Spatial pattern vision and linear systems theory

- **Receptive fields and neural images**
- Shift-invariant linear systems and convolution
- Fourier transform and frequency response
- Applications of linear systems to spatial vision
  - Contrast sensitivity
  - Spatial frequency and orientation channels
  - Spatial frequency and orientation adaptation
  - Masking

Receptive field

- In any modality: that region of the sensory apparatus that, when stimulated, causes a response

Receptive field

OBSERVATIONS ON THE SCRATCH-REFLEX IN THE SPINAL DOG. BY C. S. SHERRINGTON. (27 Figures in Text.)

VII. THE RECEPTIVE FIELD.

The whole area of skin from whose points the scratch-reflex can be elicited may be conveniently termed the *receptive field* of that reflex (Fig. 19). The receptive field may be considered as composed of receptive points. That is to say, what is referred to as 'the scratch-reflex'

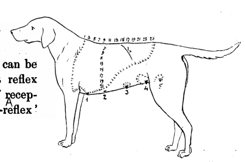

Receptive field

OBSERVATIONS ON THE SCRATCH-REFLEX IN THE SPINAL DOG. By C. S. SHERRINGTON. (27 Figures in Text.)

Within the receptive fields of the type-reflexes not all the receptive points serve with equal facility or potency to excite the reflex. The points composing certain areas of the field are most effective, those of certain other areas least effective, and from the rest of the field the reflex

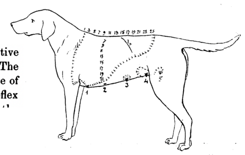

Receptive field

- In any modality: that region of the sensory apparatus that, when stimulated, can directly affect the firing rate of a given neuron
- Spatial vision: spatial receptive field can be mapped in visual space or on the retina

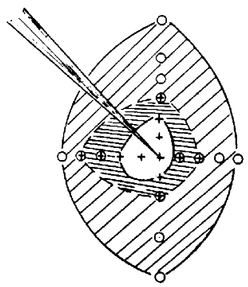

Kuffler, 1953

Receptive field

- In any modality: that region of the sensory apparatus that, when stimulated, can directly affect the firing rate of a given neuron
- Spatial vision: spatial receptive field can be mapped in visual space or on the retina

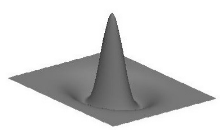

LGN

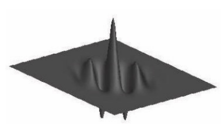

V1

Receptive field

A spatial receptive field is an image

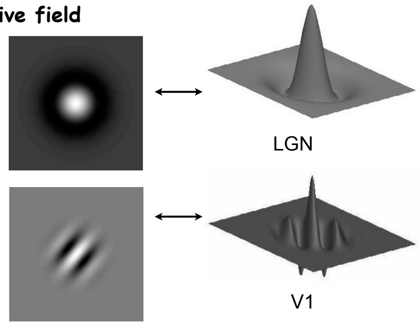

Orientation selective receptive field

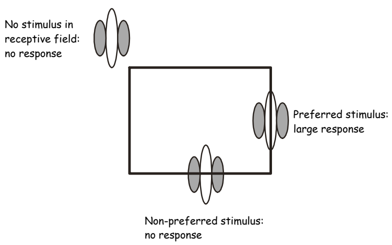

Neural image: retinal ganglion cell responses

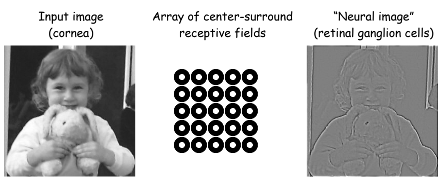

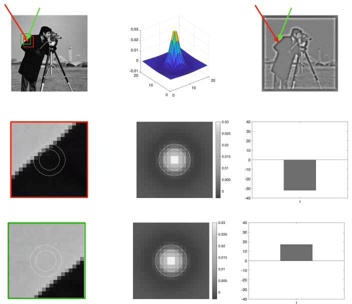

```
% Make a Difference of Gaussian receptive field using fspecial
DoG = fspecial('gaussian', 20,2) -  fspecial('gaussian', 20,5);
im = imread('cameraman.tif');

% Make a neural image by convolution
neuralim = conv2(double(im), DoG);

% Show the image, the RF, and the neural image
figure,  tiledlayout(3,3);
nexttile(); imshow(im)
nexttile(); surf(DoG); axis square
nexttile(); imshow(neuralim, DisplayRange=[-80 80]);

imcrop1= im(86+(-10:9), 56+(-10:9));
imcrop2= im(80+(-10:9), 50+(-10:9));

nexttile(); imshow(imcrop1); hold on; contour(DoG, [0 .01], 'w')
nexttile(), imshow(DoG, DisplayRange=[-.003 .03]); colorbar;
set(gca, 'Colormap', parula);
hold on; contour(DoG, [0 .01], 'w')
nexttile(), bar(dot(double(imcrop1(:)), DoG(:))); axis padded;
ylim([-40 40])
nexttile(), imshow(imcrop2); hold on; contour(DoG, [0 .01], 'w')
nexttile(), imshow(DoG, DisplayRange=[-.003 .03]); colorbar;
set(gca, 'Colormap', parula);
hold on; contour(DoG, [0 .01], 'w')
nexttile(), bar(dot(double(imcrop2(:)), DoG(:))); axis padded;
ylim([-40 40])
```

10

Neural image: simple cell responses

Input image (cornea)

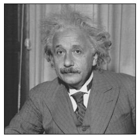

Array of orientation-selective receptive fields

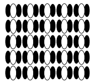

"Neural image" (V1 simple cells)

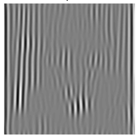

Lots of neural images: V1 simple cells

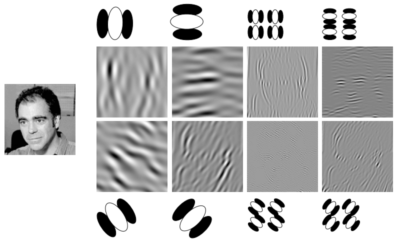

Lots of neural images

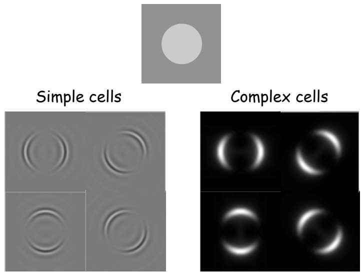

Neural image of a sine wave

For a linear, shift-invariant system such as a linear model of a receptive field, an input sine wave results in an identical output sine wave, except for a possible lateral shift and scaling.

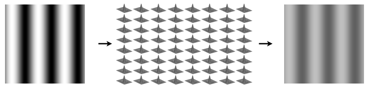

Frequency response

This scaling of contrast by a linear receptive field in the neural image is a function of spatial frequency determined by the shape of the receptive field.

