Spatiotemporal vision & filters

John Greenwood
Department of Experimental Psychology

NEUR3001

Contact: john.greenwood@ucl.ac.uk

Today

- How do you characterise a visual system?
- A useful tool: Fourier analysis
  - Relationship to natural scenes
- Contrast Sensitivity Functions
- Early stages of visual processing
  - Local/Gabor-based decomposition (LGN & V1)
  - Motion processing (space-time RFs)
- Why Fourier analysis?

The Fourier Transform

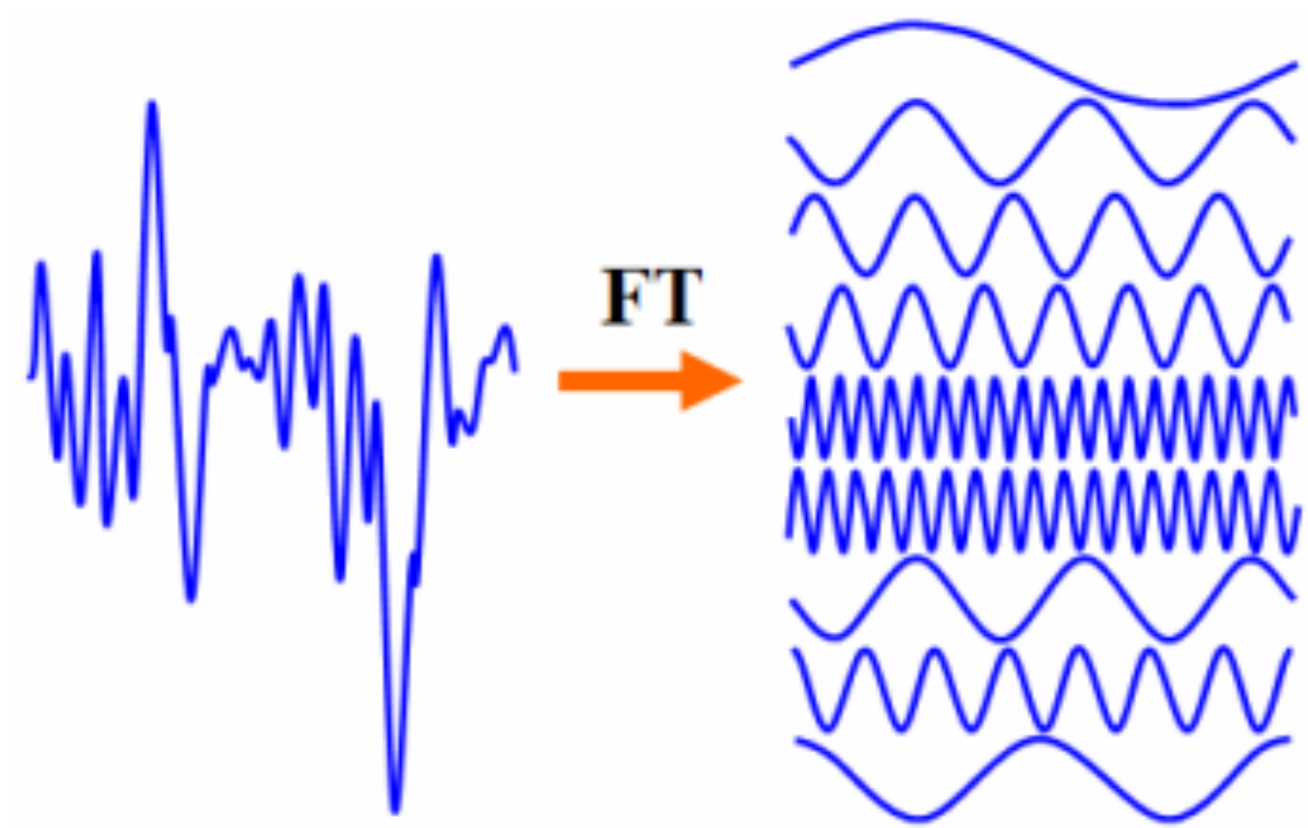

- Fourier (1822) showed that any signal can be decomposed into a sum of sine waves at different frequencies, amplitudes and phases

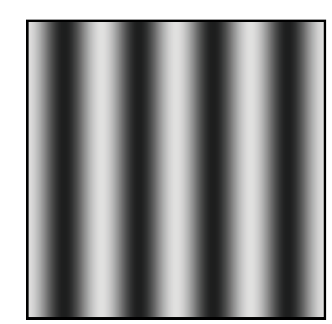

Sine wave amplitude

- Amplitude for a sine wave grating gives luminance contrast
    - Michelson contrast = $\dfrac{\text{Lmax - Lmin}}{\text{Lmin + Lmax}}$

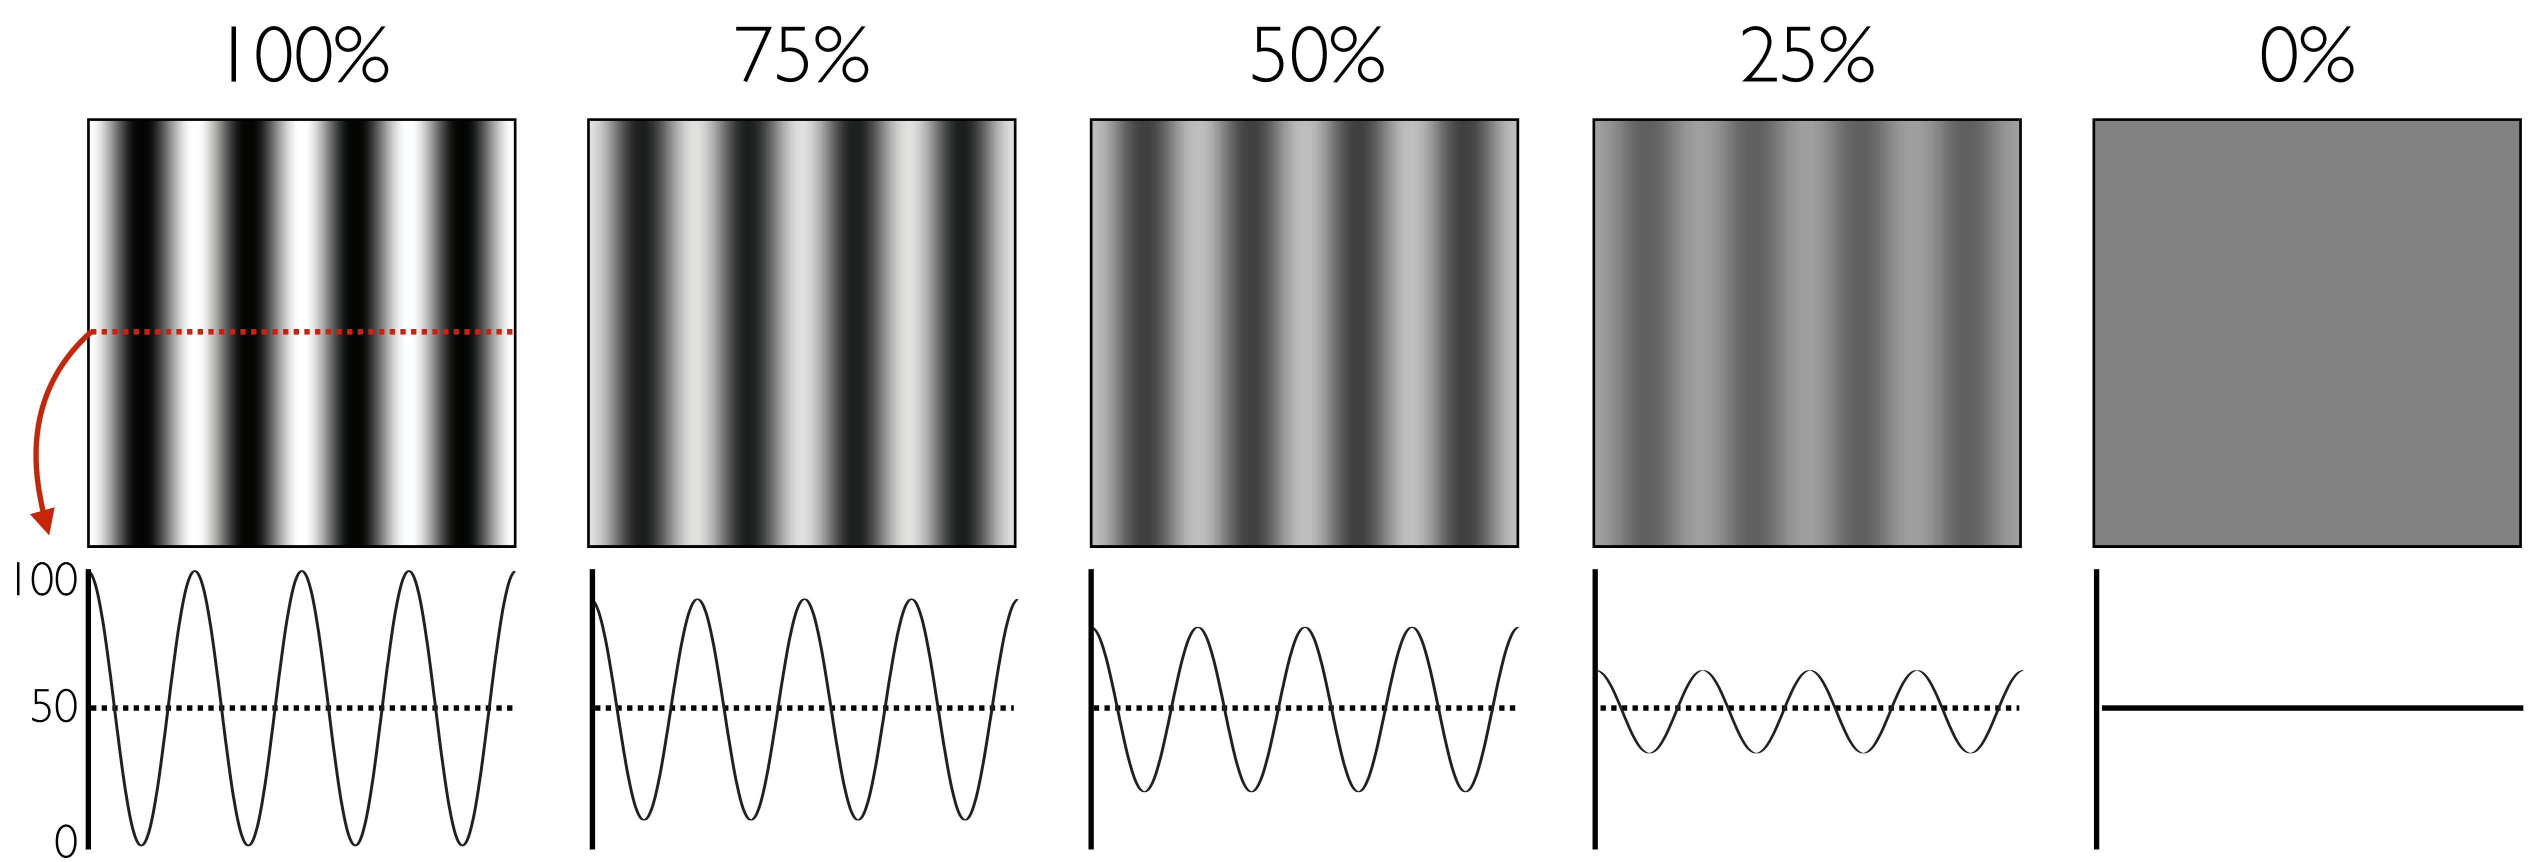

Sine wave spatial frequency

- Spatial frequency determines the variations across space
  - Reported as the number of cycles in a spatial region (peak to peak)
  - Captures the fine vs. coarse detail in an image

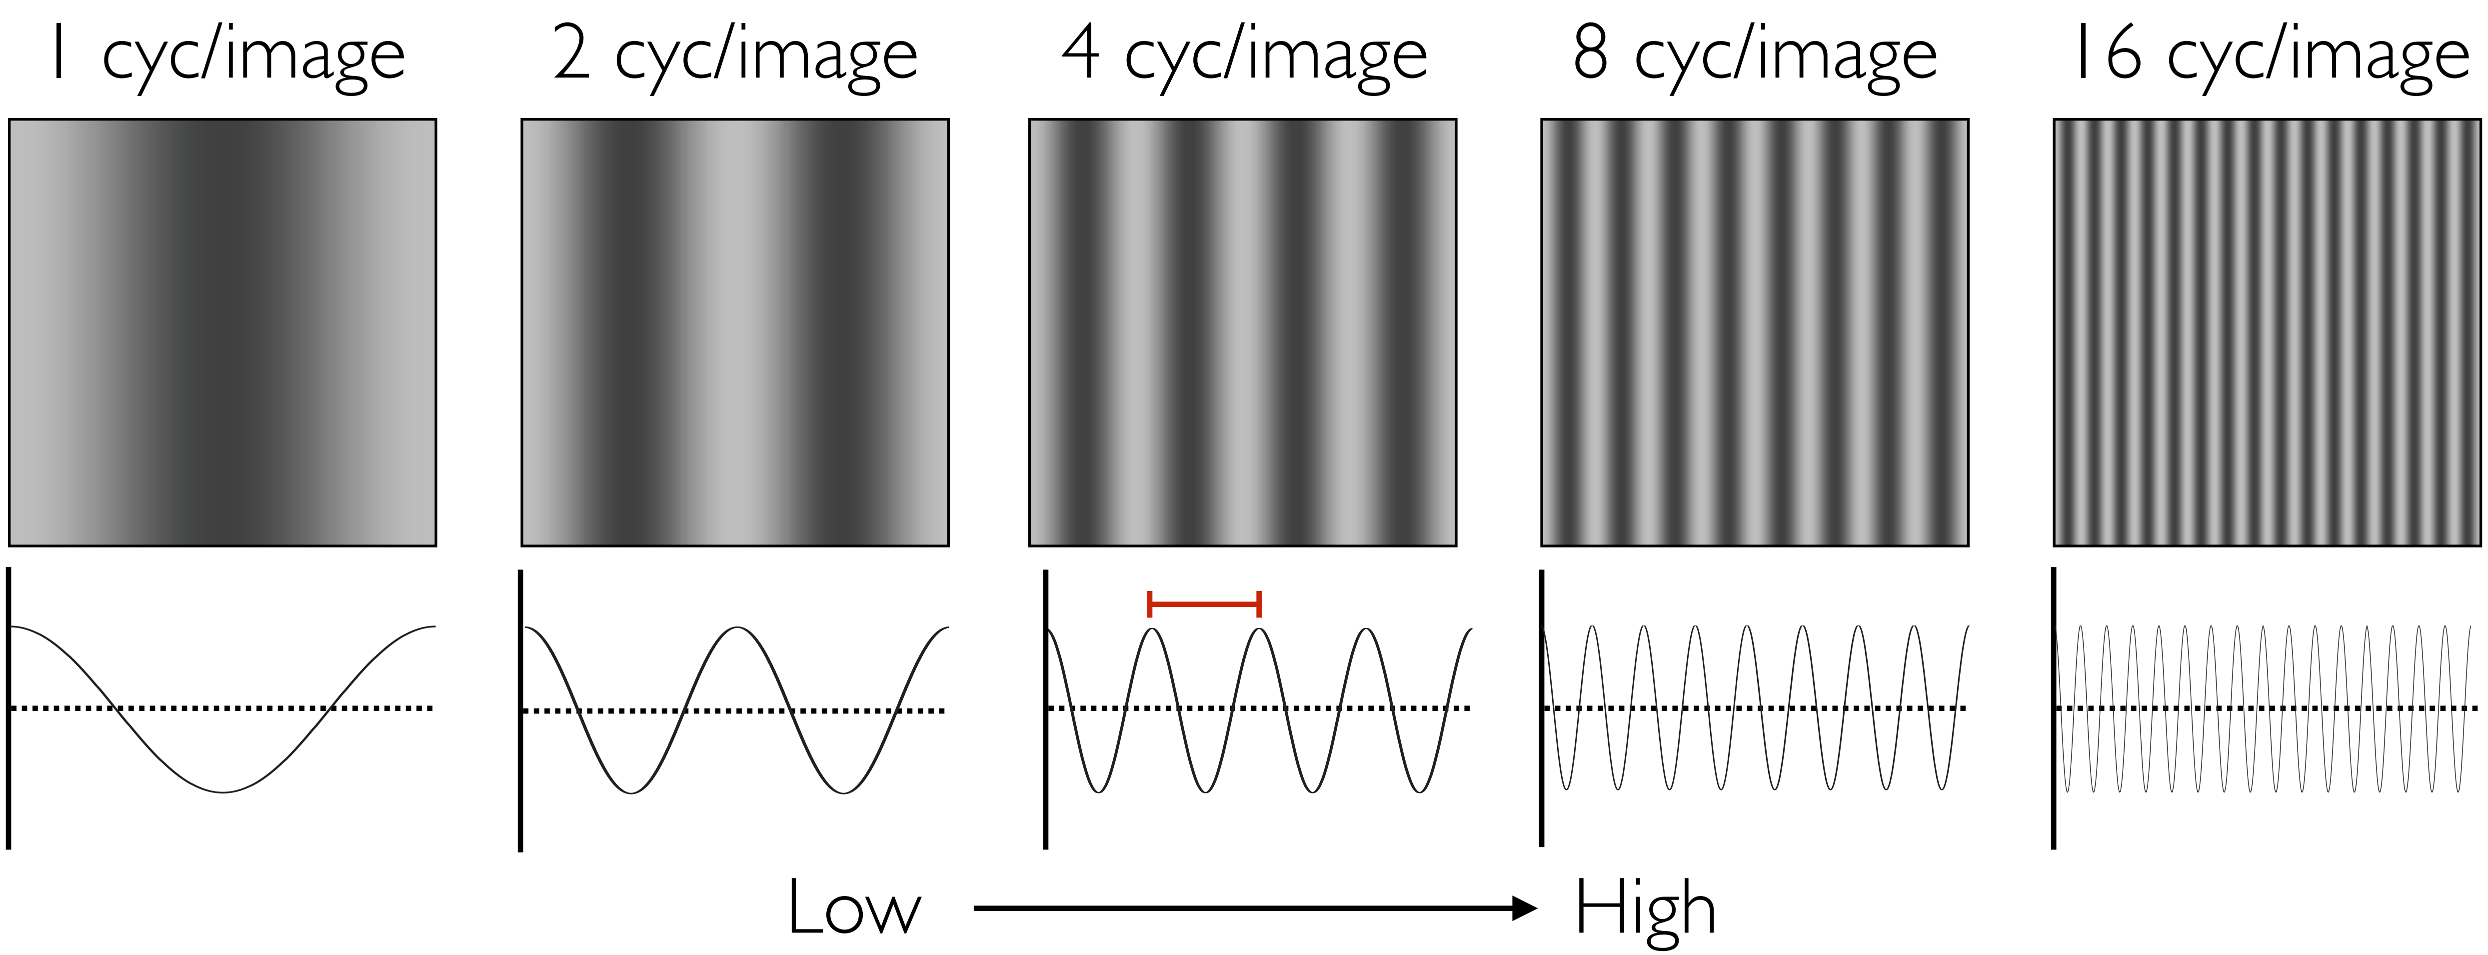

Sine wave phase

- Phase determines the point at which variations occur in space, e.g. the starting point of the cycle
  - Represented in radians with a cyclical structure

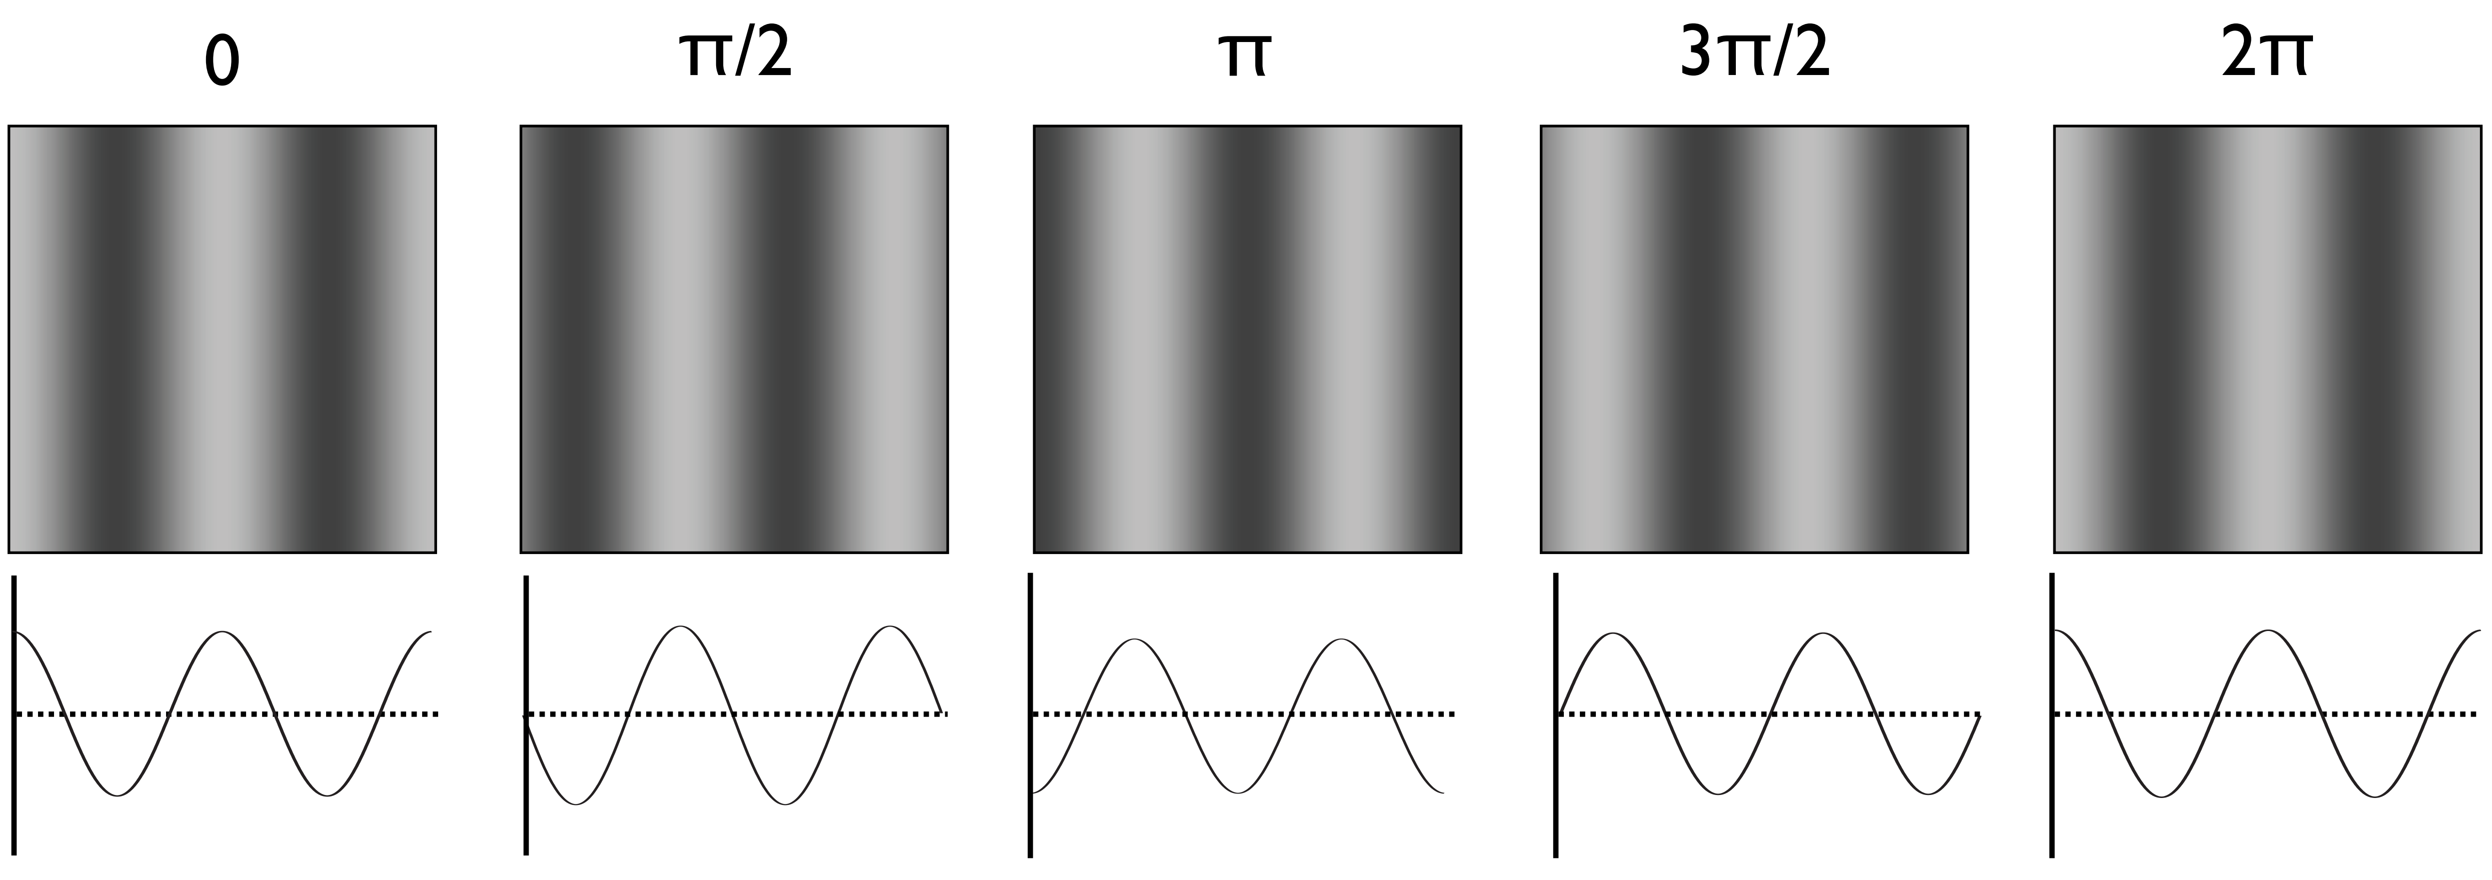

From sine to square wave

- How do we make an image using sine waves?
- Easiest example: a square wave
  - How do you get a square wave from sine wave components?

