistributions`*)
```

---

### Topographic maps

Work in monkey, and human brain, shows that the cortex is characterized by numerous distinct areas. It has been estimated that there are more than 30 visual areas alone in the macaque cortex. The earlier areas typically show a spatial topographic representation of visual space--nearby regions of visual space map to nearby regions of cortex. The so-called retinotopic map of the primary visual area (V1) is the clearest example of this. For a recent study using brain imaging techniques in humans, see: Engel et al. (1994). Other areas of the brain also show geometrical organization. For example, the auditory cortex has a tonotopic map in which the spatial order of cell responses corresponds to pitch or acoustic frequency. The motor cortex also shows a spatial organization.

In regions of the cortex with no obvious spatial map, it is quite possible that other kinds of maps wait to be discovered. Tanaka and colleagues (Tanaka, 1996) have shown that region TE of the monkey inferotemporal cortex has columns with cells that have similar visual shape preferences. Along the surface of the cortex, receptive field properties may correspond to other kinds of variation, such as rotation in depth of a face, over limited extents (on the order of 1 mm or so).

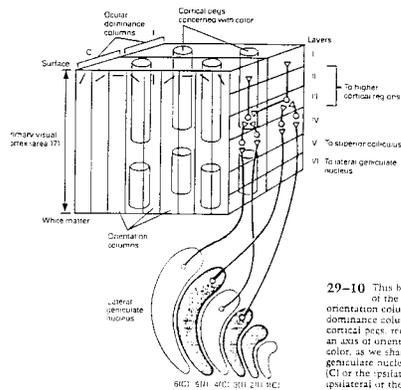
The widespread use of spatial organization in cortex suggests the possibility of a general constraint underlying the development of neural receptive field organization. We know more about primary visual cortex than any other area, so let's take a closer look at what it does.

## Dimension reduction framework for understanding cortical maps

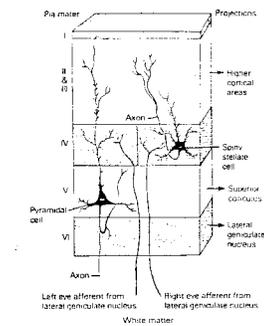
Primary visual cortex does not simply have the job of representing nearby retinal points at nearby cortical locations. Much physiological research has shown that V1 brings together information from the two eyes, along similar orientations, as well as location. Together with anatomical studies, it is now well-established that neurons with similar orientation preferences and various degrees of relative input from the two eyes are organized into "hypercolumns" (See Figure below, and Lecture 13). Hypercolumns preserve spatial contiguity and smoothness of the placement of neurons selective for features of the input.

This observation suggests that a general principle may account for the organization and development of cortical maps: Neighboring points in feature or parameter space (e.g. orientation, ocular dominance, and retinal position) should map to nearby points on the 2D cortical sheet. (See: Durbin & Mitchison, 1990)

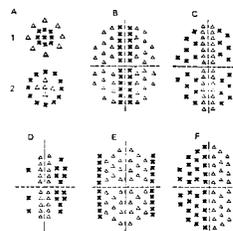
The underlying assumption is that most operations performed in the cortex are local, thus the related input for these computations should be physically near the computing units. For example, one task of vision is to go beyond the mere detection of contour segments, but to actually link contours that are likely to belong together to form a global object outline. Thus it would make sense to have the cells that signal similar orientations to be near. Visual information from the two eyes is close in the world, but separated by a great distance anatomically in the left and right eyes. This information needs to be brought physically together to process the two images binocularly, for example, to compute stereoscopic depth.



29-10 This basic cortical module (hypercolumn) in area 17 of the visual cortex contains a complete set of orientation columns representing 360° and a set of ocular dominance columns. Each hypercolumn also contains several cortical areas of cortex in which the cells do not have an axis of orientation. The cells in the pres are concerned with color, as we shall see in Chapter 40. Each layer of the lateral geniculate nucleus receives input from either the contralateral (C) or the ipsilateral (I) eye and projects in turn to the ipsilateral or the contralateral ocular dominance columns.



28-14 The afferent and efferent connections of the primary visual cortex are made in specific layers of cortex.



20-7 Comparison of the receptive fields of neurons in the retina and in the lateral geniculate nucleus with those of simple cortical cells in area 17. A. Cells of the retina and lateral geniculate fall into two classes: on-center (filled triangles) and off-center (open triangles). Neurons of the primary visual cortex also fall into two major classes: simple and complex. Each of these classes, moreover, has several subclasses. This is illustrated here for simple cells. Despite this variety, however, all simple cells are characterized by three features: (1) specific retinal position; (2) their discrete excitatory (filled triangles) and inhibitory (open triangles) zones; and (3) specific axis of orientation. For simplicity, only receptive fields with a vertical axis of orientation from 12 to 6 o'clock are shown in this figure; each has a rectangular contour. In fact, each neuron of the retina is represented in area 17, not only for this but for all axes of orientation—vertical, horizontal, and various obliques. (Adapted from Hubel and Wiesel, 1962.)

