Computational Vision
U. Minn. Psy 5036
Daniel Kersten
Lecture 9: Spatial Frequency analysis

http://vision.psych.umn.edu/www/kersten-lab/courses/Psy5036/SyllabusF2000.html)

■ **Initialize:**

```
Off[General::spell1];
```

Outline

Last time: Key ideas

1. The output or response of a linear system can be modeled as a matrix operation. If the system is shift-invariant, the matrix has the form of a convolution.

2. Some representations of images are better than others.

3. Spatial frequency representation of images good for modeling optical transformation.

E.g. Eigenfunctions of the system: Avoids convolution--one can project the image onto the appropriate basis set providing the spectrum. Then scale each eigenvector by the product of spectrum with the MTF (i.e. eigenvalues of the system), and then add them all up.

This time

Single-channel spatial filtering

Multiple channel filters

Psychophysical experiments.

->Multi-resolution, and wavelet bases

->A model of the spatial filtering properties of neurons in the primary visual cortex

Tutorials

■ **Mathematica tutorial on convolutions**

■ **Mathematica tutorial on fourier analysis of images**

Single channel spatial frequency filtering

Mach bands & perception

In the 19th century, the Austrian scientist Ernst Mach noticed that the brightness of a luminance ramp didn't look like one would predict simply from physical measurements of light intensity.The red line below is proportional to the actual intensity. The blue line shows an informal sketch of apparent brightness.Why is this? Mach advanced an explanation in terms of lateral inhibition which hasn't changed much in 100 years. We'll return to it below. But first let's take a look at what is known about retinal anatomy and physiology.

Ernst Mach was an Austrian physicist and philospher. In addition to being well-known today for a unit of speed, he is also known for several visual illusions. One illusion is called "Mach bands". Let's make some.

```
size = 120;
Clear[y];
low = 0.2; hi = 0.8;
y[x_] := low /; x<40
y[x_] :=
    ((hi-low)/(size/3)) x + (low-(hi-low)) /; x>=40 && x<80
y[x_] := hi /; x>=80
```

```
Plot[y[x],{x,0,120},PlotRange->{0,1}];
```

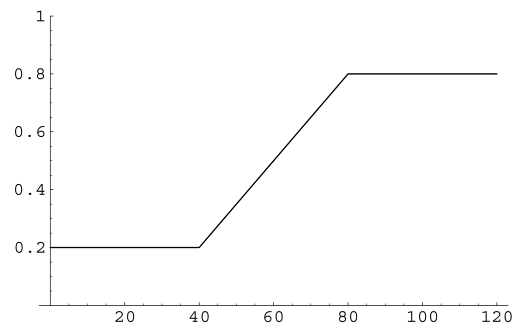

Let's make a 2D gray-level picture displayed with **ListDensityPlot** to experience the Mach bands for ourselves. **PlotRange** allows us to scale the brightness.

```
picture = Table[Table[y[i],{i,1,size}],{i,1,size/2}];
ListDensityPlot[picture,Frame->False,Mesh->False,
                PlotRange->{0,1}, AspectRatio->Automatic];
```

What Mach noticed was that the left knee of the ramp looked too dark, and the right knee looked too bright. Objective light intensity did not predict apparent brightness.

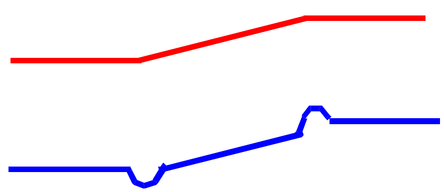

■ **Mach's explanation**

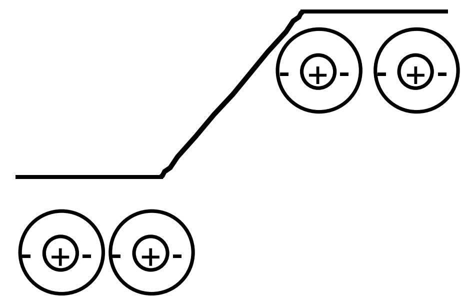

Physiological basis for spatial filtering

Neural basis?Lateral inhibitory filtering in sensory cells. Found in: Invertebrates:Limulus (horseshoe crab)--Hartline, who won the 1967 Nobel prize for this work that began in the 30's; vertebrates: Frog - Barlow, and mammals: Cat --Kuffler.