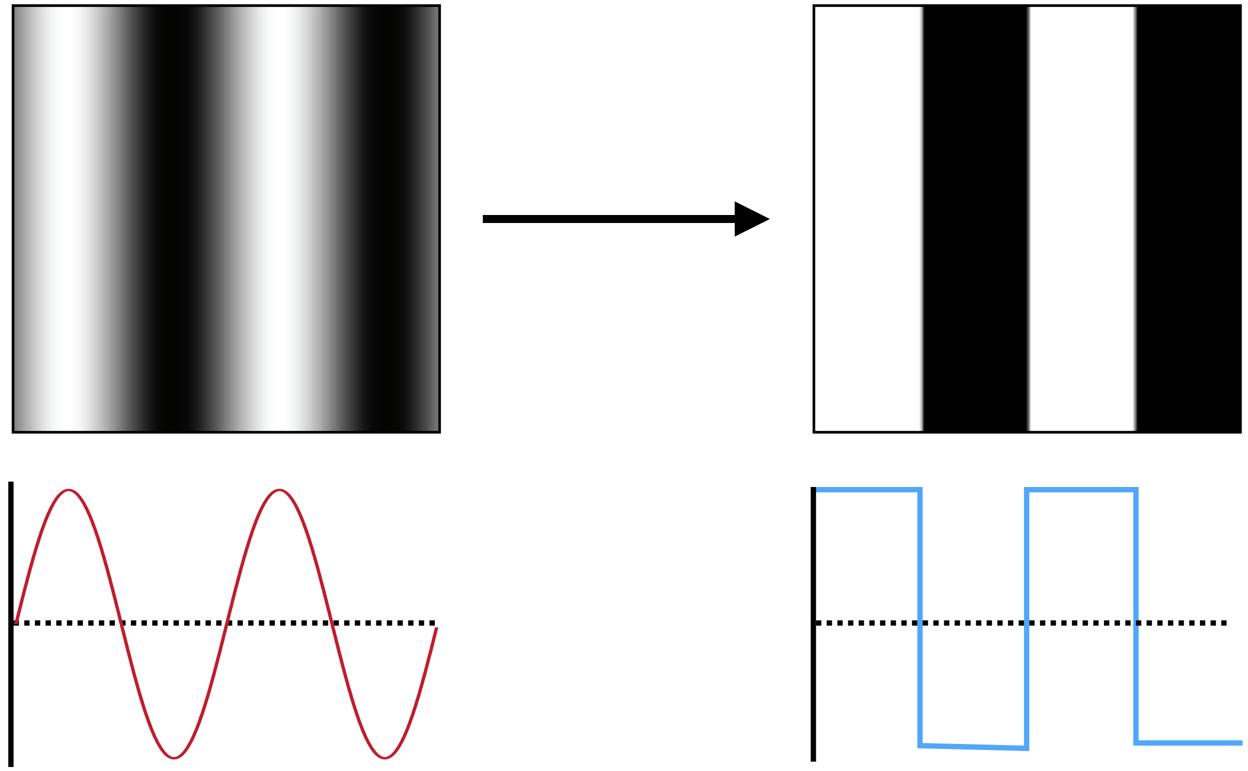

From 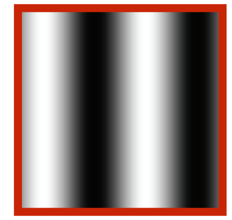 to 

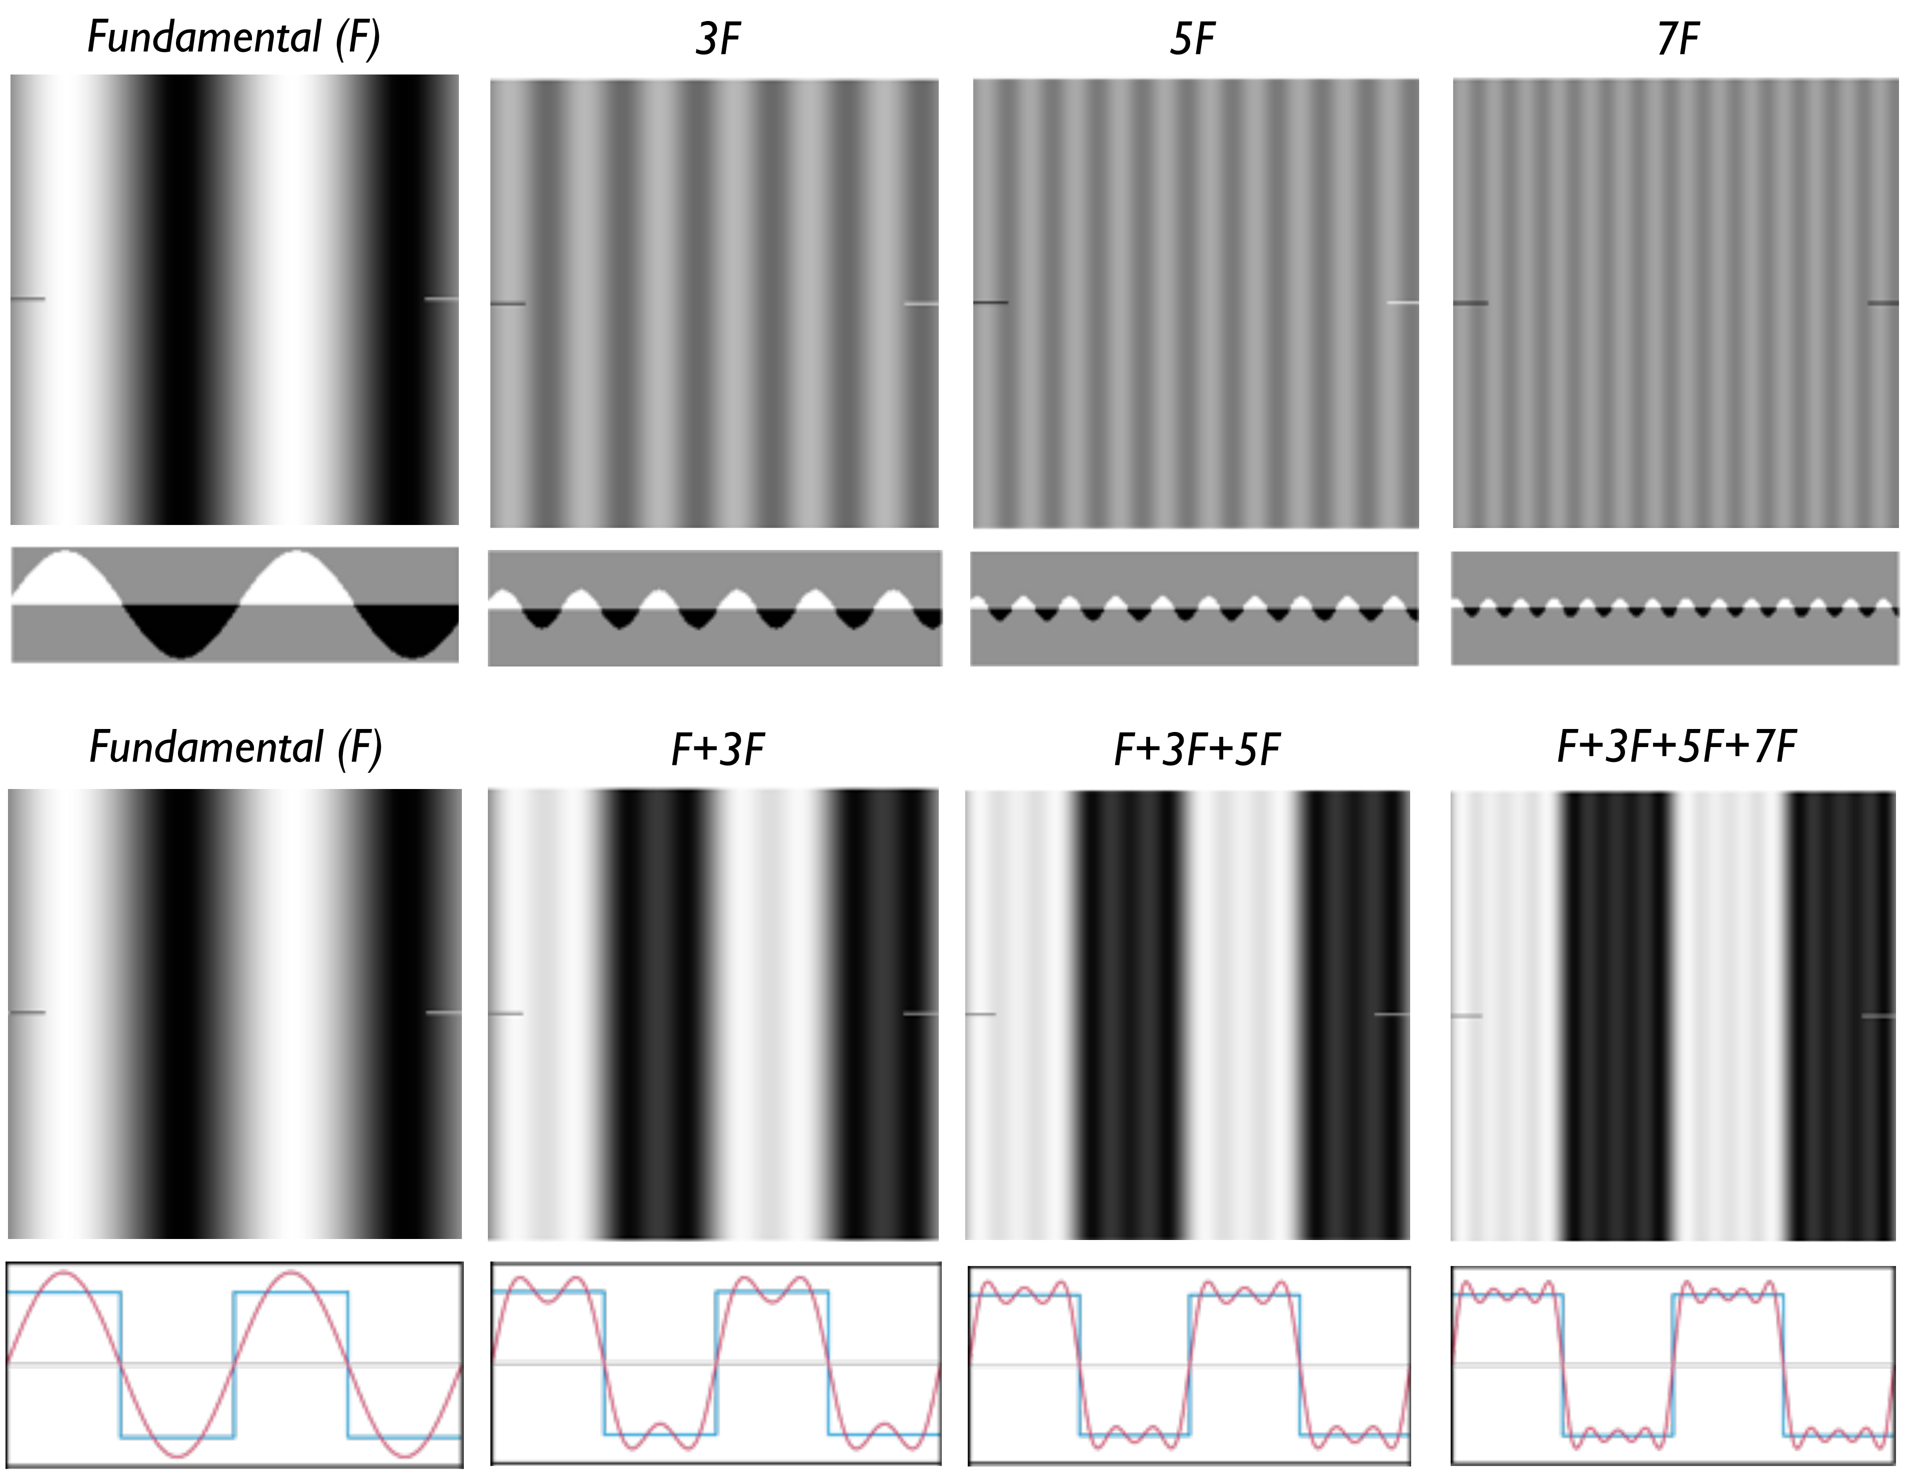

- Take a sine wave with matched spatial frequency: the fundamental
- Add the odd harmonics with decreasing amplitude

Representing the components

- Representation of a 1D square wave in terms of its components is easy
  - Plot SF against amplitude
- Gives us an amplitude spectrum for 1D stimuli

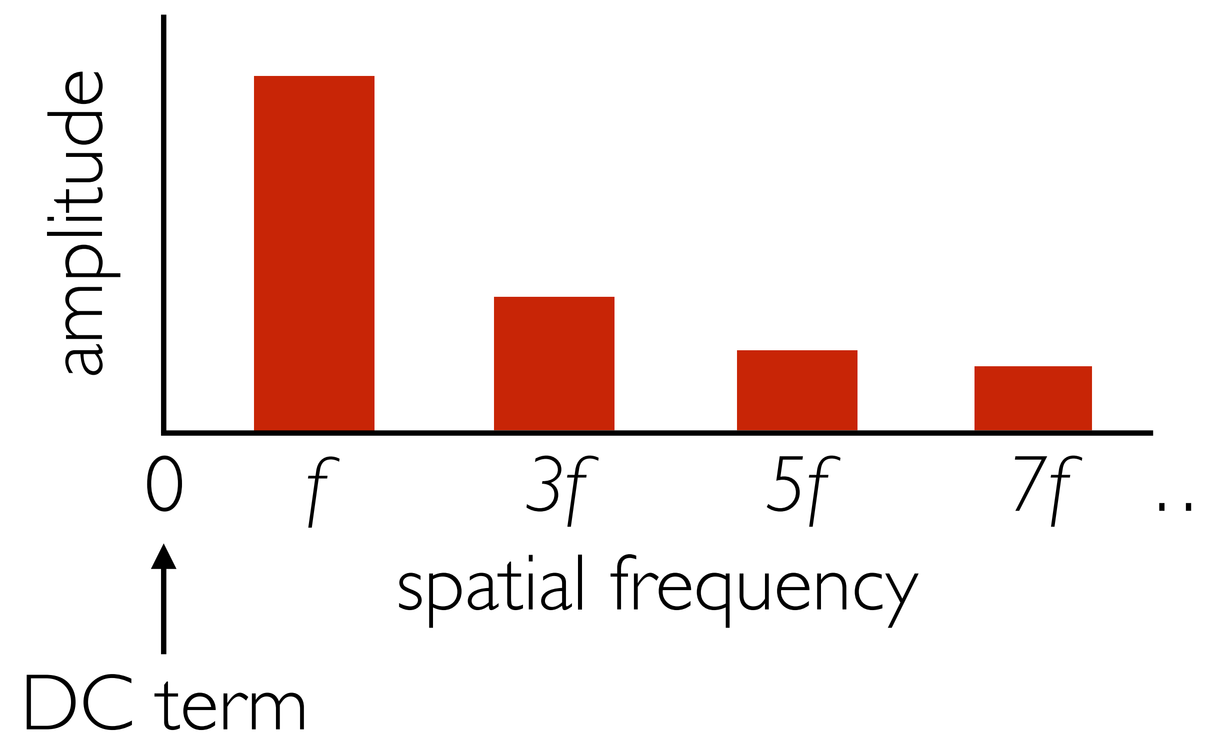

Fourier analysis in 2D

- Fourier analysis allows the decomposition of any 2D image into the sum of its components at different SF and orientations
- Returns phase and amplitude of those components

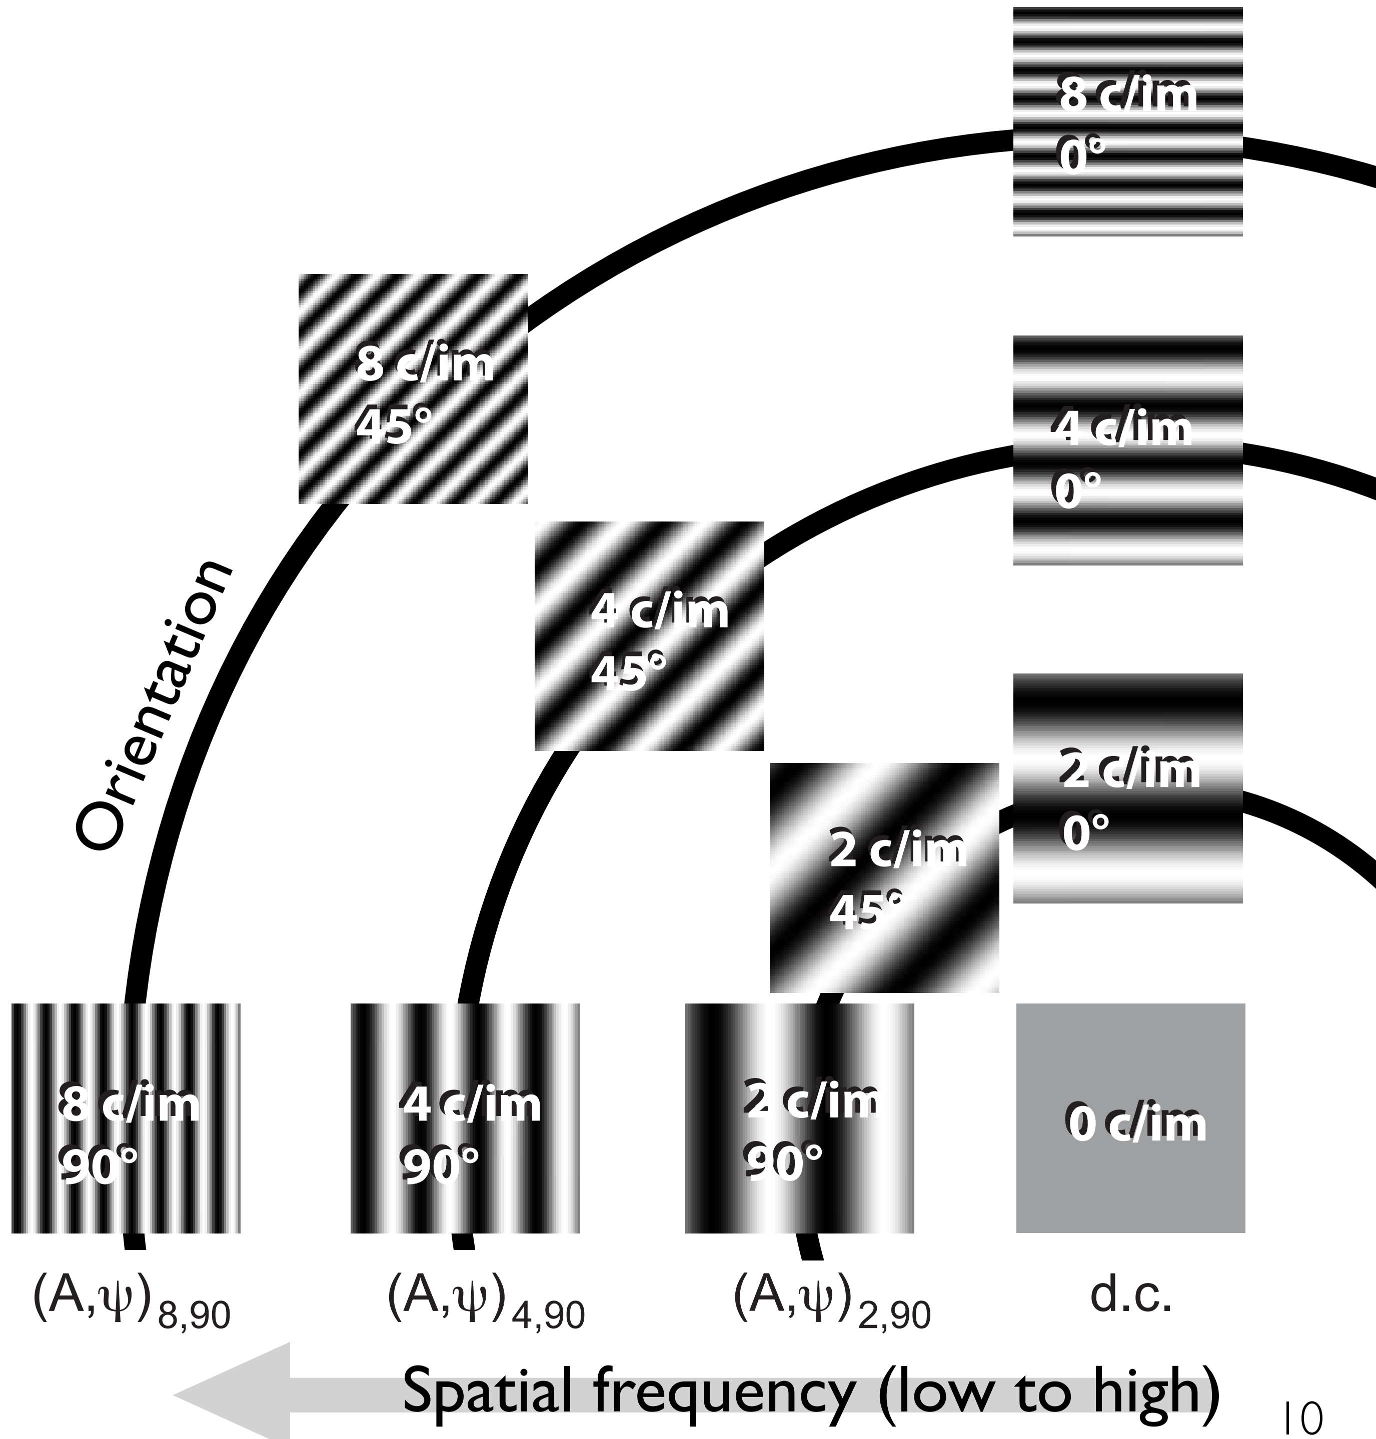

Representing a sine wave

- Simple example: a sine wave with 4cyc/im at 90deg

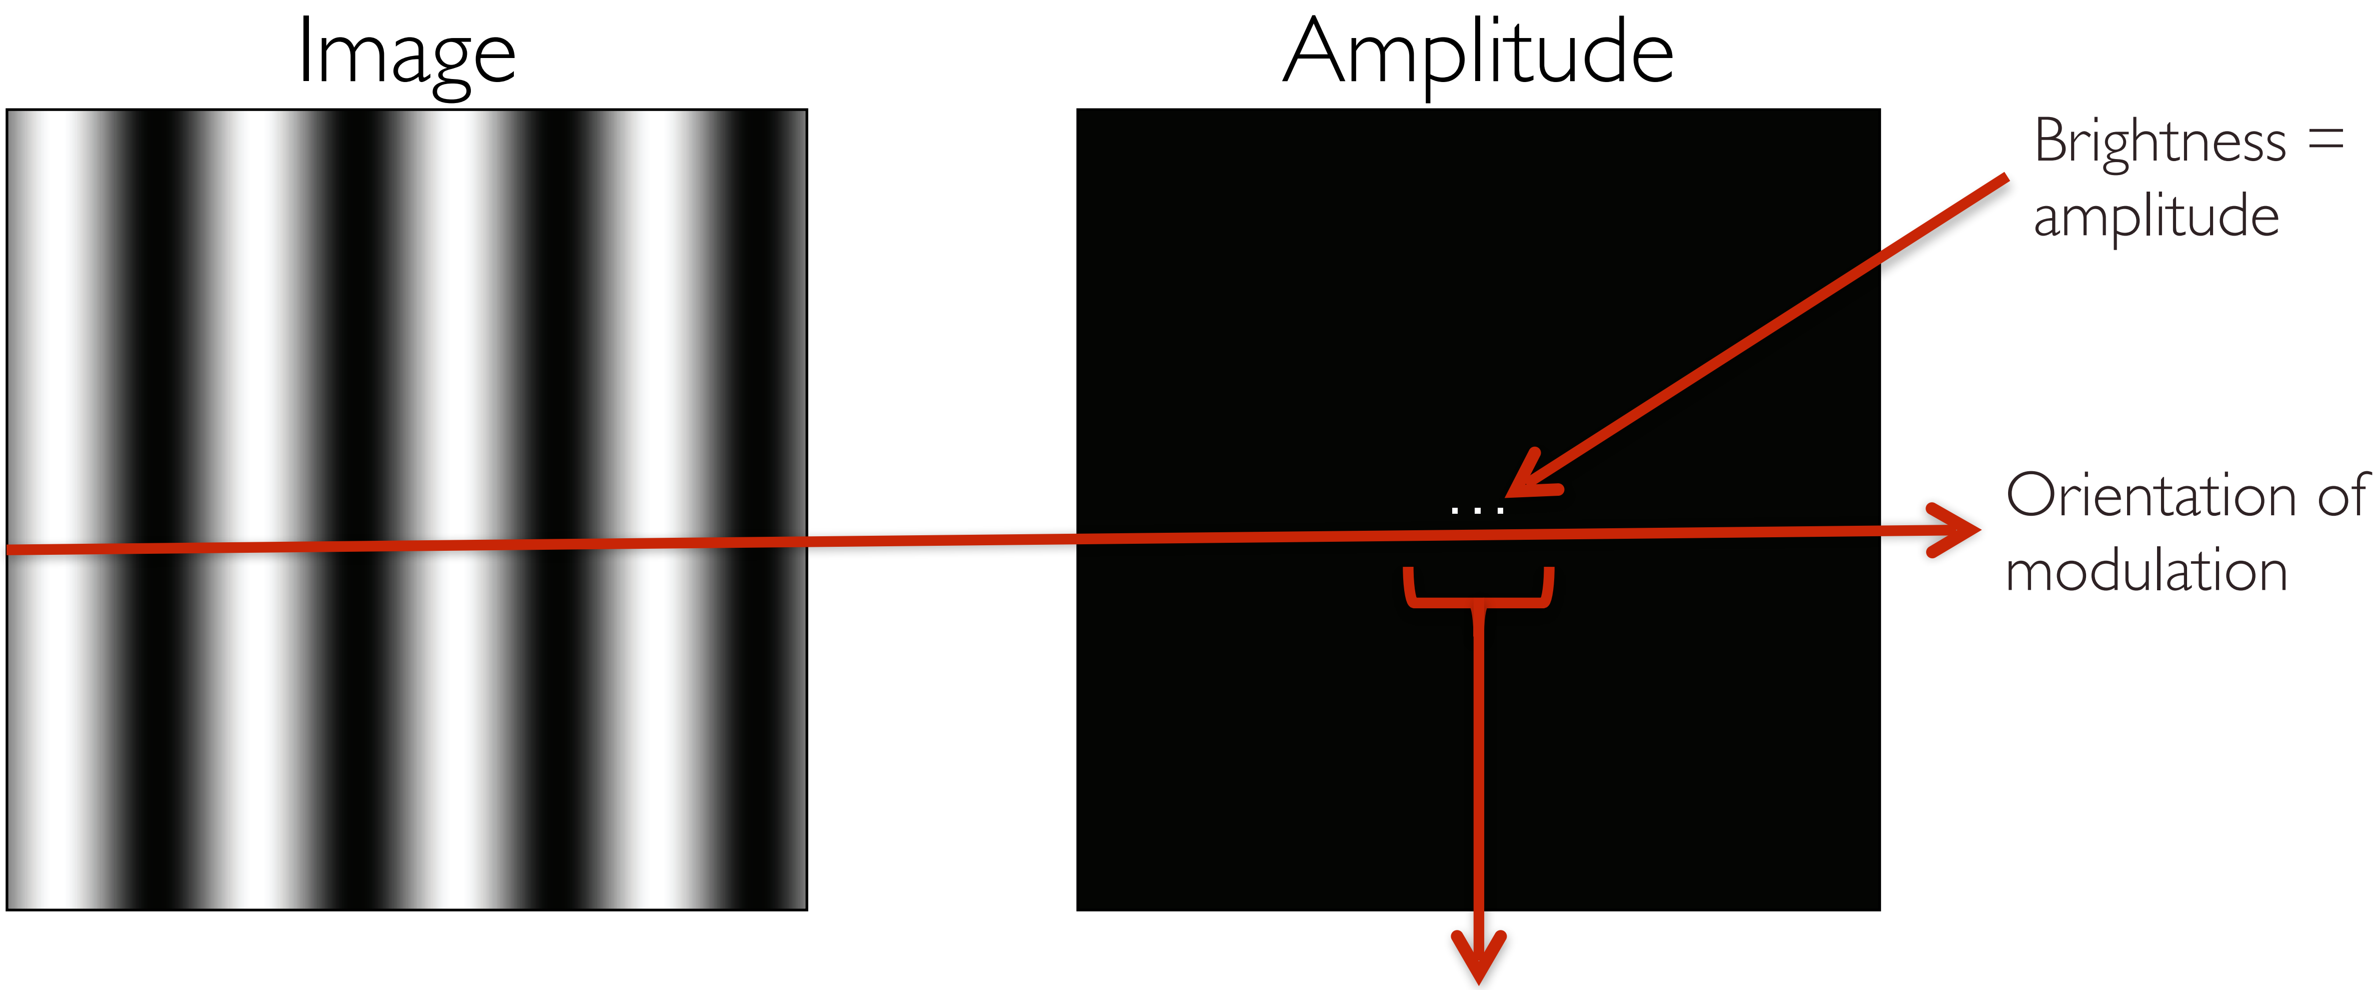

Radial separation = spatial frequency
(Further from centre = higher SF)
Note that Fourier analysis returns both +f and -f components plus the DC component

