Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Contrast normalization, 2nd order histogram analysis of filter outputs

Initialize

■ **Spell check off**

In[537]:=
```
SetOptions[ArrayPlot, ColorFunction → "GrayTones", DataReversed → True,
  Frame → False, AspectRatio → Automatic, Mesh → False,
  PixelConstrained → True, ImageSize → Small];
```

In[538]:=
```
nbinfo = NotebookInformation[EvaluationNotebook[]];
dir =
 ("FileName" /. nbinfo /. FrontEnd`FileName[d_List, nam_, ___] :→
    ToFileName[d]);
```

In[539]:=
```
Off[General::spell1];
```

Adaptive spatial filters & V1 cells

■ **Simultaneous contrast**

A classic brightness illusion demonstrates what we might expect of a spatial filter that adapts its response to the response of its neighbors. The circle on the left appears to be darker than the circle on the right; however, the intensity is the same for both (**ming**). One way to explain this is that the response of a single unit that signals intensity gets divided by a measure of the magnitude of the responses of neighboring units that also signal intensity.

In[540]:=
```
width = 256; radius = width/5;
maxg = 0.85`; ming = 0.5`; maxg2 = 0.55`;
d1 = Table[If[(i - width/2)^2 + (j - width/2)^2 < radius^2, ming, maxg],
   {i, 1, width}, {j, 1, width}];
g1 = ArrayPlot[d1, PlotRange → {0, 1}];
d2 = Table[If[(i - width/2)^2 + (j - width/2)^2 < radius^2, ming, maxg2],
   {i, 1, width}, {j, 1, width}];
g2 = ArrayPlot[d2, PlotRange → {0, 1}];
GraphicsRow[{g1, g2}, Spacings → Scaled[0.1]]
```

Out[546]=

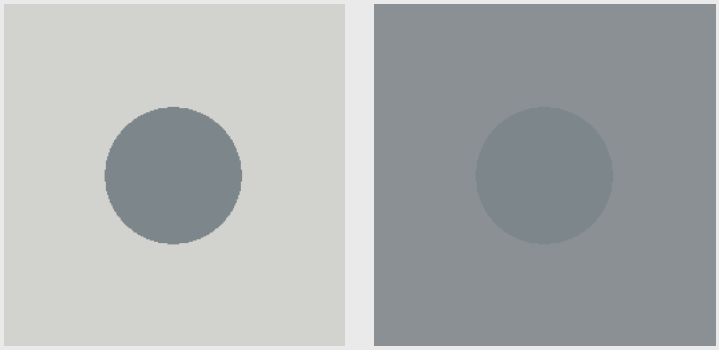

■ Simultaneous "Contrast of Contrast"

The above illusion involved manipulating mean light level. What if we manipulate contrast (or variance) and keep the means fixed?

In[547]:=
```
width = 256; radius = width/5;
maxg := 0.8` (RandomReal[] - 0.5`) + 0.5`;
ming := 0.25` (RandomReal[] - 0.5`) + 0.5`;
maxg2 := 0.15` (RandomReal[] - 0.5`) + 0.5`;
d1 = Table[If[(i - width/2)^2 + (j - width/2)^2 < radius^2, ming, maxg],
   {i, 1, width}, {j, 1, width}];
g1 = ArrayPlot[d1, Mesh → False, PlotRange → {0, 1}];
d2 = Table[If[(i - width/2)^2 + (j - width/2)^2 < radius^2, ming, maxg2],
   {i, 1, width}, {j, 1, width}];
g2 = ArrayPlot[d2, Mesh → False, PlotRange → {0, 1}];
Show[GraphicsRow[{g1, g2}, Spacings → Scaled[0.1`]]]
```