The receptors are connected to horizontal and bipolar cells which in turn are connected to amacrine cells. One overall function of this circuitry is to provide gain control, with high-pass spatial and temporal filtering, so that the retinal output signals light levels relative to mean level, rather than signalling absolute level.The function of the retina may best be seen by a close examination of what its output cells are doing. These cells are called ganglion cells and they send their axons out of the eye through the optic disk ("blind spot") to the lateral geniculate nucleus (LGN). The LGN has been crudely (and perhaps erroneously) likened to a relay station en route to cortex. Several types of ganglion cells, each with distinctive anatomy and function have been identified in cat and monkey (Enroth-Cugell and Robson, 1966; see Shapley and Perry, 1986 for a comparison with monkey ganglion cells). In cat, the two principle types are the X, Y cells. They code contrast into trains of action potentials (spikes) whose temporal frequency grows with contrast (see figure). In addition, these cells act as approximately circularly symmetric spatial-temporal band-pass filters, with small departures from linearity. What this means should become clear after we explain the idea of a receptive field.

If one measures the response of a ganglion cell to a uniformly illuminated screen, one typically finds a mean spike discharge rate (e.g. 50 spikes/sec). Then if a small light spot is positioned on the screen, the cell will increase its firing for some positions, and decrease its firing rate for other positions. A "receptive field" can be mapped out which shows the sensitivity of the cell to the light spots in various locations. One of the striking findings of the 1950's was the discovery that retinal cells almost uniformly show a concentric center-surround organization in which the center is excitatory and the surround inhibitory (or vice versa). The figure below shows a density plot where light areas represent where the cell increases firing to a spot, and the darker areas where it is inhibited

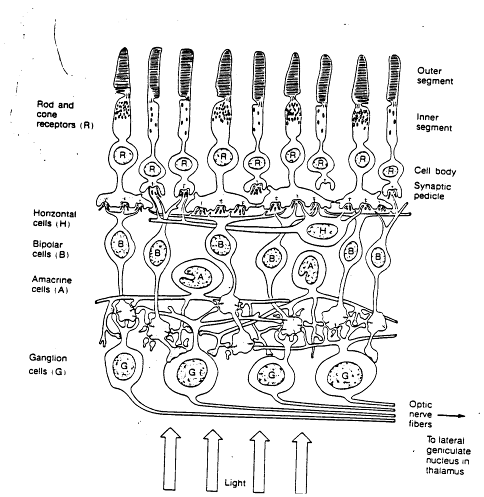

Three forms for the "mexican-hat" filter: w(x,y)

■ Difference-of-Gaussians (DOG)

```
DOG[x_, y_, s1_, s2_] := 2 Exp[(-x^2 - y^2) / s1^2] - Exp[(-x^2 - y^2) / s2^2];
```

```
Plot[DOG[x,0,1,2],{x,-4,4}];
```

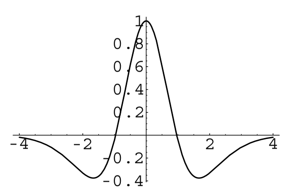


```
DensityPlot[DOG[x, y, 1, 2], {x, -4, 4}, {y, -4, 4},
  Mesh → False, PlotPoints → 64, PlotRange → {-1, 1}];
```

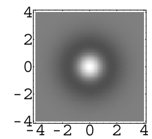

■ Radial Cosine gabor function

```
cgabor[x_, y_, fx_] := Exp[(-x^2 - y^2)] Cos[2 Pi (fx Sqrt[(x^2 + y^2)])];
DensityPlot[cgabor[x, y, .25], {x, -4, 4}, {y, -4, 4},
  Mesh → False, PlotPoints → 64, PlotRange → {-.25, .25}];
```

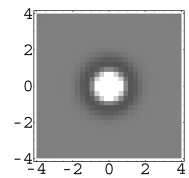

■ $\nabla^2 G$

```
blur[x_, y_] := Exp[-x^2 - y^2];
dx[x_, y_] := D[blur[u, v], u] /. u → x /. v → y
dy[x_, y_] := D[blur[u, v], v] /. u → x /. v → y
delsqG[x_, y_] := D[blur[u, v], {v, 2}] + D[blur[u, v], {u, 2}] /. u → x /. v → y

DensityPlot[-delsqG[x, y], {x, -4, 4}, {y, -4, 4},
  Mesh → False, PlotPoints → 64, PlotRange → {-2, 2}];
```