More sine waves

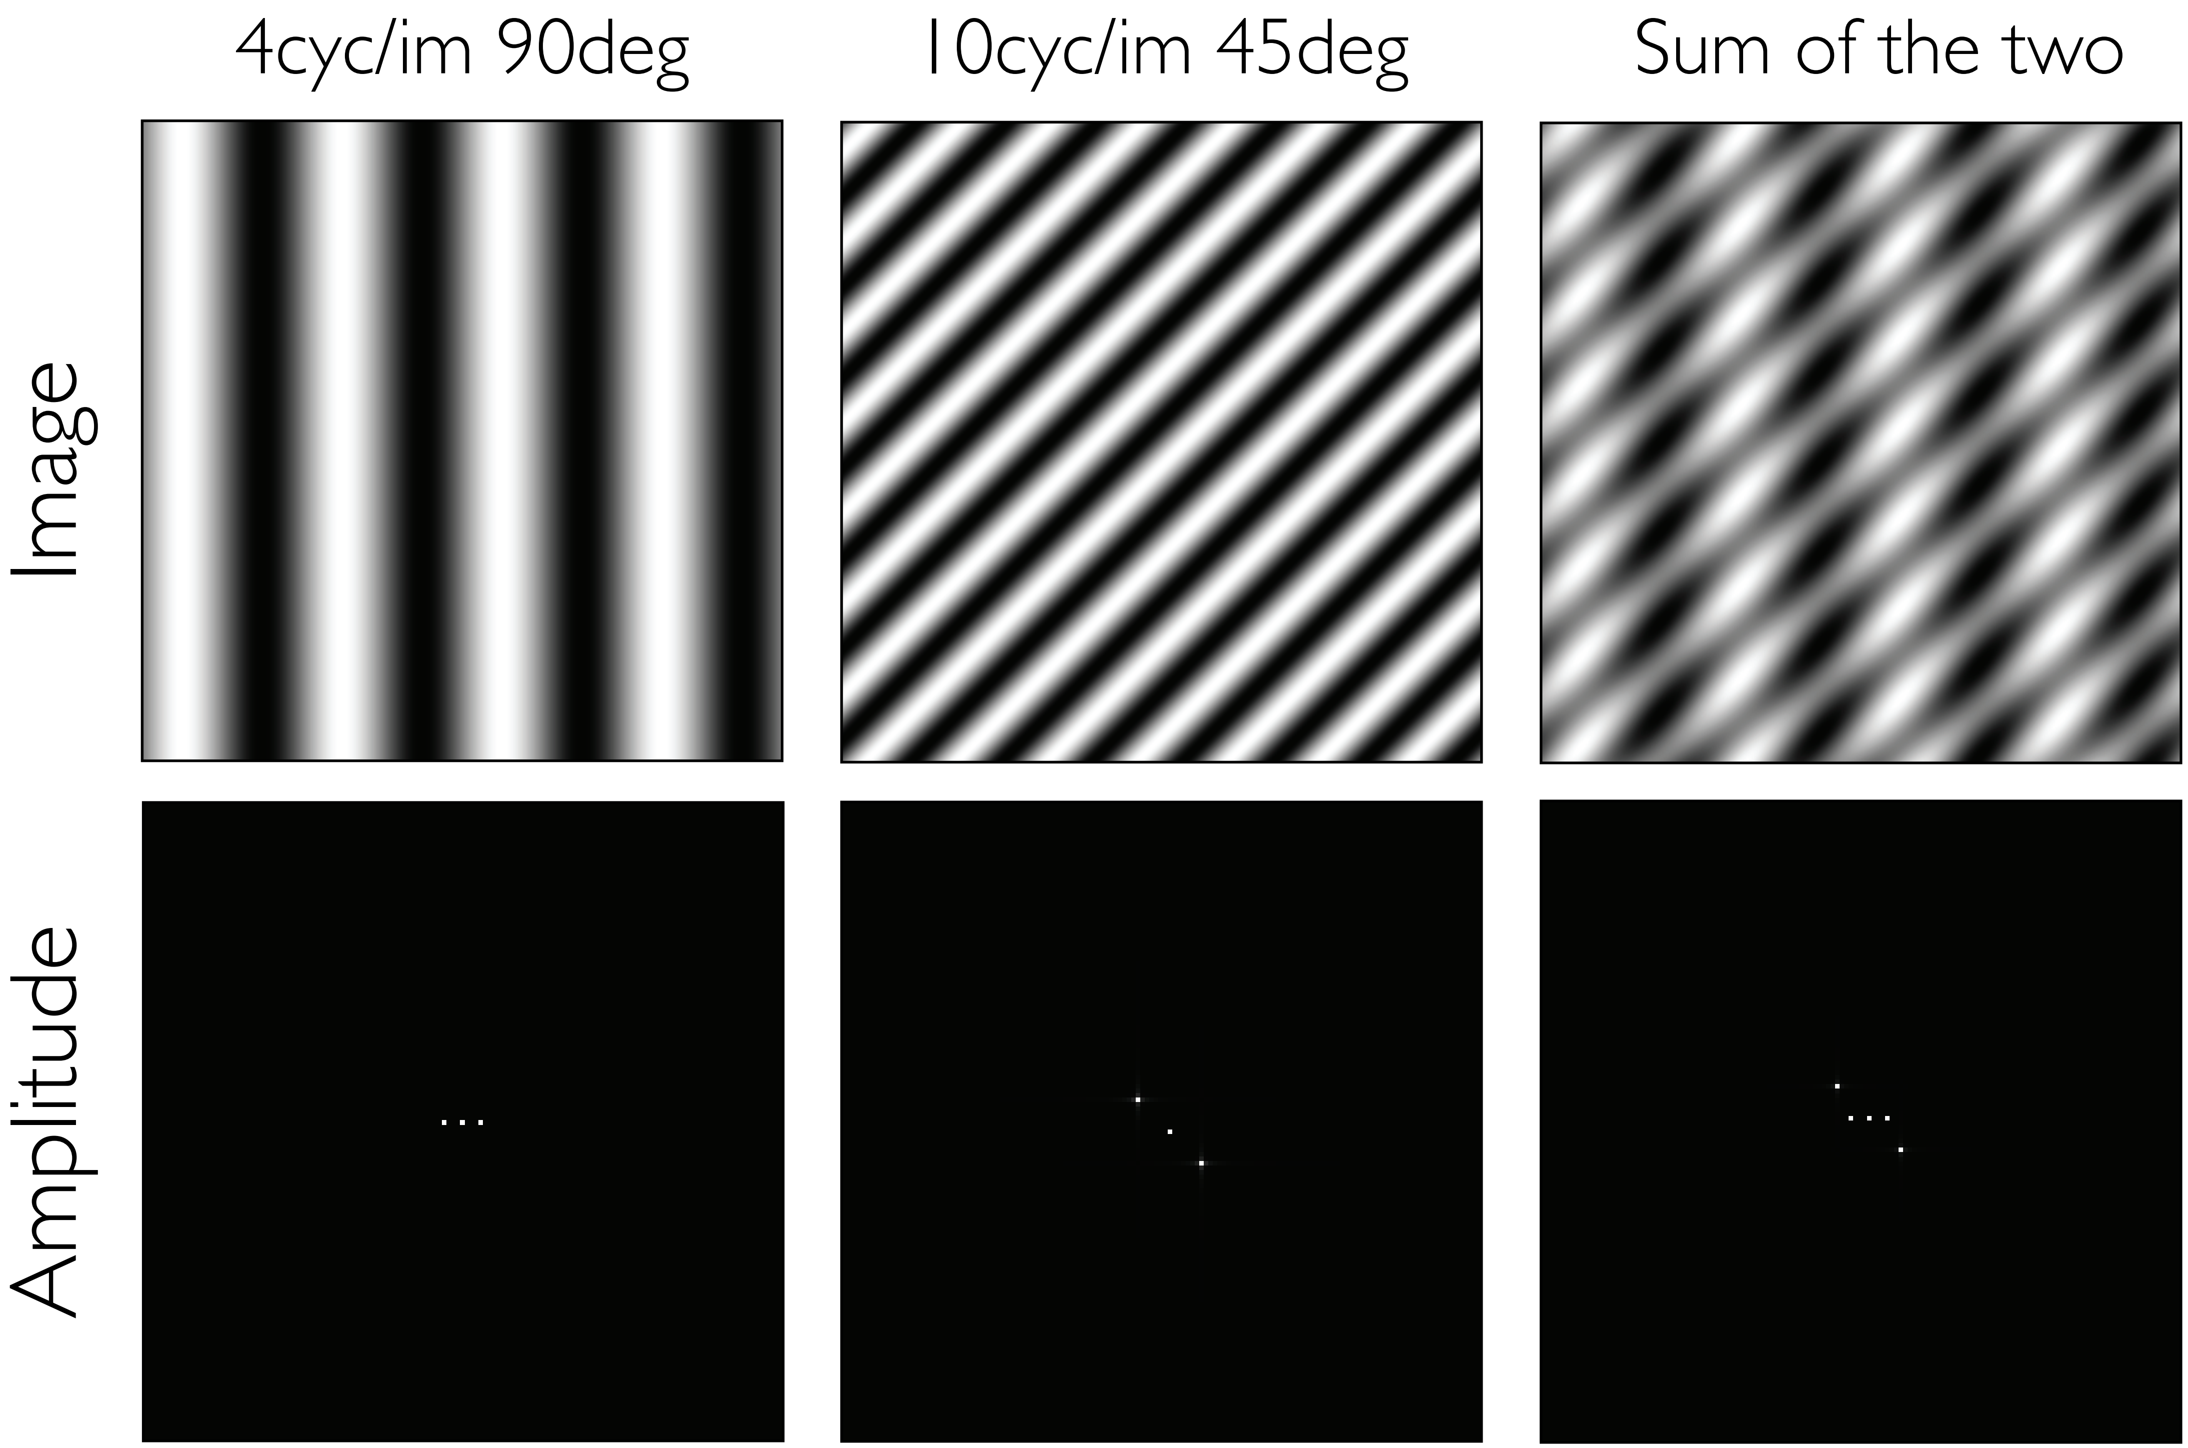

2D analysis for scenes

- Amplitude spectra of two natural images
- The house has more vertical and horizontal structure, seen in the spokes of its amplitude spectrum
- Kittens are fluffier (i.e. there is less high frequency content)

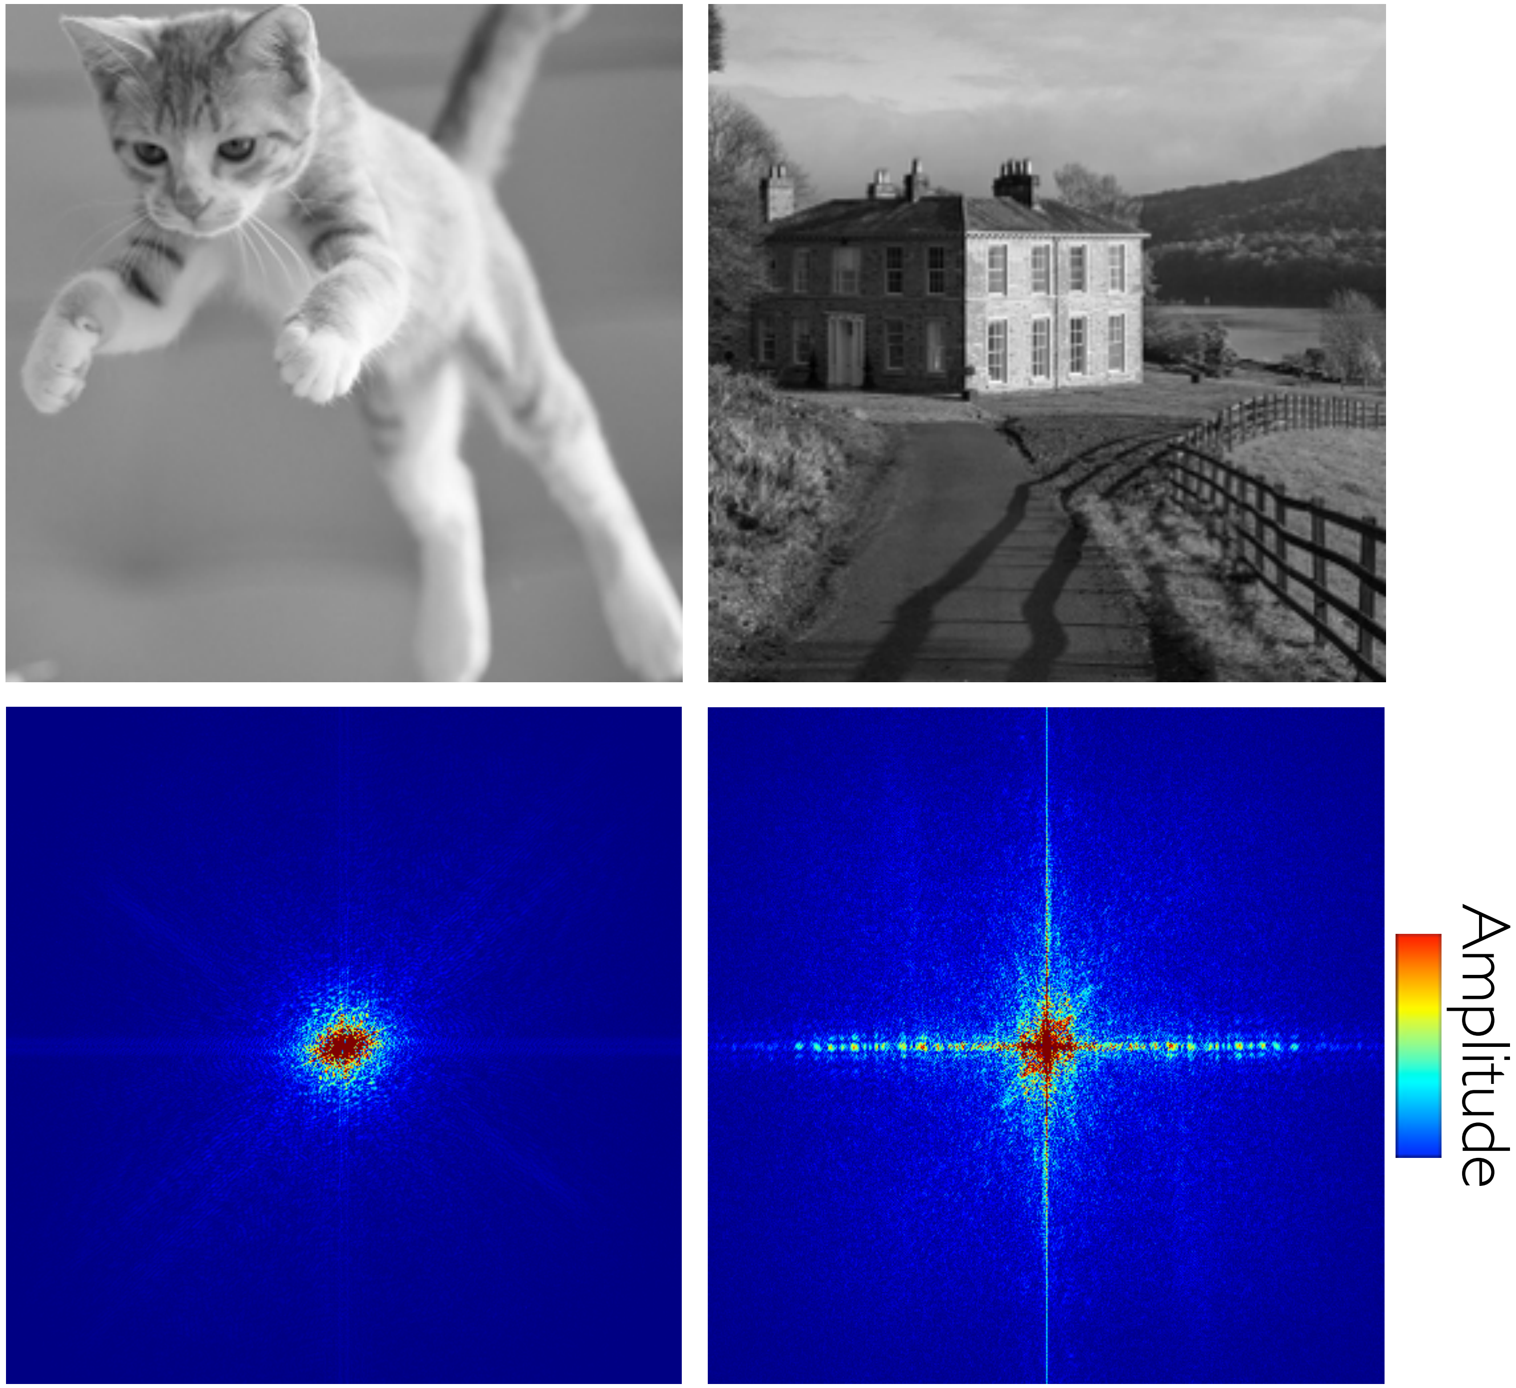

SF in natural scenes

- What does spatial frequency mean for natural scenes?
- Low-pass filtering: allow only the lowest SFs to be visible (broad blobby things)
- High-pass filtering: allow only the highest SFs to be visible (edges & fine detail)

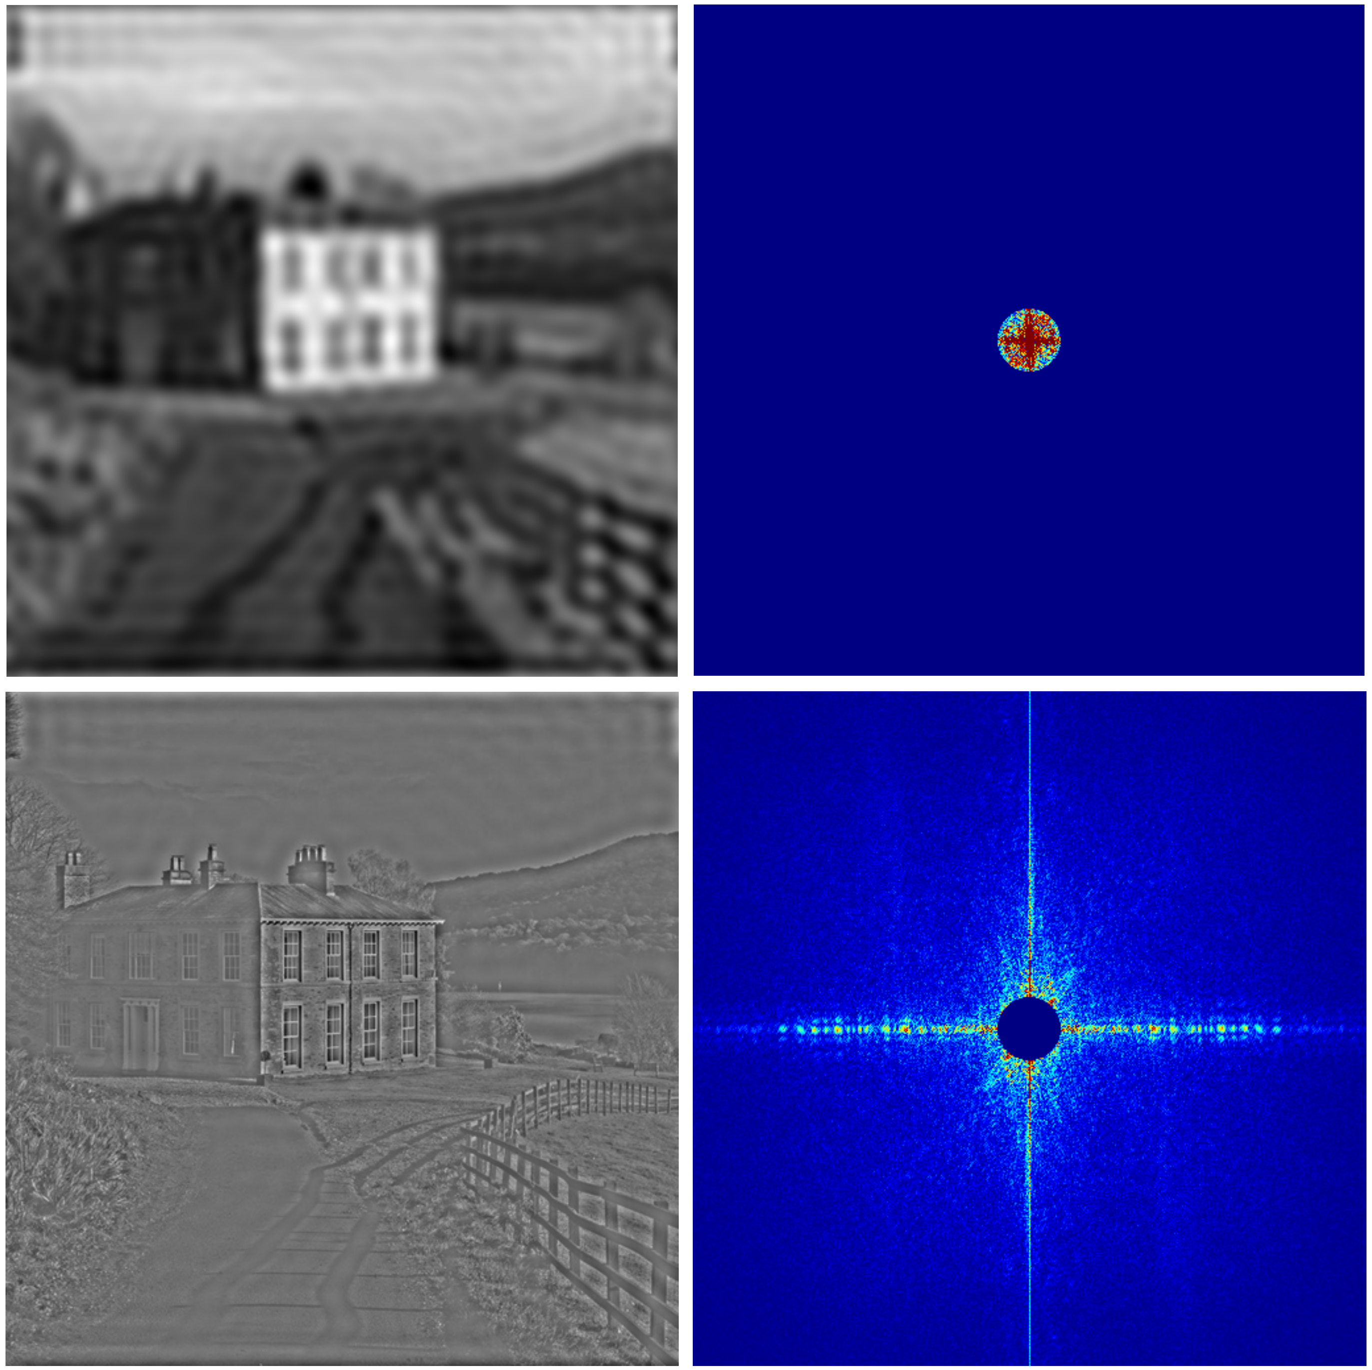

Phase in natural scenes

- Phase spectrum is plotted with the same conventions
- But for natural scenes it's hard to interpret visually - easier to understand if we modify the phase spectrum

