Introduction to Neural Networks
U. Minn. Psy 5038

Representation of visual information

Neural codes and representation

Overview

■ **Final home work**

■ **Final projects**

■ **More on representation--tie with V1 orientation selectivity**

■ **More on representation--contingent adaptation--tie with decorrelating neural networks**

Extra-striate cortical visual areas

V1 plus extra-striate areas, ~50% of primate cortex.

In lecture 1, and in Figure 10.9 in the Anderson book we see that cells in the visual cortex send their visual information to an incredibly complex, and yet structured collection of extra-striate areas. Any hypothesized function of striate cortex must eventually take into account what the information is to be used for. Two primary functions: object perception and recognition--"within object" processing, and spatial processing (between object, and view-object relations).

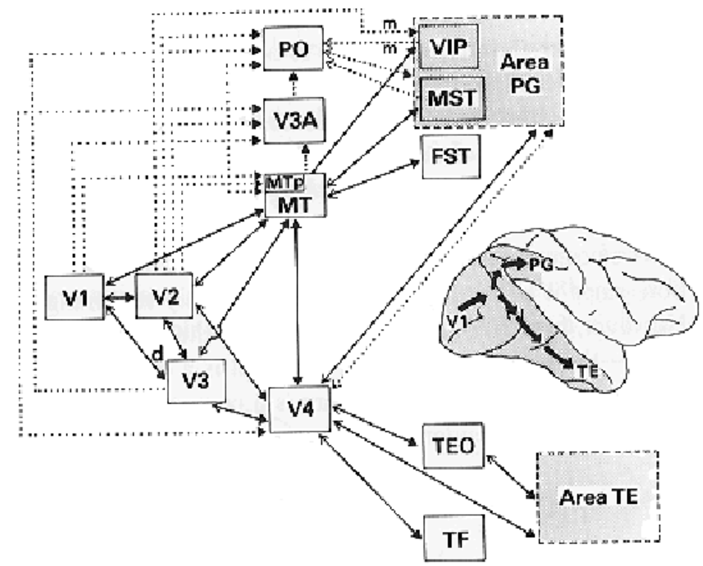

■ **Connectivity between areas**

Feedforward (ascending pathways) from mainly superficial layers (I,II, III) to layer IV of receiving areas. More diffuse feedback (descending pathways) from outside IV back to layers also outside IV (mainly I or VI). (~cm)

But there are also feedforward/feedback local circuits at a finer grain, within columns (~mm)

■ Two main large-scale functional pathways (involving mutliple areas and area-area connections)

Behavioral evidence: Double dissociation of function.

Monkeys

IT Lesions (Mishkin) showed impaired ability in matching tasks (non-match-to-sample task) but can do

spatial task.

Posterior parietal lesions: impaired ability to perform spatial task but can do matching

task.

Human patients.

Milner & Goodale's Patient D. F. (agnosic): Motor competence, good acuity, color, and motion perception. But can't recognize objects or faces. Further, DF cannot judge orientation or size. Normal intelligence and language comprehension.

Motor interactions are close to normal. She can "post" a letter-like object into a slot, adjust grasp size correctly before touching the object.

How can she “post” letter into an arbitrarily oriented slot with her hand but cannot tell you in advance whether the orientation of the slot is vertical or horizontal!

(Other patients have been studies that show the opposite pattern of deficits)

■ Spatial, action pathway

V1,MT,MST,LIP,...

Viewer-centered computations

"Where" vs. "What"

(**"where"** or **"how"** or **"now"**)

Spatial computations,such as coordinate transformations for action

■ Object perception, recognition pathway ("what")

object perception, recognition pathway

V1, V2, V4, Posterior IT, Anterior IT, ...

- generic feedforward neural network models PLUS feedback?

Receptive field properties

Size grows: V1 = .5°-2°; TE (mean) = 26°

Response properties become more difficult to characterize. Earlier RF responses often show approximate additivity--response is determined by the combination of responses to shape and orientation of parts or features. Neural RFs in later areas (TE), more difficult to characterize. Respond to highly specific features, sometimes interpretable (faces), sometimes more abstract ("toilet brush-like thing").

Some information has to be discounted, and other information selected that is diagnostic for accurate classification.

