THEORETICAL NEUROSCIENCE I

Lecture 11: Linear filters in 2D

Prof. Jochen Braun

Otto-von-Guericke-Universität Magdeburg,
Cognitive Biology Group

Outline

1. LFM in higher dimensions

2. Simple and complex cells

3. Population of LFMs: x-t motion

4. RResponse-weighted covariance (advanced)

1 Linear filter models in higher dimensions

In sensory domains with higher-dimensional stimuli, LFM must use higher-dimensional linear kernels (=*preferred stimuli*).

Reverse correlation with random stimuli produces higher-dimensional kernels.

'Reverse correlation' applies only to time dimension!

'Direct correlation' is used for all other dimensions.

Reverse correlation

We consider a spatiotemporal stimulus $s(x, y, t)$ and the associated response $r(t)$:

$$s(x, y, t) \qquad\qquad \rightarrow \qquad\qquad r(t)$$

The reverse correlation (response-weighted average) is

$$Q_{rs}(x, y, -\tau) = \frac{1}{T} \int_0^T r(t)\, s(x, y, t - \tau)\, dt$$

It can be shown to be identical to a spike-triggered average.

Optimal kernel

Given a sufficiently long record of observations of *white-noise* stimuli $s(x, y, t)$ and associated responses $r(t)$, the reverse correlation gives the optimal linear kernel $D(x, y, \tau)$:

$$D(x, y, \tau) = \frac{Q_{rs}(x, y, -\tau)}{\sigma_s^2}$$

where τ is the delay (shift) between stimulus and response and σ_s^2 is the stimulus variance.

The kernel $D(x, y, \tau)$ determines how a stimulus at $(x, y, t - \tau)$ affects the response at time t.

It is 'optimal' in the sense that it minimizes the squared prediction error for complex stimuli (such as white noise).

Separable and non-separable kernels

For some neurons, the optimal filter is *separable* in that it is the product of a purely spatial and a purely temporal filter:

$$D(x, y, \tau) = D_s(x, y)\, D_t(\tau)$$

For other neurons, the optimal filter is *non-separable* in that it is *not* the product of a purely spatial and a purely temporal filter:

$$D(x, y, \tau)$$

Linear filter model

A linear-filter model of the neuron's response has two parts. A linear part with kernel $D(x, y, \tau)$ is

$$L(t) = \int d\tau \int dx\, dy\, D(x, y, \tau)\, s(x, y, t - \tau)$$

and a static non-linearity

$$r_{est}(t) = F\,[L(t)]$$

where $F[x] \geq 0$ is a suitable non-linear function.

Together, the two parts estimate the response r_{set} to any stimulus sequence $s(x, y, t)$.

Example: center-surround receptive field

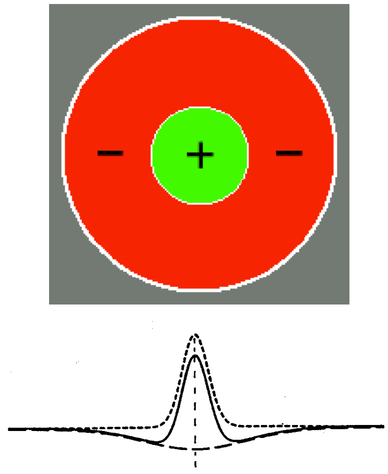

Figure 1: Differrence of Gaussian (DOG) model of retinal ganglion or lateral geniculate cells.

Spatial and temporal ‘bandpass’ filtering

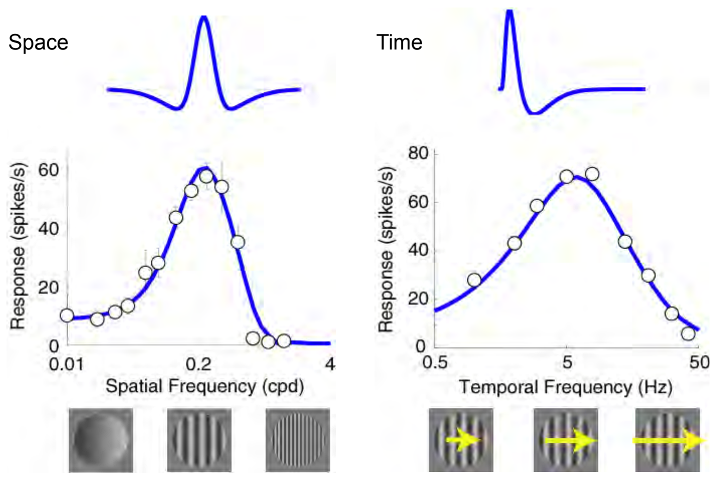

Figure 2: Emphasize a particular range of luminance gradients in space and time.