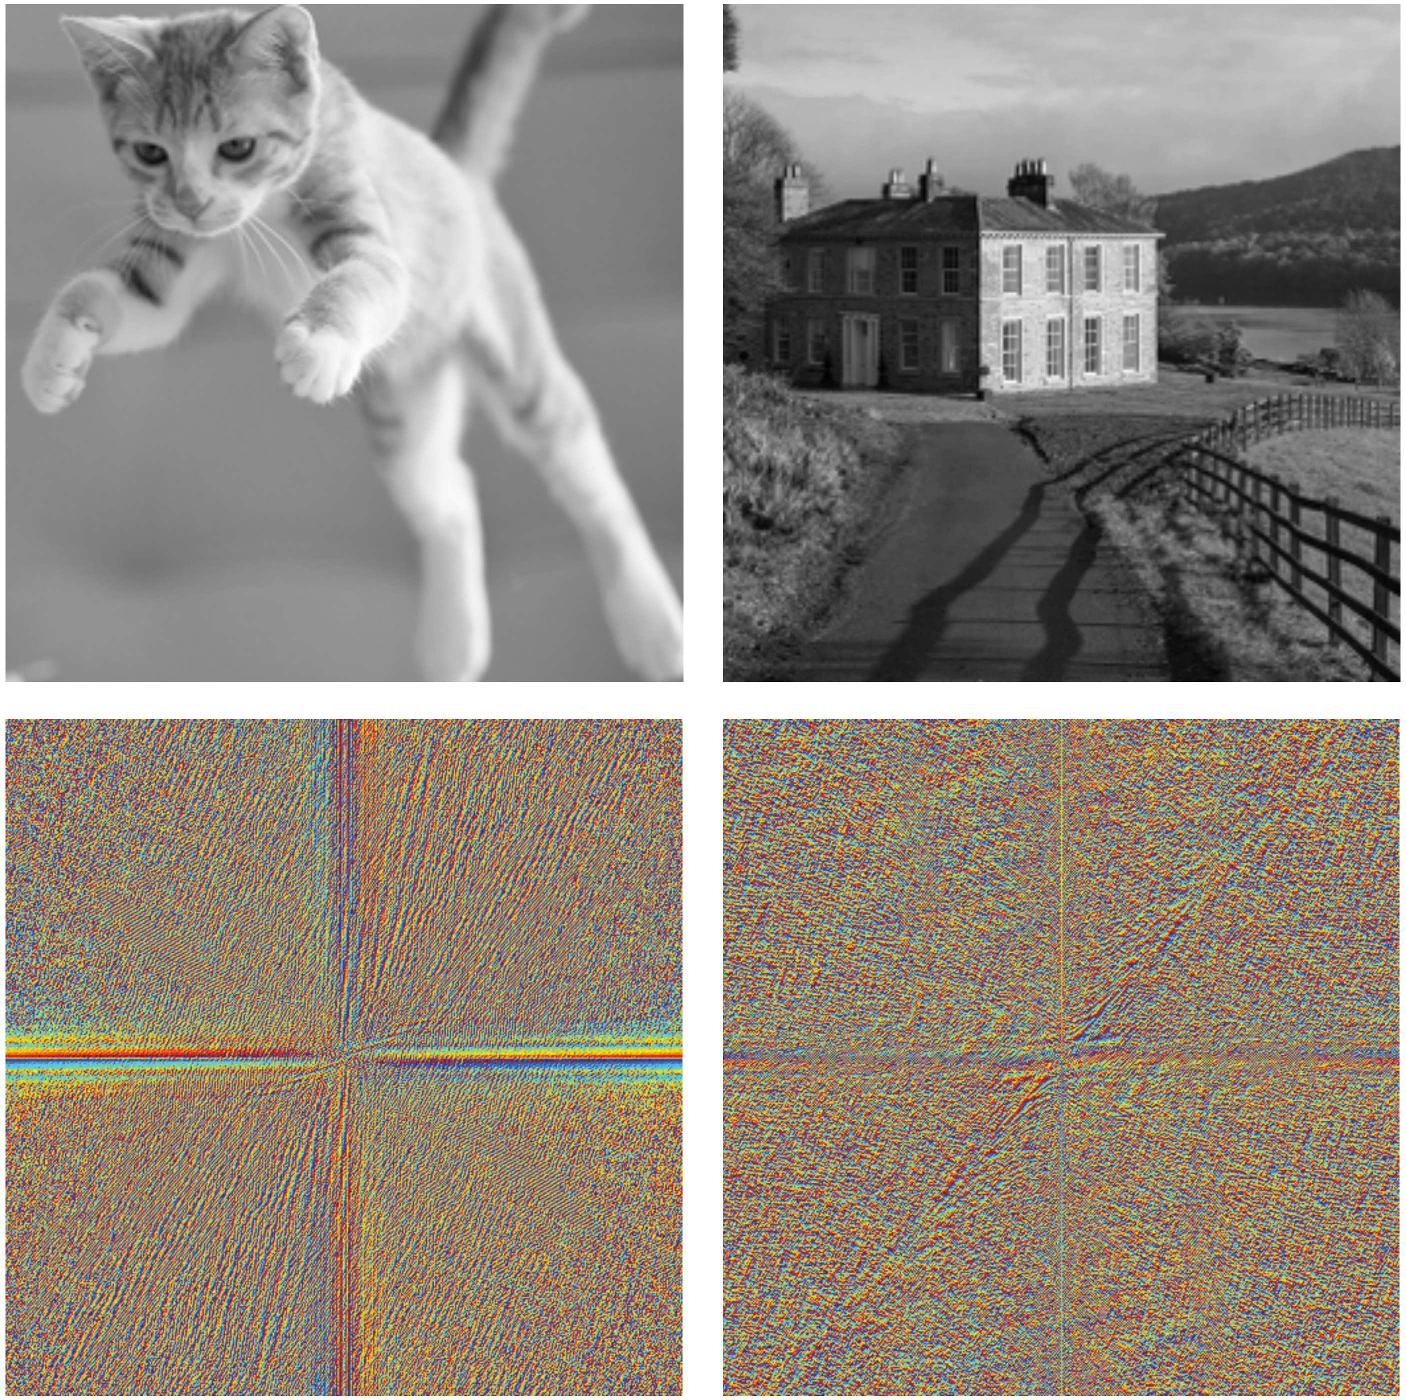

Phase scrambling

- Replace the phase spectra of these images with the phase from a white noise image and combine it with the original amplitude spectra
  - Kitty amplitude spectrum gives blobby noise
  - House amplitude gives edgy noise with a lot of cardinal orientations

Phase swapping

- Combine the kitty amplitude spectrum with the house phase spectrum
  - Makes a blobby house
  - The opposite combination is an edgy cat
- The phase of edges is very important for objects (Oppenheim & Lim, 1981)

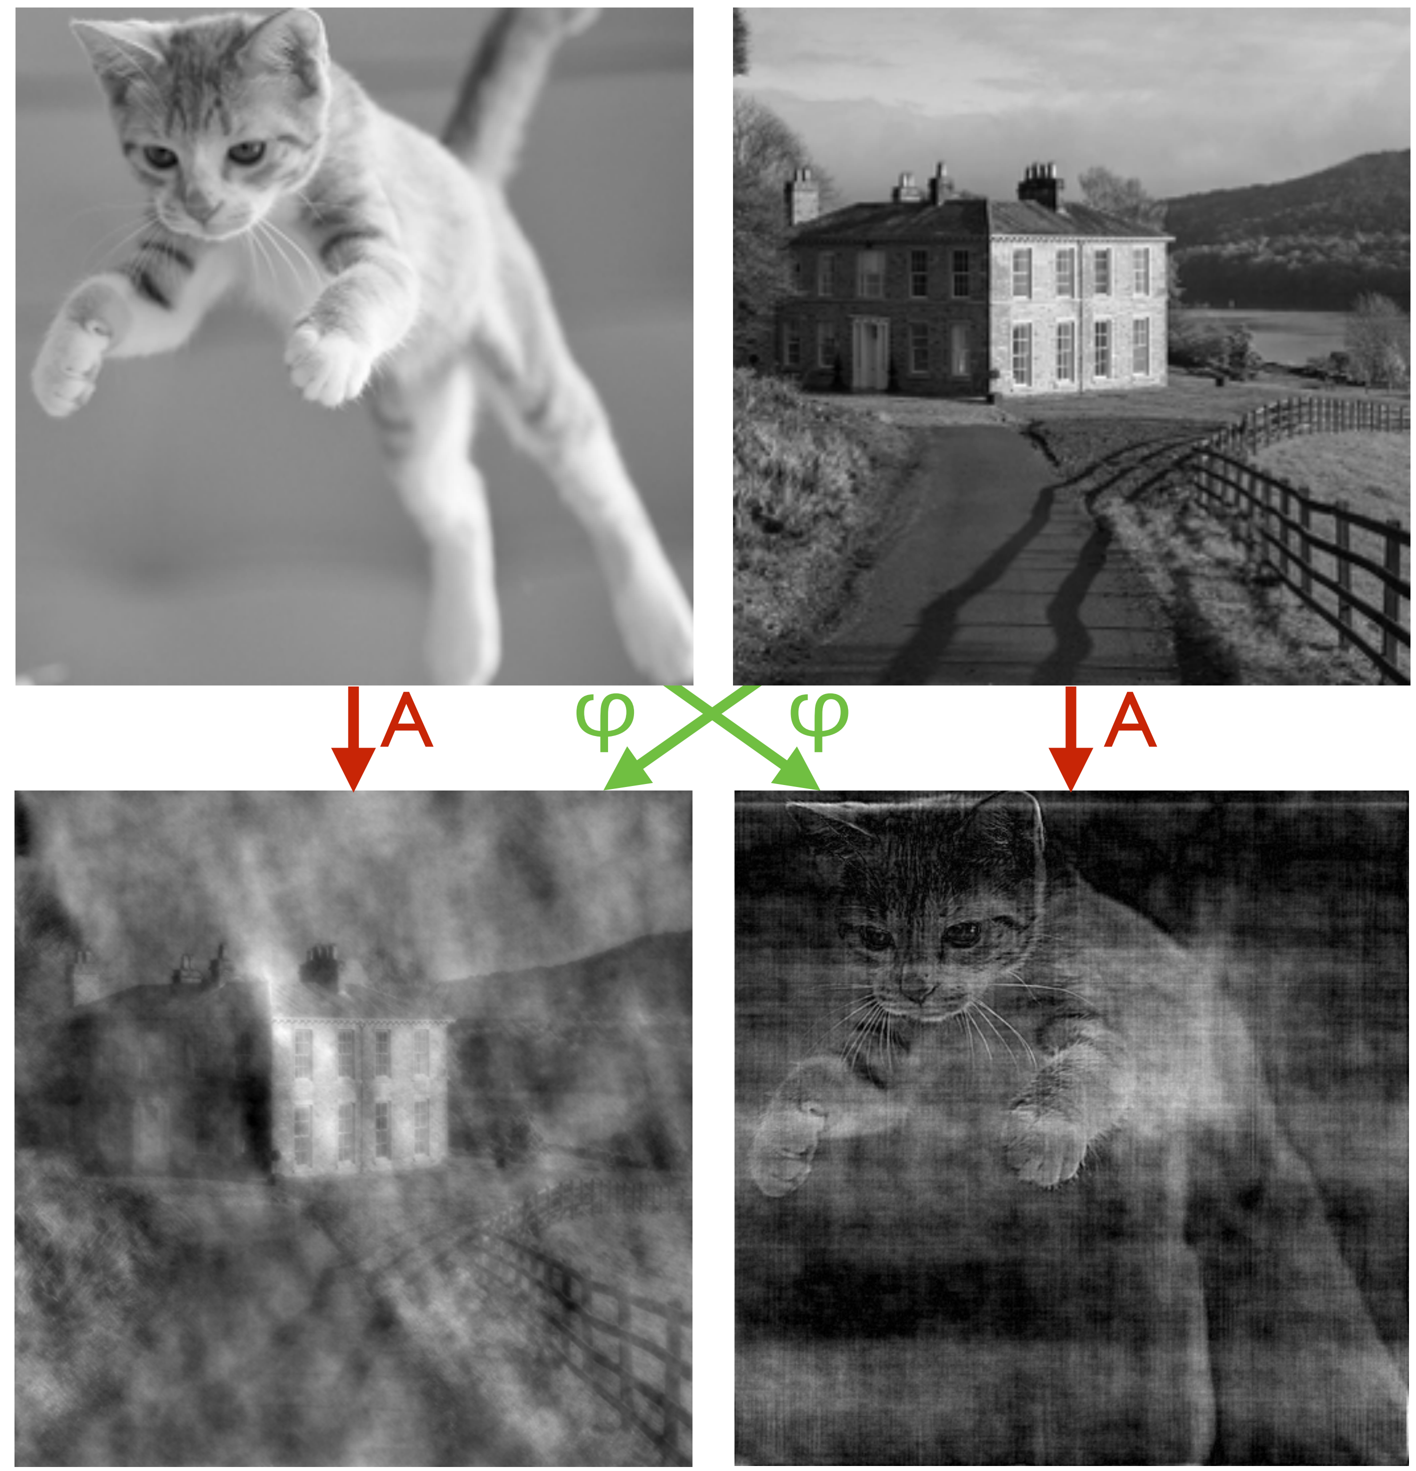

Vision via Fourier

- Fourier analysis gives us a way to think about scale
  - Images contain information at different spatial frequencies
  - Which of these components is visible to an observer?
- A common way to examine spatial vision is to test acuity, which is akin to the highest visible spatial frequency
- With Fourier analysis we can take a broader view of the image content that is visible to a given observer

A note on units

- First: how do we characterise spatial variations?
- Cycles/image is OK for theoretical Fourier analyses
- But for visual perception, image size on the retina is affected by both size and distance
- Need to measure *retinal size*
  - Calculated as degrees of visual angle, where **tan(α) = Height/Distance**
  - For SF gives cycles/degree

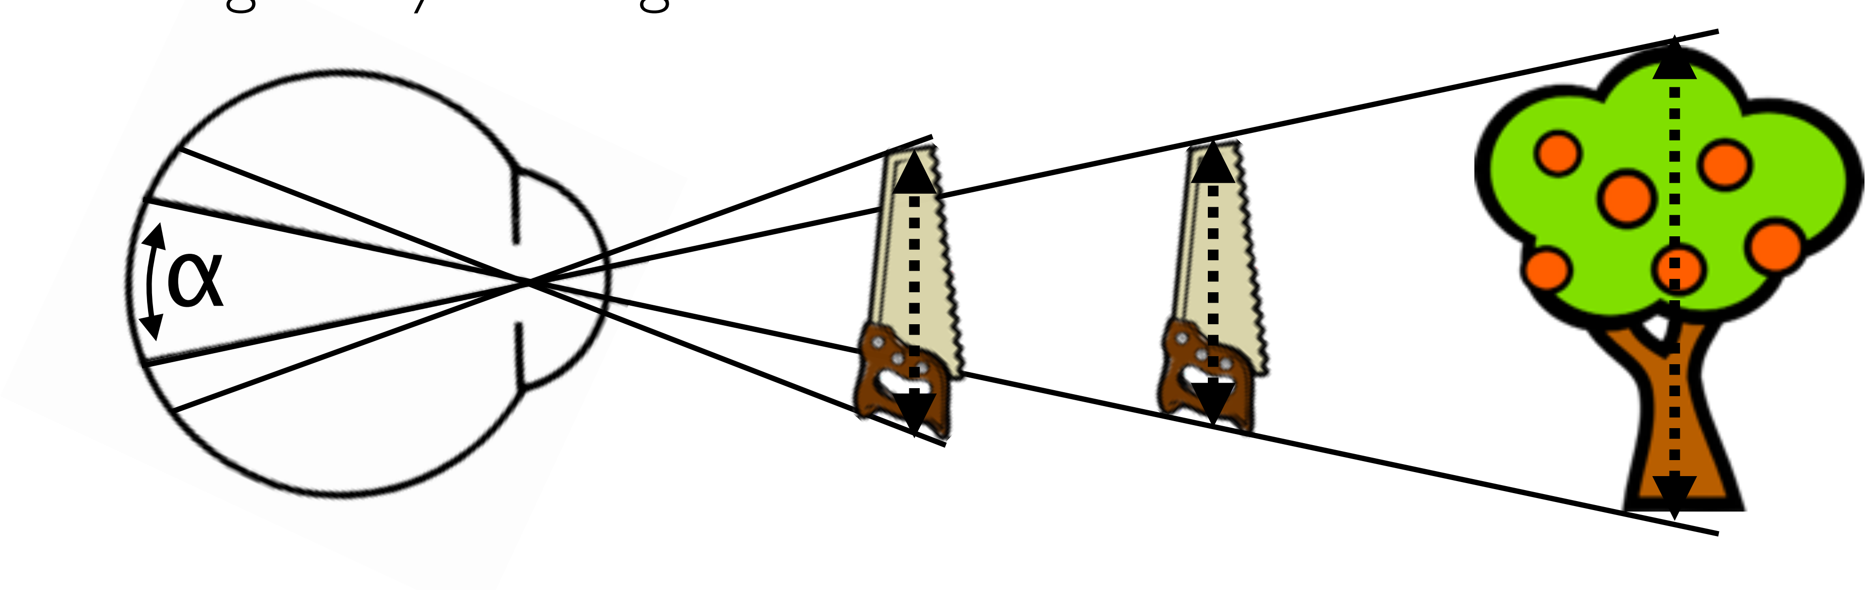

Fourier analysis & visibility

What we see is determined by the visibility of information at different spatial frequencies - different scales of analysis

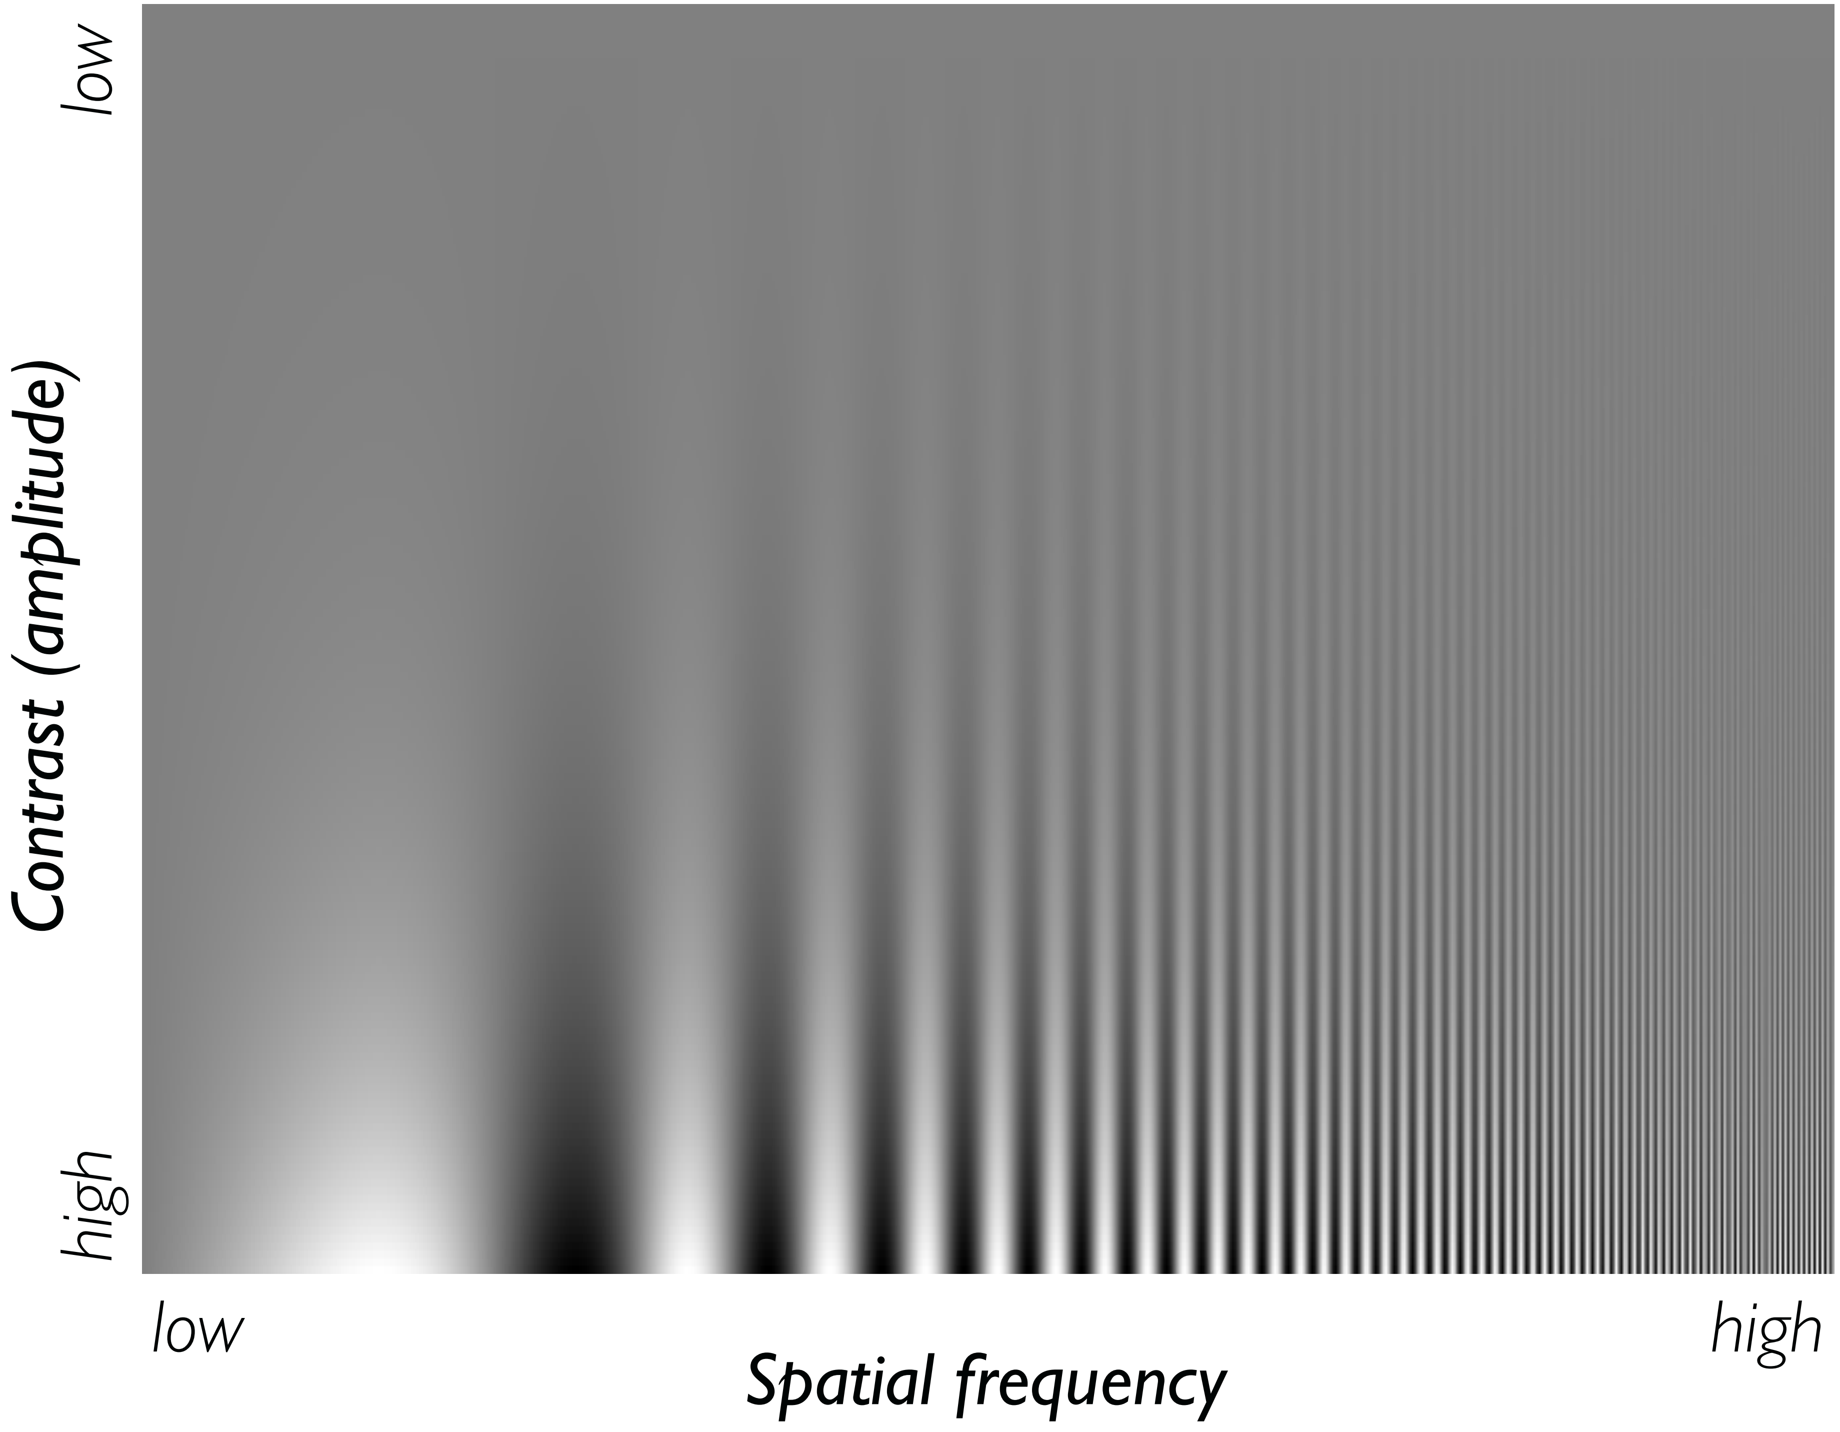

Contrast sensitivity functions

- Campbell & Robson (1968):
  - Measured contrast sensitivity (inverse of threshold!) at a range of spatial frequencies
  - Contrast sensitivity function (CSF) is band-pass & peaks around 4 c/deg
  - Defines our 'window of visibility'

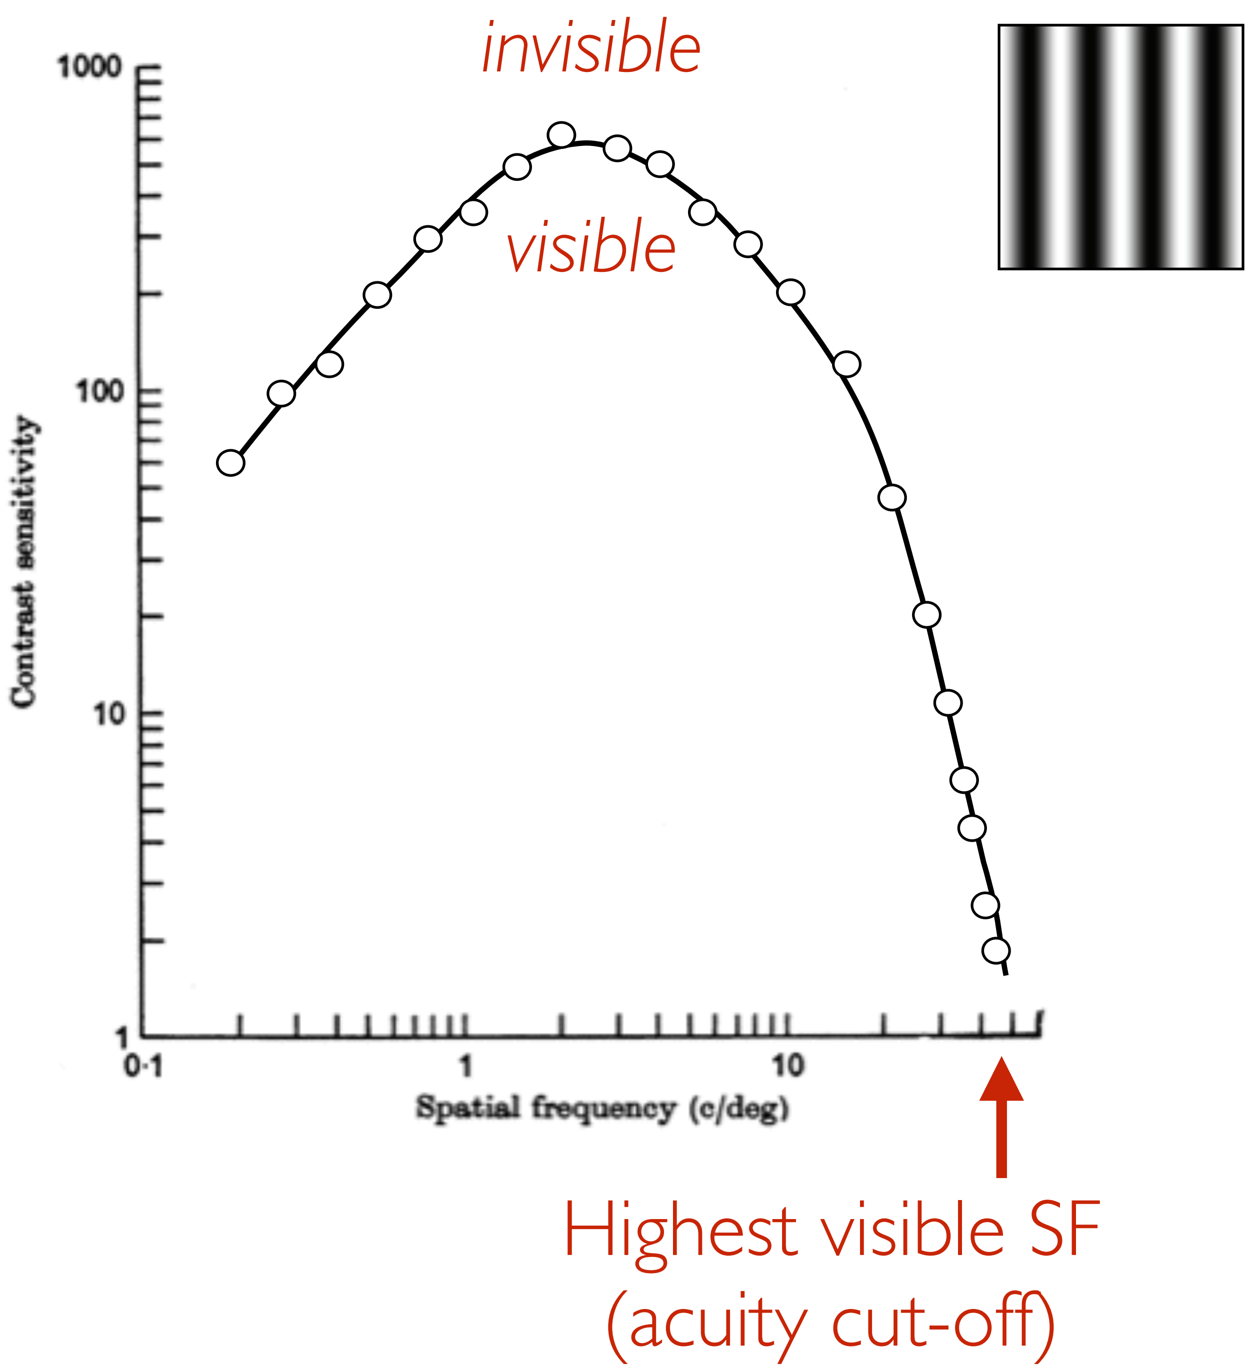

CSF: eccentricity

- CSF can capture our visual experience across the visual field
- With increasing eccentricity (moving into the periphery), the CSF shifts to lower SFs (Rovamo et al, 1978)
  - i.e. we lose sensitivity to high spatial frequencies

