THEORETICAL NEUROSCIENCE I

Lecture 10: Linear filter models

Prof. Jochen Braun

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

Content

1. Linear systems theory

2. Waveforms in time and frequency

3. Cross-correlation and convolution

4. Linear-filter models (LFM) of neural response (advanced)

5. Non-linear part of LFM: ‘static non-linearity

Story so far ...

We investigated neuronal responses in the early visual pathway of frogs and cats:

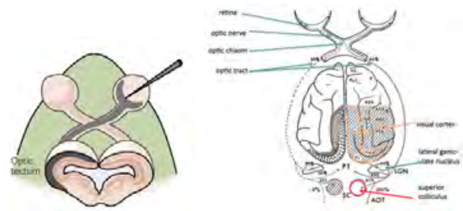

Figure 1: Visual pathway of frogs and cats. [1]

- In frog retina, neuronal responses seemed to signal to large, behaviourally relevant objects (edible prey).
- In cat LGN, neuronal responses seemed to signal tiny patterns of luminance contrast (ON/OFF venter-surround).
- In both cases, a specific part of the visual surroundings needed to be stimulated (small 'receptive field' within overall 'visual field').

In describing a neurone's response, three heuristic notions were useful:

- **Receptive field:** area in sensory space where stimulation drives neuronal response.
- **Response variability:** variability (or 'unreliability') of response limits informativeness.
- **Tuning curve:** dependence of average response on one particular stimulus attribute.

A more systematic approach ('reverse correlation') relied on stimulus sequences drawn randomly from a large ensemble:

- Spike-triggered or response-weighted average of preceding stimulus.

This was done without theoretical justification!

1 Linear systems theory

For simplicity, we consider one-dimensional, time-varying signals (e.g., sound pressure):

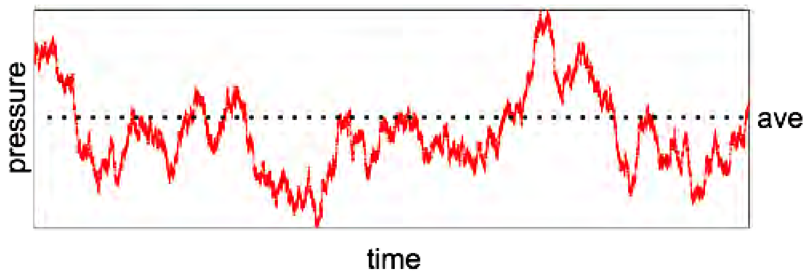

Figure 2: Sound pressure.

We would like to understand the response of a 'dynamic system' to such stimuli (e.g., the response of a sensory neuron).

A lookup table?

One conceivable approach would be to accumulate a 'look-up table', which would list the system/neuron response to every auditory stimulus observed so far.

This would be inefficient, impractical, and pointless!

A linear model!

Another approach is to develop a theory, which permits us to predict future observations.

Ideally, this theory will be determined by only a few actual measurements.

Naturally, the theory must be appropriate for the system/neuron we are studying.

Linear systems theory is an excellent theory for many sensory neurons. Not all sensory neurons are linear, but many important ones are.

Stimulus and response

We stimulate a linear system with S_1 and observe the associated response R_1:

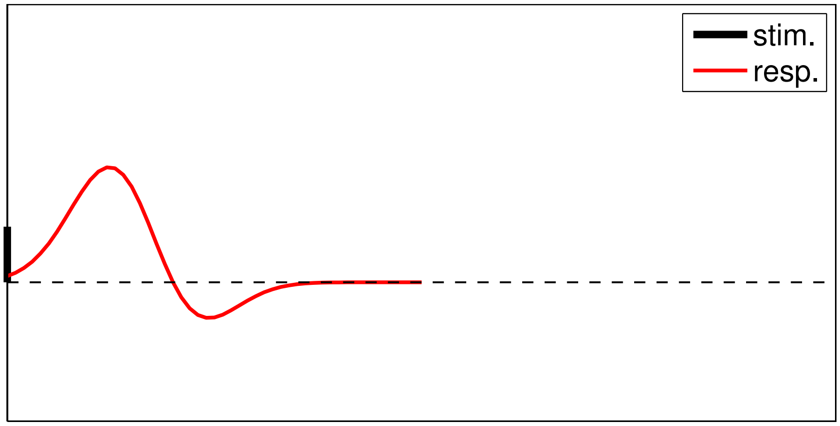

Figure 3: Assosiated response.