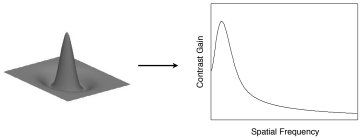

Frequency response

This scaling of contrast by a linear receptive field in the neural image is a function of spatial frequency determined by the shape of the receptive field.

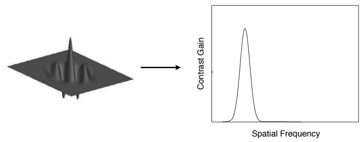

Orientation tuning

If a receptive field is not circularly symmetric, the scaling of contrast is also a function of orientation (for a given spatial frequency) determined by the shape of the receptive field.

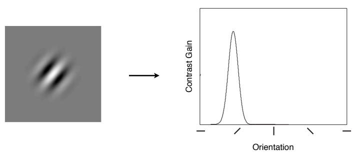

Linear systems analysis

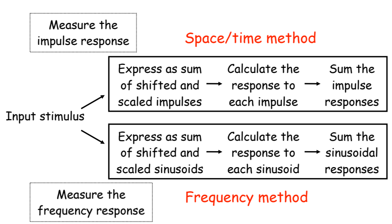

Spatial pattern vision and linear systems theory

- Receptive fields and neural images
- **Shift-invariant linear systems and convolution**
- Fourier transform and frequency response
- Applications of linear systems to spatial vision
  - Contrast sensitivity
  - Spatial frequency and orientation channels
  - Spatial frequency and orientation adaptation
  - Masking

Functions and Systems

A **function**: a relation between a set of inputs and a set of permissible outputs with the property that each input is related to exactly one output.

Examples

- $f(x) = x^2$
- $f(x) = \cos(x)$
- $f(x,y) = \mathrm{sqrt}(x^2 + y^2)$

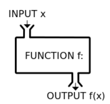

wikipedia

Functions and Systems

A **system**: a generalization of a function whereby inputs and outputs can be numbers or objects, discrete or continuous, any dimensionality

Examples

- Physiological optics: Input: scene spectral irradiance; Output: retinal image
- Eye: Input: scene spectral irradiance; Output: retinal ganglion cell firing rate
- Human: Input: scene spectral irradiance; Output: knob adjustment (eg color match)
- fMRI: Input: neuronal activity Output: T2*w image

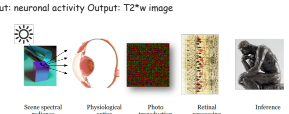

Linear Systems Analysis

Systems with signals as input and output

- 1-d: low- and high-pass filters in electronic equipment, fMRI data analysis, or in sound production (articulators) or audition (the ear as a filter)

  $y(t) = T\{x(t)\}$
- 2-d: optical blur, spatial receptive field

  $g(x,y) = T\{f(x,y)\}$
- 3-d: spatio-temporal receptive field

  $g(x,y,t) = T\{f(x,y,t)\}$

Note that inputs and outputs do not have to have the same dimensionality

Homogeneity (scaling)

$a*T[x(t)] = T[a*x(t)]$

Visual stimulus RGC firing rate

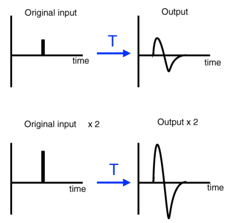

Additivity

$T[x_1(t)] + T[x_2(t)] = T[x_1(t) + x_2(t)]$

Visual stimulus RGC firing rate

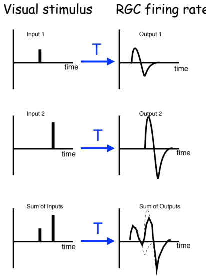

Summary: Linear systems requirements

A system (or transform) converts (or maps) an input signal into an output signal:

$y(t) = T[x(t)]$

A linear system satisfies the following properties:

1) Homogeneity (scalar rule):

$T[a\, x(t)] = a\, y(t)$

2) Additivity:

$T[x_1(t) + x_2(t)] = y_1(t) + y_2(t)$

Often, these two properties are written together and called superposition:

$T[a\, x_1(t) + b\, x_2(t)] = a\, y_1(t) + b\, y_2(t)$

Shift invariance

$y(t) = T[x(t)] \longrightarrow y(t\text{-}s) = T[x(t\text{-}s)]$

Visual stimulus RGC firing rate

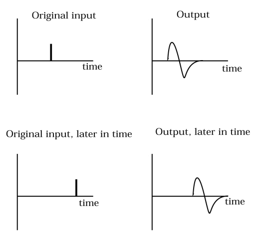

Shift-invariant linear systems and impulses

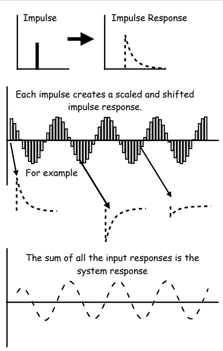

Convolution

Discrete-time signal: x[n] = [x1, x2, x3, ...]

A system or transform maps an input signal into an output signal:

$$y[n] = T\{x[n]\}$$

A shift-invariant linear system can always be expressed as a convolution:

$$y[n] = \sum_m x[m]\, h[n-m]$$

where h[n] is the impulse response.

Convolution as sum of impulse responses

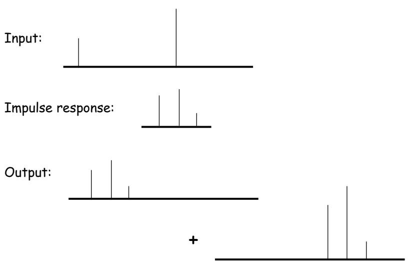

Convolution as sequence of weighted sums

past		present				future			
0	0	0	1	0	0	0	0	0	input (impulse)
	1/8	1/4	1/2						weights
0	0	0	1/2	1/4	1/8	0	0	0	output (impulse response)

past		present				future			
0	0	0	1	1	1	1	1	1	input (step)
	1/8	1/4	1/2						weights
0	0	0	1/2	3/4	7/8	7/8	7/8	7/8	output (step response)

Convolution as matrix multiplication

Linear system <=> matrix multiplication

Shift-invariant linear system <=> Toeplitz matrix

B = Ax

$$\begin{bmatrix} \vdots \\ 5 \\ 2 \\ -3 \\ 4 \\ \vdots \end{bmatrix} = \begin{bmatrix} & & & \vdots & & & \\ \cdots & 1 & 2 & 3 & 0 & 0 & 0 & \cdots \\ \cdots & 0 & 1 & 2 & 3 & 0 & 0 & \cdots \\ \cdots & 0 & 0 & 1 & 2 & 3 & 0 & \cdots \\ \cdots & 0 & 0 & 0 & 1 & 2 & 3 & \cdots \\ & & & \vdots & & & \end{bmatrix} \begin{bmatrix} \vdots \\ 1 \\ 2 \\ 0 \\ 0 \\ -1 \\ 2 \\ \vdots \end{bmatrix}$$

Columns contain shifted copies of the impulse response.