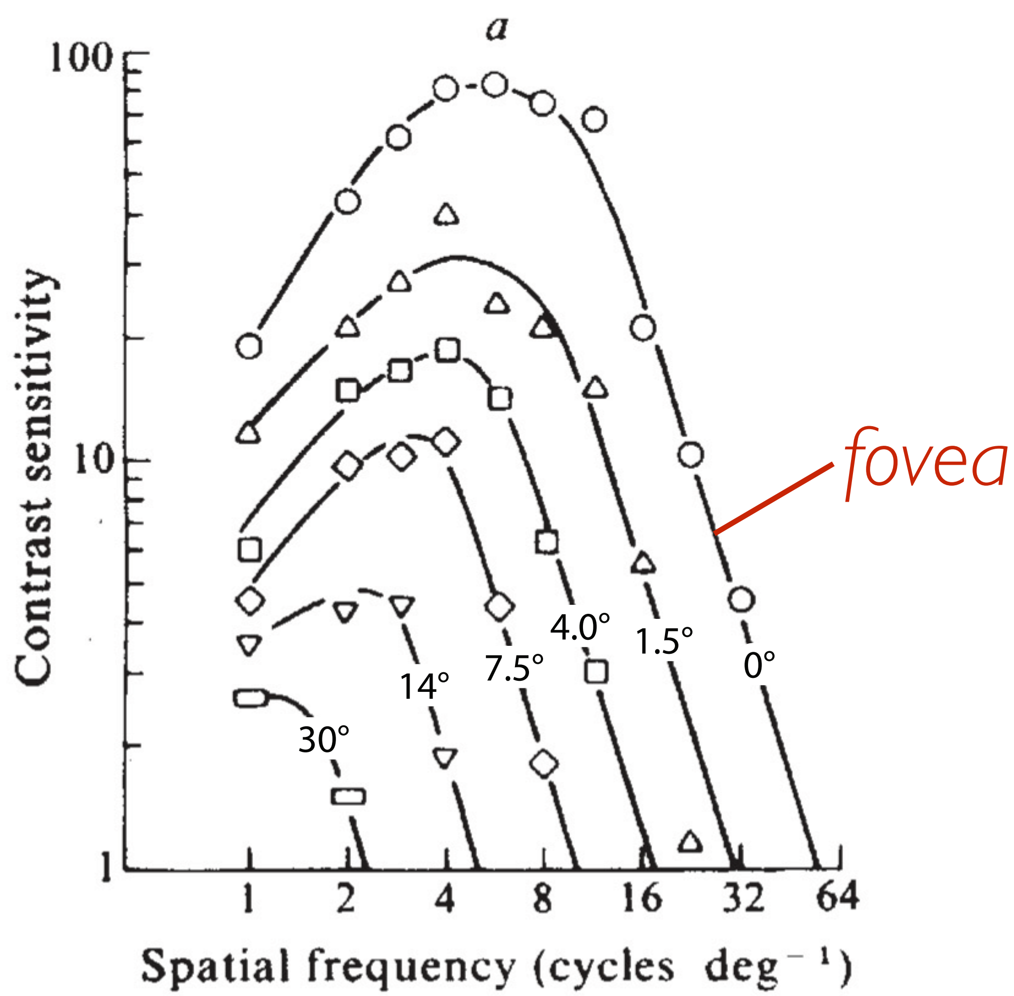

Visibility in other species

- We can use this to think about the vision of other species
- What does a cat see? (Bisti & Maffei, 1974)
  - Present gratings on a monitor
  - Present/absent task
  - Cats trained to press a lever when grating is seen, with a milk reward when correct
  - Grating contrast and SF varied along Method of Constant Stimuli

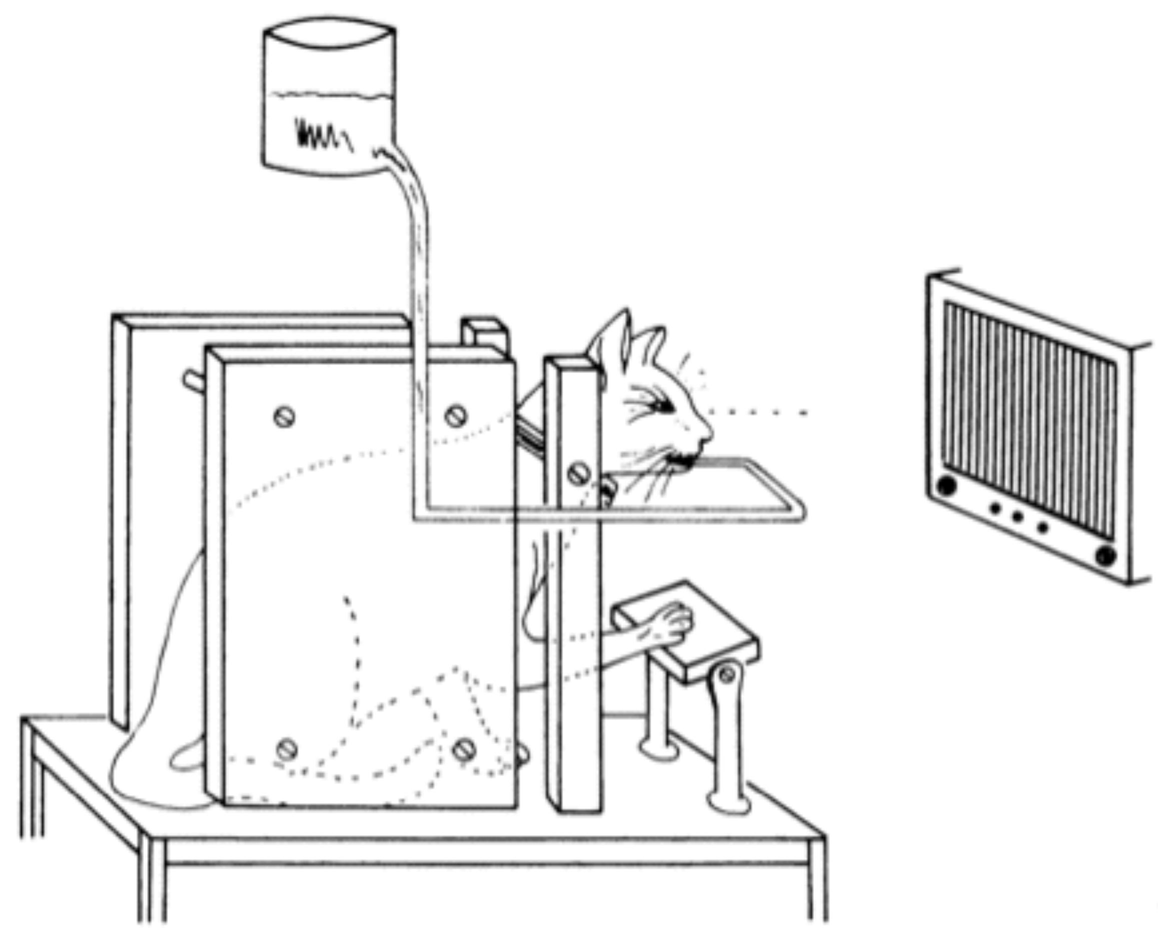

Fig. 1. Schematic diagram of the experimental set up (for explanations see text).

CSF of cats

- Feline sensitivity:
  - Has a lower cutoff point (i.e. worse acuity limit)
  - Peaks at a lower SF (0.3-0.4 cyc/deg)
  - But sensitivity much higher than ours in the low SF range
  - Do cats see strange shadows on the wall?

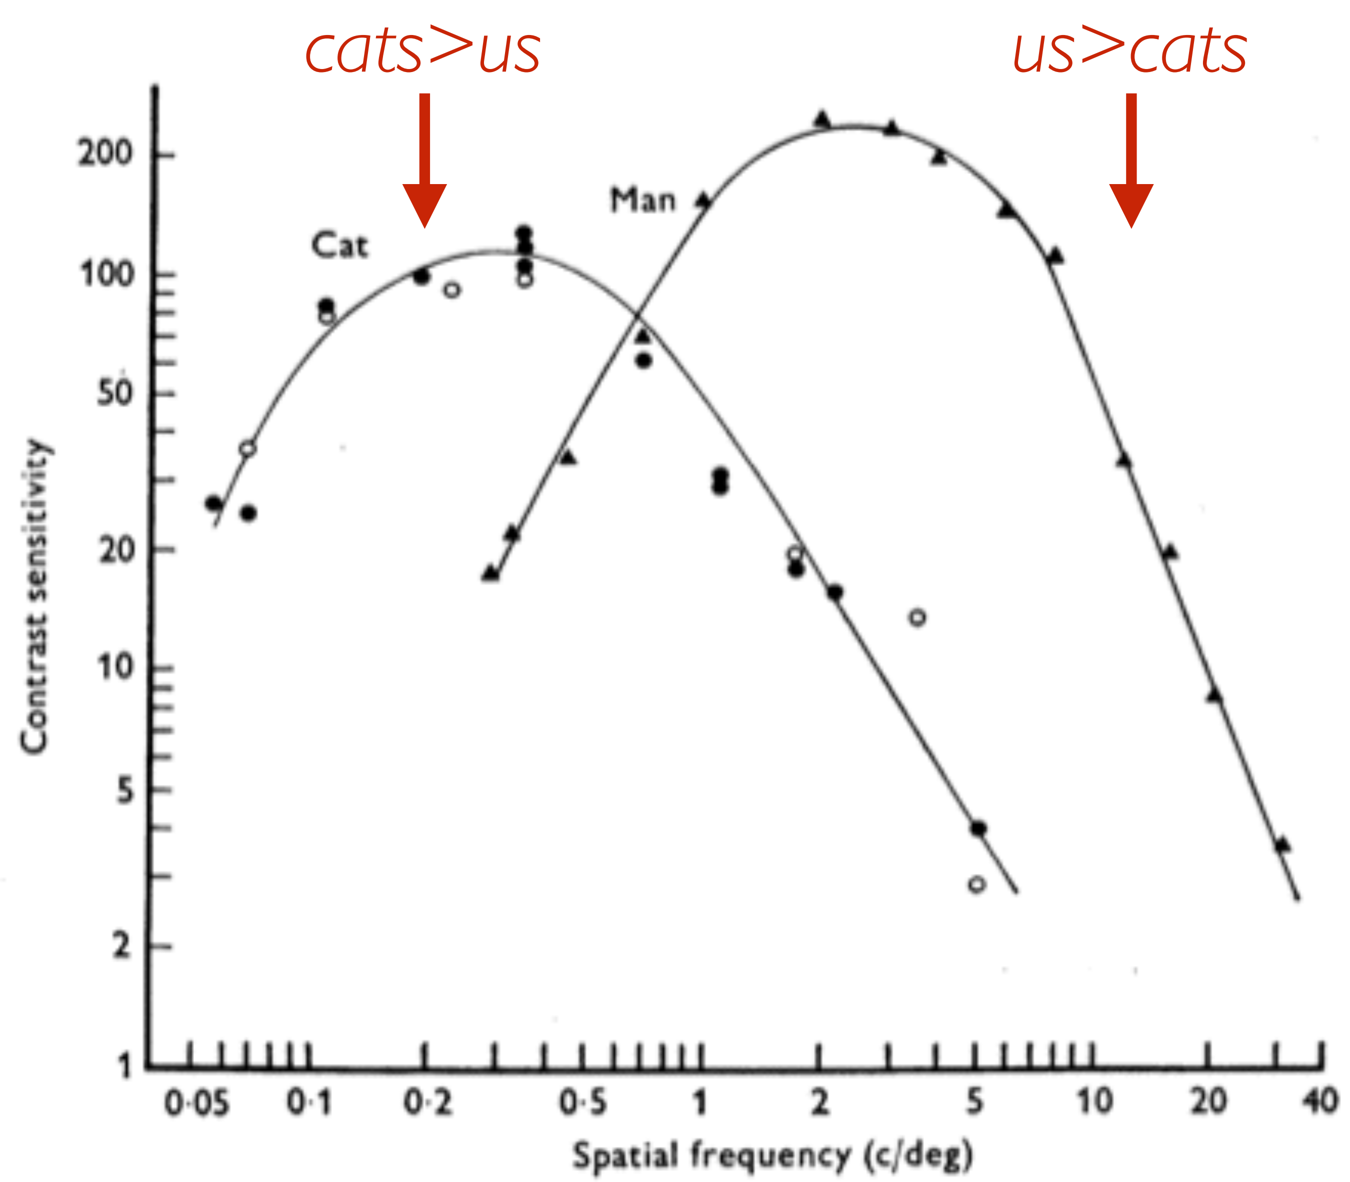

Fig. 4. Comparison of the contrast sensitivity curves of the cat and of a human subject (L.M.) in similar experimental conditions.

What produces the CSF?

- Why do we show this pattern of sensitivity?
- Campbell & Robson (1968) hypothesised that the visual system is composed of spatial frequency channels - each sensitive to a restricted range of SFs
- Blakemore & Campbell (1969) tested this using adaptation

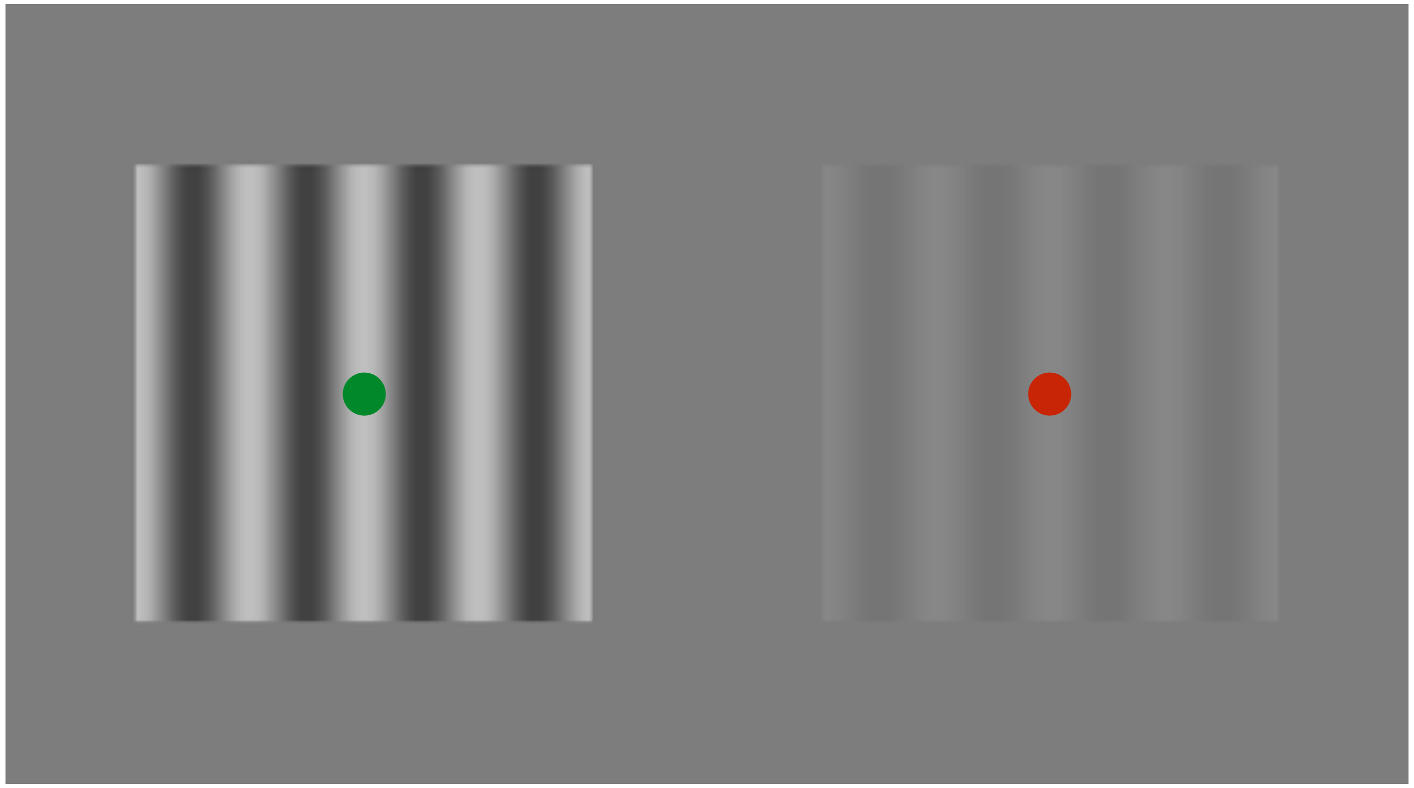

Adaptation: predictions

- Adaptation reduces sensitivity to contrast
- But does it affect all SFs or just those of the adaptor?

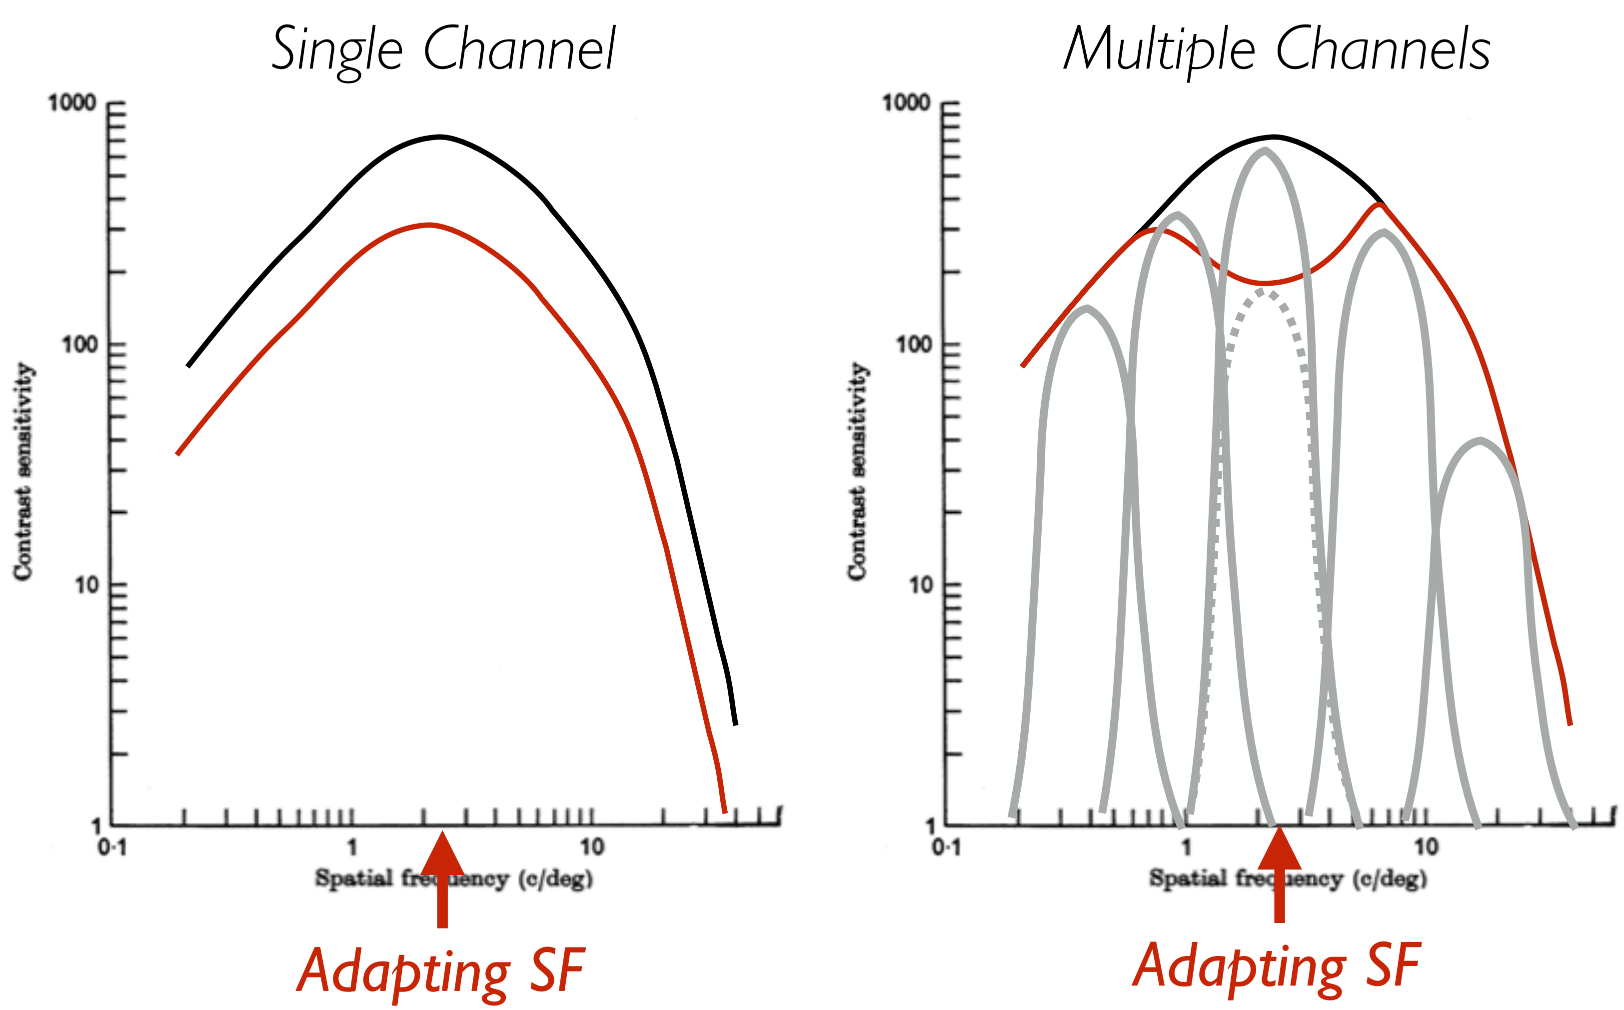

Multiple SF channels

- Adaptation to a sine grating with 7.1 cycles per degree
  - Sensitivity is strongly reduced at the adapted spatial frequency
  - Decreased effect for adjacent SFs
- Consistent with multiple channels for spatial frequency
  - Evidence that we separate the visual scene into its Fourier components (at least for SF)

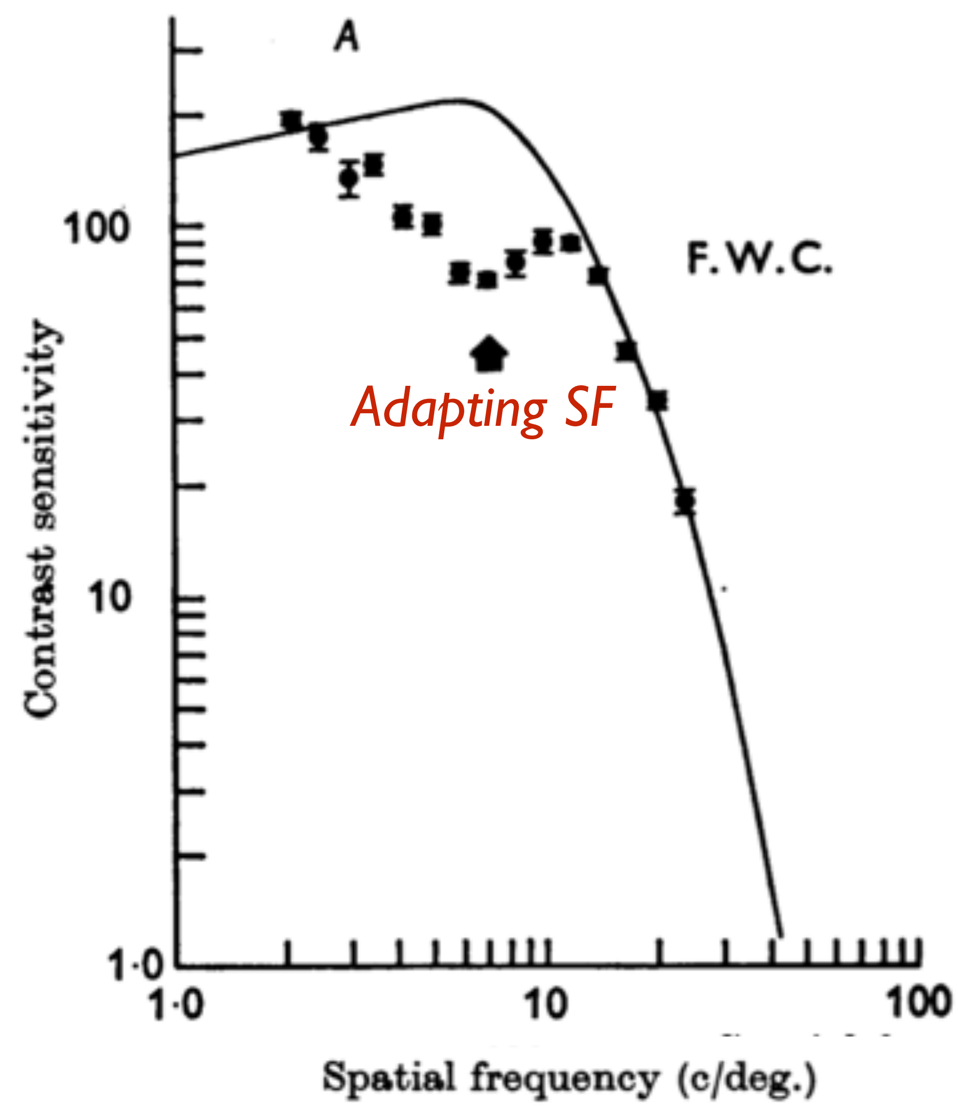

Independent access?