## Minimum wiring length constraint

### ■ Nematode

A number of people have sought a simple organizational principle that would predict the spatial layout of neurons. One such principle is that the layout of nervous system components minimizes total connection cost. Christopher Cherniak, a philosopher at the University of Maryland calculated the total wiring length for the 40,000,000 possible layouts of the 10 ganglia in the nematode worm. Remarkably, he reported that the layout the worm actually has is indeed the one with the shortest total connection length (Cherniak, 1991, 1995). (For criticisms of this result, see: <http://york37.ncl.ac.uk/replicate.html>) Similar arguments have been made by Cherniak and others for the layout of the multiple areas of cortex. The problem of minimizing connection lengths is also encountered in VLSI component layout in the design of computer chips.

### ■ Minimum wiring length & dimensionality reduction in cortical maps

Let's look at a simple and small version of the problem that Durbin and Mitchison addressed, that of mapping a higher dimensional parameter space to one of lower dimension. Suppose we have an  $N \times N$  2D feature space that we wish to map to a "1D cortex". Points in this 2D feature space can be represented by an index number,  $f_{ij}$ , assigned to the  $(i,j)$ th coordinate. There are several ways of assigning costs for various orderings. Mitchison and Durbin analysed the following connection cost:

$$C(f) = \sum \Delta_{ij}$$

$$\Delta_{ij} = |f_{i,j+1} - f_{i,j}|^q + |f_{i+1,j} - f_{i,j}|^q$$

If  $q = 1$ , then the standard "raster" ordering, the index,

$$f_{ij} = (i - 1)N + j$$

gives a cost of:

$$C(f) = N^3 - N$$

(In standard raster ordering, matrix rows are laid out one after the next in one long vector).

One interesting biological application of this principle was published by Durbin and Mitchison in Nature (1990). It is not computationally feasible to find the minimum cost for dimensional reduction mappings of higher dimension, for example from  $(x,y,r,\theta) \rightarrow (x',y')$  as one would like to do for the formation of retinotopic and orientation maps in V1. The alternative is to see whether some biologically plausible rules could act to accomplish an efficient mapping of the higher dimensional feature space onto the 2D cortex.

Two biologically plausible rules are: 1) there are competitive winner-take-all interactions selective for distinct inputs 2) the units also strengthen their responses to those stimuli that their neighbors respond to. The first rule divides up the input

domain, and the second rule imposes a continuity constraint on the formation of a map. Durbin and Mitchison developed an algorithm which applied these rules and showed that the kind of 2D maps which developed looked very much like the visual cortical maps, revealed from photo-sensitive dye studies (e.g. T'so et al., 1990).

## Kohonen's algorithm for topology-preserving mappings

### ■ Theory

The Finnish scientist, Teuvo Kohonen, was the first to develop topology-preserving adaptive maps for neural networks (Kohonen, T. (1984). Let's look at the basic structure of a simple adaptive map.

Kohonen boiled the essential features of self-organizing topology-preserving maps down to two basic processes:

1. Find the activity maximum of a set of neurons to an input. We assume that maximum activity occurs for patterns which match the receptive field (i.e. as with a cross-correlator).
2. Define a set of neighbors around this maximum, and make these neighbors more likely to respond to that input in the future by making their weights more like the input. Typically the neighborhood  $N_c$  starts off large, and is gradually reduced over time.

Let  $x$  be an  $n$ -dimensional vector representing a feature sample. Let  $m_i$  be a  $k$ -dimensional vector representing the weights of the  $i^{\text{th}}$  unit. Specifically, we will use the following rules:

**Similarity matching.** Find unit  $c$  such that:

$$\left| x(t_k) - m_c(t_k) \right| = \min_i \left| x(t_k) - m_i(t_k) \right|$$

**Updating.** Update the weights for unit  $c$ , and all the units within  $c$ 's neighborhood:

$$m_i(t_{k+1}) = \begin{cases} m_i(t_k) + \alpha(t_k)[x(t_k) - m_i(t_k)] & \text{for } i \in N_c \\ m_i(t_k) & \text{otherwise} \end{cases}$$

## Demonstration of Kohonen's algorithm for mapping 2D features to a 1D line

### ■ Define functions