Rows contain time-reversed copies of impulse response.

Convolution as matrix multiplication

$$\begin{bmatrix} \vdots \\ 3 \\ 2 \\ 1 \\ 0 \\ \vdots \end{bmatrix} = \begin{bmatrix} & & & \vdots & & & \\ \cdots & 1 & 2 & 3 & 0 & 0 & 0 & \cdots \\ \cdots & 0 & 1 & 2 & 3 & 0 & 0 & \cdots \\ \cdots & 0 & 0 & 1 & 2 & 3 & 0 & \cdots \\ \cdots & 0 & 0 & 0 & 1 & 2 & 3 & \cdots \\ & & & \vdots & & & \end{bmatrix} \begin{bmatrix} \vdots \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix}$$

Columns contain shifted copies of the impulse response.

Rows contain time-reversed copies of impulse response.

Convolution as matrix multiplication

$$\begin{bmatrix} \vdots \\ 3 \\ 2 \\ 1 \\ 0 \\ \vdots \end{bmatrix} = \begin{bmatrix} & & & \vdots & & & \\ \cdots & 1 & 2 & 3 & 0 & 0 & 0 & \cdots \\ \cdots & 0 & 1 & 2 & 3 & 0 & 0 & \cdots \\ \cdots & 0 & 0 & 1 & 2 & 3 & 0 & \cdots \\ \cdots & 0 & 0 & 0 & 1 & 2 & 3 & \cdots \\ & & & \vdots & & & \end{bmatrix} \begin{bmatrix} \vdots \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix}$$

Columns contain shifted copies of the impulse response.

Rows contain time-reversed copies of impulse response.

Derivation:
shift-invariant linear system => convolution

Homogeneity:

$T\{a\,x[n]\} = a\,T\{x[n]\}$

Additivity:

$T\{x_1[n] + x_2[n]\} = T\{x_1[n]\} + T\{x_2[n]\}$

Superposition:

$T\{a\,x_1[n] + b\,x_2[n]\} = a\,T\{x_1[n]\} + b\,T\{x_2[n]\}$

Shift-invariance:

$y[n] = T\{x[n]\} \Rightarrow y[n-m] = T\{x[n-m]\}$

Convolution derivation (cont)

Impulse sequence:

$d[n] = 1$ for $n = 0$, $d[n] = 0$ otherwise

Any sequence can be expressed as a sum of impulses:

$$x[n] = \sum_m x[m]\, d[n-m]$$

where

$d[n-m]$ is impulse shifted to sample m

$x[m]$ is the height of that impulse

Example:

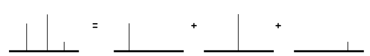

Convolution derivation (cont)

$x[n]$: input

$y[n] = T\{x[n]\}$: output

$h[n] = T\{d[n]\}$: impulse response

1) Represent input as sum of impulses:

$$y[n] = T\{x[n]\}$$

$$y[n] = T\left\{\sum_m x[m]\, d[n-m]\right\}$$

2) Use superposition:

$$y[n] = \sum_m x[m]\, T\{d[n-m]\}$$

3) Use shift-invariance:

$$y[n] = \sum_m x[m]\, h[n-m]$$

Matrix multiplication => scaling

$$\begin{array}{ccccc} \text{Nx1 vector} & & \text{NxN Toeplitz matrix} & & \text{Nx1 vector} \\ \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_N \end{bmatrix} & = & \begin{bmatrix} \\ \text{Impulse} \\ \text{response} \\ \\ \end{bmatrix} & & \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_N \end{bmatrix} \end{array}$$

$$\begin{bmatrix} ay_1 \\ ay_2 \\ ay_3 \\ \vdots \\ ay_N \end{bmatrix} = \begin{bmatrix} \\ \text{Impulse} \\ \text{response} \\ \\ \end{bmatrix} \begin{bmatrix} ax_1 \\ ax_2 \\ ax_3 \\ \vdots \\ ax_N \end{bmatrix}$$

Matrix multiplication => additivity

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} \\ \text{Impulse} \\ \text{response} \\ \\ \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ \vdots \\ x_N \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} z_1 \\ z_2 \\ z_3 \\ \vdots \\ z_N \end{bmatrix} = \begin{bmatrix} \\ \text{Impulse} \\ \text{response} \\ \\ \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ \vdots \\ w_N \end{bmatrix}$$

$$\begin{bmatrix} y_1 & & z_1 \\ y_2 & & z_2 \\ y_3 & + & z_3 \\ \vdots & & \vdots \\ y_N & & z_N \end{bmatrix} = \begin{bmatrix} \\ \text{Impulse} \\ \text{response} \\ \\ \end{bmatrix} \begin{bmatrix} x_1 & & w_1 \\ x_2 & & w_2 \\ x_3 & + & w_3 \\ \vdots & & \vdots \\ x_N & & w_N \end{bmatrix}$$

Toeplitz matrix => shift invariance

$$\begin{bmatrix} \vdots \\ 3 \\ 2 \\ 1 \\ 0 \\ \vdots \end{bmatrix} = \begin{bmatrix} & & & \vdots & & & & \\ & 1 & 2 & 3 & 0 & 0 & 0 & \\ \cdots & 0 & 1 & 2 & 3 & 0 & 0 & \cdots \\ & 0 & 0 & 1 & 2 & 3 & 0 & \\ & 0 & 0 & 0 & 1 & 2 & 3 & \\ & & & \vdots & & & & \end{bmatrix} \begin{bmatrix} \vdots \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ \vdots \end{bmatrix}$$

Neural image and convolution

A spatial receptive field may also be treated as a linear system, by assuming a dense collection of neurons with the same receptive field translated to different locations in the visual field. In this view, it is a linear, shift-invariant system*§.

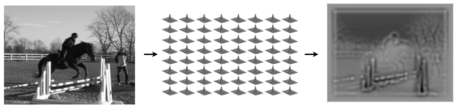

* This is the basis of CNNs (convolutional neural networks).
§ The nervous system doesn't actually work like this. (It's not linear and it's not shift invariant!)