Invariances required for recognition:

photometric: illumination level, direction, shadows

geometrical: translation, size, orientation in depth

category-related: levels of abstraction

Combinatorial problem of object representation and classification

grandmother cells, distributed codes, sparse codes

Modeling large-scale neural systems & systems analysis

Much of the modeling of visual processing has been built on the tools that we've learned about. But there are many aspects of functional brain modeling that require additional tools and ways of thinking.

■ Modeling information processing functions

*Neural representation and coding:

grandmother cells, distributed codes, sparse codes, coarse codes

binding problem, "binding by synchrony"

Large-scale architectures (e.g. Inter-area processing) -- ?

Feedback

Information processing roles of feedback

*Dynamical behavior

Timing and sequences (e.g. speech, motor sequences)

Dynamical issues for real-time control, visuo-motor control

*Efficient coding, dimensionality reduction

*Handling uncertainty - Probabilistic models

■ Measuring and characterizing neural systems

Linear and non-linear systems analysis, statistical and stochastic processes analysis (time series),...

An adaptation demonstration: Orientation

Motivate concepts of coarse coding, population codes.

Make stimuli

In[140]:=
```
width = 64;
grating[x_,y_,xfreq_,yfreq_] := Cos[(2. Pi)*(xfreq*x + yfreq*y)];
```

■ **Left-slanted adapting grating**

In[142]:=
```
xfreq = 4; theta = 0.8 * Pi / 2;
yfreq = xfreq / Tan[theta];

gleft = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1}, {y, 0, 1},
   PlotPoints → 64, Mesh → False, Frame → False, DisplayFunction → Identity];
```

■ **Right-slanted adapting grating**

In[145]:=
```
xfreq = 4; theta = 1.2 * Pi / 2;
yfreq = xfreq / Tan[theta];

gright = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1}, {y, 0, 1},
   PlotPoints → 64, Mesh → False, Frame → False, DisplayFunction → Identity];
```

■ **Vertical test grating**

In[148]:=
```
xfreq = 4; theta = Pi / 2;
yfreq = xfreq / Tan[theta];

gvertical = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1}, {y, 0, 1},
   PlotPoints → 64, Mesh → False, Frame → False, DisplayFunction → Identity];
```


■ **Gray fixation bar**

In[151]:=
```
gbar =
  Graphics[{GrayLevel[.77], Rectangle[{0, .45}, {1, .5}]}, AspectRatio → 1 / 4];
```

Test: Try it

In[152]:=
```
Show[GraphicsArray[{{gleft, gvertical}, {gbar, gbar}, {gright, gvertical}},
   DisplayFunction → $DisplayFunction ]];
```

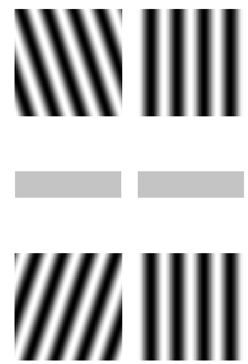

What happens if you adapt with the left eye and test with the right eye?

Can this effect be explained in terms of changes to neurons in the retina? LGN?

Neural codes and representation

What are the "languages" of the brain? How is information represented? What is the information in a train of action potentials?

Can we explain orientation adaptation?

First some background concepts.

Firing rate

Firing rate correlates well with subjective intensity in sensory systems. What does it mean elsewhere in the brain?

"Labeled lines"

Suppose that when a particular cell fires it means something in particular. A ganglion cell normally fires when stimulated by light coming from the upper left visual field. If the cell fires for any reason at all (e.g. you press on your eyeball), the fact that information is coming from this cell means "bright spot in upper left visual field". Similarly, for a pressure-sensitive cell on your finger tip. The identity of the cell that is firing represents information. Assuming a neuron has a "label" doesn't say much about how that label information gets passed around in the brain (other than by virtue of connectivity), but is useful for comparing behavioral/perceptual responses to neural measurements.

We noted that there are about 12 visual cortical areas with topographic maps of the visual field. Spatial location is represented on the cortical surface. If cells are labeled lines for position, then excitation of a cell signals information about location. But extra-striate areas have increasingly larger receptive fields with a coarser representation of space (see below).

Over 100 years ago, William James (1890) proposed the thought experiment that if we could splice the nerves so that the excitation of the ear sends input to the brain's visual area, and vice versa, we would “hear the lightning and see the thunder”. A recent study has re-wired visual input to the auditory cortex of the ferret, but contrary to James, the auditory cortex of the ferret adjusted to become like the visual area "V1", and further the behavior was consistent with the experience of sight being derived from visual inputs to their re-wired auditory cortex (von Melchner et al., 2000).