‘Filtered images’

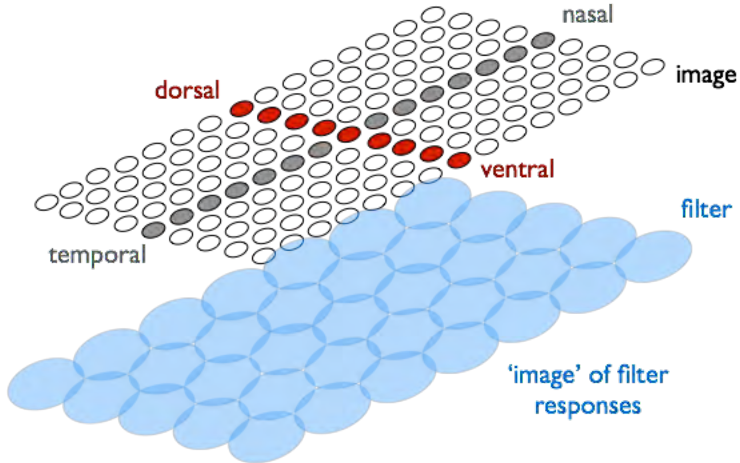

Figure 3: Apply identical filter to every image position, collect linear responses into 'filtered image'. [1]

'Bandpass filtering' of an image

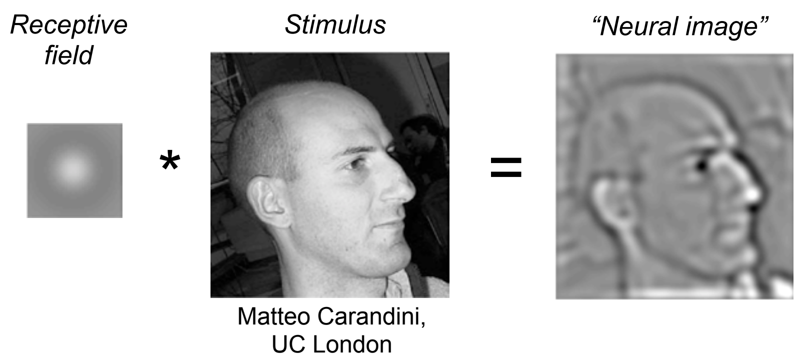

Figure 4: Apply identical center-surround filters to every image position, collect linear responses into 'filtered image'. [2]

Summary high-dimensional filters

- **Some sensory neurons represent high-dimensional stimuli (*e.g.*, time-frequency, time-space).**
- **Their responsiveness may also be described by linear-filter models.**

- **Optimal linear kernels may be determined by reverse correlation with white noise stimuli.**
- **The optimal linear kernel equals the** *preferred* **higher-dimensional stimulus.**
- **Example: center-surround space-time receptive field responds to luminance gradients over space and time.**
- **Acts as 'bandpass filter' for both spatial and temporal frequency.**
- **Each filter type provides 'filtered' version of original image/movie.**

2 Simple and complex cells

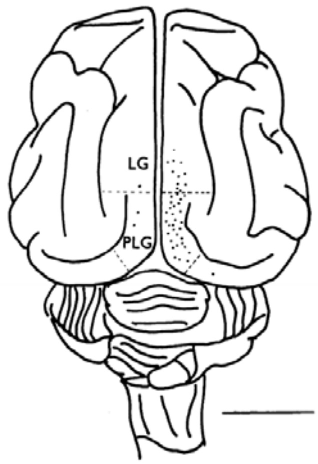

Figure 5: Hubel and Wiesel (1962) described two distinct cell types in visual cortex of cat (and, later, monkey). [3]

Simple cells

Receptive field comprises distinct ON and OFF regions:

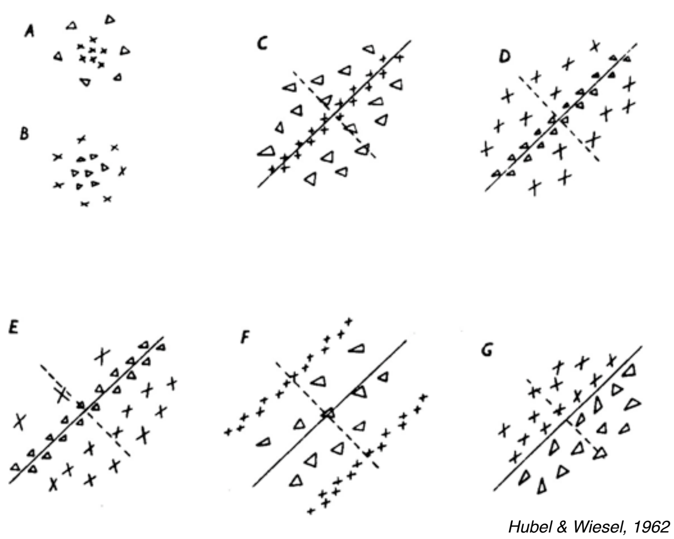

Figure 6: A,B: LGN cells, C to G: simple cells [4]

Complex cells

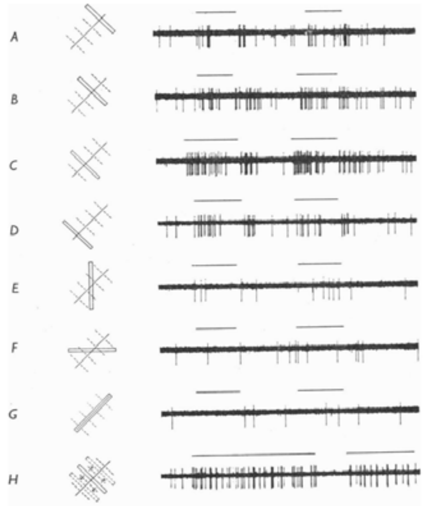

Figure 7: Receptive field shows no such structure. [5]