- Does independent channels mean independent access to each frequency band?
- No: Pairing high SF image with low SF image shows high SFs are difficult to ignore
  - Low SFs can appear by squinting/defocusing/shrinking the image
- High spatial frequency channels dominate the lower SFs in object recognition

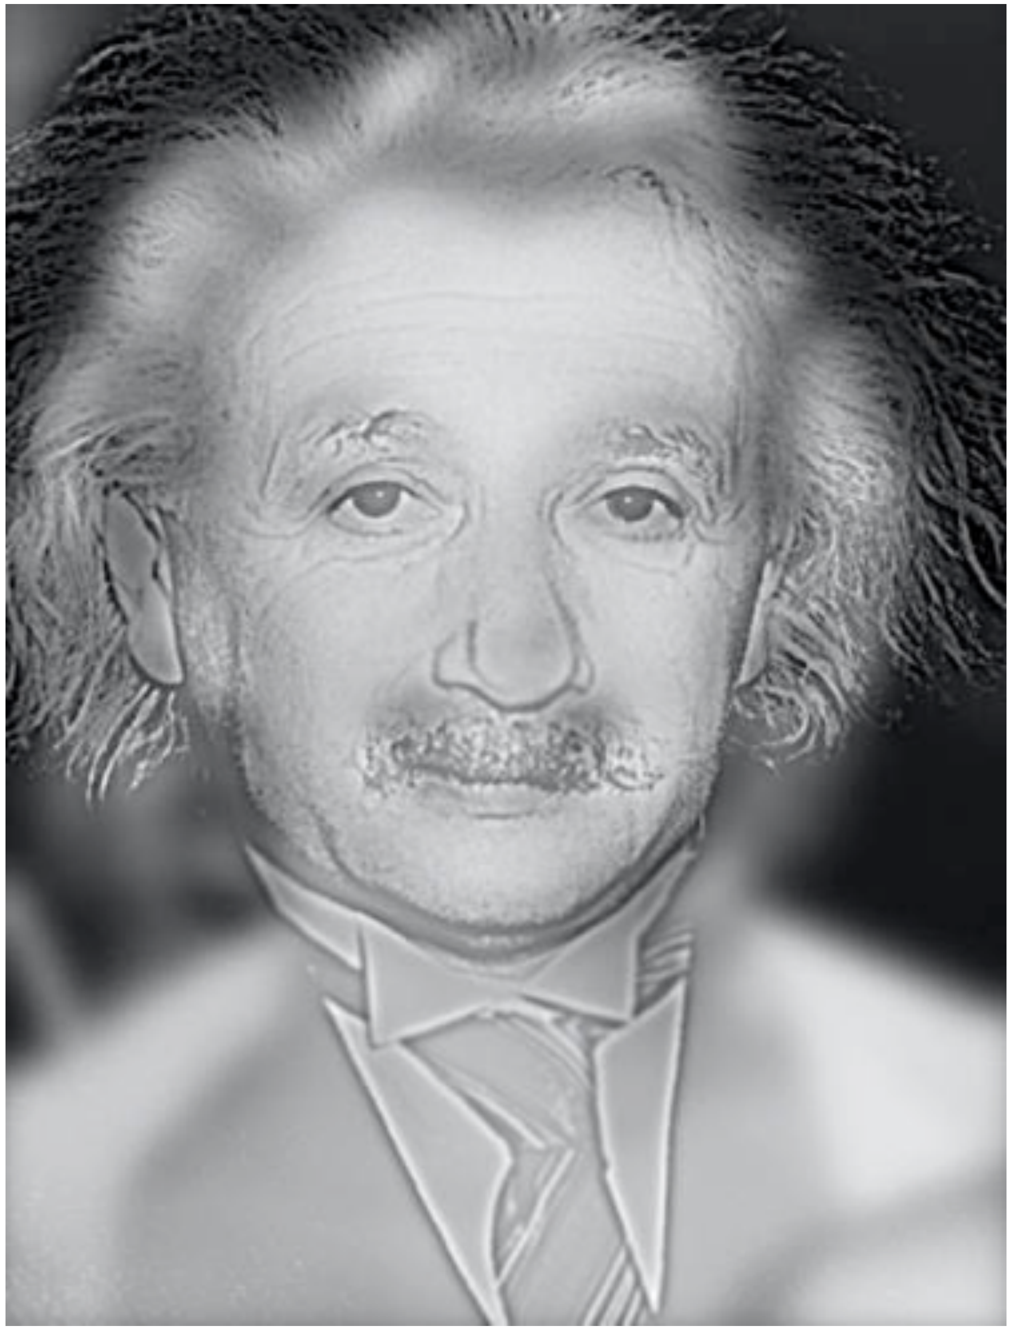

Schyns & Oliva (1999)

Neuro-Fourier?

- Fourier analysis:
  - Gives us a useful metric to consider the aspects of a scene that will be visible to a person (or animal)
  - Describes meaningful aspects of our visual experience
- But do neurons also analyse visual input in this way?
  - Receptive field properties in the early stages of visual recognition:
    - Lateral Geniculate Nucleus (LGN) of the thalamus
    - Cortical area V1 (primary visual cortex)

Visual pathways

- Consider the earliest stages of vision:

Retina
↓
Thalamus (LGN)
↓
Primary visual cortex (V1)

- Can Fourier analysis characterise the selectivity of neurons in these areas?

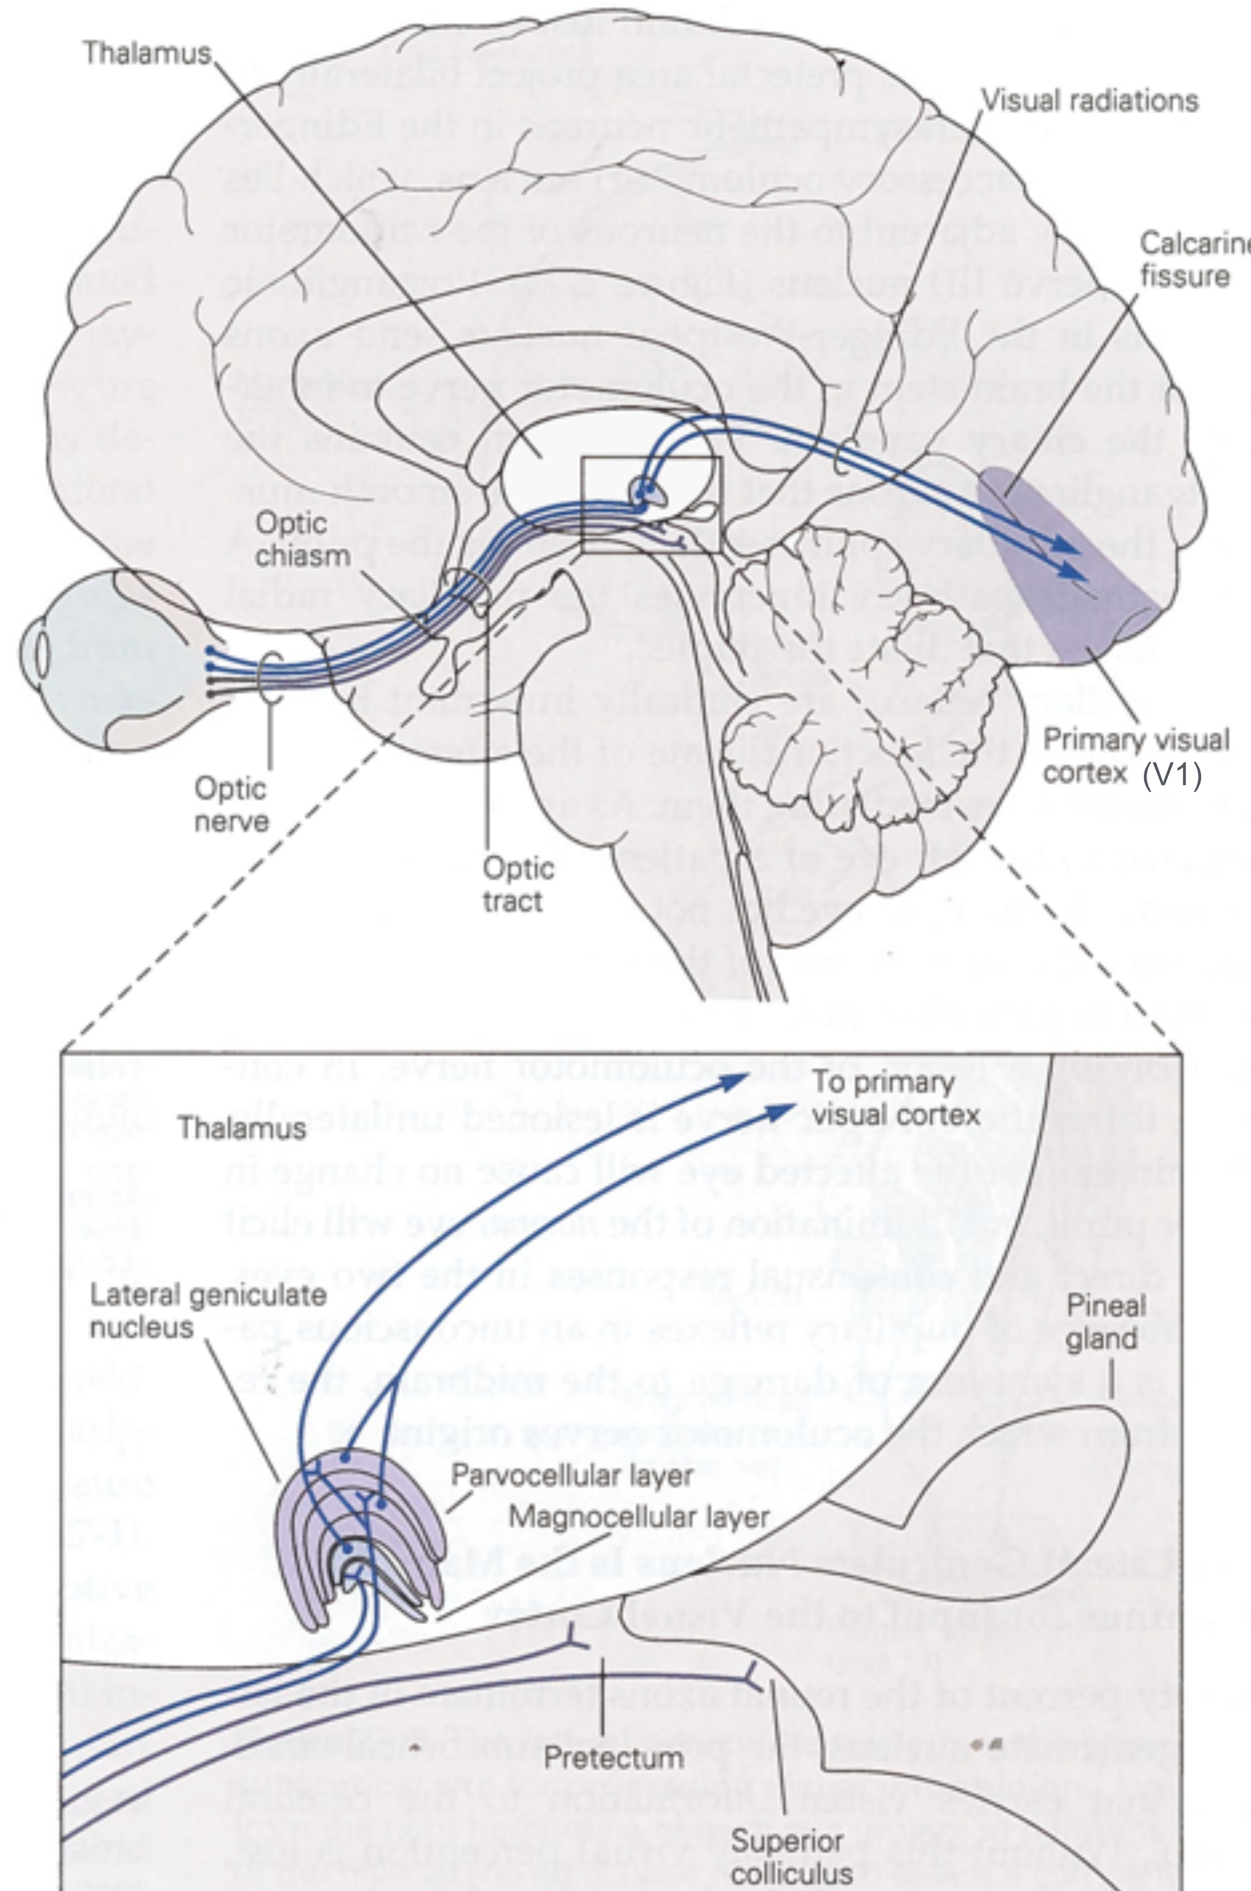

Centre-surround RFs

- Both retinal ganglion cells and neurons in the Lateral Geniculate Nucleus show centre-surround receptive fields
  - Either on-centre or off-centre
- What is the purpose of this RF organisation?
  - We can picture this through "spatial filtering"
    - Convolve an image with a filter matching these RFs
    - Simulates how a population of neurons with centre-surround RFs would "see" the world

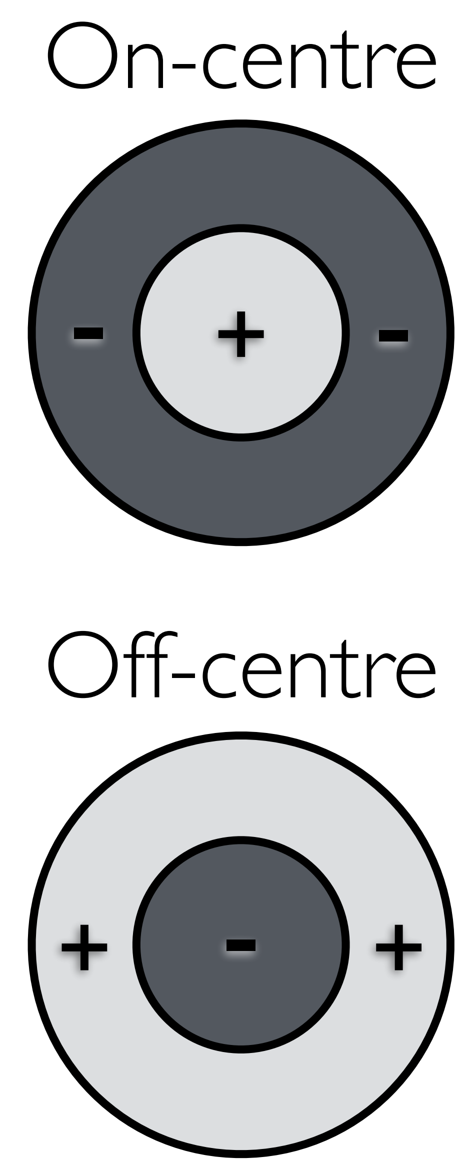

Centre-surround filters

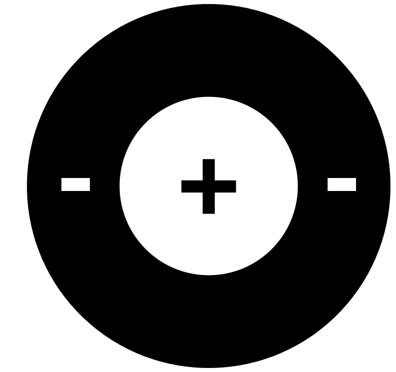

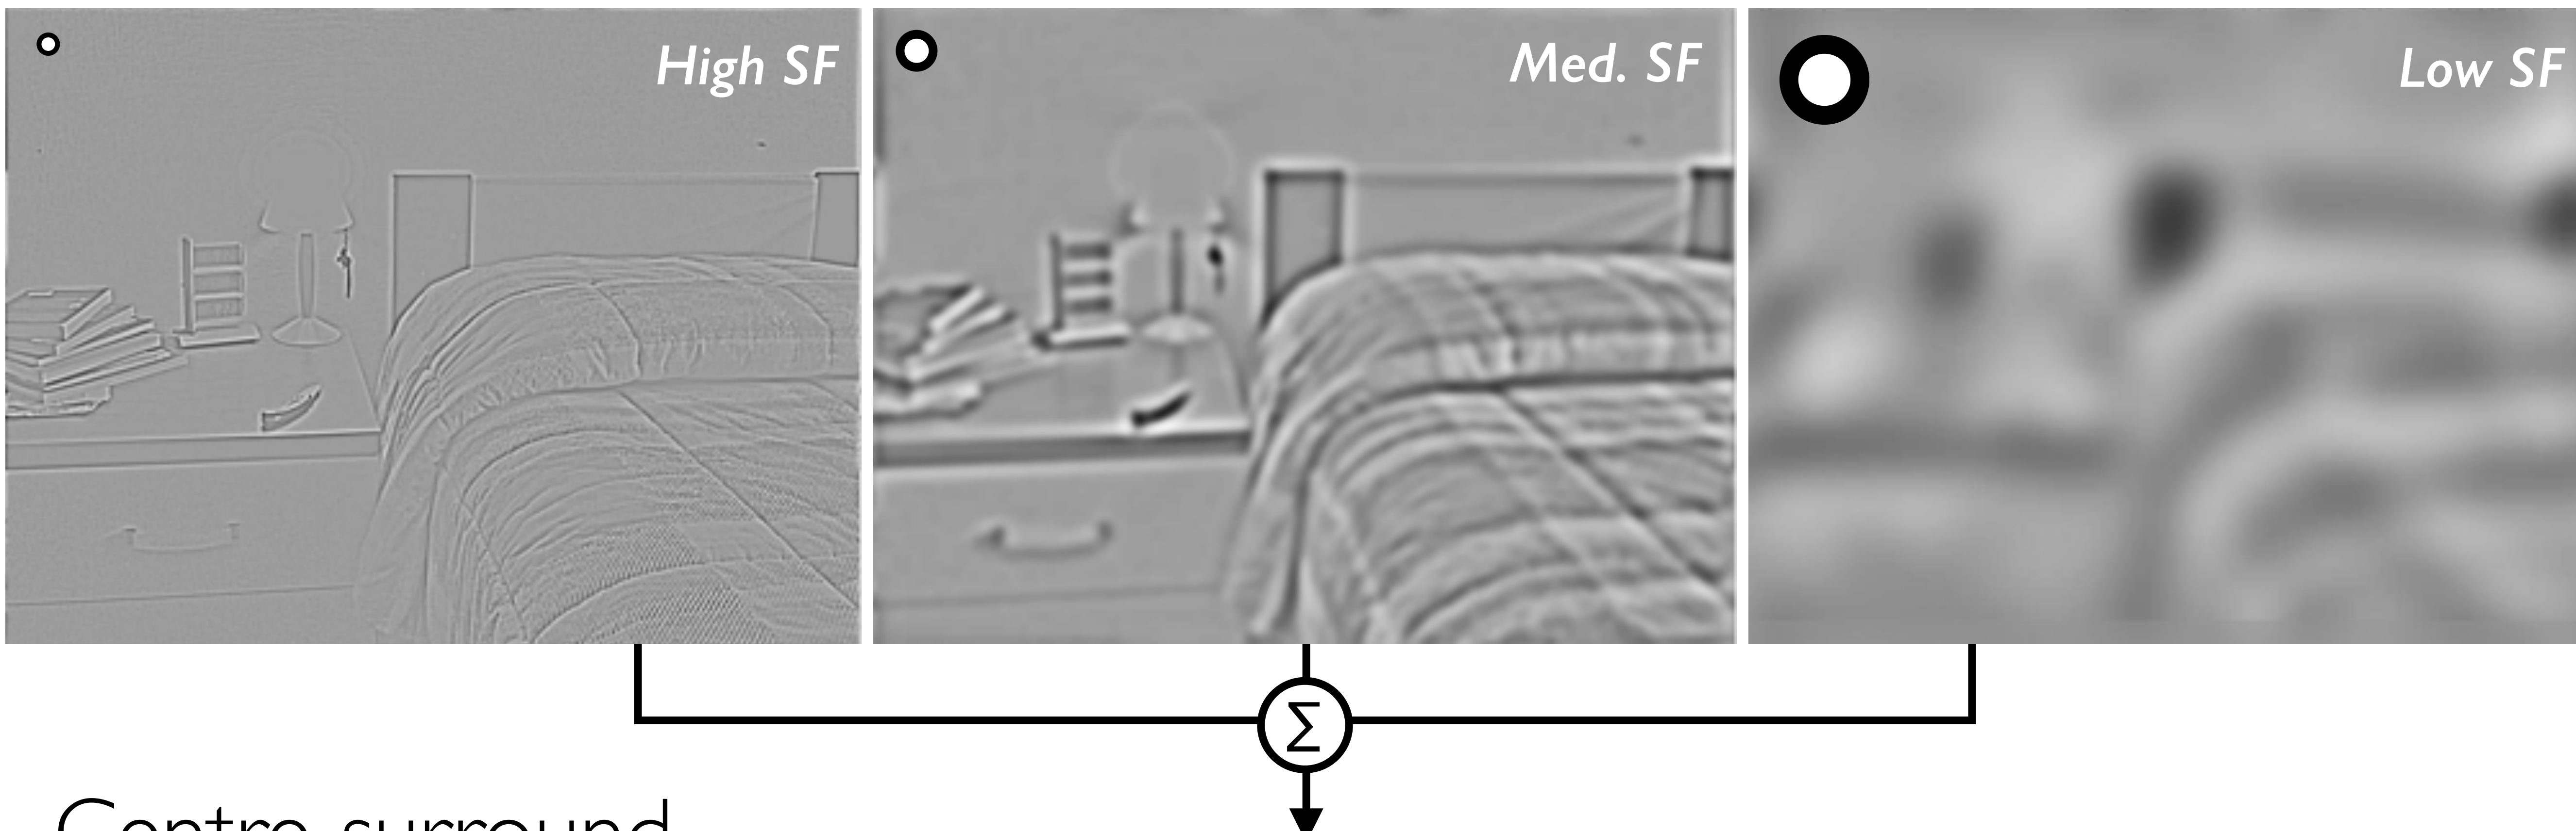

Centre-surround filtering highlights the edges and *changes* in the scene, see e.g. Shapley & Lennie (1985)

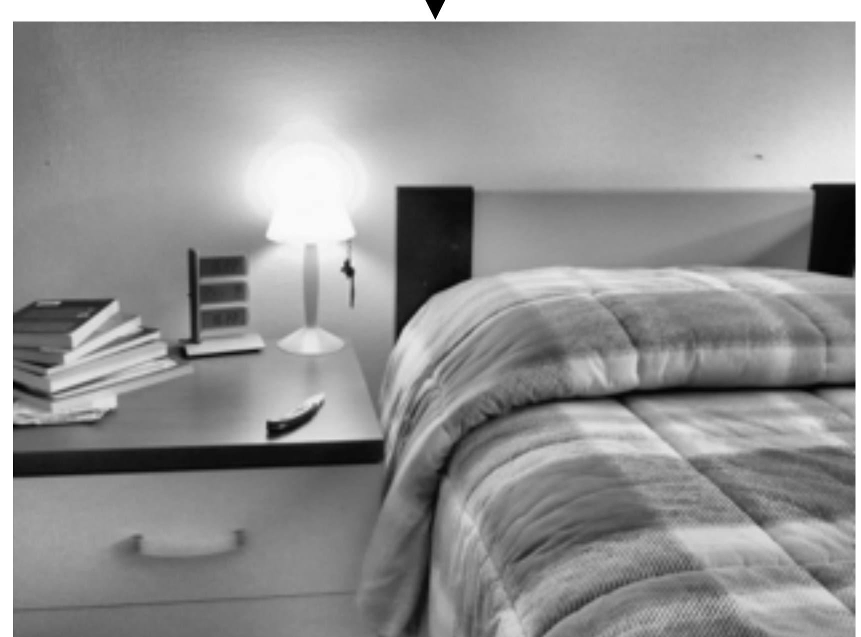

Varied sizes of RF centre/surround allows separation into spatial frequency bands

Functional subdivision

- Lateral Geniculate Nucleus has clearly defined layers
- Magnocellular layers:
  - Large cells
  - Rapid response
  - Colour blind
- Parvocellular layers:
  - Smaller cells
  - Slower response properties
  - Colour selective
- Koniocellular layers (far less understood)

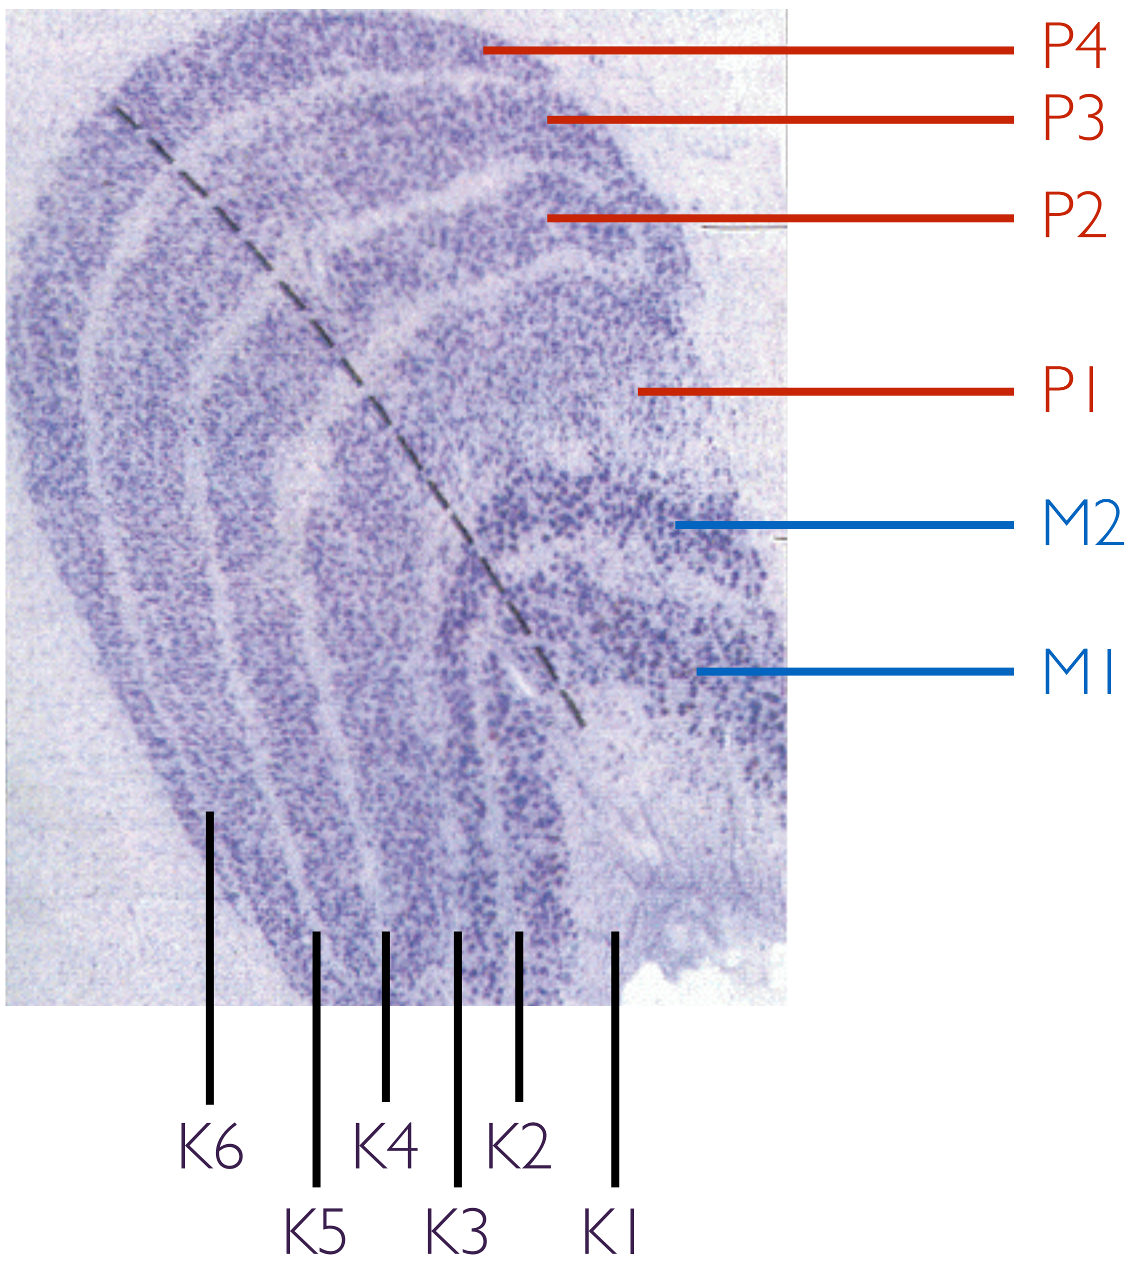

M & P selectivity

- Distinct patterns of spatial frequency selectivity:
  - P cells: High contrast sensitivity that matches overall sensitivity of the monkey
  - M cells: Much lower contrast sensitivity overall
- Likewise for temporal variations:
  - P cells: sensitive to slow flicker
  - M cells: sensitive to fast flicker