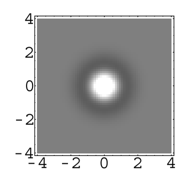

Spatial filtering

If one thinks of the retina (or more precisely a small patch) as a having a whole array of identical linear ganglion cells, one could also interpret the receptive field as the analog of the point spread function of the optics. So now the above figure would represent the strength of the response of the array of ganglion cells at each location to a small spot of light.

This concentric antagonistic spatial organization is called lateral inhibition. There are both ON-center and OFF-center types of ganglion cells. How can we quantitatively model the ganglion cell response? Let $r_{k,l}$ be the response (in spikes/sec) of a ganglion cell at x-y location (k,l). The average response, to a first approximation, is determined by the weighted sum of the input light intensities, $g_{i,j}$ at spatial location (i,j)

$$r_{k,l}=\sum_{i,j} w_{k,l;\,i,j}\, g_{i,j}$$

As with the PSF, we assume spatial homogeneity, and thus shift-invariance:

$r_{k,l}=\sum_{i,j} w_{k-i;\,l-j}\, g_{i,j}$. Or by suitable arrangement of rows and columns as matrix operation, **r = W.g**

In the continuous case,by the convolution integral:

$$r(x,y) = w*g = \int\int w(x-x', y-y')\, g(x', y')\, dx'\, dy'$$

Neural signals often show small-signal suppression, and high-signal saturation. We can describe this as a simple *point-wise* non-linearity, $\sigma()$:

As we have seen earlier, one can describe a linear system either in the spatial or spatial frequency domain. Thinking in frequency terms, a center-surround filter (as in the above figure) can also be interpreted as a band-pass filter. One way to see this is to note that the amplitude spectrum of the Fourier transform of a DOG shaped filter has the shape of an inverted U -- that is, the contrast amplitudes of low and high spatial frequencies are attenuated relative to the middle frequencies. The optimal frequency, in fact, would have a periodicity close to that of the receptive field. If one uses the simple linear model, then the integral of the receptive field has to be zero; otherwise, the filter will generate a "d.c." offset that is proportional to the mean light level. Recall that one of the striking properties of X and Y cat ganglion cells is that they show no sensitivity to absolute light level.

■ So what about Mach bands?

We've seen Mach's qualitative explanation. In Mathematica, it is straightforward to convolve a DOG shaped filter with Mach’s luminance ramp to predict a response.


```
LDOG = 0.25 Table[DOG[x, 0, 1 / 1.9, 2 / 2], {x, -4, 4, .25}];
ListPlot[LDOG , PlotJoined → True, PlotRange → {-.5, .5}];
r = ListConvolve[LDOG, picture[[1]]];
ListPlot[r, PlotJoined → True, PlotRange → {-.1, .1}];
```

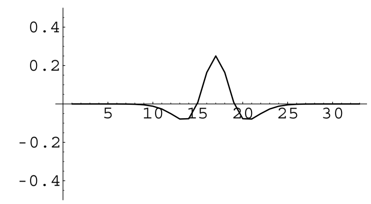

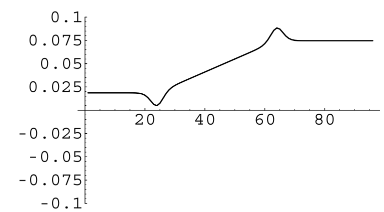

■ Exercise: What happens (to perception and to the model) when the slope of the ramp increases, approachin that of a step function?

■ And what about the shape of the Contrast Sensitivity Function (CSF)?

The figure (`CSF.gif`) is a demonstration of how your own contrast sensitivity function depends on spatial frequency.

With a well – calibrated computer screen, the contrast along any horizontal line is the same, independent of spatial frequency. The contrast declines going up the figure. If your contrast sensivity was independent of spatial frequency, you'd see the transition from visible to invisible as a straight horizontal boundary. But you don't. You see the middle range better than high or low frequencies. Further, you can demonstrate to yourself that this is a property of spatial frequency measured in retinal units (e.g. in cycles per degree of visual angle), because the apparent peak will shift with viewing distance.Peak sensitivity is typically between 3 and 5 cycles per degree.