Shared properties

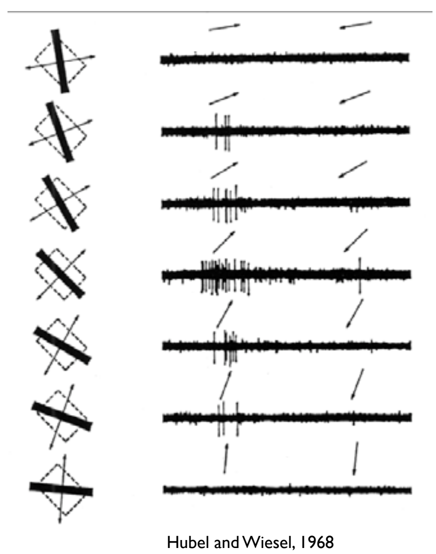

Figure 8: Both types prefer visual stimuli of some particular *orientation* (e.g., bars). Some cells also prefer a particular *direction* of motion. [5]

LFM of simple cell

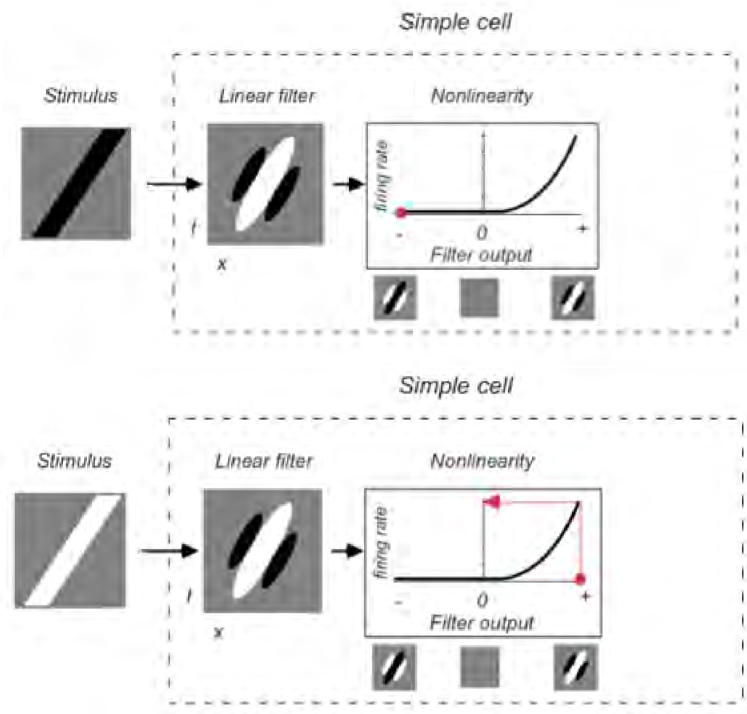

Figure 9: LFM of simple cell. [6]

stimulus $\rightarrow$ linear filter $\rightarrow$ static non-linearity

LFM of complex cell

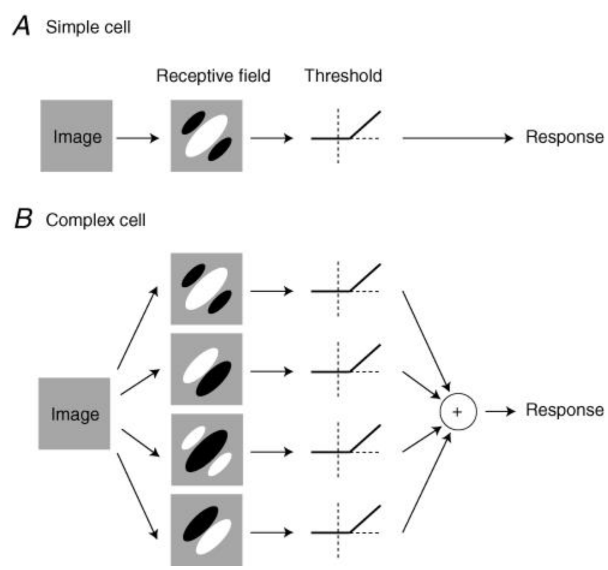

Figure 10: LFM of complex cell. [6]

stimulus $\rightarrow$ linear filter $\rightarrow$ static non-linearity

LFM of complex cell

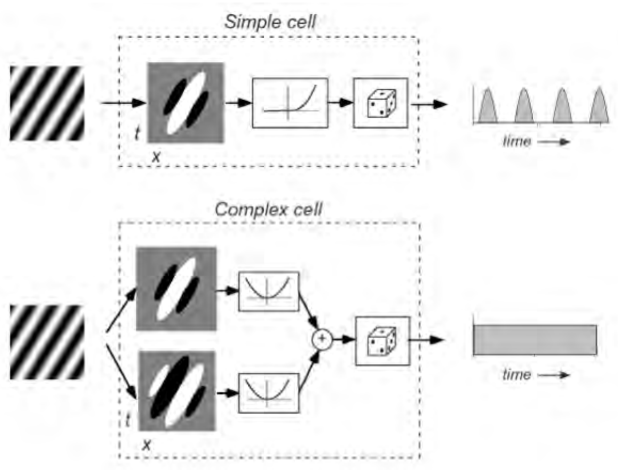

Figure 11: LFM of complex cell with noise. [6]

$$\text{stimulus} \rightarrow \text{linear filter} \rightarrow \text{static non-linearity} \rightarrow \text{noise}$$

Summary simple and complex cells

- Simple and complex cells in the visual cortex of cat (or monkey) illustrate general attributes of sensory neurons.
- Receptive fields are small and cover only a (tiny) fraction of visual space.
- Responses are 'tuned' to several stimulus properties (contrast, orientation, spatial frequency, ...).
- **Simple cell** receptive fields show distinct ON and OFF regions arranged in parallel stripes.
- This structure explains certain aspects of 'tuning' (orientation, spatial frequency).
- **Complex cells** do not show distinct ON and OFF regions, but exhibit similar 'tuning'.
- Complex cell may combine several different linear receptive fields.