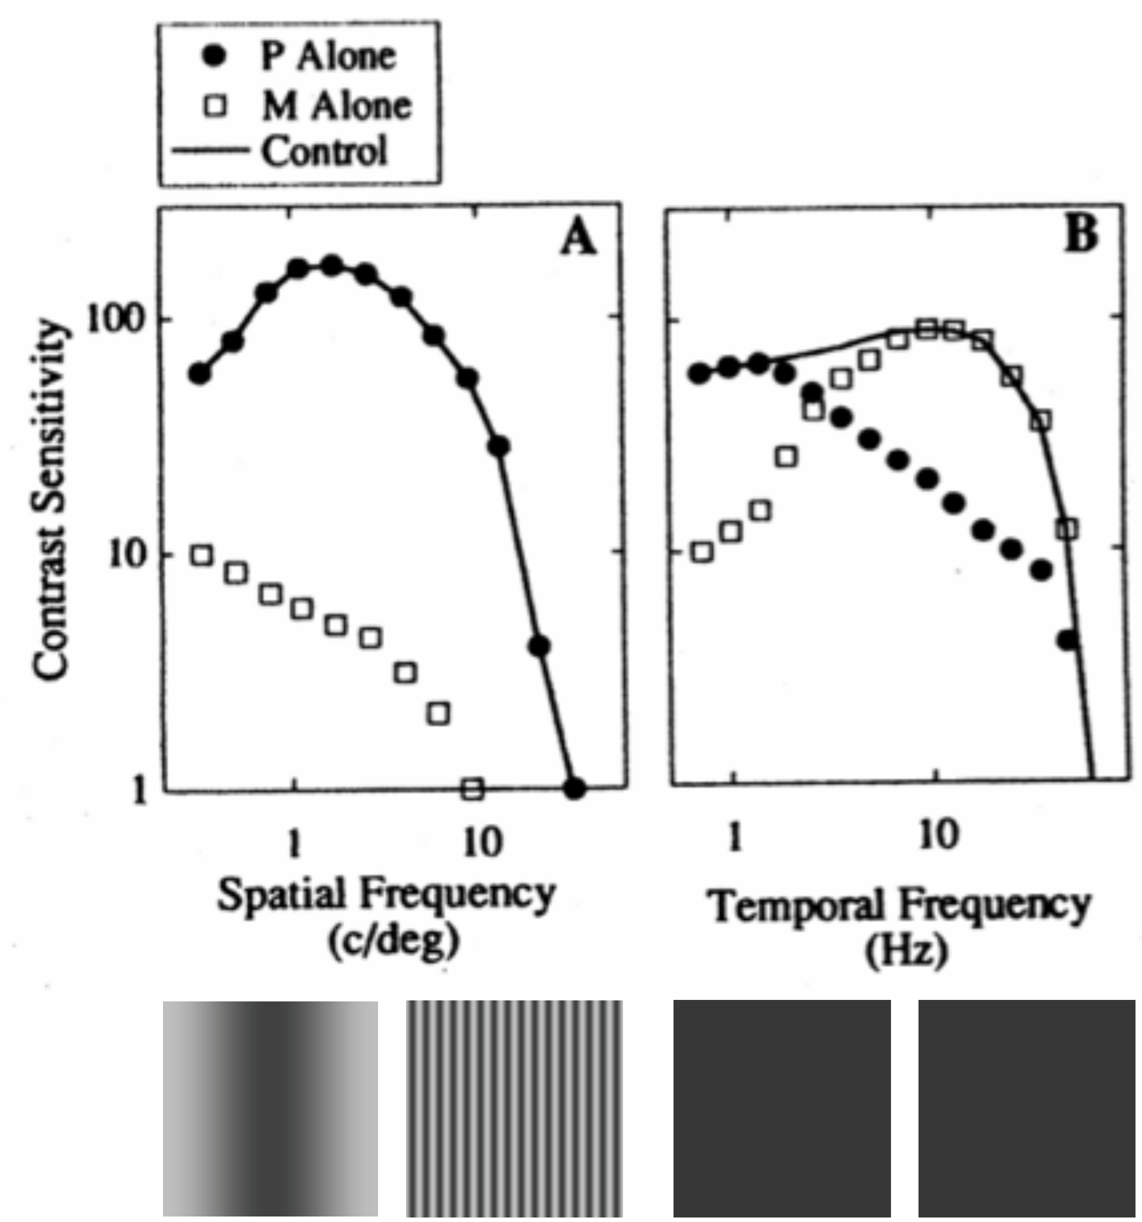

from Merigan & Maunsell (1993)

Building towards form

- Hierarchical organisation of the visual pathways:
  - RF sizes increase along the visual pathway
  - Cells show selectivity to particular stimulus properties (e.g. dark vs. light)
  - Selectivity becomes increasingly complex along the visual pathway
  - Hubel & Wiesel (1968): V1 cells are a linear combination of LGN inputs, allowing sensitivity to orientation

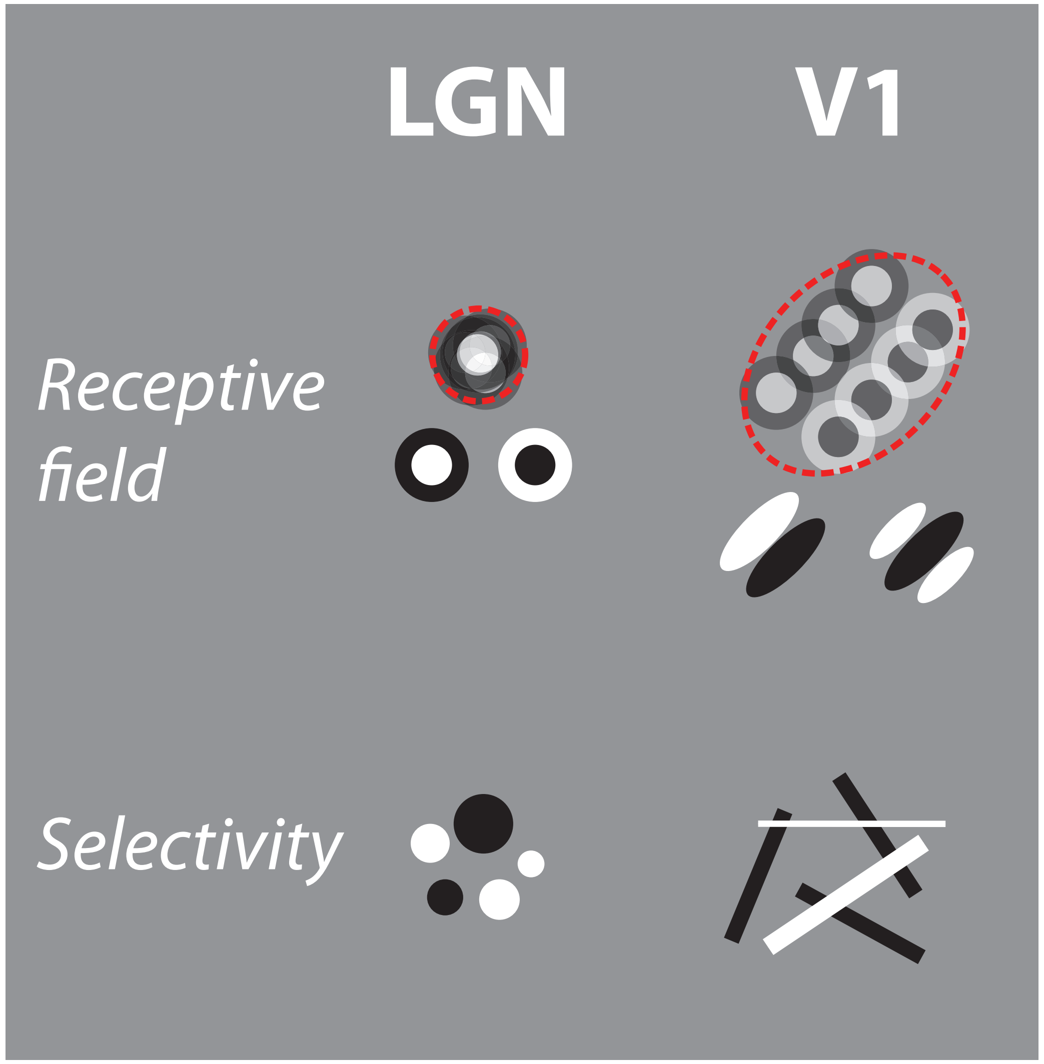

V1 filters

- Present white noise in the receptive fields of V1 cells and correlate cell responses with pixel values
- Produces many V1 receptive fields that look like Gabors
  - The multiplication of a sine wave with a 2D Gaussian profile

V1 receptive field

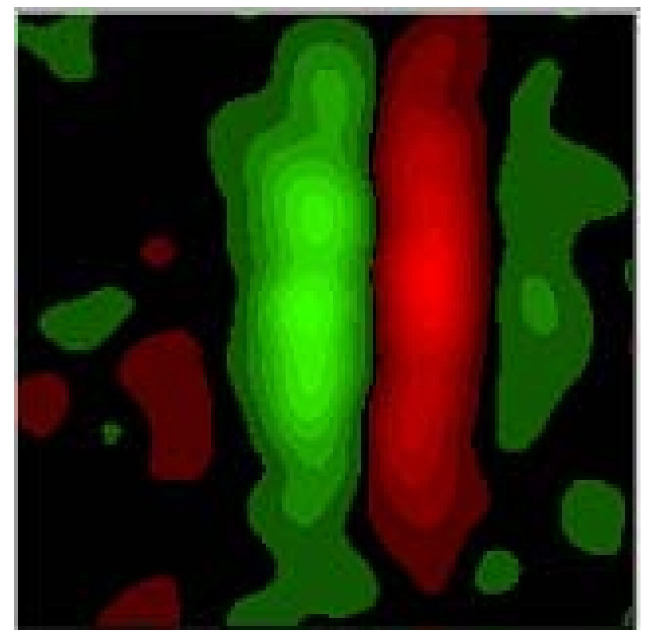

Gabor model

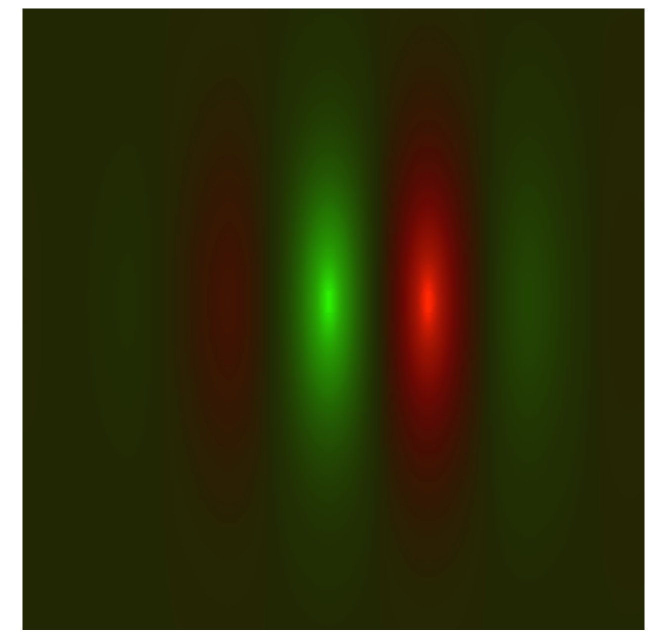

De Angelis, Ohzawa & Freeman (1995)

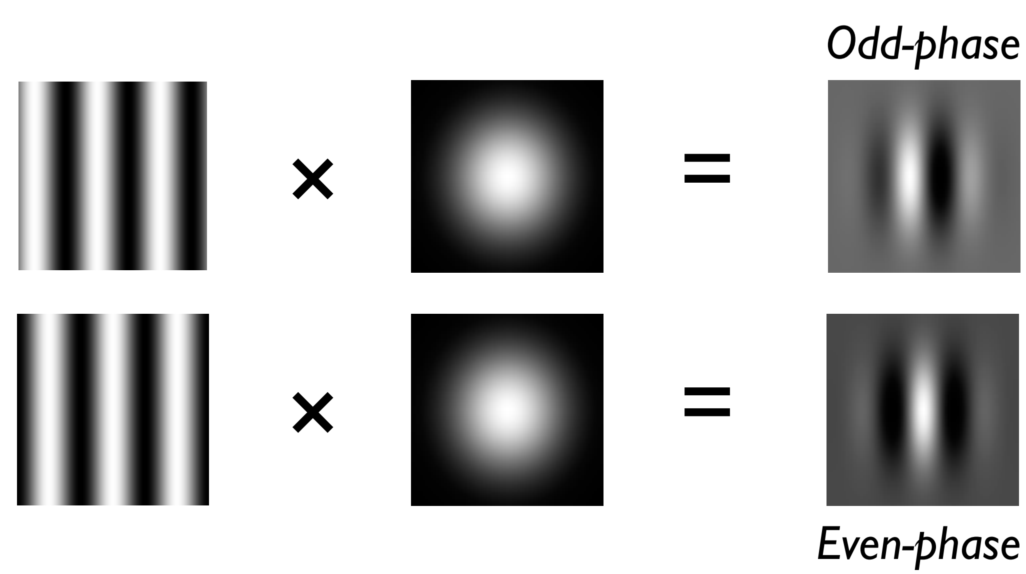

V1 cells as Gabor filters

- Shape of V1 receptive fields could allow a local Fourier analysis (vs. a global analysis over the whole image)
- Gabors are an optimal trade-off for uncertainty in space/SF

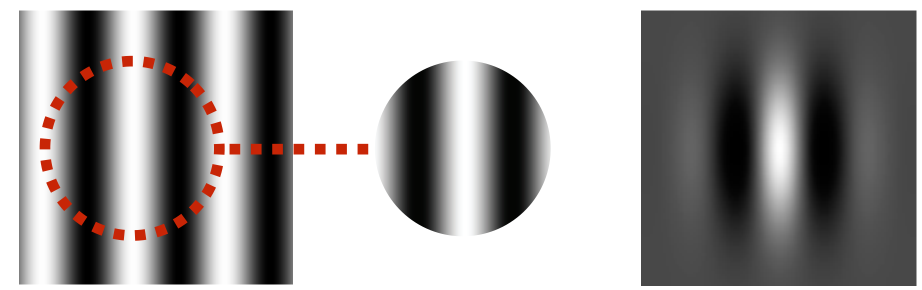

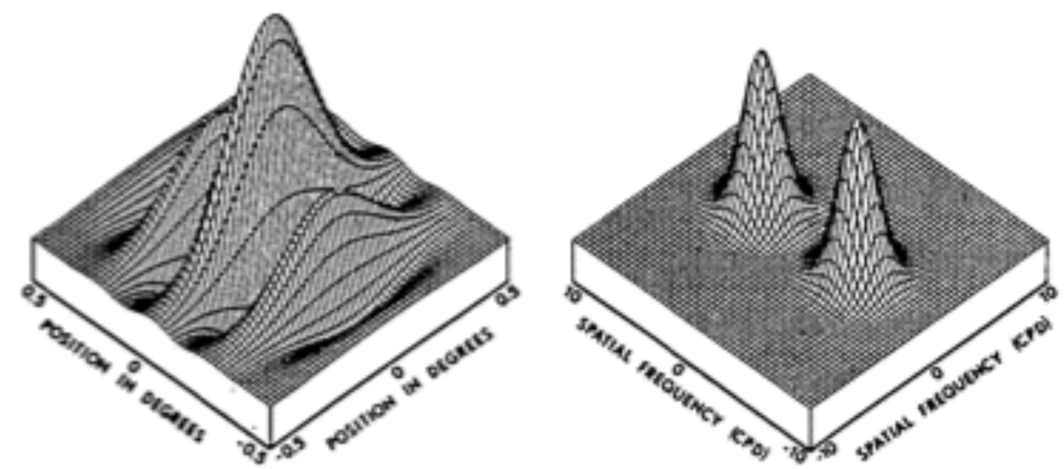

(Daugman, 1985)

- The response of a bank of Gabor filters to a natural scene
  - Different orientations at a range of spatial frequencies

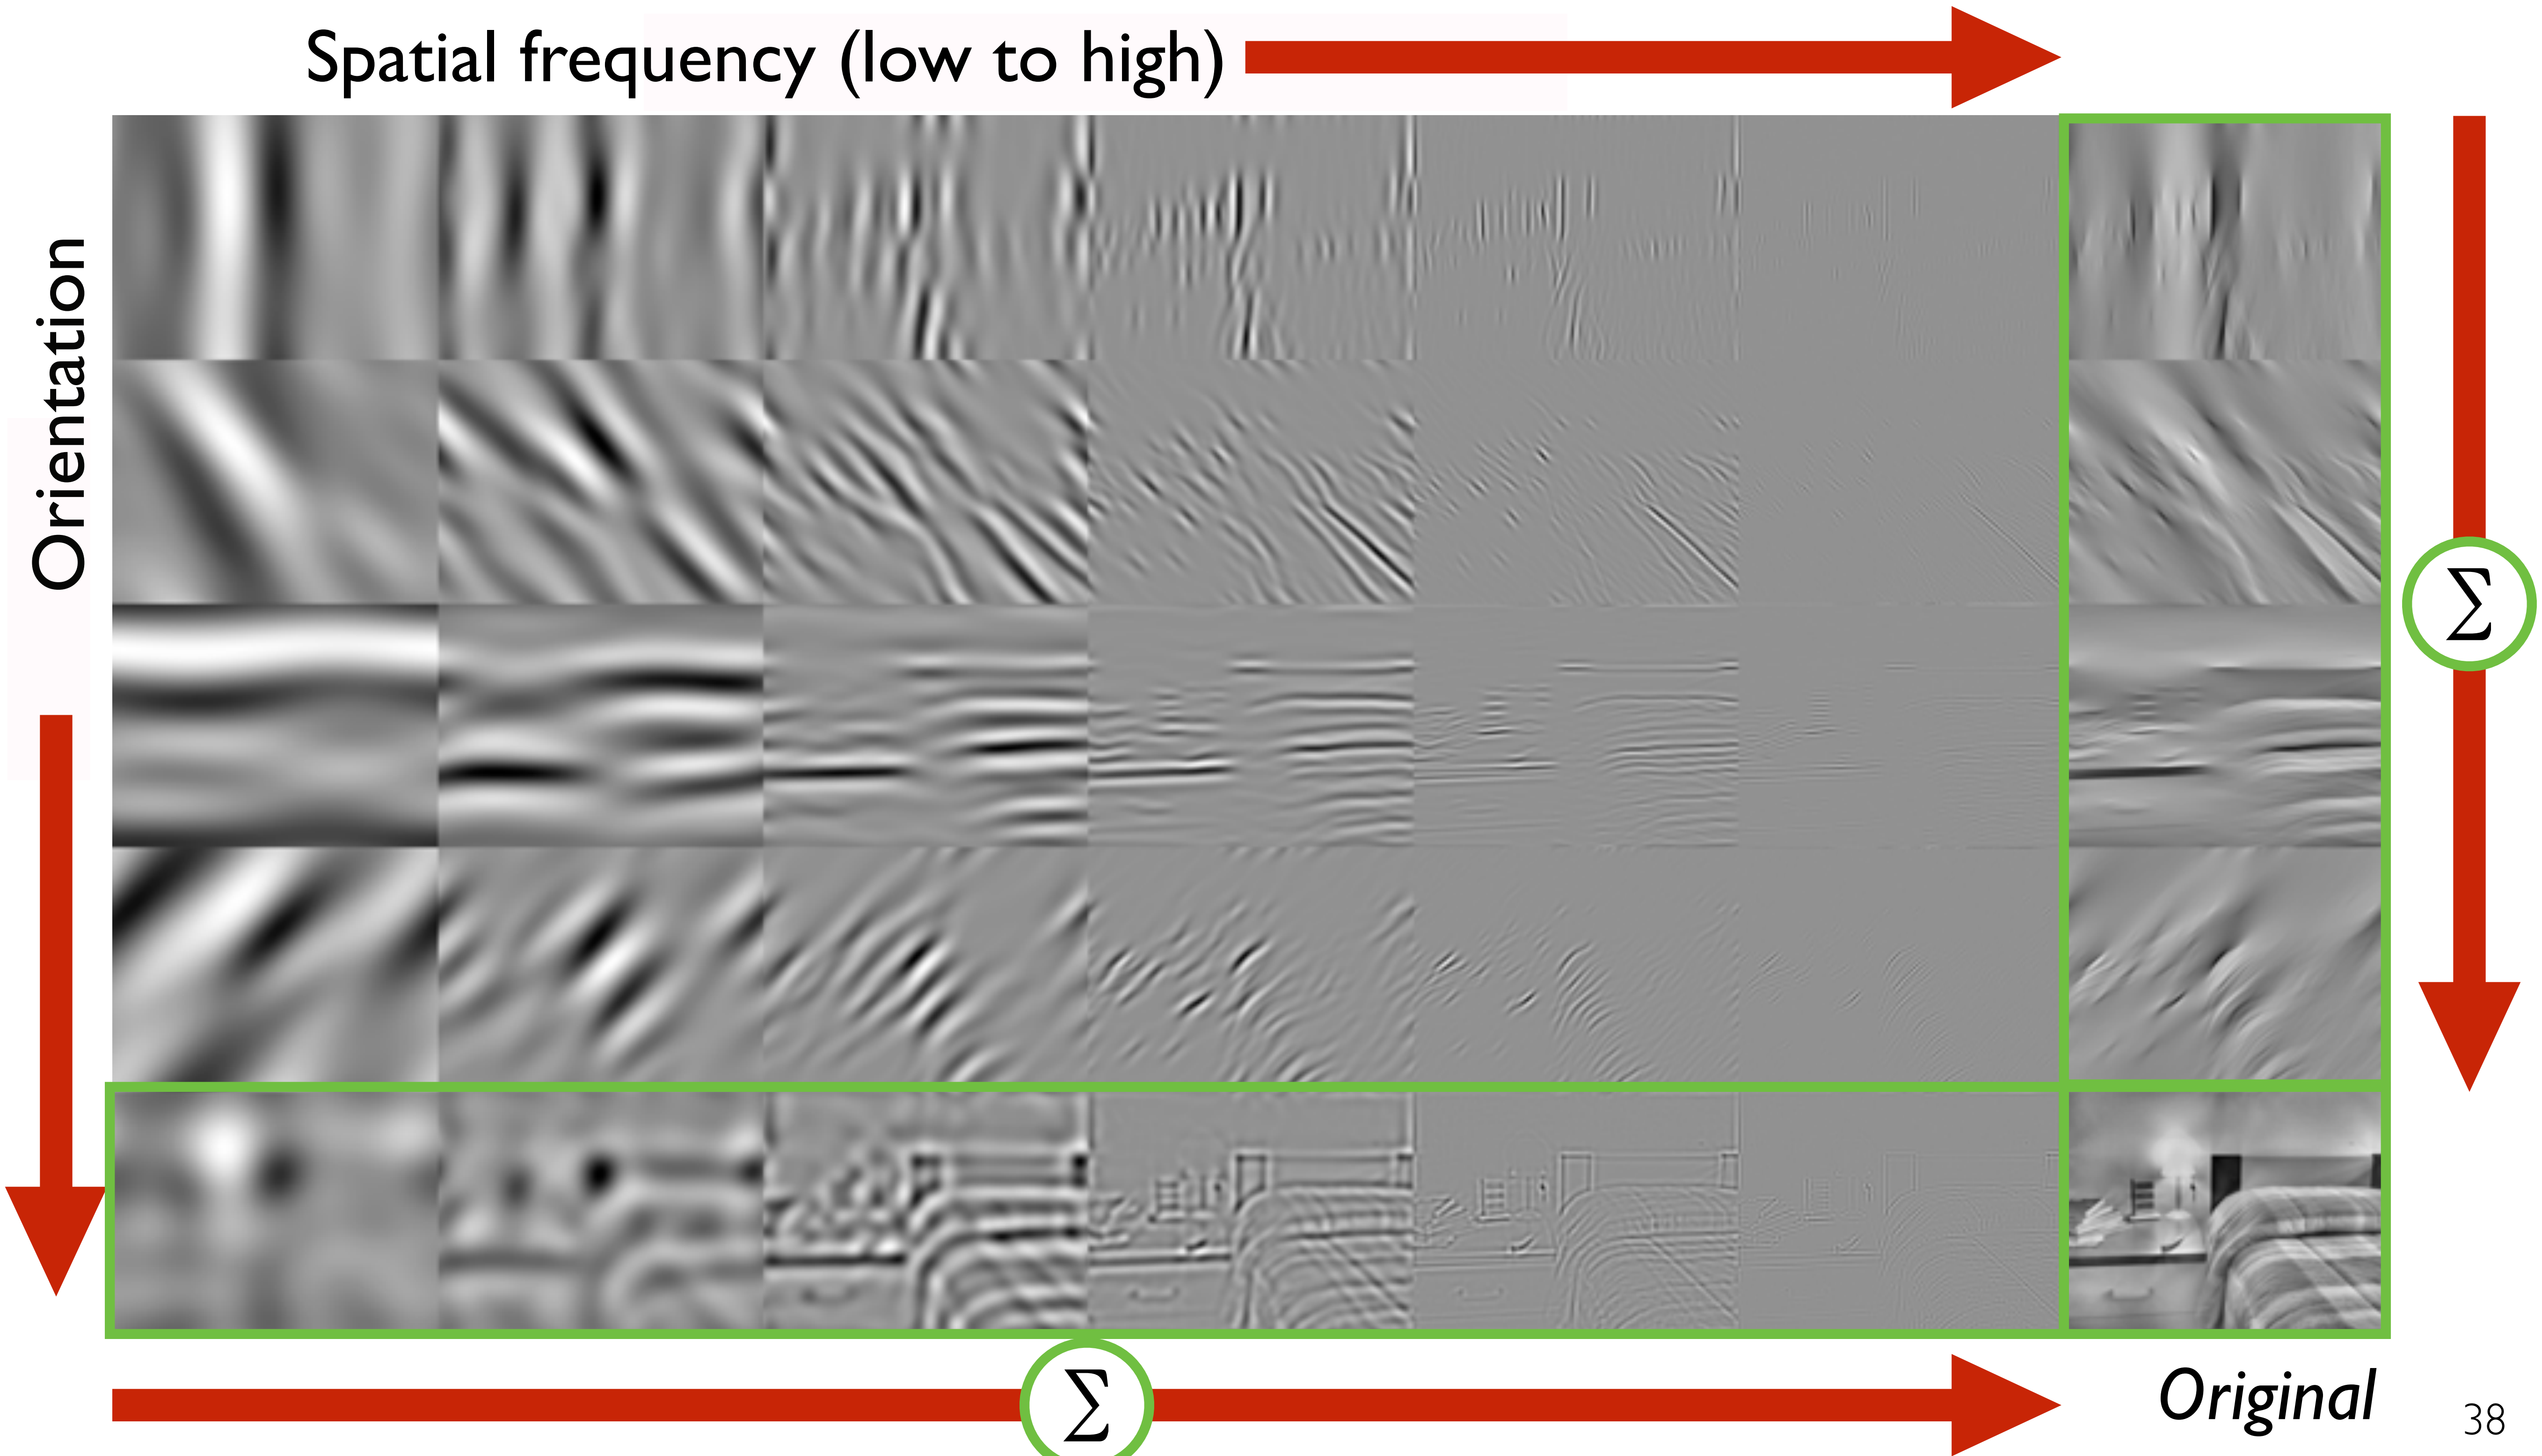

From spatial to temporal

- V1 cells are also selective for direction of motion
  - Show a preferred direction of motion with little response to the opposite direction
- Can we understand the temporal domain in terms of Fourier components as well?