What if the visual system’s spatial contrast processing of spatial frequency was linear throughout (i.e. not just the optics)? Further suppose the system’s contrast threshold was determined by the ganglion cell receptive field. How could you calculate a quantitative prediction for the receptive field of the ganglion cell? Knowing the spatial receptive field property of a retinal ganglion cell would lead naturally to a prediction of what the contrast sensitivity function should look like. But one has to assume spatial homongeneity.

Contrast sensitivity as a function of spatial frequency has an inverted U-shape:

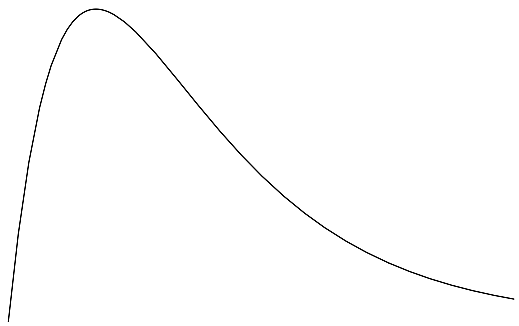

The visual system responds well to contrasts over a middle range, but not as well to low or high spatial frequencies.

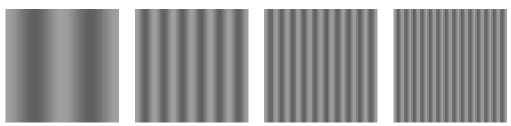

One theory (the "single-channel") theory can account for this if the size and spacings of the excitatory and inhibitory regions of the receptive field weighting function,

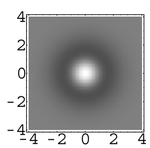

are well-matched to a middle frequency. Although the single-channel theory works here, it doesn't account for other results that we will get to later. One theory is that human CSF is probably determined by the envelope of the contrast sensitivities of a family of cell types in primary visual cortex of the brain.

Multiple channel spatial filtering

Campbell and Robson -- 1968

■ Contrast sensitivity for a cosine grating and for a square wave grating

Given a cosine grating and a square-wave grating of equal physical contrasts, which one is more detectable near threshold? Or would contrast sensitivity be the same?

```
lowcontrast = 0.0125; mediumcontrast = 0.125;
contrast = mediumcontrast;
Grating[x_, y_, fx_, fy_] := Sin[2 Pi (fx x + fy y)];
Square[x_, y_, fx_, fy_] := Sign[Grating[x, y, fx, fy]];
```

```
gsine = DensityPlot[contrast * Grating[x, y, 3, 0],
   {x, -1, 1}, {y, -1, 1}, PlotPoints → 64, Mesh → False,
   Frame → False, PlotRange → {-1, 1}, DisplayFunction → Identity];
gsquare = DensityPlot[contrast * Square[x, y, 3, 0], {x, -1, 1},
   {y, -1, 1}, PlotPoints → 64, Mesh → False, Frame → False,
   PlotRange → {-1, 1}, DisplayFunction → Identity];
Show[GraphicsArray[{gsine, gsquare}], DisplayFunction → $DisplayFunction ];
```

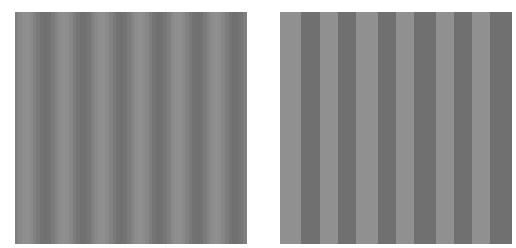

■ Answer: Human contrast sensitivity is higher for a square than for a sine wave.

Campbell and Robson measured the CSF for sinusoidal and square-wave gratings. They showed that to be seen equally well (at threshold), one needed about 27% more contrast for a sine wave than for a square wave. (This was for spatial frequencies higher than the peak frequency of the CSF). Or, putting it another way--if they are both at the same physical contrast, when a square-wave can just be seen, the sine-wave is quite invisible. Why?