Out[555]=

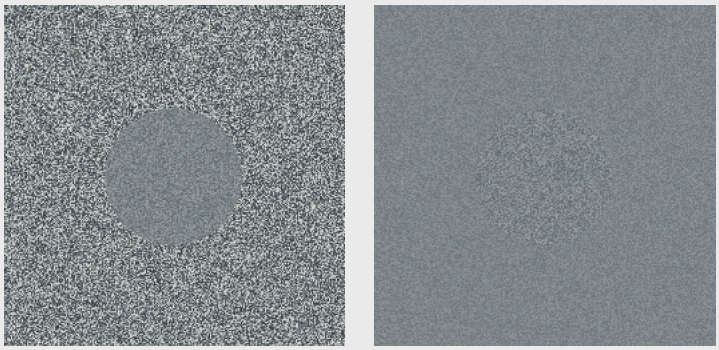

■ Contrast normalization

It turns out that neurons in V1 show an analogous response to your own perception of contrast. One way to model this is to assume that the response of a single unit that signals contrast for a particular location, spatial frequency and orientation preference, gets divided by the average of a measure of the magnitude of the responses of neighboring units that also signal contrast over a range of spatial frequencies and orientations.

The linear spatial receptive field model for a V1 neuron says that that response should scale linearly with contrast. But simple cells don't show this property--instead, the response begins to saturate at high input contrasts (e.g. for a drifting sinusoidal grating matching the orientation, spatial frequency and motion direction preferences of the cell). Time-wise, the response also begins to occur sooner as the stimulus contrast is increased. Another break-down is seen in the response of a cell to the combination of a horizontal and vertical sinusoidal gratings. Linearity predicts the response to the sum should be the sum of the responses, but it isn't. Instead neurophysiologists find "cross-orientation inhibition". Interestingly enough, a cell that prefers say the vertical grating will typically show zero response to the horizontal one; yet, the presence of the horizontal one still inhibits the cell's firing to the vertical.

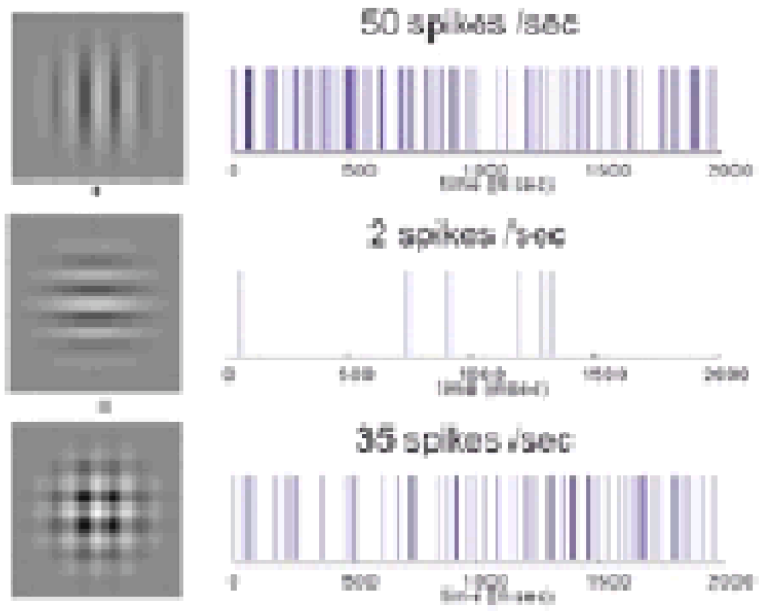

A model that quantitatively accounts for the responses of simple cells is called the "contrast normalization" model (Heeger, 1997 ; Carandini, M., & Heeger, D. J. (1994)).