Distributed representations

Earlier in the coarse we discussed distributed vs. "grandmother" cell (or "local") representations of an object or event. Consider object memory. Suppose we have n neurons that can be active or inactive. In a grandmother cell representation, the activity of a single unit signals a unique object. There are strong theoretical arguments against a grandmother cell representation for objects--one needs an new neuron for every new object, i.e. there would be a single neuron whose firing would uniquely signal your "grandmother", hence the name. ("yellow volkswagen" detectors is another phrase of historical interest.). Representational capacity is n.

In a distributed representation, object identity is represented by the pattern.

Advantages to distributed coding?

- Capacity: If there are m distinguishable levels for each neuron, the system can represent m^n objects.
- Similarity between two patterns can be represented in the correlation (e.g. dot product, or angle between them, or cosine).

...but how are decisions made? i.e. "this is or is not my grandma".

Fully vs. sparse distributed representations

The latter case, m^n, would be a fully distributed system. But we noticed before that the cortex seems to be quiet on average--e.g. most V1 cells at any given moment are not firing. So maybe the truth is somewhere between, and an object is coded by a small population that is active for an event. This is sparse coding. We have n units, but an object is represented by the firing of p units, where 1<p<<n.

We discussed coding in terms of objects, but the issue is relevant for any kind of information. Later we'll talk about sparse coding of images. I.e. for any given image, say 256x256 8-bit pixels, what are the properties of V1 coding schemes that represent the image? The type of coding interacts with the statistical structure of the set of events to be encoded. If images were arbitrary, i.e. any image was equally likely (even TV snow), then we'd require representational space equal to the task (e.g. max representational capacity is 256x256x8 bits, or 2^(256x256x8) possible signals). But if there is statistical structure, we could get by with less. The space of natural images is much much smaller than 256x256x8 bits. Further, it turns out that the Gabor set described in the previous lecture is a sparse distributed code for natural images. There is more than one cell activated for any given image, but the number is relatively small.

Coarse vs. fine coding

Features are coarsely sampled (few detectors to span the range), and the receptive fields broad (so no empty regions). Broadly tuned cells mean that similar inputs to the cell's preferred input also fire the cell. Coarse coding means that the (broad) receptive fields overlap.

Fine coding means that the neurons finely sample the feature space, with correspondingly narrower receptive fields. Neurons are more closely tuned to the exact feature, and show little or no response to similar features.

How can information be represented given a coarse code?

Let's construct an hypothetical tuning function for some feature s with tuning width w:

```
truncCos[s_, w_] := If[-w / 2 < s < w / 2, Cos[Pi * (1 / w) * s], 0];
Plot[truncCos[s, 2], {s, -3, 3}];
```

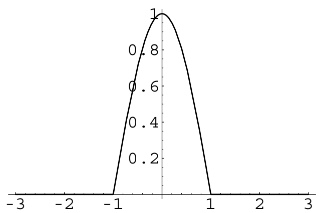

Try the above cell with various values of w.

Coarse coding

```
w = 4;
R[i_, x_] := truncCos[x - i, w];
Plot[{R[-3, x], R[0, x], R[3, x]}, {x, -10, 10}];
```

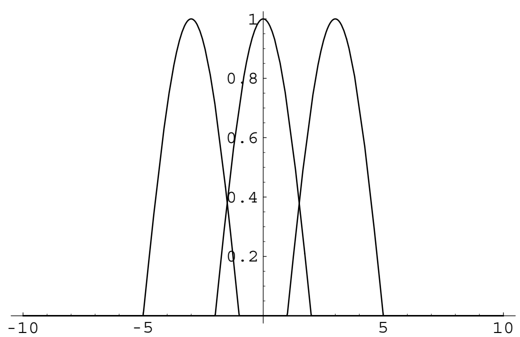

E.g. wavelength coding in the retina

Finer coding


```
w = 1;
R[i_, x_] := truncCos[x - i, w];
Plot[{R[-3, x], R[-2, x], R[-1, x], R[0, x], R[1, x],
   R[2, x], R[3, x], R[4, x], R[5, x], R[6, x]}, {x, -10, 10}];
```