■ Fourier series for a square-wave grating

Here is a plot of the relative intensity as a function of distance across the screen for the sine and square-wave gratings:

```
square1[x_] := Sign[Sin[x]]; Plot[{Sin[x], square1[x]}, {x, 0, 7}];
```

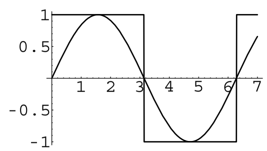

Recall that Fourier showed that (almost) any periodic function could be synthesized from a sum of sine and cosine waves. The fourier series for a square-wave is: $\sum_{n=0}^{\infty} \frac{1}{2n+1} \sin[(2n+1)x]$. The first term in the series (n=0) is sin(x) and is called the fundamental component of the series. Let's plot up the sum of the first 13 terms together with a plot of the fundamental:

```
approxsquare[x_] := Sin[x] + (1 / 3) Sin[3 x] + (1 / 5) Sin[5 x] + (1 / 7) Sin[7 x] +
   (1 / 9) Sin[9 x] + (1 / 11) Sin[11 x] + (1 / 13) Sin[13 x] + (1 / 15) Sin[15 x] +
   (1 / 17) Sin[17 x] + (1 / 19) Sin[19 x] + (1 / 21) Sin[21 x] + (1 / 23) Sin[23 x]
```

```
Plot[{Sin[x], approxsquare[x]}, {x, 0, 7}];
```

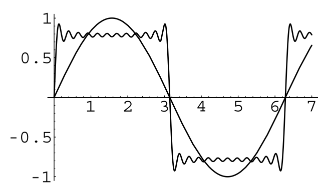

In fact, Fourier analysis tells us that the contrast amplitude of the fundamental is $\pi/4$ (=1.27) greater than that of the square-wave:


```
realsquare[x_] := (Pi / 4) Sign[Sin[x]];
Plot[{Sin[x], realsquare[x]}, {x, 0, 7}];
```

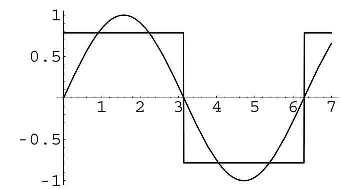

What does the psychophysics mean for neurophysiology?

Campbell and Robson's conclusion was that the human visual system was analyzing the square-wave grating in terms of its spatial frequency components. Campbell and Robson were a bit more conservative that this, but the irresistable grand conjecture was that:

The visual cortex represents complex images in terms of their projections onto the set of sinusoidal grating basis functions

I.e. dot products with $\{\cos[2\pi(f_x x + f_y y) + \phi]\}$, or using complex variable notation, with members of the set $\{e^{i2\pi(f_x x + f_y y)}\}$. (Recall that these are eigenfunctions of a linear shift-invariant system).

What does this mean in terms of neurons? One interpretation would be that there are neurons in the visual cortex whose receptive field weights match those of sinusoidal gratings, and whose activities represent the magnitude of the projections. So rather than a circularly symmetric center-surround receptive field (as for the ganglion cell), the receptive field weights (for a range of cells of one orientation) would look like one of these:

```
grating[x_, y_, fx_, fy_, ϕ_] := Sin[2 Pi (fx x + fy y) + ϕ];
fx = 1 / 1; fy = 1 / 3; ϕ1 = 0;
DensityPlot[grating[x, y, fx, fy, ϕ1], {x, -2, 2}, {y, -2, 2},
  PlotPoints → 64, Mesh → False, Frame → False, PlotRange → {-1, 1}];
```

where various combinations of $\{f_x, f_y, \phi\}$ produce gratings of any spatial frequency, orientation, and phase. Try some variations on the above.

■ **Exercise: What is orientation θ as a function of f_x and f_y?**

These would be the weights for a collection of cells at just one location. What image pattern would produce the biggest response? Assuming that the patterns are all normalized, the pattern that matches the receptive field weight would give the biggest response.

You might argue that this is a pretty big stretch from a simple psychophysical result. You'd be right. What are the neurophysiological data?