Neural image and convolution

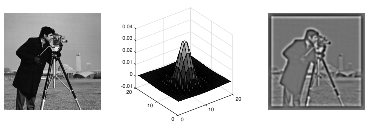

```
% Make a Difference of Gaussian receptive field using fspecial
DoG = fspecial('gaussian', 20,2) -  fspecial('gaussian', 20,5);
im = imread('cameraman.tif');

% Make a neural image by convolution
neuralim = conv2(double(im), DoG);

% Show the image, the RF, and the neural image
figure,  subplot(1,3,1), imshow(im), subplot(1,3,2), surf(DoG)
subplot(1,3,3), imagesc(neuralim); colormap gray, axis image off
```

Neural image and convolution

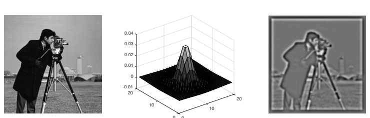

```
% Shift the image and replot
im2 = im(:, [n:end 1:n]);
neuralim2 = conv2(double(im2), DoG);
```

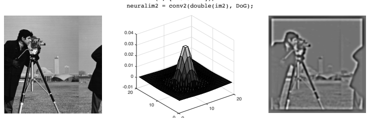

Continuous-time derivation of convolution

Linear systems requirements

A system (or transform) converts (or maps) an input signal into an output signal:

$y(t) = \mathrm{T}[x(t)]$

A linear system satisfies the following properties.

1) Homogeneity (scalar rule):

$\mathrm{T}[\mathrm{a}\, x(t)] = \mathrm{a}\, y(t)$

2) Additivity:

$\mathrm{T}[x_1(t) + x_2(t)] = y_1(t) + y_2(t)$

Often, these two properties are written together and called superposition:

$\mathrm{T}[\mathrm{a}\, x_1(t) + \mathrm{b}\, x_2(t)] = \mathrm{a}\, y_1(t) + \mathrm{b}\, y_2(t)$

Pulses and impulses

$$\delta(t) = \begin{cases} \infty & \text{if} \quad t = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$\delta_\Delta(t) = \begin{cases} \frac{1}{\Delta} & \text{if} \quad 0 < t < \Delta \\ 0 & \text{otherwise} \end{cases}$$

$$\delta(t) = \lim_{\Delta \to 0} \delta_\Delta(t).$$

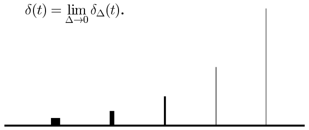

Staircase approximation to continuous time signal

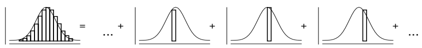

$$\tilde{x}(t) = \sum_{k=-\infty}^{\infty} x(k\Delta)\,\delta_\Delta(t - k\Delta)\,\Delta.$$

$$x(t) = \lim_{\Delta\to 0} \sum_{k=-\infty}^{\infty} x(k\Delta)\,\delta_\Delta(t - k\Delta)\,\Delta.$$

$$x(t) = \int_{-\infty}^{\infty} x(s)\,\delta(t-s)\,ds.$$

Shift invariance

For a system to be shift-invariant (or time-invariant) means that a time-shifted version of the input yields a time-shifted version of the output:

$$y(t) = \mathrm{T}[x(t)]$$

$$y(t - s) = \mathrm{T}[x(t - s)]$$

The response $y(t - s)$ is identical to the response $y(t)$, except that it is shifted in time.

Convolution

Representing the input signal as a sum of pulses:

$$\begin{aligned} y(t) = T[x(t)] &= T\left[\int_{-\infty}^{\infty} x(s)\,\delta(t-s)\,ds\right] \\ &= T\left[\lim_{\Delta\to 0} \sum_{k=-\infty}^{\infty} x(k\Delta)\,\delta_\Delta(t-k\Delta)\,\Delta\right]. \end{aligned}$$

Using additivity,

$$y(t) = \lim_{\Delta\to 0} \sum_{k=-\infty}^{\infty} T[x(k\Delta)\,\delta_\Delta(t-k\Delta)\,\Delta].$$

Taking the limit,

$$y(t) = \int_{-\infty}^{\infty} T[x(s)\,\delta(t-s)\,ds].$$

Using homogeneity (scalar rule),

$$y(t) = \int_{-\infty}^{\infty} x(s)\,T[\delta(t-s)]\,ds.$$

Defining $h(t)$ as the impulse response,

$$y(t) = \int_{-\infty}^{\infty} x(s)\,h(t-s)\,ds.$$

Linear systems analysis

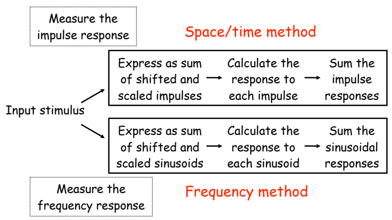

Spatial pattern vision and linear systems theory

- Receptive fields and neural images
- Shift-invariant linear systems and convolution
- **Fourier transform and frequency response**
- Applications of linear systems to spatial vision
  - Contrast sensitivity
  - Spatial frequency and orientation channels
  - Spatial frequency and orientation adaptation
  - Masking

Fourier transform and frequency response: summary

- Signals can be represented as sums of sine waves
- Linear, shift-invariant systems operate “independently” on each sine wave, and merely scale and shift them.
- A simplified model of neurons in the visual system, the linear receptive field, results in a neural image that is linear and shift-invariant.
- Psychophysical models of the visual system might be built of such mechanisms.
- It is therefore important to understand visual stimuli in terms of their spatial frequency content.
- The same tools can be applied to other modalities (e.g., audition) and other signals (EEG, MRI, MEG, etc.).

Temporal frequency and Fourier decomposition

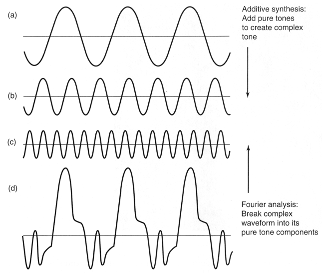

Figure 10.7
Additive synthesis. (a) Pressure changes for a pure tone with frequency of 440 Hz; (b) the second harmonic of (a), with a frequency of 880 Hz; (c) the third harmonic, with a frequency of 1,320 Hz. (d) The sum of the three harmonics above creates the waveform for a complex tone. Working in the other direction, applying Fourier analysis to the complex tone reveals its pure tone components.