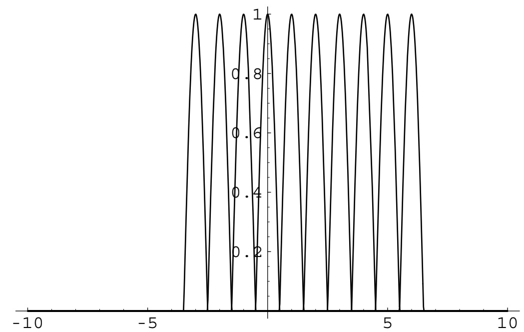

E.g. positional coding in retina. V1 shows finer coding for position than extra-striate topographic areas, like V2.

Population vector coding

■ Definition

Receptive fields typically overlap (e.g. a bright spot at one location creates a neural point spread function, the "projective field"). E.g. a bar at one orientation will more or less activate cells within a certain feature range (+/- 15 deg in V1).

Suppose we have access to the responses of a bunch of neurons all "seeing" the same stimulus bar. How can information be extracted from this pattern of activity? We aren't going to answer this with a neural mechanism, but rather with an interpretive measure that an experimenter could employ.

One can combine information across a population in terms of a "population vector".

Let x_k be a vector representing a stimulus feature (e.g. the kth 2-D position, or kth motion direction, or kth orientation, etc.). Let the firing rate of the ith cell be $R_i(x_k)$ in response to input x_k . Let the feature that produces the peak response of the ith cell be x_i^p --i.e. x_i^p is the ith cell's preferred feature, the one that fires the cell the most.

Given an input feature x_k , the firing rate of the ith cell can be interpreted as the strength of its "vote" for its preferred feature.

So for this example, position could be represented by the weighted average over the population of cells, each responding with various firing rates to x_k :

$$\mathbf{x} = \sum_i R_i(x_k)\, x_i^P \qquad (1)$$

Analogous to computing the center of mass, or an average, we can normalize the estimate by the total activity:

$$x = \sum_i R_i (x_k)\, x_i^p \Big/ \sum_i R_i (x_k) \tag{2}$$

This measure has not only been applied to modeling sensory coding, but also in the motor system and cognitive processes involving direction of movement (Georgopoulos et al., 1993). In certain tasks the population vector can be measured in real neuronal ensembles and be seen to evolve in time consistent with behavioral measures (mental rotation, reach planning).

■ Illustration of population coding

Simple example. Population vector will be 1-D, i.e. just a scalar. Coarse coding with only 3 units, and broad tuning with width = 10.

Suppose the input to the population is xinput = 5.

```
xinput = 5; (*stimulus feature*)

w = 10;
R[i_, x_] := N[truncCos[x - i, w]];
input[x_] := If[Abs[x] < .1, 1, 0];
(* just for plotting purposes-- a narrow deltafunction-like pulse *)

prefx = {1, 4, 7}; (*preferred features for the 3 cells*)

responses = Table[R[prefx[[i]], xinput], {i, 1, Length[prefx]}];

Plot[{R[prefx[[1]], x], R[prefx[[2]], x], R[prefx[[3]], x], input[x - xinput]},
  {x, 0, 10}, PlotStyle → {RGBColor[0, 0, 0],
    RGBColor[0, 0, 0], RGBColor[0, 0, 0], RGBColor[1, 0, 0]}];
```

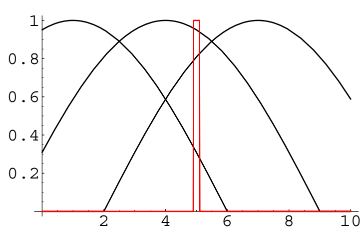


```
Plot[
  {responses[[1]] input[x - prefx[[1]]], responses[[2]] input[x - prefx[[2]]],
   responses[[3]] input[x - prefx[[3]]]}, {x, 0, 10}, PlotPoints → 100];
```

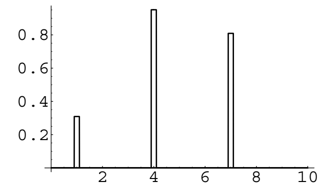

Population vector response

```
xpop = responses.prefx / Apply[Plus, responses]
```

```
4.72496
```

How could accuracy be improved?