Hubel & Wiesel

■ **Bar detectors**

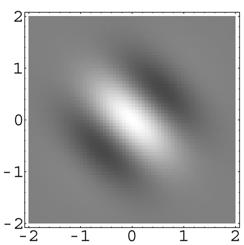

■ **Edge detectors**

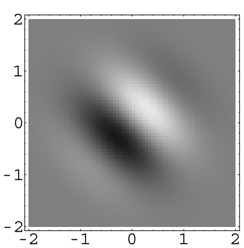

■ **Conclusion-"global" fourier representation of images NOT consistent with neurophysiology**

Multiresolution

Localization in frequency vs. localization in space

■ **Fourier transform**

Often it is more convenient to work with the general continuous Fourier case (rather than the discrete). Then, we need to use complex numbers to represent the sinusoidal gratings, their amplitudes and phases. Complex numbers can be written in polar notation where the complex number has a length r = **Abs[z]** and an angle θ = **Arg[z]** that it makes with the real axis:

$z = r(\cos\theta + i\sin\theta)$.

Using Euler's formula, $e^{i\theta} = \cos\theta + i\sin\theta$,

we can put the length and angles back together:

z = Abs[z] Exp[I Arg[z]], (where i=I).

Given a function f(x), the fourier transform is given by:

Fourier transform : $F(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x)\, e^{i\omega x}\, dx$,

where $F(\omega)$ is in general a complex number. Arg[$F(\omega)$] gives the phase spectrum,
which specifies how the gratings should be offset relative to each other,
and Abs[$F(\omega)$] gives the amplitude spectrum, which specifies the amplitude of each grating.

Inverse fourier transform : $f(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} F(\omega)\, e^{-i\omega x}\, d\omega$

■ **Localization in space: Fourier transform of a delta function is constant**

The Dirac delta function is precisely localized in space,

```
FourierTransform[DiracDelta[x], x, ω]
```

but is spread out over the whole frequency spectrum:

$$\frac{1}{\sqrt{2\pi}}$$

■ Localization in frequency: What is the Fourier transform of Cos[ω_0x]? ($\omega = 2\pi f$)

Cosine is spread out in space,

```
FourierTransform[Cos[ω0 x], x, ω]
```

but is precisely located in frequency:

$$\sqrt{\frac{\pi}{2}}\ \text{DiracDelta}[\omega-\omega_0]+\sqrt{\frac{\pi}{2}}\ \text{DiracDelta}[\omega+\omega_0]$$

■ Is there some compromise?

The "Gabor functions": {Exp[-(x / σ) ^2] Cos[ω_0 x], Exp[-(x / σ) ^2] Sin[ω_0 x]} are said to be "well-localized" in both space and time. Heisenberg had shown the uncertainty principle--simultaneous exact localization in space and momentum was not possible. Dennis Gabor showed that of all basis functions, "Gabor functions" achieved an optimal compromise (according to criteria he specified) in achieving simultaneous compaction in both frequency and space (or time).

Plot a gabor function:

```
σ = 1; ω0 = 4;
Plot[Exp[-(x / σ)^2] Cos[ω0 x], {x, -4, 4},
  PlotRange → {-1, 1}, AxesLabel → {"x", "Contrast"}];
```

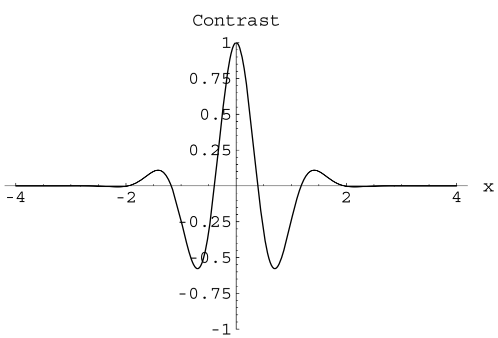

Plot the frequency spectrum of a gabor function:


```
Clear[σ, ω, ω0];
FourierTransform[Exp[-(x / σ)^2] Cos[ω0 x], x, ω]
```

$$\frac{e^{-\frac{1}{4}\sigma^2(\omega+\omega_0)^2}\left(1+e^{\sigma^2\omega\omega_0}\right)\sqrt{\frac{1}{\sigma^2}}\ \sigma^2}{2\sqrt{2}}$$

```
σ = 1; ω0 = 4;
Plot[ e^(-1/4 σ^2 (ω+ω0)^2) (1 + e^(σ^2 ω ω0)) Sqrt[1/σ^2] σ^2 / (2 Sqrt[2]), {ω, -8, 8},
  PlotRange → {{-8, 8}, {0, .5}}, AxesLabel → {"ω", "Amplitude"}];
```