Shift-invariant linear systems & sinusoids

Measure the scaling and shifting for each sinusoid

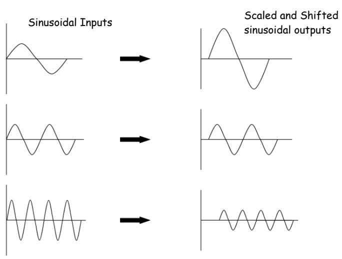

See Matlab demo

Shift-invariant linear systems & sinusoids

Frequency response

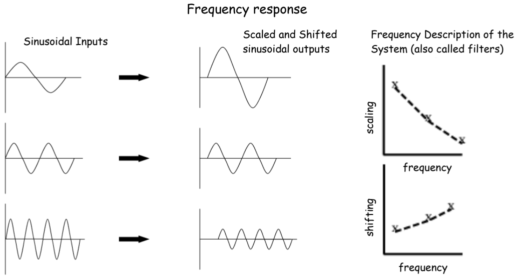

Shift-invariant linear systems & sinusoids

```
% Define a sine wave and a square wave at 10 Hz
t          = (1:100)/100;
sinewave   = sin(2*pi*t*10);
squarewave = square(2*pi*t*10);

% Define an arbitrary impulse response function
h = randn(1,10);

% Compute the output
r_sine   = conv(sinewave,h, 'same');
r_square = conv(squarewave,h, 'same');

figure, set(gcf, 'Color', 'w')
tiledlayout(3,1)

linewidth = 2;
fontsize = 15;

nexttile();
plot(h, 'o-', LineWidth=linewidth);
title('Impulse Response Function')
yline(0)
set(gca, FontSize=fontsize);

nexttile();
plot(t,sinewave, t, r_sine, LineWidth=linewidth);
title('Sine wave')
legend('Input', 'Output', 'Location','eastoutside')
set(gca, FontSize=fontsize);


nexttile();
plot(t,squarewave, t, r_square, LineWidth=linewidth);
title('Square wave')
legend('Input', 'Output', 'Location','eastoutside')
set(gca, FontSize=fontsize);
```

https://foundationsofvision.stanford.edu/appendix/#Convolution_and_Harmonic_Functions

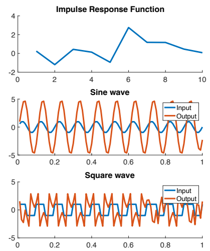

Linear systems analysis

Measure the impulse response

Space/time method

Input stimulus

Express as sum of shifted and scaled impulses → Calculate the response to each impulse → Sum the impulse responses

Express as sum of shifted and scaled sinusoids → Calculate the response to each sinusoid → Sum the sinusoidal responses

Measure the frequency response

Frequency method

Convolution and multiplication

time sample

Fourier transform

frequency sample

Convolution:

$$x_1[n] * x_2[n] \longleftrightarrow X_1[k]\, X_2[k]$$

convolution

Multiplication:

$$x_1[n]\, x_2[n] \longleftrightarrow (1/N)\, X_1[k] * X_2[k]$$

Linear filter example

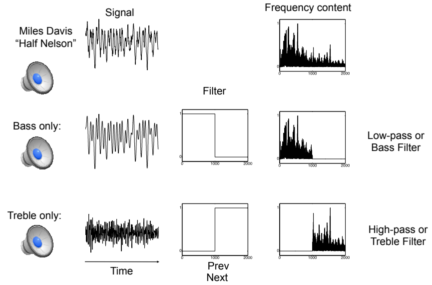

Linear filter example

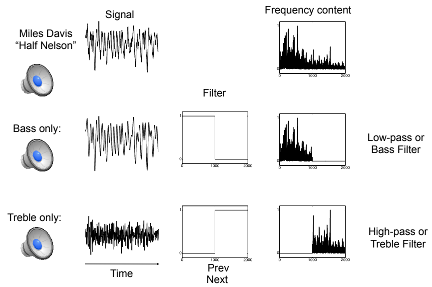

Linear filter example

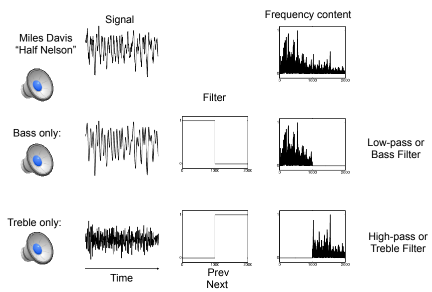

Linear filter example (1-D)

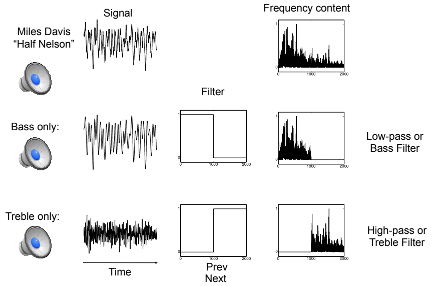

Linear filter (2-D)

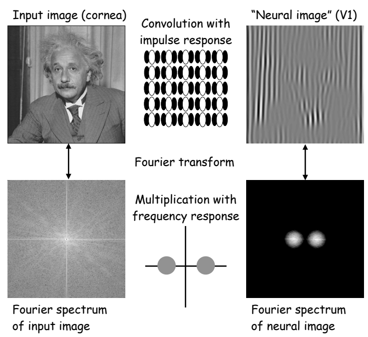

Fourier Analysis

Signals as sums of sine waves

- 1d: time series
  - fMRI signal from a voxel or ROI
  - mean firing rate of a neuron over time
  - auditory stimuli
- 2d: static visual image, neural image
- 3d: visual motion analysis

Auditory example: Pure tones

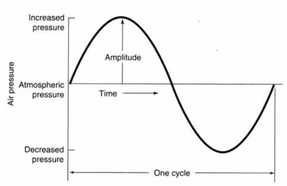

Pure tones can be described by 3 numbers:
Frequency = rate of air pressure modulation (related to pitch)
Amplitude = sound pressure level (related to loudness)
Phase = sin vs. cosine vs. another horizontal shift

Frequency and amplitude

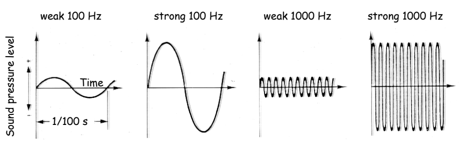

Fourier spectrum representation of sound

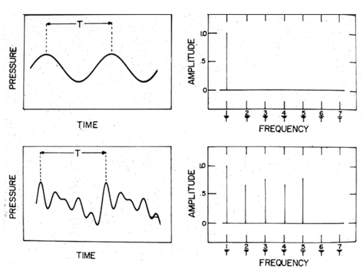

Fourier spectra of some sounds

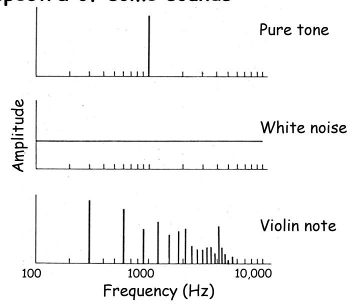

Fundamental frequency and harmonics

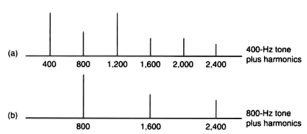

Lots of Fourier transforms

Name	Time domain	Freq domain
Fourier transform	continuous, infinite	continuous, infinite
Fourier series	continuous, periodic	discrete, infinite
DTFT	discrete, infinite	continuous, periodic
DFS	discrete, periodic	discrete, periodic
DFT	discrete, finite	discrete, finite

FFT algorithm

- Computes DFT of finite length input.
- Efficient for inputs of length N = m^n.
- Produces 2 outputs, each of size/length equal to that of the input: real part (cosine coeffs), imaginary part (sine coeffs).