```

VectorLength[x_]:=Sqrt[x.x];

(*ramp[] is used below to define both the  $\alpha()$  term,
and the rate of change of the neighborhood size *)

ramp[x_,yint_,end_,plat_] :=
  If[(x>=end) || (-2*x*yint/end+yint<plat),plat,
  -2*x*yint/end+yint];

(*rv samples the 2D "feature space"*)
rv := {Random[Real],Random[Real]};

(*These are alternative samplings to try
rvdiscrete := 1/8 + (Floor[(4 rv)]/4);
rvline := {xx=Random[Real],xx};*)

```

### ■ Neighborhood function

The neighborhood function determines the neighbors, and thus the topology of the connections between the neurons. In our example, the neighborhood is 1-D and is defined along a line.

**neigh[]** is a neighborhood function that produces a list of indices for the neighbors of unit **c**. We will not use a toroidal geometry here. Instead, **neigh[]** generates shorter lists of indices near the borders, so the **min\_** and **max\_** of the range need to be specified. This neighborhood function only defines neighbors along a line, i.e. in one dimension. You could elaborate this algorithm to find maps from 2D to 2D, allowing neighbors to be nearby regions of 2D space.

```

neigh[c_,numneigh_,min_,max_] :=Module[{i,nn,temp},
  temp = numneigh/2;
  nn={};
  For[i=c-temp,(i<=c+temp)&&(i<=max),i++,
    If[i>=min,nn=Join[nn,{i}]];
  ];
  Return[nn];
];

```

### ■ Initializing the simulation parameters

**n** is the number of nodes in the 1D line. **mu** is the matrix with the weights that will get updated according to the above update rule. **niter** is the number of iterations.

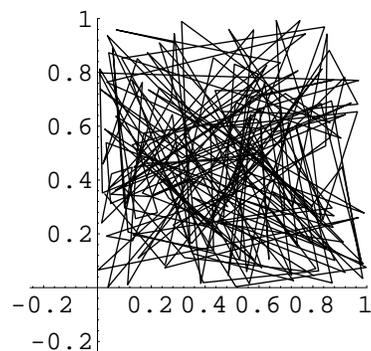
**numneigh0** is the initial neighborhood size. If this is too small, the topography map can get tangled. We will start off with a neighborhood size that is 60-80% of the total size, **n**.