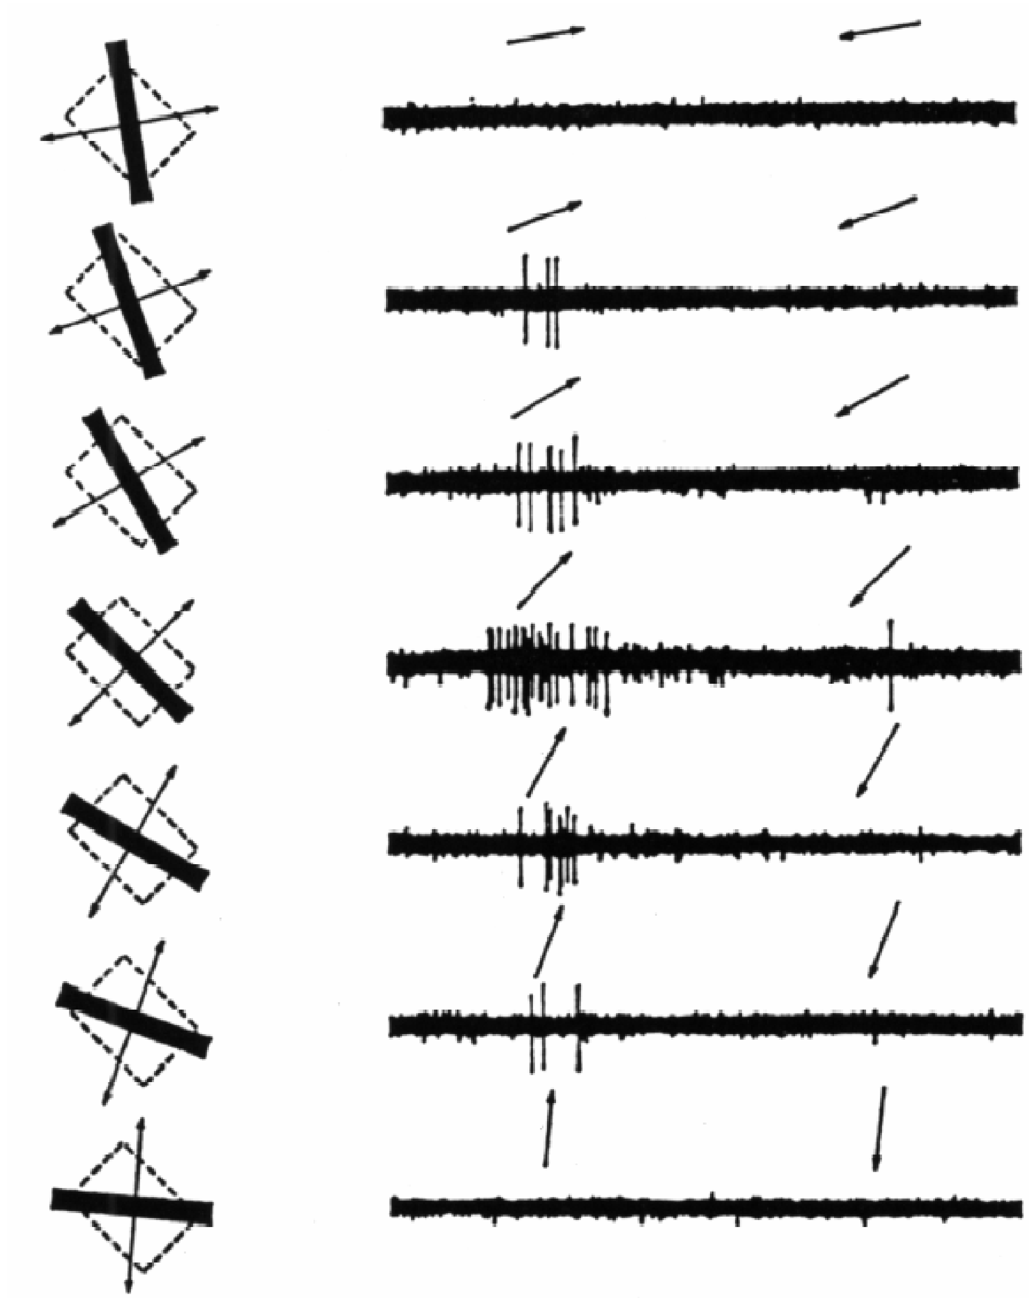

Hubel & Wiesel (1959)

Temporal filters

- Neurons' direction selectivity can be modelled with filters
- e.g. space-time inseparable RFs
  - In these plots, direction is an orientation in (x,t) space

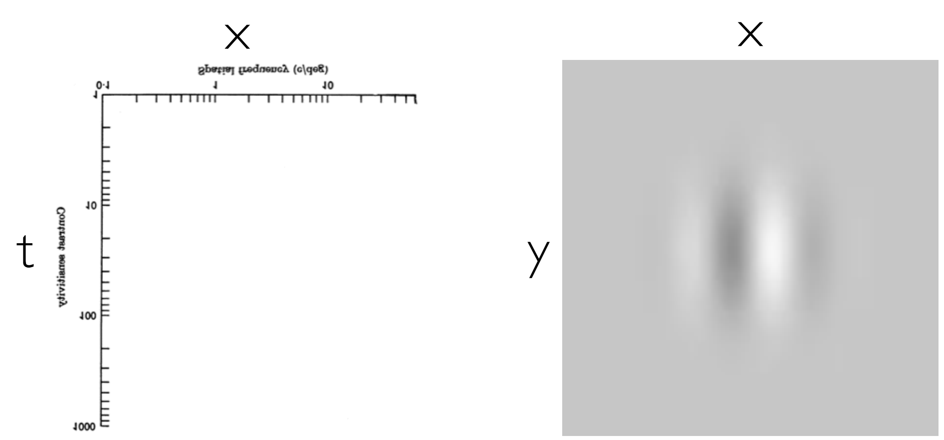

De Angelis, Ohzawa & Freeman (1995)

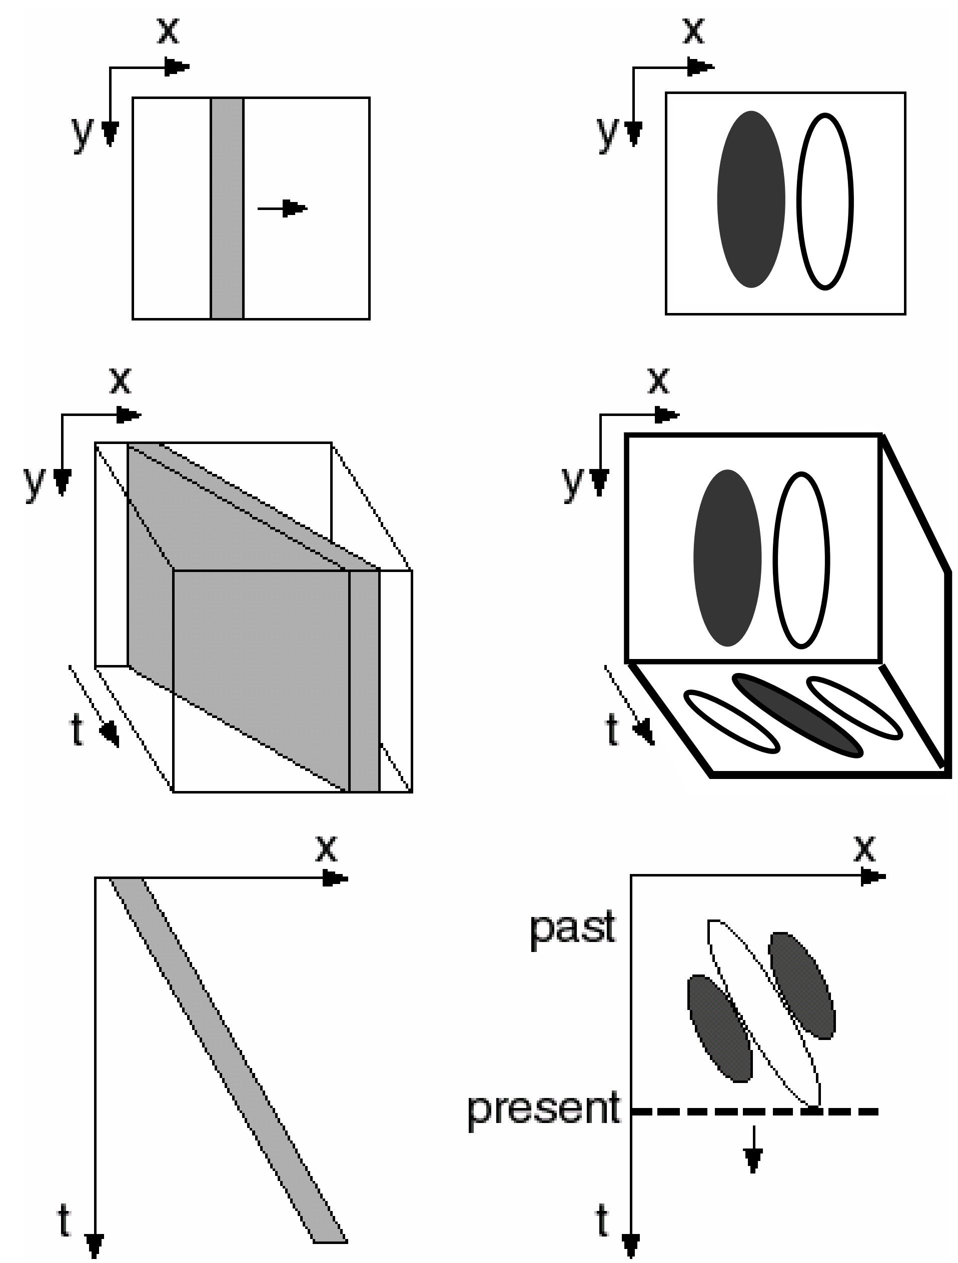

(Adelson & Bergen, 1985;
Carandini, Heeger & Movshon , 1999)

Temporal filtering

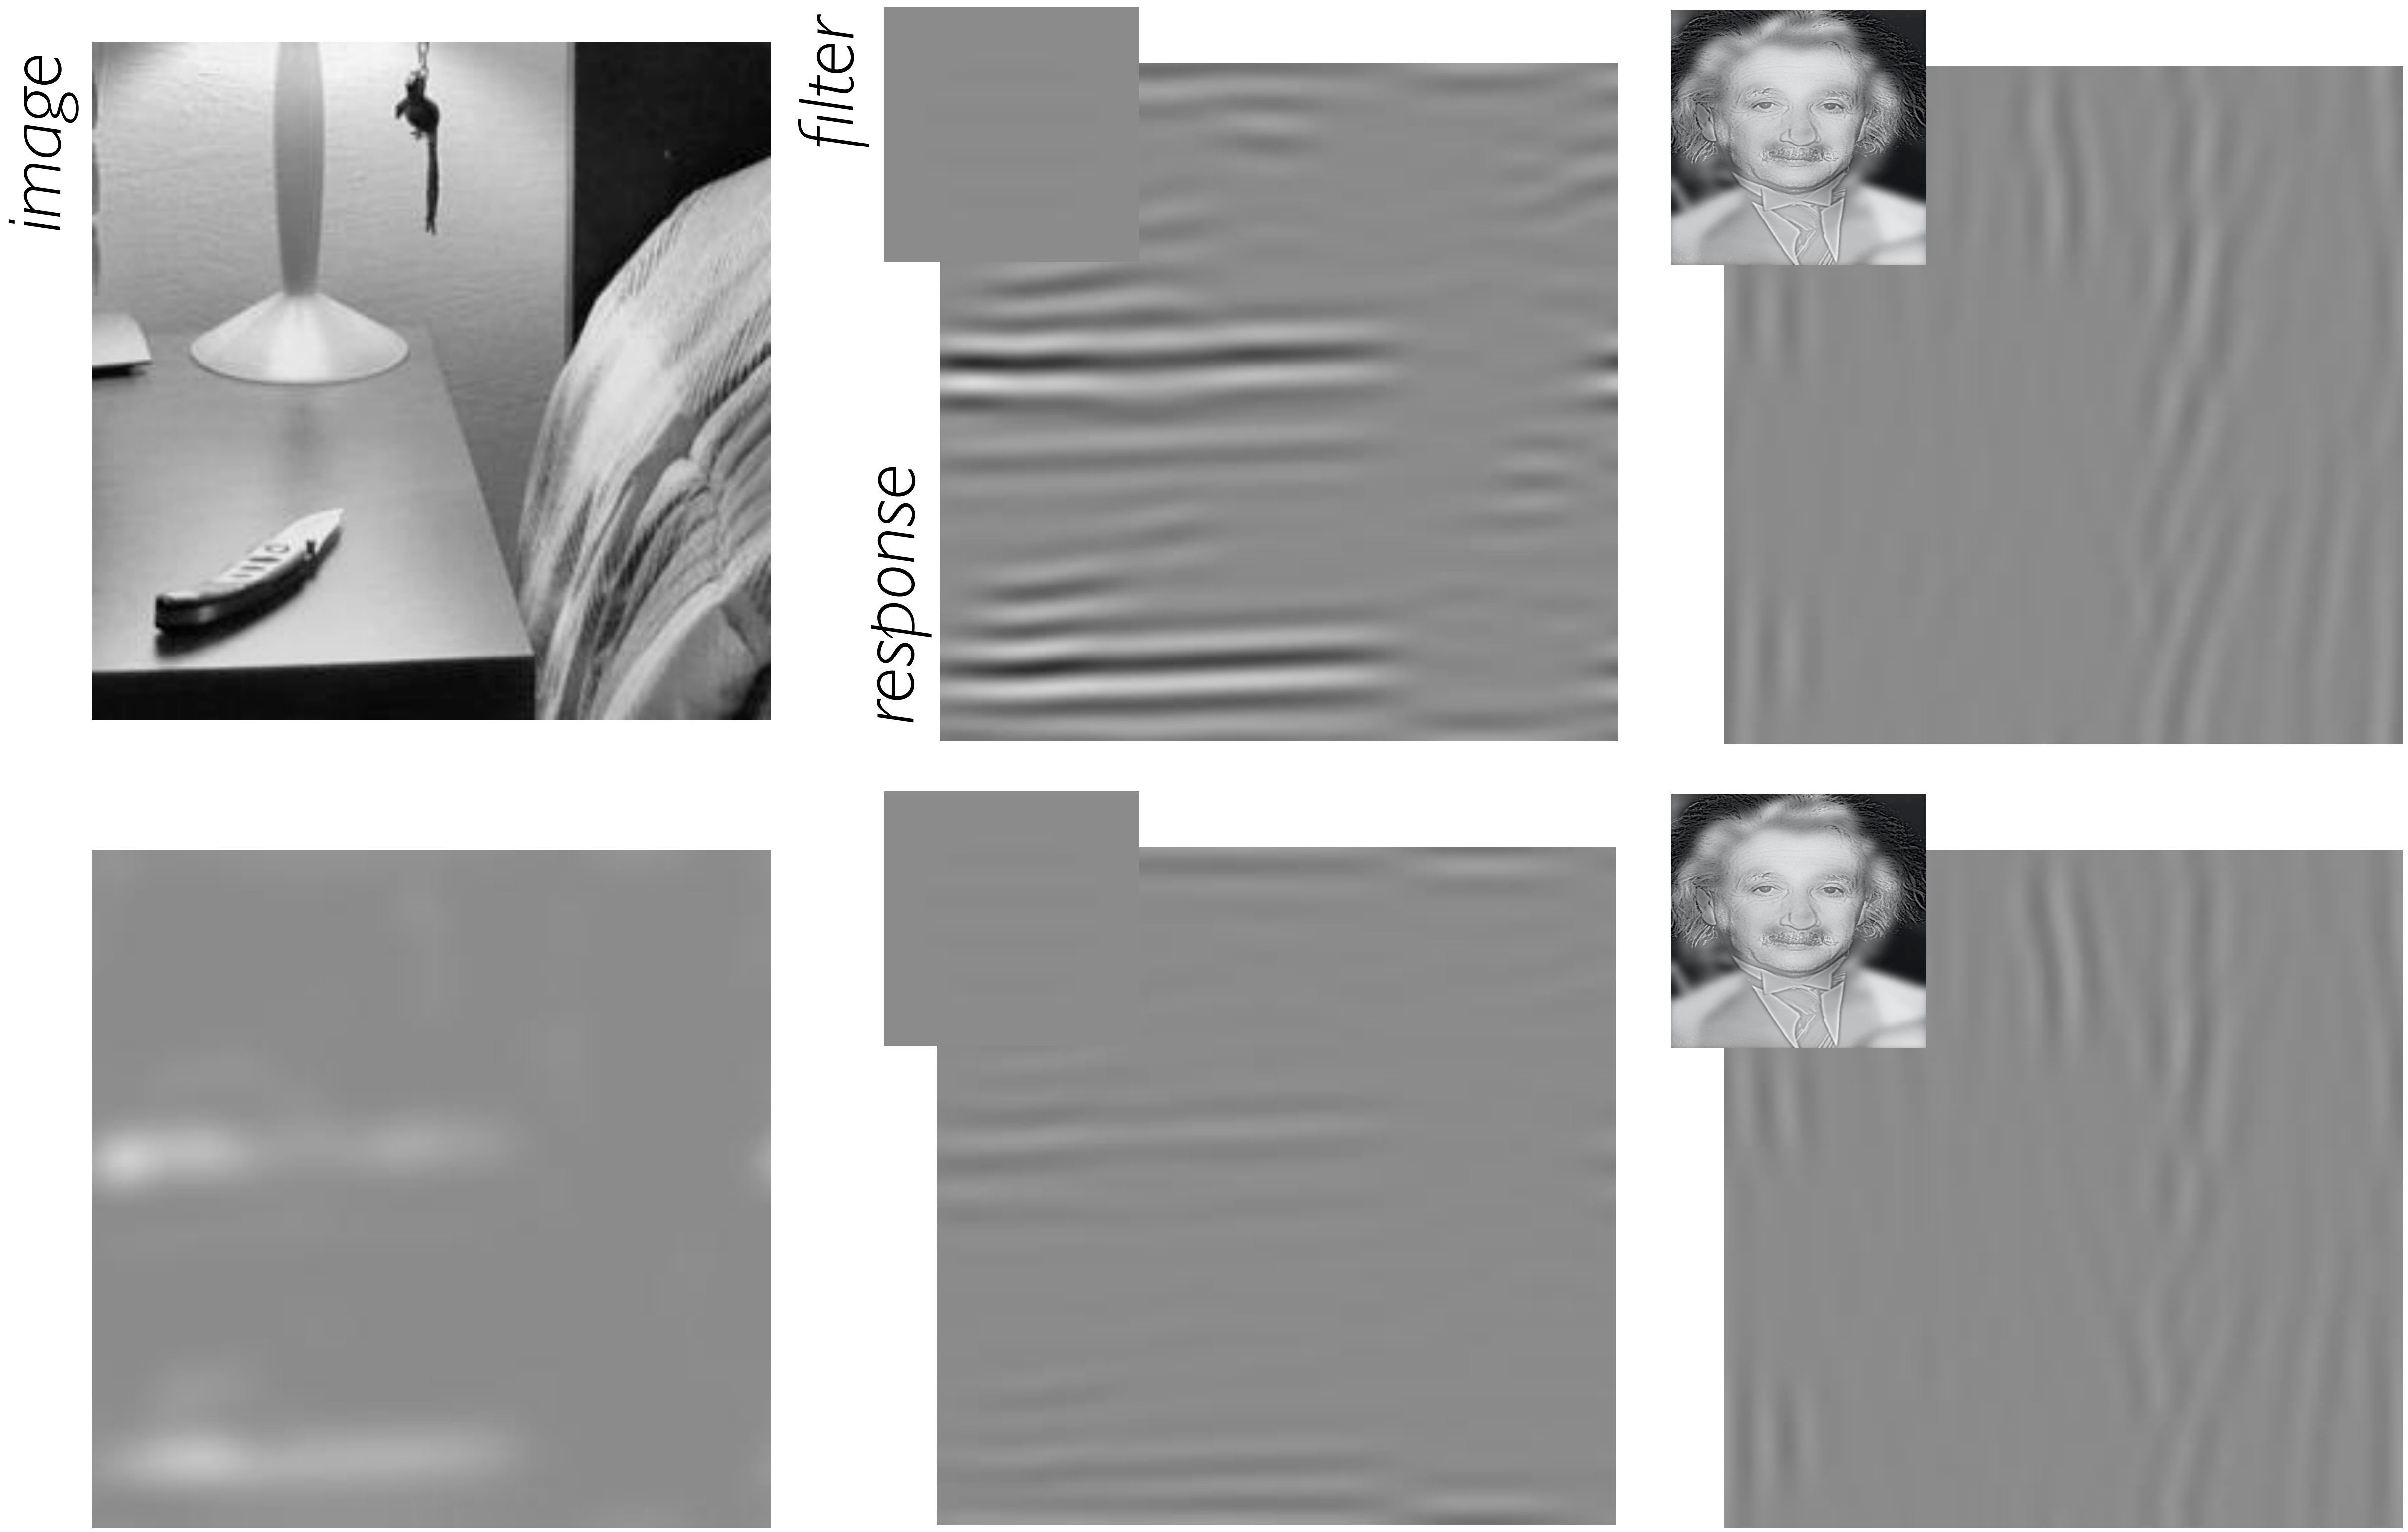

Local energy across all directions

Is there an advantage here?

- Action potentials are metabolically demanding (Laughlin et al., 1998)
- Solution: the visual system is optimised for the natural environment
    - Want efficient codes to minimise redundancy (Barlow, 1961) by adapting cell selectivities to the statistics of their input
    - Natural scenes show a consistent spatiotemporal structure
    - Gabor-like filters allow a balance between requirements of efficient coding and energy cost (e.g. Rao & Ballard, 1999)

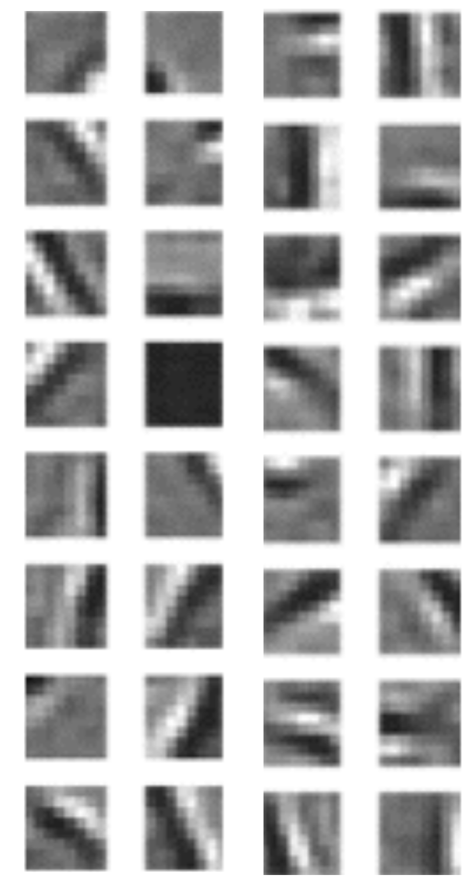

(Rao & Ballard, 1999)

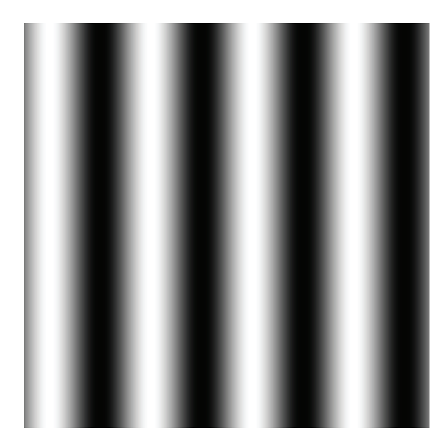

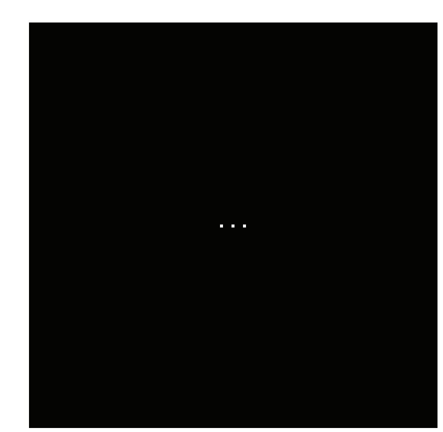

Summary

- Fourier analysis allows us to separate out properties like orientation, phase, contrast, and direction in natural stimuli
- We see evidence for spatial frequency channels that may perform these analyses in the mammalian visual system
  - Allows us to characterise vision via the Contrast Sensitivity Function
- Receptive fields in the early visual system can be thought of as filters that perform a local Fourier analysis of scenes
- This arrangement constitutes an efficient scheme for coding structure in the natural visual environment that minimises metabolic cost

Reading

- Chapter 3 of either Wolfe et al. or Goldstein et al. *Sensation & Perception* gives an overview of these ideas
- Further reading (if interested / confused):
  - Papers referenced in the lecture, e.g.
    - Campbell & Robson (1968). Application of Fourier analysis to the visibility of gratings.
    - Blakemore & Campbell (1969). On the existence of neurones in the human visual system selectively sensitive to the orientation and size of retinal images.
    - Schyns & Oliva (1999). Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations.
    - DeAngelis et al (1995). Receptive-field dynamics in the central visual pathways.