In 1994, Gilbert Strang described the FFT as "the most important numerical algorithm of our lifetime",[3][4] and it was included in Top 10 Algorithms of 20th Century by the IEEE magazine *Computing in Science & Engineering*.[5]

FFT algorithm

```
% Random signal with N time points, random IRF with .1*N time points
N = 1000000;
n = .1*N;
t  = 0:N-1;
x = randn(N,1);
h = zeros(size(x));
h(1:n) = randn(n,1);

% Compare Convolution and FFT multiply
disp('Convolution time')
tic,  y = conv(x,h, 'full'); toc, y = y(1:N);

disp('FFT time')
tic, y2 = ifft(fft(x) .* fft(h)); toc

% Show that they produce the same output
figure(1), clf
subplot(2,1,1), set(gca, 'FontSize', 20), hold on
plot(t,x, t,h), legend('x (signal)', 'h (irf)'), axis off
subplot(2,1,2), set(gca, 'FontSize', 20), hold on
plot(t, y, 'r', t, y2, 'go'); xlim([n n+100])
legend('conv(x, h)', ' ifft(fft(x) .* fft(h))'), axis off
```

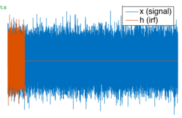

Convolution time
Elapsed time is 7.762392 seconds.

FFT time
Elapsed time is 0.049519 seconds.

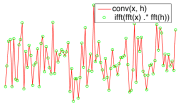

DFT of a cosine

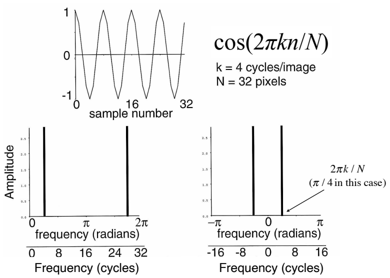

Why do 4, 28, and -4 cycles per image have the same cosine?

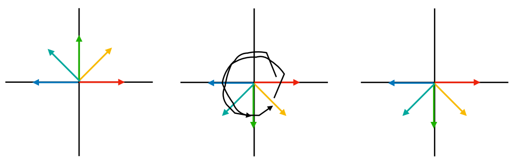

4 cycles, 32 samples =>
1/8 rotation per sample

28 cycles, 32 samples =>
7/8 rotation per sample

-4 cycles, 32 samples =>
-1/8 rotation per sample

Why do 4, 28, and -4 cycles per image have the same cosine?

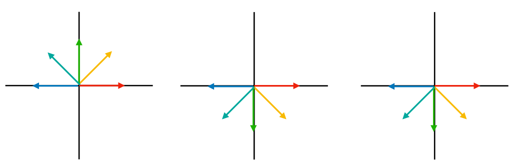

4 cycles, 32 samples =>
1/8 rotation per sample

28 cycles, 32 samples =>
7/8 rotation per sample

-4 cycles, 32 samples =>
-1/8 rotation per sample

Real and imaginary parts

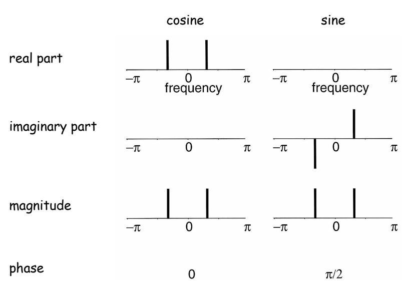

DFT of impulse signal and constant signal

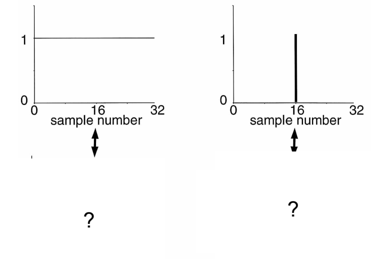

Uncertainty principle

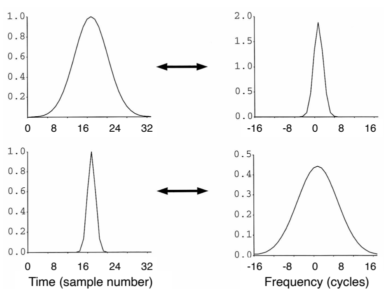

Multiplication and convolution

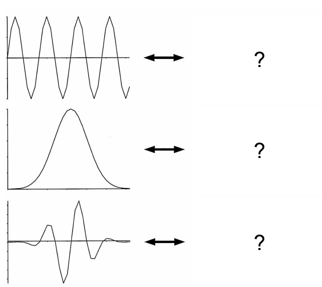

Discrete Fourier Transform (DFT)

Analysis:

$$X[k] = \begin{cases} \sum_{n=0}^{N-1} x[n]\, e^{-j(2\pi/N)kn} & 0 \le k \le N-1 \\ 0 & \text{otherwise} \end{cases}$$

Synthesis:

$$x[n] = \begin{cases} \frac{1}{N}\sum_{k=0}^{N-1} X[k]\, e^{j(2\pi/N)kn} & 0 \le n \le N-1 \\ 0 & \text{otherwise} \end{cases}$$

$x[n]$: discrete, finite

$X[K]$: discrete, finite

Complex numbers and complex exponentials

$$z = \underbrace{a}_{\text{real part}} + j\underbrace{b}_{\text{imaginary part}} = \overbrace{A}^{\text{amplitude}} e^{j\overbrace{\phi}^{\text{phase}}} = A\left[\cos(\phi) + j\sin(\phi)\right]$$

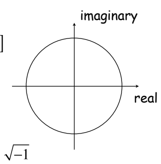

where:

$a = A\cos(\phi)$ $\quad A = \sqrt{a^2 + b^2}$ $\quad j = \sqrt{-1}$

$b = A\sin(\phi)$ $\quad \phi = \tan^{-1}(b/a)$

Why bother with complex exponentials?

$$(A_1 e^{j\phi_1})(A_2 e^{j\phi_2}) = \underbrace{A_1 A_2}_{\text{amplitudes multiply}} e^{\underbrace{j(\phi_1 + \phi_2)}_{\text{phases add}}}$$

Discrete Fourier transform matrix

Analysis: $X[k] = \sum_n x[n] \exp(-j2\pi kn/N)$

For real valued inputs:

$$X_c[k] = \sum_n x[n]\cos(\ldots) \qquad X_s[k] = \sum_n -x[n]\sin(\ldots)$$

$$\begin{pmatrix} X_c[k] \\ X_s[k] \end{pmatrix} = \underbrace{\begin{pmatrix} \text{cosines} \\ \cdots\cdots \\ \text{sines} \end{pmatrix}}_{\mathbf{P}} \begin{pmatrix} x[n] \end{pmatrix}$$