Noise and quantization

Suppose position is coarsely coded by just two neurons, with peak sensitivity say at 0 and the other at position 1, but with broad receptive fields spanning both. Does this mean that position can't be represented accurately and reliably?

No. There are only 3 cone photoreceptor types for daytime vision, but we can distinguish thousands of colors.

Discussion or project question: Where might wavelength discrimination be best? At the peak sensitivities? At the cross-over points?

But is firing rate the "code"? Timing

Information in the detailed timing of spikes?See F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek (1996).

Information in the temporal coherence across spiking ensembles? "Binding by synchrony", see Shadlen & Movshon.

Adaptation and population coding

```
adaptstrength = {1, 1, 0.3};
xinput = 5; (*stimulus feature*)

w = 10;
R[i_, x_] := N[truncCos[x - i, w]];
input[x_] := If[Abs[x] < .1, 1, 0];
(* just for plotting purposes-- a narrow deltafunction-like pulse *)

prefx = {1, 4, 7}; (*preferred features for the 3 cells*)

responses = Table[R[prefx[[i]], xinput], {i, 1, Length[prefx]}];
responses = adaptstrength responses

Plot[{adaptstrength[[1]] * R[prefx[[1]], x],
   adaptstrength[[2]] * R[prefx[[2]], x],
   adaptstrength[[3]] * R[prefx[[3]], x], input[x - xinput]},
  {x, 0, 10}, PlotStyle → {RGBColor[0, 0, 0], RGBColor[0, 0, 0],
    RGBColor[0, 0, 0], RGBColor[1, 0, 0]}];
```

```
{0.309017, 0.951057, 0.242705}
```

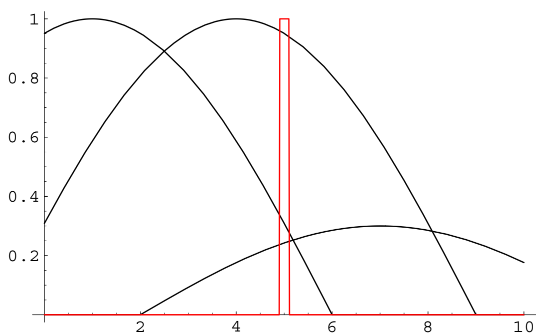


```
xpopadapt = responses.prefx / Apply[Plus, responses]
```

```
3.86762
```



```
Plot[{R[1, x], R[4, x], 0.3 * R[7, x], input[x - xpop], input[x - xpopadapt]},
  {x, 0, 10}, PlotStyle → {RGBColor[0, 0, 0], RGBColor[0, 0, 0],
    RGBColor[0, 0, 0], RGBColor[1, 0, 0], RGBColor[1, 0, 1]}, PlotPoints → 100];
```

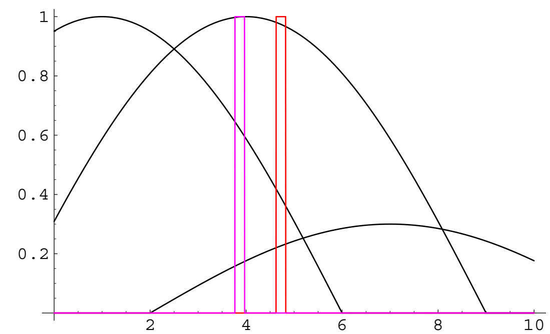

Explanation of orientation after-effect?

Contingent Adaptation: McCollough effect

Celeste McCollough (1965).

Make stimulus

In[169]:=
```
width = 64;
grating[x_,y_,xfreq_,yfreq_] := Cos[(2. Pi)*(xfreq*x + yfreq*y)];
```

■ Vertical red adapting grating

In[171]:=
```
xfreq = 4; theta = Pi / 2;
yfreq = xfreq / Tan[theta];

gvertred = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1},
   {y, 0, 1}, PlotPoints → width, Mesh → False, Frame → False,
   ColorFunction → (RGBColor[#, 0, 0] &), DisplayFunction → Identity];
```

■ **Horizontal green adapting grating**

In[174]:=
```
xfreq = 0; theta = 0;
yfreq = 4;

ghorizgreen = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1},
   {y, 0, 1}, PlotPoints → width, Mesh → False, Frame → False,
   ColorFunction → (RGBColor[0, #, 0] &), DisplayFunction → Identity];
```