```
n = 200;
mu = Table[rv, {j, 1, n}];
g1 = ListPlot[mu, PlotRange -> {{-0.25, 1}, {-0.25, 1}},
  AspectRatio -> 1,
  PlotJoined -> True, DisplayFunction -> Identity];

niter = 2000;
numneigh0 = Floor[.6 * n / 2];

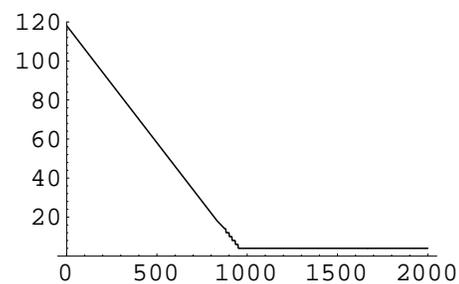
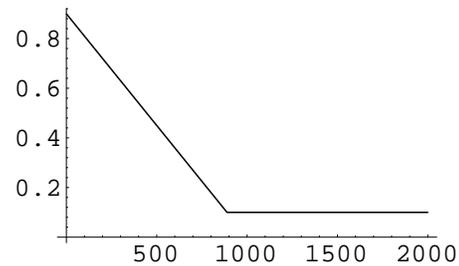
(*Set the learning rate, and neighborhood size*)
eta[t_] := ramp[t, 0.9, niter, 0.1];
numneigh[t_] := 2 Floor[ramp[t, numneigh0, niter, 2]];
```

```
Show[g1, DisplayFunction -> $DisplayFunction];
```



Let's take a look at how the neighborhood size and  $\alpha = \mathbf{eta}$ , decrease with the number of iterations, **t** :

```
Plot[eta[t],{t,1,niter}];  
Plot[numneigh[t],{t,1,niter},PlotRange->{0,2 numneigh0}];
```



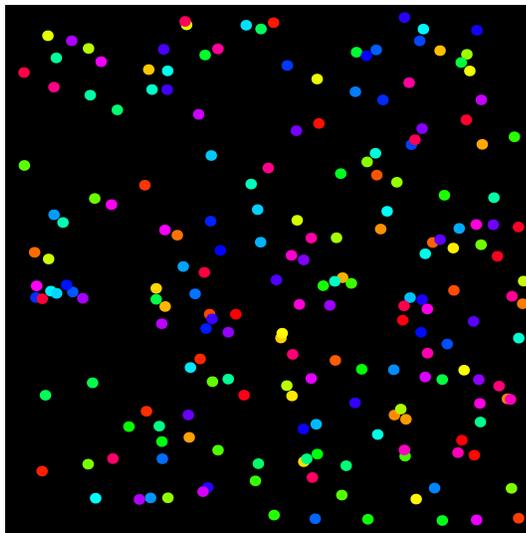
### ■ The algorithm

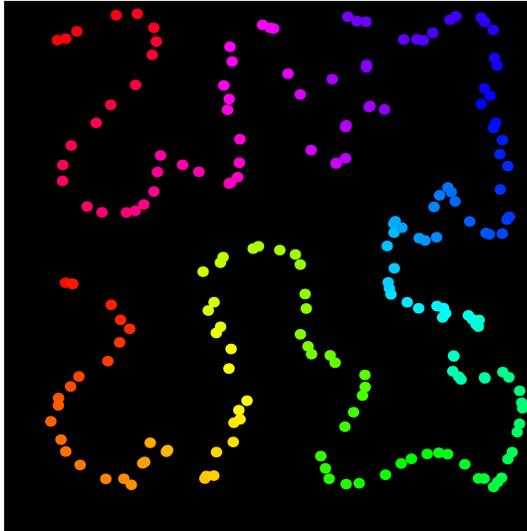
We will make a series of plots, showing the first ten iterations, and after that sampling every 200. The plots will show how the matrix **mu** (which has the topography of a 1D line, because of the way we defined the **neigh[]** function) gradually fits itself to the geometry of the 2D input space.