Shift-invariance:

If we repeat stimulus S_1 at a different time, we expect an identical response (only shifted in time). A system with shift-invariant responses is also called **stationary**.

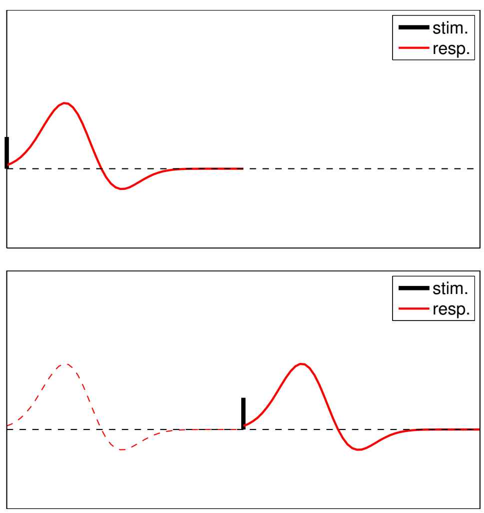

Figure 4: Identical response shifted in time

Homogeneity:

If we double the stimulus $S_2 = 2 \cdot S_1$, we expect double the response $R_2 = 2 \cdot R_1$.

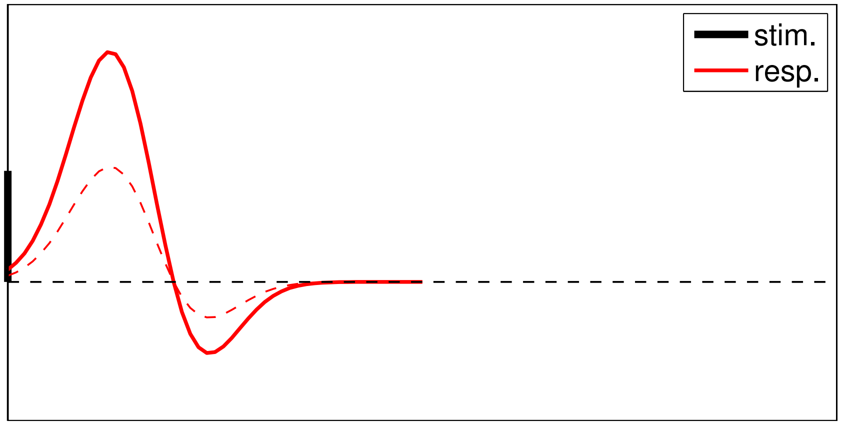

Figure 5: Doubled response

Homogeneity, ctd:

If we invert the stimulus $S_2 = -S_1$, we expect an inverted response $R_2 = -R_1$:

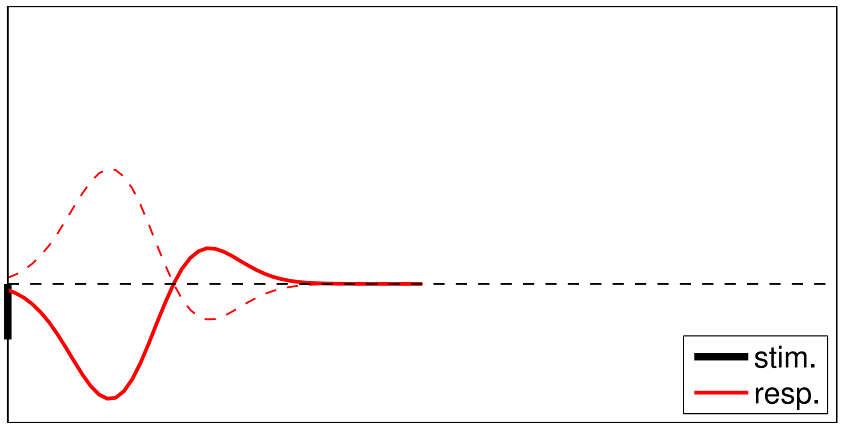

Figure 6: Inverted response

Additivity: Presented individually, two stimuli S_1 and S_2 produce responses R_1 and R_2, respectively.