3 Populations of LFMs: x-t motion

Applying identical filters to every image locations produces a 'filtered' image, emphasizing the locations at which the image matches the filter (positively or negatively).

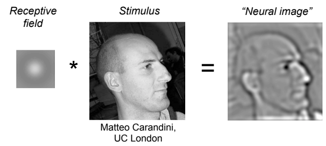

Figure 12: Filtered image. [2]

Sensory systems apply a heterogeneous set of filters to every image location. Thus, they produce a corresponding set of 'filtered images' (one image per filter). It is on the basis of this image set that sensory systems achieve their astonishing sensitivity in challenging sensory environments.

Heterogeneous sets of filters

Every image location is analysed by a set of *different* filters:

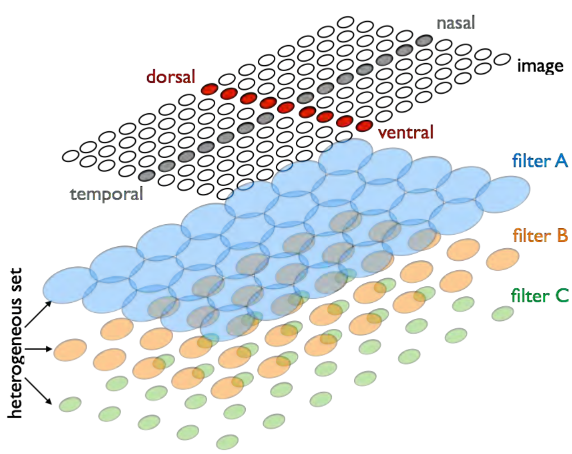

Figure 13: Layers of filters. [7]

To illustrate this principle in a challenging sensory situation, we consider visual motion of a 1D spatial pattern.

X-T motion

We consider the horizontal (x-direction) motion of vertical stripes (y-direction). Negative to positive contrast is represented on a color scale (dark orange representing zero).

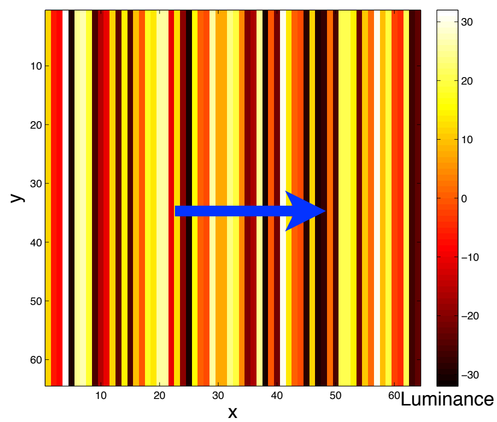

Figure 14: Such stimuli vary only in the x- and the t-direction.

Example stimulus with signal plus noise

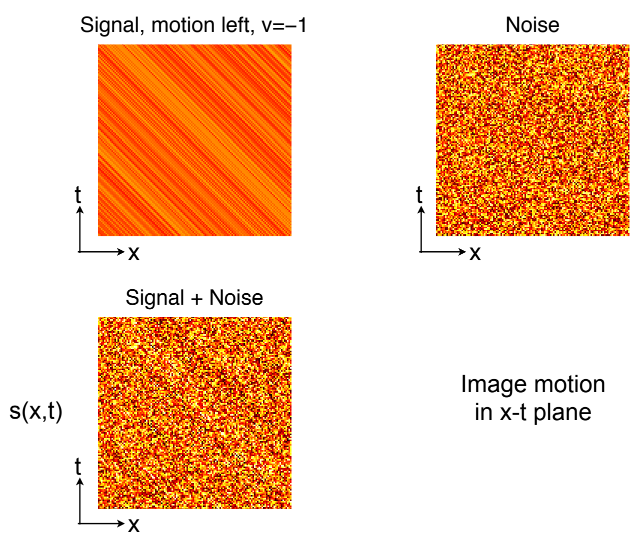

Figure 15: The signal $s(x,t)$ consists of diagonal motion to the left ($v = -1$).

Filters selective for rightward motion

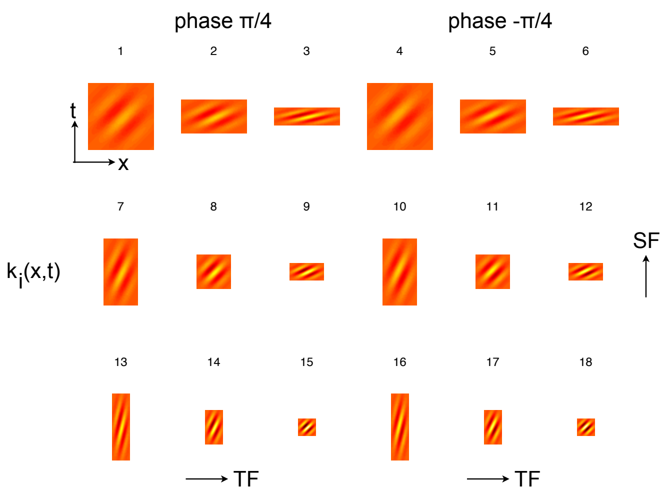

Figure 16: Filters selective for rightward motion.