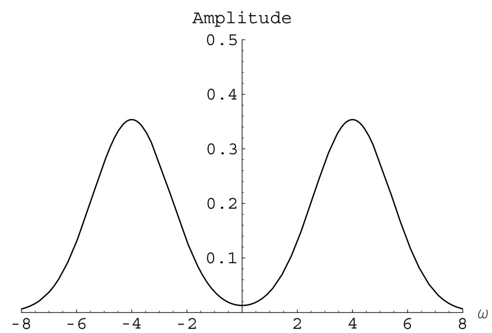

Psychophysical clues to human spatial image representations: What does the eye see best?

■ Burgess, Wagner, Jennings and Barlow (1981)

Burgess et al. combined the SKE observer and spatial frequency analysis of human vision to find out how efficiently humans detected patterns.

Burgess, Wagner, Jennings, and Barlow showed in a 1981 Science article that narrowly windowed sinusoids were detected with high efficiency (>70%) when added to static visual noise. Further, these targets were detected more efficiently than disks of light.

You basically have all the tools to replicate the experiment of Burgess et al. You can compute d' for the ideal observer for signal-known-exactly patterns. And you can generate Gaussian-windowed sinusoids and add them to gaussian white noise. If you measure the percent correct, and convert that to d' for the human observer, you can calculate the absolute efficiency for human detection--and contribute to answer the question of what the eye sees best.

■ **Watson, Barlow & Robson (1983)**

Watson et al. found that that a 7 c/deg grating drifting at 4 Hz, (with a narrow gaussian envelope in space and time) was detected more efficiently than other patterns. Further, the quantum efficiency was very low (<0.05%).

■ **Kersten, 1984**

Kersten measured efficiency for 1-d gratings (i.e. vertical) in temporal visual noise for various spatial frequencies and widths. Here i

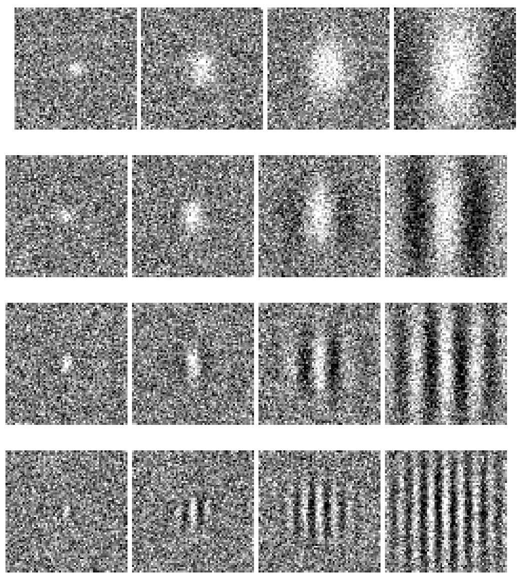

Peak efficiency was found for patterns of about the same shape, regardless of spatial frequency. The cross-sectional profiles looked like:

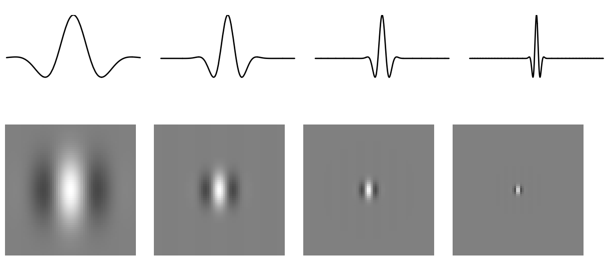

When the filters have the same shape except for a change of scale ($x \to \alpha x$), they are called *self-similar*.