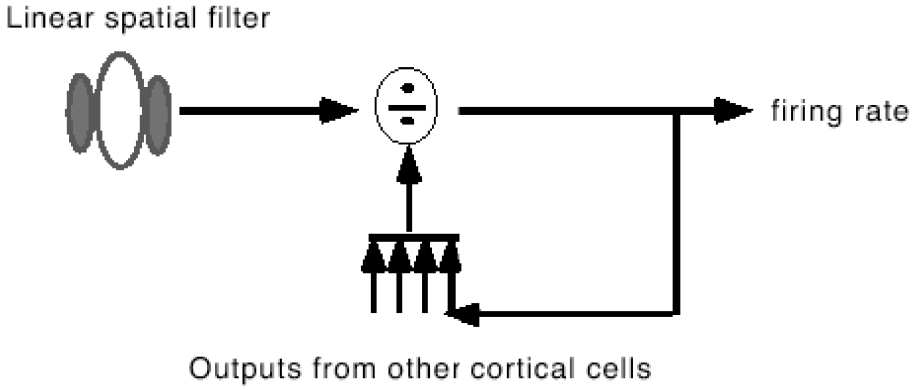

■ Contrast normalization & image statistics

We've noted that a number of image basis sets yield image coefficients that are essentially uncorrelated for natural images. Fourier, and PCA bases all tend to produce uncorrelated outputs for natural images. Wavelets and gabor-filters also tend to produce uncorrelated outputs for natural images. Wavelets may be close to providing an independent components representation for natural images. Independent Components Analysis (ICA) seeks basis sets for which p(a,b) = p(a)p(b), where a and b are coefficients. This is in contrast to PCA which seeks bases for which E(ab) = E(a)E(b) (Recall that independence implies no correlation, but not the reverse).

Consider a standard model for V1 cells consisting of a set of gabor-like filters, or more specifically a wavelet basis set selective for several orientations and spatial scales, spread across the image domain. Despite the fact that the responses are decorrelated, gathering image statistics shows that the wavelet coefficients are not statistically independent.

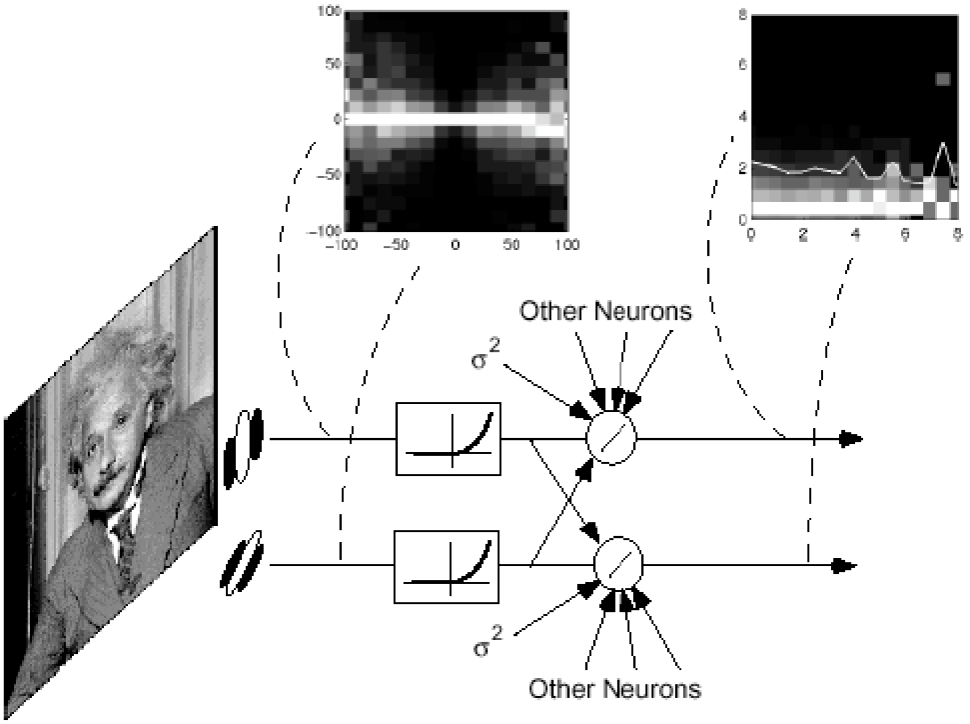

Figure 1: Illustration of image statistics as seen through two neighboring receptive fields. Left image: Joint conditional histogram of two linear coefficients. Pixel intensity corresponds to frequency of occurrence of a given pair of values, except that each column has been independently rescaled to fill the full intensity range. Right image: Joint histogram of divisively normalized coefficients (see text).