A heterogeneous set of 18 filters $k_i(x,t)$, selective for **rightward** motion at different spatial and temporal frequencies (SF and TF).

Filters selective for leftward motion

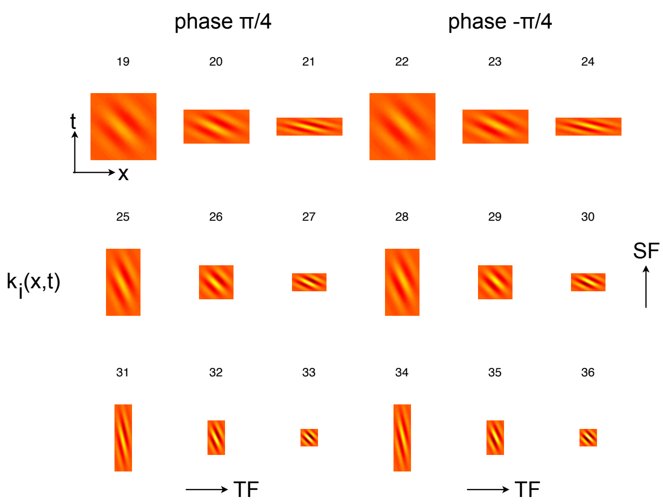

Figure 17: Filters selective for leftward motion.

A heterogeneous set of 18 filters $k_i(x,t)$, selective for **leftward** motion at different spatial and temporal frequencies (SF and TF).

One filtered image

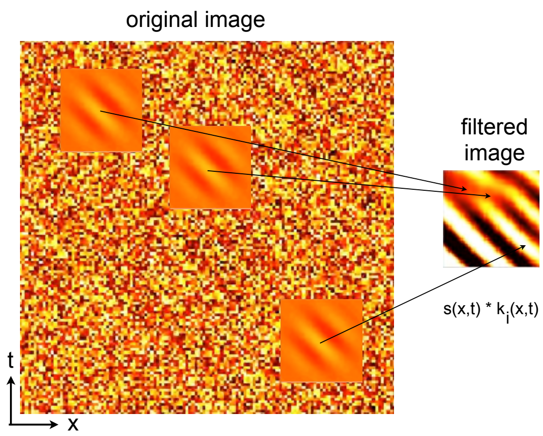

Figure 18: Filtered image.

The original stimulus (in x-t plane) is convolved with (filtered by) one particular linear filter ($k_{19}(x,t)$). This produces a filtered image, representing the degree to which each x-t location of the stimulus matches this particular filter (positively or negatively).

Many filtered images ('simple' cells)

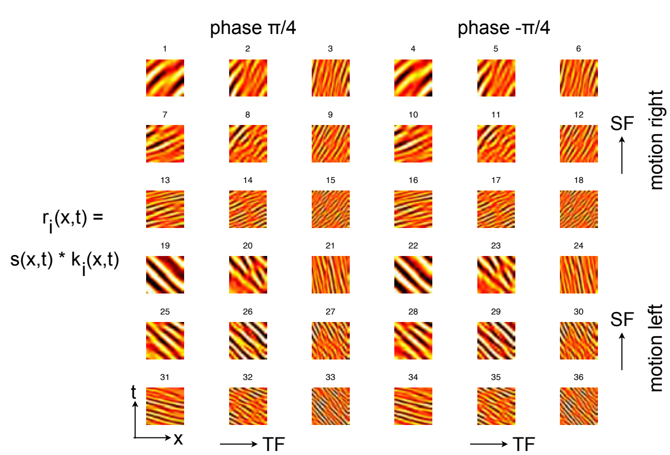

Figure 19: Many filtered images ('simple' cells)

Convolution of the original image with the entire set of 36 filters produces 36 filtered images $r_i(x,t)$. Each represents how each x-t location matches the filter in question. Due to the high level of noise, all x-t locations match all filters to some degree.

Many filtered images ('complex cells')

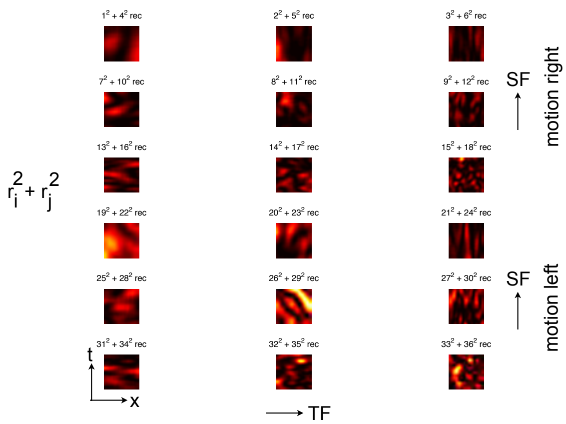

Figure 20: Many filtered images ('complex cells')

Combining the responses of filter pairs (same SF, same TF, complementary phase) reveals which filters are **consistently** matched by the stimulus: filters representing diagonal motion to the left ($v=-1$, $r_{19}^2+r_{22}^2$, $r_{26}^2+r_{29}^2$, $r_{33}^2+r_{36}^2$).

Summary heterogeneous filter sets

- Sensory systems analyze stimuli with a heterogeneous set of filters.
- Conceptually, one stimulus image is represented by a set of 'filtered images' (one per filter).
- Typically, in the presence of noise, all stimulus locations (x,t) match all filters to some extent.