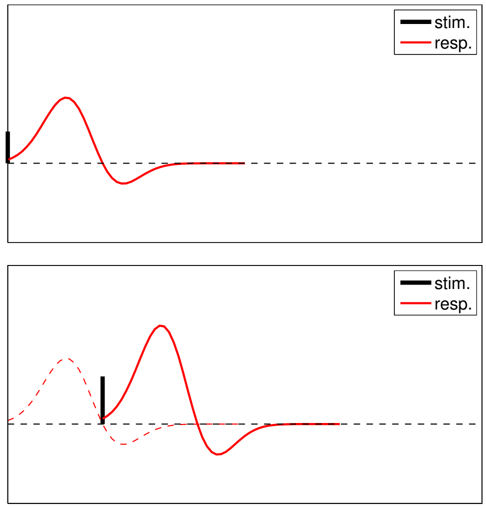

Figure 7: Additivity

Additivity, ctd:

Presented together as $S = S_1 + S_2$, they produce the response $R = R_1 + R_2$:

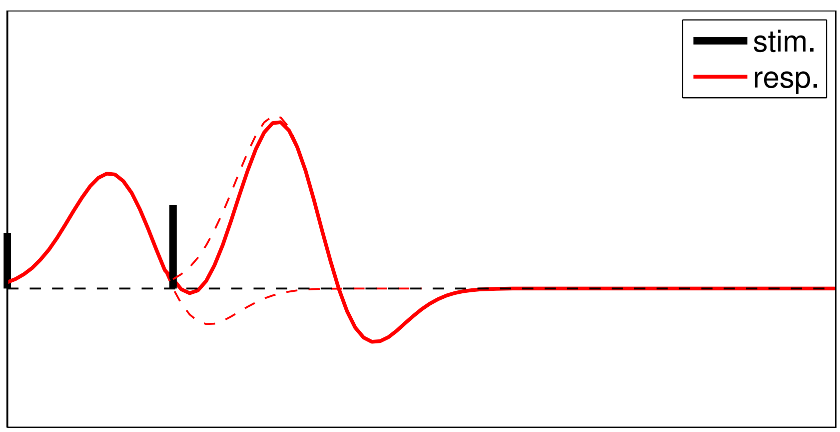

Figure 8: Response

Conditions for linearity

- Responses satisfying both *homogeneity* and *additivity* are considered to be linear.
- *Homogeneity* and *additivity* constitute the "principle of superposition".
- Responses satisfying *shift-invariance* are considered to be stationary.
- Not all linear systems are stationary, but ours are.
- Linear-system models of neuronal responses exhibit all three properties: *homogeneity*, *additivity*, *shift-invariance.*

Space-time method for predicting responses

With a linear-systems model, responses to arbitrary stimuli can be predicted from the response to just one stimulus, the *impulse response function.*

Arbitrary stimuli can be 'decomposed' into a sequence of scaled and shifted versions of a standard impulse, $S = \Sigma_k \; a_k \, S_1$.

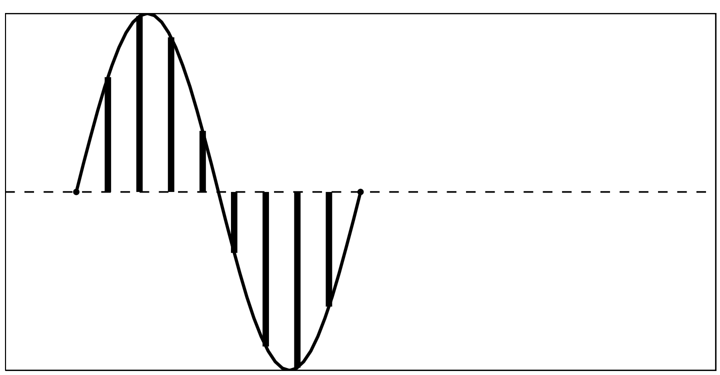

Figure 9: Impulse response function

Superposition of impulse responses

The predicted response is the sum of scaled and shifted versions of the impulse response, $R = \Sigma_k \; a_k \, R_1$.

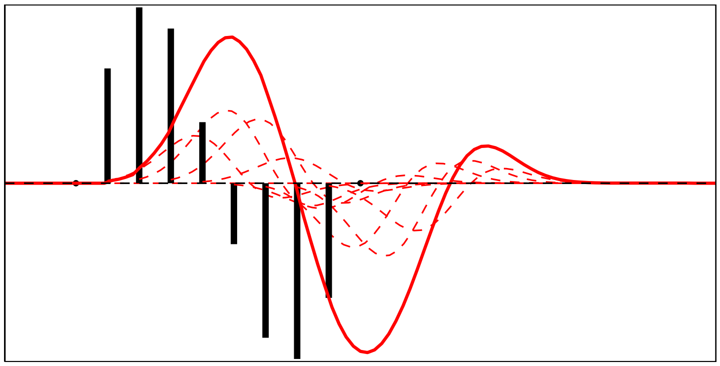

Figure 10: The predicted response is the sum of scaled and shifted versions of the impulse response

Frequency method for predicting responses

With a linear-systems model, responses to arbitrary stimuli can be predicted from the response to *harmonic stimuli.*

Arbitrary stimuli can be 'decomposed' into a sum of scaled and shifted harmonic stimuli ("Fourier series").