Figure from: Eero P Simoncelli and Odelia Schwartz (1998) Modeling Surround Suppression in V1 Neurons with a Statistically-Derived Normalization Model . Advances in Neural Information Processing Systems 11. ed. M.S. Kearns, S.A. Solla and D.A. Cohn, pp. 153-159, May 1999. © MIT Press, Cambridge, MA. See: http://www.cns.nyu.edu/~eero/-ABSTRACTS/simoncelli98d-abstract.html

The figure shows that the expected value of the ordinate coefficient is about zero regardless of the abscissa value, i.e. the correlation is about zero. However, the variance of the ordinate increases with the absolute value of the abscissa--the responses are not independent of each other. Simoncelli & Schwartz (1998) report this pattern for pairs of coefficients at neighboring spatial positions, orientation and scales. Remarkably, the statistical dependence goes away if the image filter responses are normalized by the sum of responses of nearby filters. In other words, the contrast normalization model (above) could be the consequence of an efficient coding principle.

(The Appendix of **12. SpatialCodingEfficiency.nb** has code for calculating a 2D histogram plot of joint spatial filter statistics for an image.)

2nd order statistics spatial filter

- **Joint histogram of two overlapping, possibly orthogonal, filters**

Some useful functions: scale256, histogram, argmax

In[556]:=
```
scale256[image_] := Module[{α, β},
   α = 255 / (Max[image] - Min[image]);
   β = -α Min[image];
   Return[α image + β];];
```

In[557]:=
```
myhistogram[image_] := Module[{histx},
   histx = BinCounts[Flatten[image], {0, 255, 1}];
   Return[N[histx / Plus @@ histx]];
  ];
```

In[558]:=
```
argmax[x_] := Position[x, Max[x]][[1, 1]];
```

Input image: " alpine.jpg" is a 256x256 array of gray-levels

In[559]:=
```
alpine = Import["GrayLonesome256x256.jpg", Path → dir][[1, 1]];
width = Dimensions[ alpine][[1]];
alpine256 = N[scale256[ alpine]];
```

■ Set up kernels for the filters

In[562]:=
```
sgabor[x_, y_, fx_, fy_, sig_] :=
  N[Exp[(-x^2 - y^2) / (2 sig * sig)] Sin[2 Pi (fx x + fy y)]];
fsize = 16;
filter2 = Table[sgabor[(i - fsize / 2), (j - fsize / 2), 1 / 8, -1 / 8, 4],
  {i, 0, fsize}, {j, 0, fsize}];
filter2 = Chop[filter2];
ArrayPlot[filter2, PlotRange → {-1, 1}, ImageSize → Tiny]
```

Out[565]=

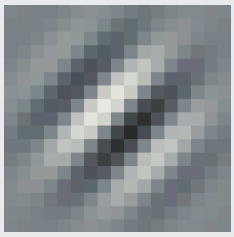

In[566]:=
```
filter = Table[sgabor[(i - fsize / 2), (j - fsize / 2), 1 / 8, 1 / 8, 4],
   {i, 0, fsize}, {j, 0, fsize}];
filter = Chop[filter];
ArrayPlot[filter, Mesh → False, PlotRange → {-1, 1}, ImageSize → Tiny]
```

Out[568]=

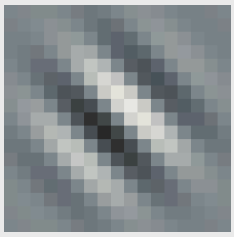

■ Convolution

In[569]:=
```
falpine = ListConvolve[filter, alpine256];
falpine256 = scale256[falpine];

falpineB = ListConvolve[filter2, alpine256];
falpine256B = scale256[falpineB];

GraphicsRow[{ArrayPlot[falpine256], ArrayPlot[falpine256B]}]
```

Out[573]=

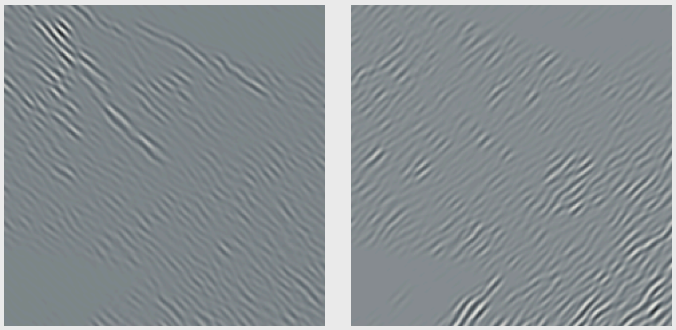

■ 2D histogram

The above filters are centered at the place, but orthogonal to each other. We can also shift the filtered ouputs relative to each other with the variable

offset below. offset=0 corresponds to no shift. binsize controls the bin size.