- However, comparing between ‘filtered images’ reveals which filters are *consistently* matched by the stimulus.

- The true power of linear filtering becomes evident with heterogeneous filter sets, where each filter response can be assessed in the context of other filter responses.

4 Response-weighted covariance (advanced)

The *spike-triggered average* fails when a neuron combines several linear receptive fields (complex cells). If we have enough data, we can use the *spike-triggered covariance* instead. This method identifies the visual patterns that account for most of the response variance.

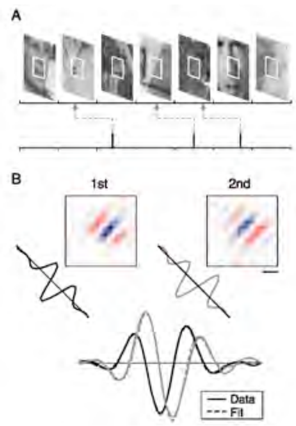

Figure 21: Response-weighted covariance. [8]

Simple cell

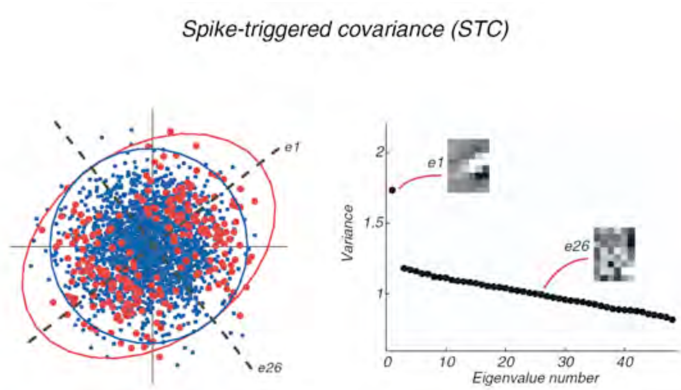

Figure 22: Spike-triggered covariance in simple cell.

In a simple cell, a single pattern accounts for most of the response variance.

Complex cell

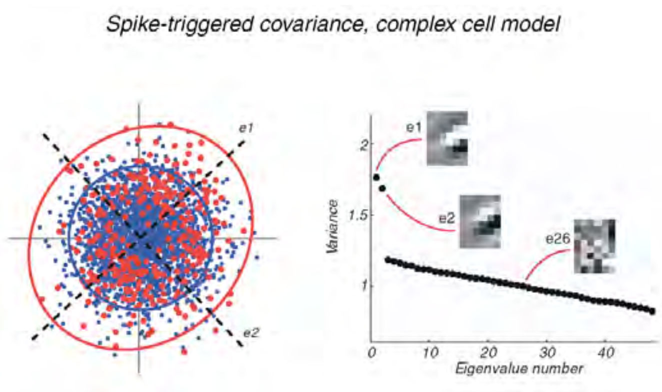

Figure 23: Spike-triggered covariance in complex cell.

In a complex cell, two patterns account for most of the response variance.

Math aside: vectors and covariances

For zero-mean vectors x_i, define average $\langle x_i \rangle$, variance $\langle x_i^2 \rangle$, and covariance $\langle x_i \cdot x_j \rangle$:

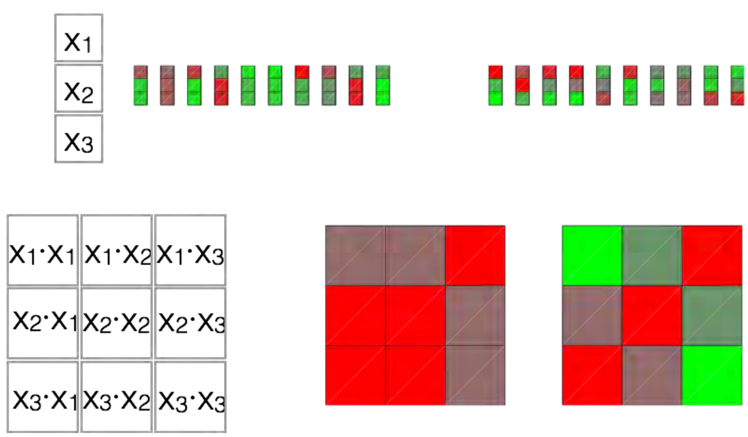

Figure 24: Vectors and covariances.

How does it work?

Formally, we compute the response-weighted average *covariance matrix* and identify its eigenvectors and eigenvalues.

The eigenvectors with the largest eigenvalues account for most of the response variance.

The linear receptive fields of the cell are linear combinations of the largest eigenvalues.

Thus, the result is the *family* of linear receptive fields that underlies the neuron's response.

Complex cell model

We begin by constructing a complex cell model with two linear kernels, $D_1(x, y)$ and $D_2(x, y)$:

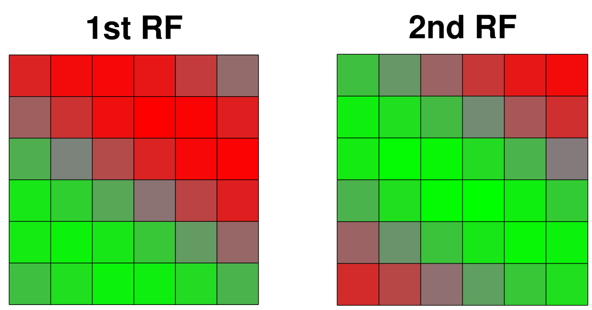

Figure 25: Two linear kernels.