```

For[t=1,t<=niter,t++,
  If[(Mod[t,200]==1) || t<=10,
    gcolor = Graphics[Table[{Hue[i/n],PointSize[0.02],
      Point[mu[[i]]]},{i,1,n}]];
    Show[gcolor,Background->RGBColor[0,0,0],AspectRatio->1 ];
  ];
  (*Pick a uniformly distributed "feature" sample
  from a 2D array*)
  x=rv;
  (*Do the similarity matching. mini is the unit whose
  weights best match the input *)
  diffs = Map[VectorLength,Transpose[Transpose[mu]-x]];
  minarg = Min[diffs];
  mini = Part[Position[diffs,minarg],1,1];
  (*Make a list, j, of the neighbors for this index,
  at this tth iteration *)
  j=neigh[mini,numneigh[t],1,n];
  (*Update the weights in the neighborhood of i to move them
  towards feature x, by eta proportion of the difference*)
  For[s=1,s<=Length[j],s++,
    mu[[ j[[s]] ]] = mu[[ j[[s]] ]] +
    eta[t] (x-mu[[j[[s]] ]])
  ];
];

```





Let's summarize what we have done. There are 200 "neurons", each with 2 input weights, represented by matrix  $\mathbf{mu}$ . We can imagine representing the weights of these 200 neurons by a location in weight space. The point for each neuron is represented by a different hue in the graph above--neurons with similar hues are neighbors. We randomly sampled a location in the 2-dimensional parameter space defined by the unit square. We then surveyed the 200 neurons to see which one had input weights that were closest to the sampled location. Note that we don't actually have to calculate the response of the unit--which if linear would be the cross-correlation of the input with its weights. Then we adjusted the 2 weights of that neuron to move them closer to the sampled point. Further, we adjusted the weights of all of the neighbors of that neuron (i.e. those with similar hues) to be closer too. We reduced the size of the neighborhood as the number of iterations increased. So a unit's weights are less affected by distant activity as time goes on. The end result is that nearby points in feature space tend to activate nearby neurons that are arrayed in a 1-D line. For other interesting examples, and for a discussion of the relationship of Kohonen maps to space-filling curves, see Kohonen (1984).

### Optional exercises

---

Compute the connection cost for the Kohonen adaptive map in the above example. Compare it to a raster scheme.

Try sampling from `rvline`, and watch how the algorithm learns the topology of the 1D input space.

Try playing with the initial neighborhood size, `numneigh0`. What happens if it starts off small, (e.g. let the number of neighbors be fixed at 2 throughout the simulation).

Define a 2D feature input space which is not rectangular. For example, `rv` could sample from a triangular or circular region within the unit square.

## References

- Cherniak, 1991, Component placement optimization in the brain, UMIACS-TR-91-98 or CS-TR-2711)
- Cherniak, C. (1995). Neural component placement. *TINS*, 18(12), 522-527.
- Durbin, R., & Mitchison, G. (1990). A dimension reduction framework for understanding cortical maps. *Nature*, 343, 644-647.
- Engel, S. A., Rumelhart, D. E., Wandell, B. A., Lee, A. T., Glover, G. H., Chichilnisky, E.-J., & Shadlen, M. N. (1994). fMRI of human visual cortex. *Nature*, 369, 525.
- Grinvald, A., Bonhoeffer, T., Malonek, D., Shoham, D., Bartfeld, E., Arieli, A., Hildesheim, R., & Ratzlaff, E. (1991). Optical Imaging of Architecture and Function in the Living Brain. In L. R. Squire, N. M. Weinberg, G. Lynch, & J. L. McGaugh (Eds.), *Memory: Organization & Locus of Change*, (pp. 49-85): Oxford University Press.
- Kohonen, T. (1984). *Self-Organization and Associative Memory*. Berlin, New York: Springer-Verlag.
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience*, 19, 109-139.
- Ts'o, D. Y., Frostig, R. D., Lieke, E. E., & Grinvald, A. (1990, 27 July 1990). Functional Organization of Primate Visual Cortex Revealed by High Resolution Optical Imaging. *Science*, 249, 417-420.