Bottom-line: image coding in terms of scale and orientation: Amodel for human spatial image representation

At each spatial location, project the image onto a collection of basis vectors (i.e. compute the dot product) that span a range of *spatial scales* and *orientations*:

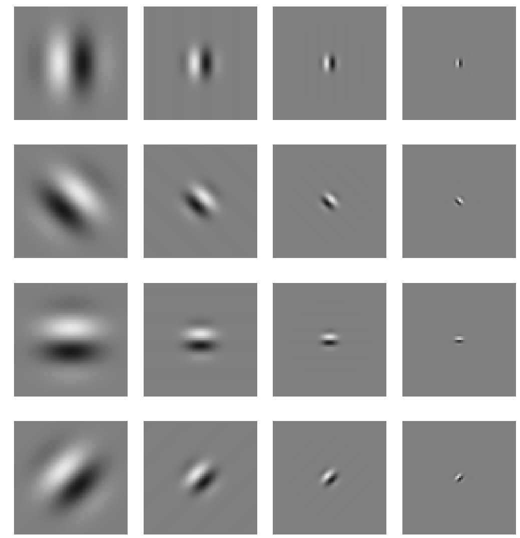

In general, these neural models of basis functions may be over-complete, and non-orthogonal. And there may be a range of phases. Above we show only the "sine-phase" or "edge-detectors" of Hubel and Wiesel.

The self-similar idea is important to vision because of the need for some kind of scale-invariance. Further, the self-similar aspect of these neural models bore a close resemblance to the emerging mathematical field of wavelet analysis. The emphases are different--over-completeness may be important and vision does the projections in parallel (the serial algorithmic component of wavelet computation is integral to the mathematical interest).

Neural image? Or neural image representation?

We can view the response activities of a family of receptive fields of neurons as representing a filtered neural image of the input image. Although useful, this view can be misleading when we start to think about function, for "who is looking at the image"?

Alternatively, thinking in terms of basis functions gives us another perspective. We can view the response activities of a family of receptive fields as a representation of the input image. If linear, an activity is the result of a projection of an image on to a basis function (receptive field weights). Given such a representation we can begin to ask questions like:

1. Is the neural basis set complete? Can any image be represented?

2. A closely related question is: Is any information lost? I.e. we do the inverse transformation, can the original input be reconstructed?

3. Maybe the neural basis set is "over-complete"?

4. Are the neural basis functions orthogonal? Are they normal?

Next time

What is the computational significance of a wavelet-like decomposition?

Efficient coding?

-> analysis of natural image statistics

Edge detection?

-> analysis of what vision needs to recognize objects, etc..

Appendices

References

Adelson, E. H., Simoncelli, E., & Hingorani, R. (1987). *Orthogonal Pyramid Transforms for Image Coding*. Paper presented at the Proc. SPIE - Visual Communication & Image Proc. II, Cambridge, MA.

Burgess, A. E., Wagner, R. F., Jennings, R. J., & Barlow, H. B. (1981). Efficiency of human visual signal discrimination. *Science, 214*, 93-94.

Daugman, J. G. (1988). An information-theoretic view of analog representation in striate cortex, *Computational Neuroscience*. Cambridge, Massachusetts: M.I.T. Press.

Enroth-Cugell, C., & Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. Journal of Physiology (London), 187, 517-552.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive Fields of Single Neurons in the Cat's Striate Cortex. *J. Physiol., 148*, 574-591.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive Fields and Functional Architecture of Monkey Striate Cortex. *J. Physiol*., 215-243.

Gerstein, G. L., & Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. Biophysical Journal, 4, 41-68.

Kersten, D. (1984). Spatial summation in visual noise. *Vision Research, 24*, 1977-1990.

Shapley, R., & Perry, H. H. (1986). Cat and monkey retinal ganglion cells and their visual functional roles. Trends in Neuroscience, 9(5), 229-235.

Silverman, M. S., Grosof, D. H., DeValois, R. L., & Elfar, S. D. (1989). Spatial-frequency organization in primate striate cortex.*, 86*, 711-715.

Simoncelli, E. P., Freeman, W. T., Adelson, E. H., & Heeger, D. J. (1992). Shiftable Multi-scale Transforms. *IEEE Trans. Information Theory, 38*(2), 587--607.

Watson, A. B., Barlow, H. B., & Robson, J. G. (1983). What does the eye see best? *Nature, 31*,, 419-422.

Watson, A. B. (1987). Efficiency of a model human image code. *Journal of the Optical Society of America, A, 4*(12), 2401-2417.

http://www.cis.upenn.edu/~eero/ABSTRACTS/simoncelli90-abstract.html