Rows of **P** called projection functions: $\frac{1}{N}\mathbf{P}^T\mathbf{P} = \mathbf{I}$

Discrete Fourier transform matrix

Synthesis: $x[n] = \sum_k X[k] \exp(j2\pi kn/N)$

$$\begin{pmatrix} x[n] \end{pmatrix} = \frac{1}{N}\underbrace{\begin{pmatrix} \text{cosines} & \vdots & \text{sines}\end{pmatrix}}_{\mathbf{B}} \begin{pmatrix} X_c[k] \\ X_s[k]\end{pmatrix}$$

Cols of **B** called basis functions. $\mathbf{B} = \mathbf{P}^T$, $\frac{1}{N}\mathbf{B}\,\mathbf{B}^T = \mathbf{I}$

Sine wave gratings and spatial frequency

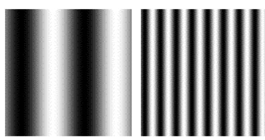

Low SF 1 cpd   High SF 4 cpd

Measured in cycles per degree (cpd or c/deg or c/°) of visual angle.

Contrast

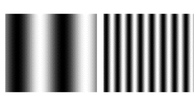

Contrast = 1

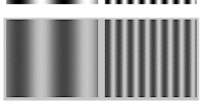

Contrast = 0.5

Two-dimensional Fourier spectra

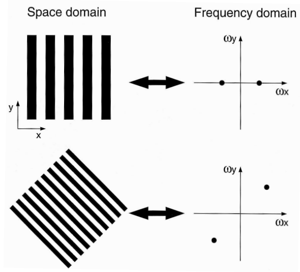

Two-dimensional Fourier spectrum

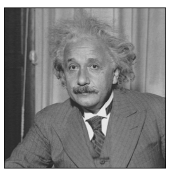

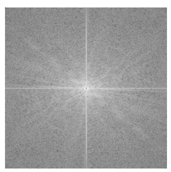

Intensity in the Fourier spectrum at each location indicates amount of contrast (in the original image) for each spatial frequency and orientation.

Linear systems analysis

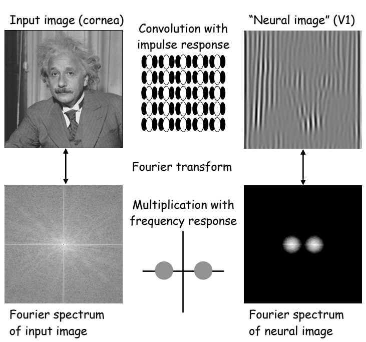

Spatial pattern vision and linear systems theory

- Receptive fields and neural images
- Shift-invariant linear systems and convolution
- Fourier transform and frequency response
- **Applications of linear systems to spatial vision**
  – Contrast sensitivity
  – Spatial frequency and orientation channels
  – Spatial frequency and orientation adaptation
  – Masking

Spatial contrast sensitivity

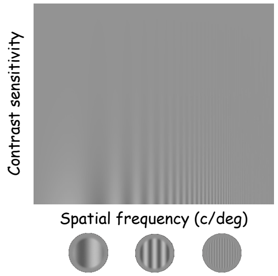

Spatial contrast sensitivity

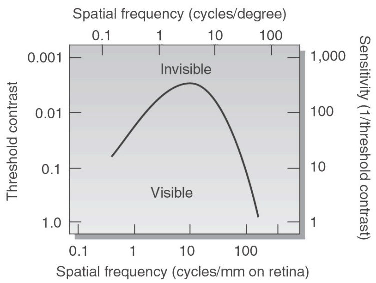

Spatial frequency channels

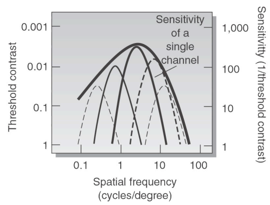

Each channel is sensitive to a narrow range of frequencies. Overall contrast sensitivity depends on all of them together.

Theory of spatial pattern analysis by the visual system

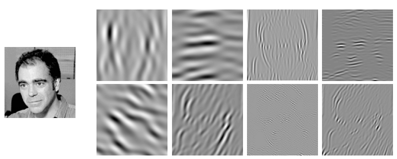

Low sf filters encode coarse-scale information (large objects, overall shape)

High sf filters encode fine-scale information (small objects, detail)

Multi-resolution model

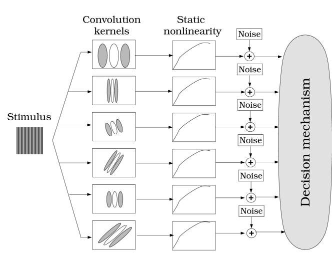

Psychophysical/perceptual evidence for spatial-frequency and orientation selective channels