$P_i(x, y)$

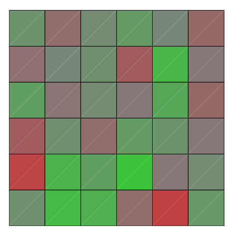

Figure 26: Input patern.

The linear responses $L_i^{(1)}$ and $L_i^{(2)}$ to an input pattern $P_i(xy)$ are

$$L_i^{(1)} \;=\; \int\int D_1(x,y)\, P_i(x,y)\, dx\, dy$$

$$L_i^{(2)} \;=\; \int\int D_2(x,y)\, P_i(x,y)\, dx\, dy$$

and the non-linear response of our complex cell is

$$R_i \;=\; \left[L_i^{(1)}\right]^2 + \left[L_i^{(2)}\right]^2$$

White noise stimuli

Next, we generate many (thousands) of white noise stimuli $P_i(x,y)$ and obtain for each stimulus the associated complex cell response R_i.

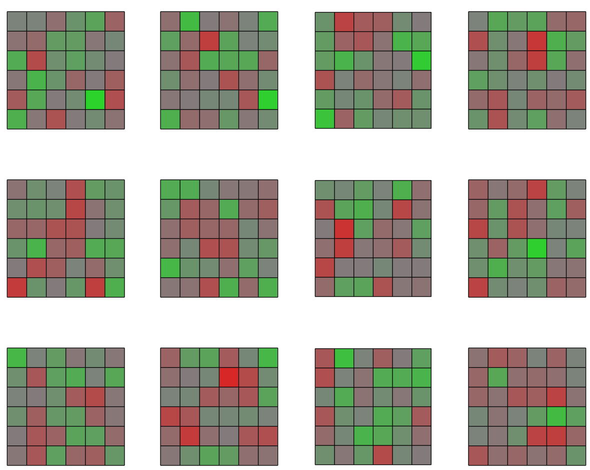

Figure 27: Many white noise stimuli.

Associated simple and complex cell responses

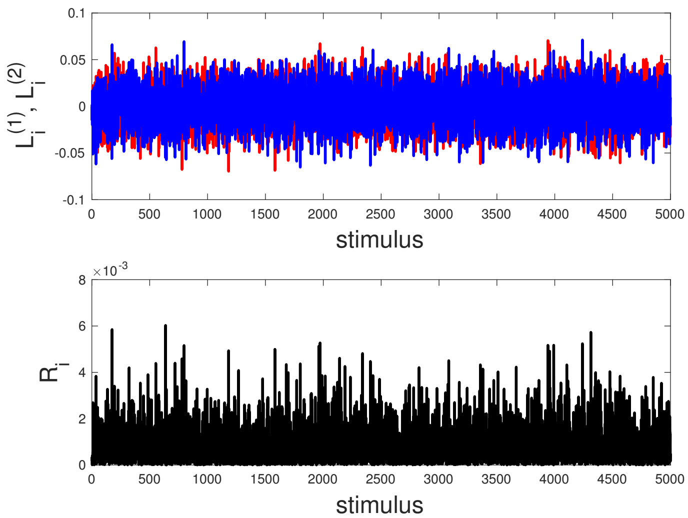

Figure 28: Associated simple and complex cell responses

STA *versus* STC

For STA, we computed the response-weighted average over patterns P: $\langle P \rangle = \frac{1}{N}\ \Sigma_i\ R_i\, P_i$

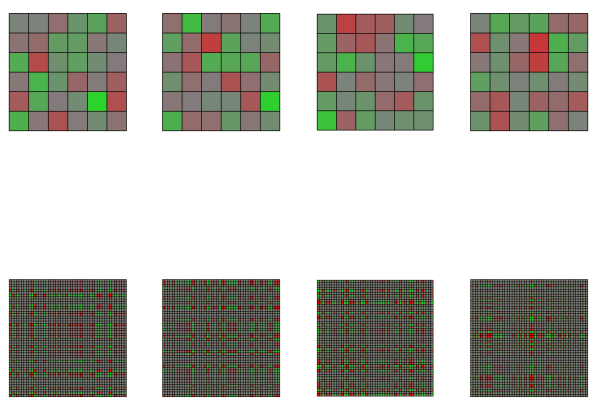

Figure 29: Spike-triggered average vs. spike-triggered covariance.

Now, we compute instead the response weighted average over covariances C: $\langle C \rangle = \frac{1}{N}\ \Sigma_i\ R_i\, C_i$

Patterns, vectors, and covariances

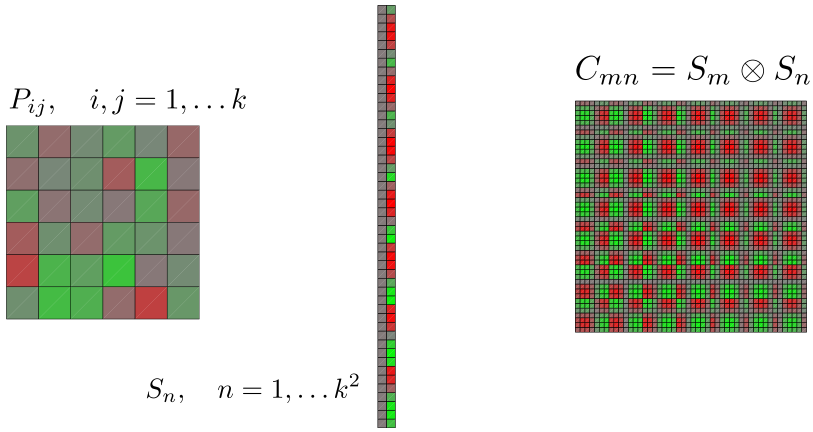

Figure 30: A 6×6 image P_{ij} can be written as a 36 dimensional vector S_n. Its covariance C_{nm} is a 36×36 matrix.