■ **Horizontal gray test grating**

In[177]:=
```
xfreq = 0;
yfreq = 4;

ghorizgray = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1},
   {y, 0, 1}, PlotPoints → width, Mesh → False, Frame → False,
   ColorFunction → (RGBColor[#, #, #] &), DisplayFunction → Identity];
```

■ **Vertical gray test grating**

In[180]:=
```
xfreq = 4;
yfreq = 0;

gvertgray = DensityPlot[grating[x, y, xfreq, yfreq], {x, 0, 1},
   {y, 0, 1}, PlotPoints → width, Mesh → False, Frame → False,
   ColorFunction → (RGBColor[#, #, #] &), DisplayFunction → Identity];
```

■ **Gray fixation bars**

In[183]:=
```
gbar =
  Graphics[{GrayLevel[.77], Rectangle[{0, .45}, {1, .5}]}, AspectRatio → 1 / 4];
```


Test: Try it

In[184]:=
```
Show[GraphicsArray[
   {{gvertred, ghorizgray}, {gbar, gbar}, {ghorizgreen, gvertgray}},
   DisplayFunction → $DisplayFunction , GraphicsSpacing → {.5, .05}]];
```

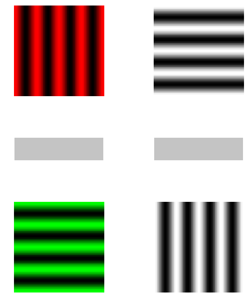

■ **Or display to a new notebook:**

```
displayToNotebook[graphic_, opts___] :=
 (nb = NotebookCreate[WindowSize → {350, 350}, WindowFrame → "Palette"];
  NotebookWrite[nb, Cell[GraphicsData["PostScript", DisplayString[graphic]],
    "Graphics", DisplayFunction → $DisplayFunction, opts]]; graphic)
opts1 = Sequence /@ {Frame → False, Mesh → False,
   DisplayFunction → $DisplayFunction}
```

In[167]:=
```
displayToNotebook[GraphicsArray[
   {{gvertred, ghorizgray}, {gbar, gbar}, {ghorizgreen, gvertgray}}], opts1];
```

Is there a functional explanation for what is going on? Is there a neural network explanation?

Barlow suggested that the adaptation is the result of the visual system adjusting to new statistical dependencies between features. It isn't just that neurons are "getting tired". Specifically, he suggested that perceptual systems adjust their representations of features to be independent or uncorrelated with each other (technically, independence is a stronger condition). This is a problem of neural self-organization, perceptual learning based on experience. Next we'll look at several neural networks that re-organize to decrease the correlations between their firing. One network may provide an explanation for McCulloch's color-contingent affter-effect.

Bibliography

Barlow, H. B., & Foldiak, P. (1989). Adaptation and decorrelation in the cortex. In C. Miall, R. M. Durban, & G. J. Mitchison (Ed.), The Computing Neuron Addison-Wesley.

Barlow, H. (1990). Conditions for versatile learning, Helmholtz's unconscious inference, and the task of perception. Vision Research, 30(11), 1561-1572.

Field, D. J. (1994). What is the goal of sensory coding? Nueral Computation, 6, 559-601.

Georgopoulos, A. P., Lurito, J. T., Petrides, M., Schwartz, A. B., & Massey, J. T. (1989). Mental Rotation of the Neuronal Population Vector. *Science, 243*, 234-236.

Lee, C., Rohrer, W. H., & Sparks, D. L. (1988). Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature, 332*(6162), 357-360.

McCollough, C. (1965, 3 September 1965). Color Adaptation of Edge-Detectors in the Human Visual System. *Science, 149,* 1115-1116.

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. Nature, 381, 607-609.

F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek (1996). Spikes: Exploring the neural code (MIT Press, Cambridge).

Shadlen, M. N., & Movshon, J. A. (1999). Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron, 24*(1), 67-77, 111-125.

Sillito, A. M., Jones, H. E., Gerstein, G. L., & West, D. C. (1994). Feature-linked synchronization of thalamic relay cell firing induced by feedback from the visual cortex. Nature, 369, N. 9, 479-482.

von Melchner, L., Pallas, S. L., & Sur, M. (2000). Visual behaviour mediated by retinal projections directed to the auditory pathway. *Nature, 404*(6780), 871-876.