In[574]:=
```
offset = 4;
binsize = 4;

falpine256n = Round[(falpine256 - Mean[falpine256])];
falpine256Bn = Round[(falpine256B - Mean[falpine256])];
falpine256Bn = RotateLeft[falpine256Bn, offset];

temp = Transpose[{Flatten[falpine256n], Flatten[falpine256Bn]}];
twoDhist = BinCounts[temp, {-130, 130, binsize}, {-130, 130, binsize}];
```

In[581]:=
```
ArrayPlot[Log[0.0001 + twoDhist]]
```

Out[581]=

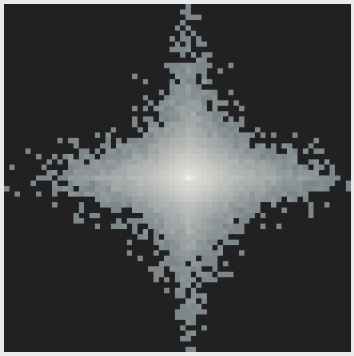

Now normalize the columns to represent conditional probabilities--i.e. elements in a column sum to one.

In[582]:=
```
twoDhistn = Table[(twoDhist[[i, j]]) / (.00001 + Total[twoDhist[[All, j]]]),
   {i, 1, Dimensions[twoDhist][[1]]}, {j, 1, Dimensions[twoDhist][[1]]}];
```

In[583]:=
```
ArrayPlot[twoDhistn, PlotRange → {0, .3}]
```

Out[583]=

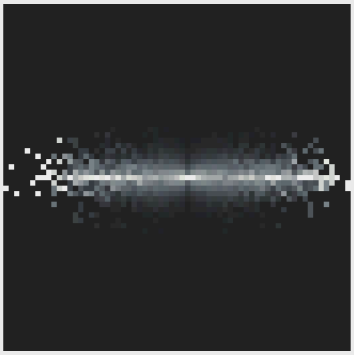

References

Carandini, M., Heeger, D. J., & Movshon, J. A. (1997). Linearity and normalization in simple cells of the macaque primary visual cortex. *J Neurosci, 17*(21), 8621-8644.

Carandini, M., & Heeger, D. J. (1994). Summation and division by neurons in primate visual cortex. *Science, 264*(5163), 1333-1336.

Friston, K. (2005). A theory of cortical responses. Philos Trans R Soc Lond B Biol Sci, 360(1456), 815-836.

Heeger, D. J., Simoncelli, E. P., & Movshon, J. A. (1996). Computational models of cortical visual processing. *Proc Natl Acad Sci U S A, 93*(2), 623-627.

Eero P Simoncelli and Odelia Schwartz (1998) Modeling Surround Suppression in V1 Neurons with a Statistically-Derived Normalization Model . Advances in Neural Information Processing Systems 11. ed. M.S. Kearns, S.A. Solla and D.A. Cohn, pp. 153-159, May 1999. © MIT Press, Cambridge, MA.

E P Simoncelli. Statistical models for images: Compression, restora-tion and synthesis. In 31st Asilomar Conf Signals, Systems and Com-puters, pages 673-678, Pacific Grove, CA, November 1997. Available from http://www.cns.nyu.edu/~eero/publications.html

Tolhurst, D. J., & Heeger, D. J. (1997). Comparison of contrast-normalization and threshold models of the responses of simple cells in cat striate cortex. *Vis Neurosci, 14*(2), 293-309.

See: http://www.cns.nyu.edu/~eero/ABSTRACTS/simoncelli98d-abstract.html