Orientation and spatial frequency selective adaptation: appearance

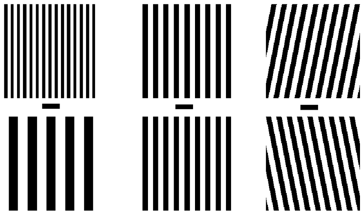

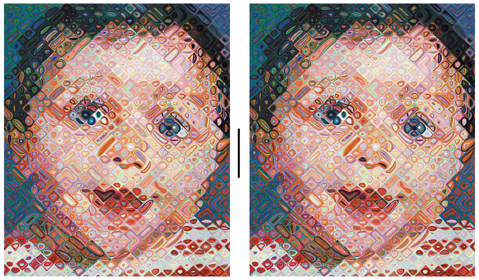

Emma by Chuck Close

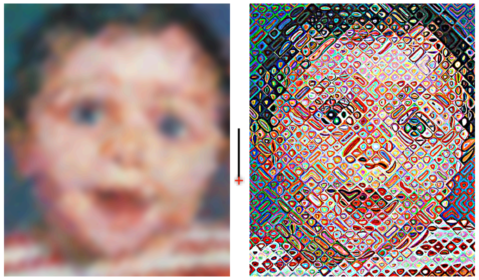

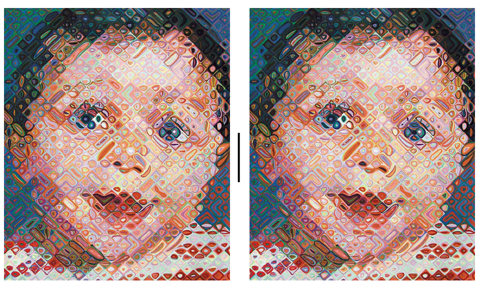

Spatial frequency selective adaptation: Appearance

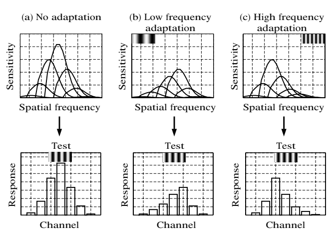

Spatial frequency selective adaptation: Sensitivity

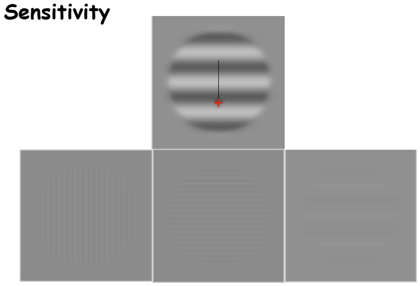

Spatial frequency selective adaptation: sensitivity

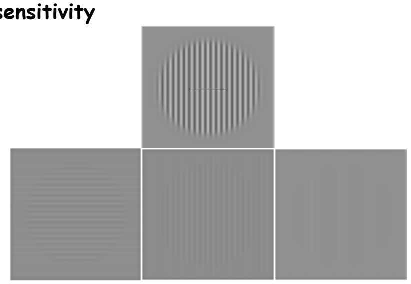

Contrast sensitivity before & after adaptation

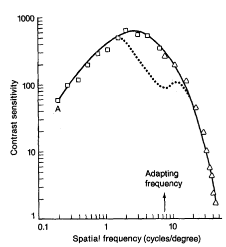

Figure 3.22
Squares and solid curve: Contrast sensitivity function for a sine-wave grating. (From Campbell & Robson, 1968.) Dotted curve: Contrast sensitivity measured after adaptation to a 7.5 cycles/degree grating.

Contrast sensitivity before & after adaptation

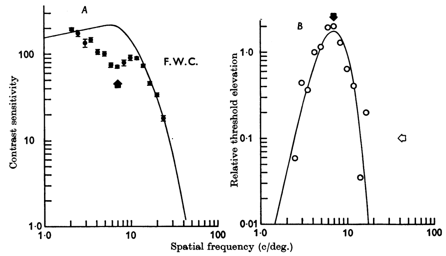

Contrast sensitivity before & after adaptation

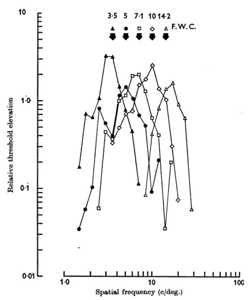

Orientation selective adaptation

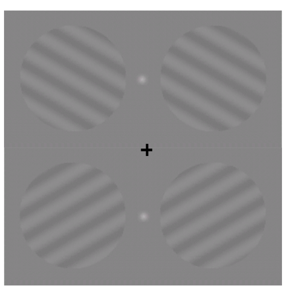

Orientation selective adaptation

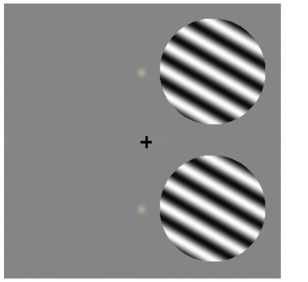

Orientation selective adaptation

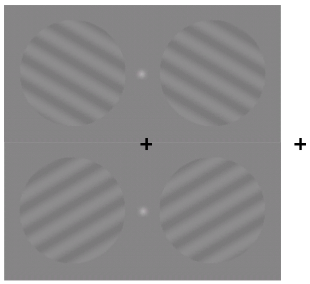

Applications: Optical Line Spread Function (Campbell & Gubisch, 1966)

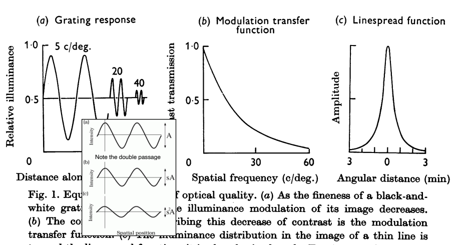

Fig. 1. Equ… optical quality. (a) As the fineness of a black-and-white grat… e illuminance modulation of its image decreases. (b) The co… ibing this decrease of contrast is the modulation transfer fu… luminance distribution in the image of a thin line is termed the linespread function; it is also obtained as the Fourier transform of the modulation transfer function.

Applications: Optical Line Spread Function (Campbell & Gubisch, 1966)

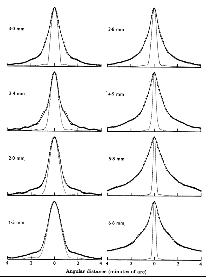

Applications: Optical Line Spread Function (Campbell & Gubisch, 1966)

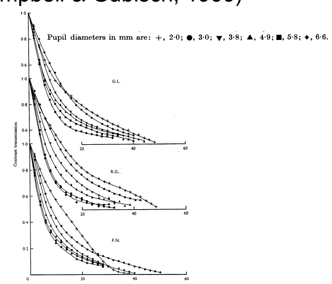

Applications: Optical Line Spread Function (Westheimer, 1986)

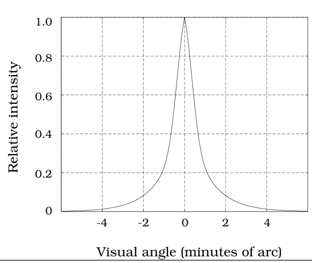

Westheimer's Linespread Function. Analytic approximation of the human linespread function for an eye with a 3.0mm diameter pupil (Westheimer, 1986).

Applications: Optical Line Spread Function (Westheimer, 1986)

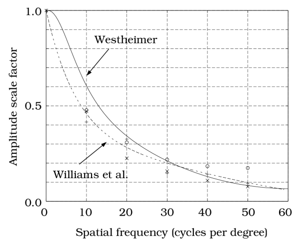

Modulation transfer function measurements of the optical quality of the lens made using visual interferometry (Williams et al., 1995; described in Chapter 3). The data are compared with the predictions from the linespread suggested by Westheimer (1986) and a curve fit through the data by Williams et al. (1995).

Applications: Optical Line Spread Function, Wavelength Dependency

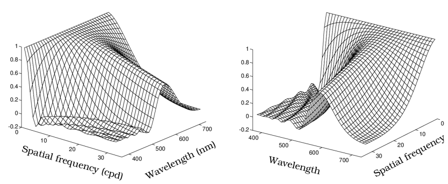

OTF of Chromatic Aberration: Two views of the modulation transfer function of a model eye at various wavelengths. The model eye has the same chromatic aberration as the human eye (see Figure 2.23) and a 3.0mm pupil diameter. The eye is in focus at 580nm; the curve at 580nm is diffraction limited. The retinal image has no contrast beyond four cycles per degree at short wavelengths. (From Marimont and Wandell, 1993).

Applications: Optical Line Spread Function, Wavelength Dependency