In the case of our complex cell, the STA shows little structure. This is because the cell responds to patterns of opposite sign (bright *and* dark spots at the same location), which cancel in the average.

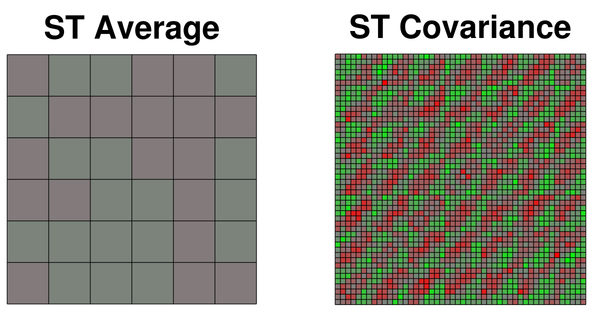

Figure 31: STA vs. STC

In contrast, the STC shows a highly consistent structure. This is because responses correlate with bright spots at one location and dark spots at another (but not *vice versa*). In fact, responses correlate with two distinct arrangements of dot pairs (because the cell has two receptive fields).

Eigenvectors and -values of STC

The STC matrix combines the covariances (possibly several!) that were consistently associated with large responses. The eigenvectors and eigenvalues of the STC matrix reveal theses covariances.

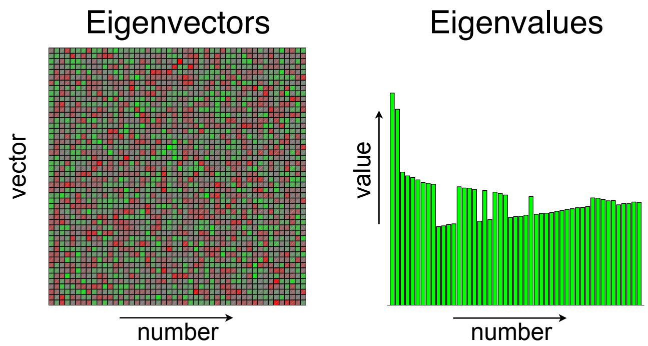

Figure 32: Eigenvectors and eigenvalues.

In the present case, two eigenvectors (covariances) have far larger eigenvalues than all others. Accordingly, the two principal eigenvectors account for most of the response variance.

Compare principal eigenvectors to linear kernels

The two principal eigenvectors are related to the two linear kernels of our model complex cell: the linear kernels are linear combinations of the principal eigenvectors.

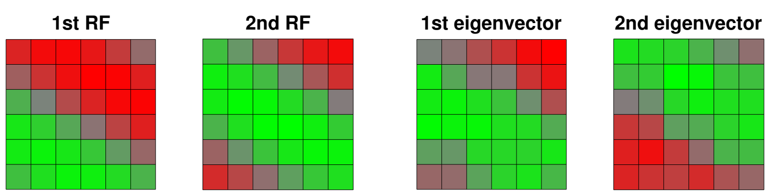

Figure 33: Comparison between eigenvectors and linear kernels.

Thus, we succeeded in identifying the family of linear kernels that drives responses.

Complex cell

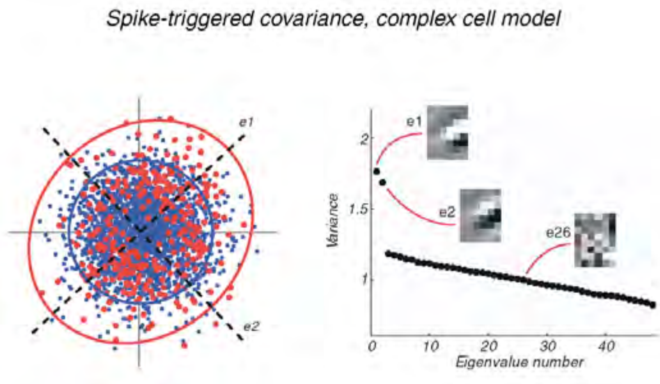

Figure 34: Spike-triggered covariance in complex cell.

In a complex cell, two patterns account for most of the response variance.

Summary spike-triggered covariance

- Spike-triggered averaging fails when neurons combine multiple linear filters (e.g., complex cells).
- An alternative approach (requiring quadratically larger data sets) is spike-triggered covariance.
- This method was applied successfully to visual cortical neurons by Touryan, Felsen, and Dan (2005)

5 Bibliography

1. Matteo Carandini (London). Neural Image. Personal communication.
2. Matteo Carandini (London). Neural Image. Personal communication.
3. Hubel & Wiesel, 1962
4. Charles Frye, What were the fundamental findings of Hubel and Wiesel, J. Physiol. 1962?, Anonymous, 2016, Image 1 Ref http://charlesfrye.github.io/FoundationalNeuroscience/02/
5. Hubel & Wiesel, 1962
6. McMillan & Creelman, Detection Theory, 1991 Anthony Movshon, New York
7. Matteo Carandini (London). Neural Image. Personal communication.