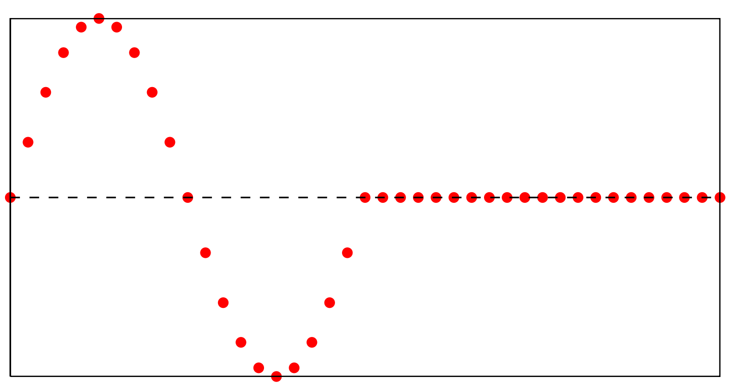

Figure 11: Harmonic stimuli

Harmonic functions

Harmonic functions of different frequency and phase (shifted vertically for clarity)

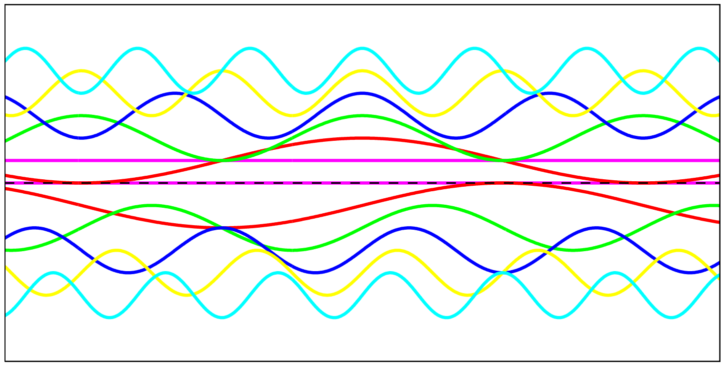

Figure 12: Harmonic functions of different frequency and phase

Square wave function

Fourier decomposition of square wave function:

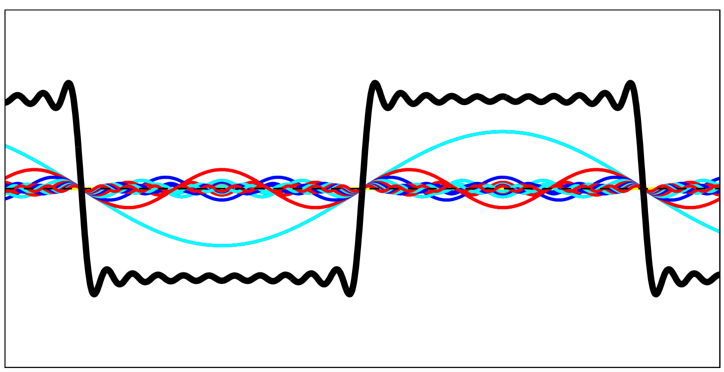

Figure 13: Fourier decomposition of square wave function

Sawtooth function

Fourier decomposition of saw tooth function:

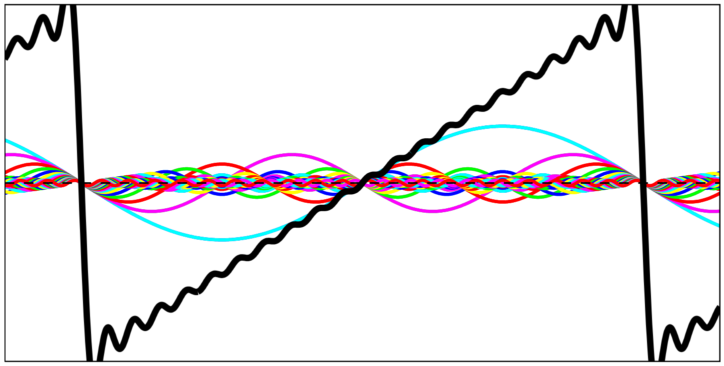

Figure 14: Fourier decomposition of saw tooth function

Impulse function

Fourier decomposition of impulse function:

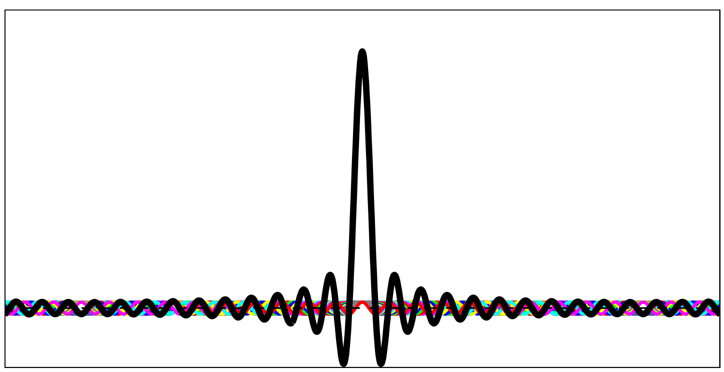

Figure 15: Fourier decomposition of impulse function

Fourier decomposition of impulse function:

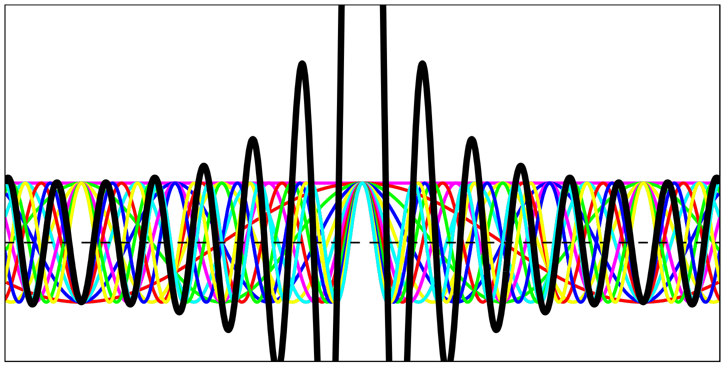

Figure 16: Fourier decomposition of impulse function

Frequency decomposition of stimulus

Stimulus is decomposed into harmonic functions of different amplitude and phase.

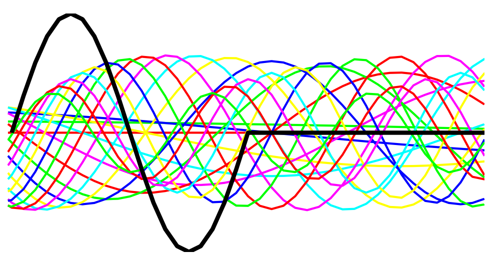

Figure 17: Frequency decomposition of stimulus

Frequency decomposition

Amplitudes and Phases of individual stimulus components:

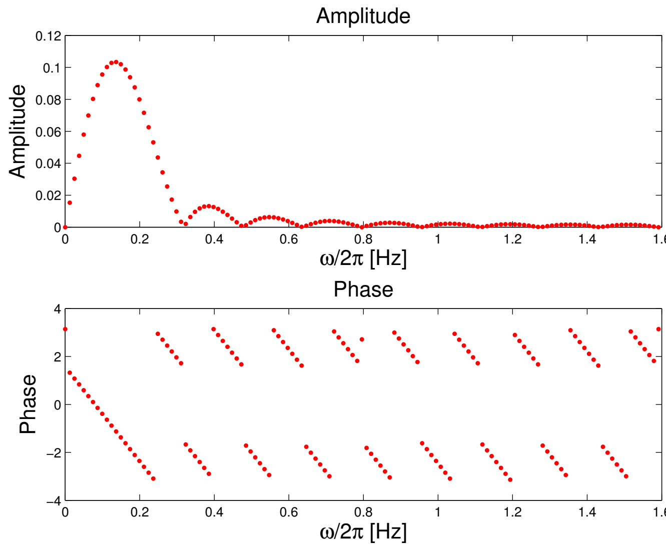

Figure 18: Amplitudes and Phases of individual stimulus components

Two ways of characterising a linear system

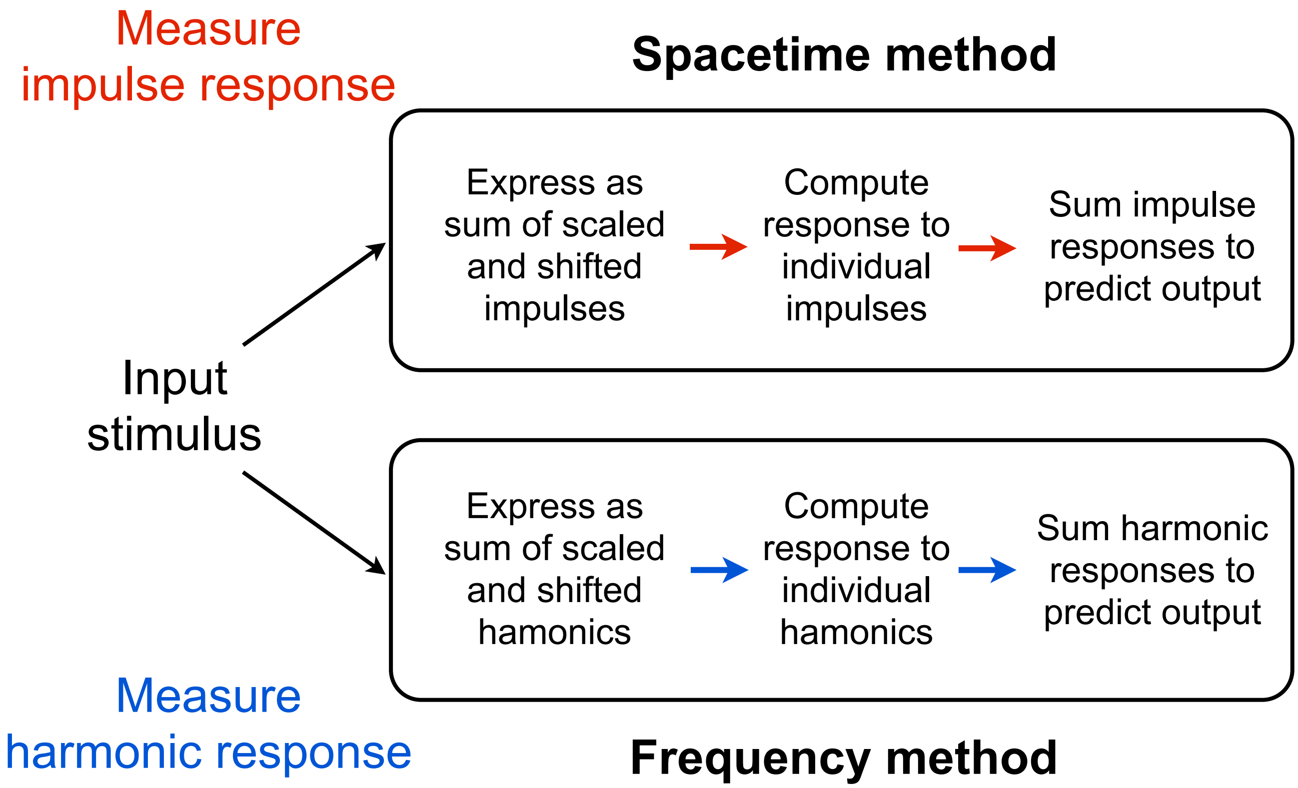

Figure 19: Characterising a linear system

Summary linear systems

- Stationary systems exhibit *shift-invariance*.

- Linear systems exhibit *homgeneity* and *additivity*.

- In stationary linear systems, responses to arbitrary stimuli can be predicted from the measured responses to a few stimuli.

- *Spacetime method:* predict responses from measured response to impulse stimulus.

- *Frequency method:* predict responses from measured responses to harmonic stimuli.

2 Waveforms in time and frequency space

An example of a sensory signal that varies in only one dimension, namely time, is the sound received by one ear. For experimental purposes, we are free to choose the waveforms that best suit our purpose.

We introduce several useful waveforms and their decomposition into harmonic functions (Fourier series)

- Delta function
- Harmonic function
- Gaussian function
- Gabor function
- White noise

Delta functions (Dirac or Kronecker function)

Consider a dry clap in an anechoic chamber. Sound pressure increases sharply and immediately returns to zero.

This waveform is approximated by a **delta function**

$$\delta(x) = \lim_{\epsilon \to 0} \begin{cases} 1/\epsilon, & |x| < \epsilon/2 \\ 0, & |x| \geq \epsilon/2 \end{cases} \qquad\qquad \int_{-\infty}^{+\infty} \delta(x)\, dx = 1$$

A **delta function** is the sum of cosine functions with equal amplitude at all frequencies. Accordingly, its Fourier decomposition is a constant.

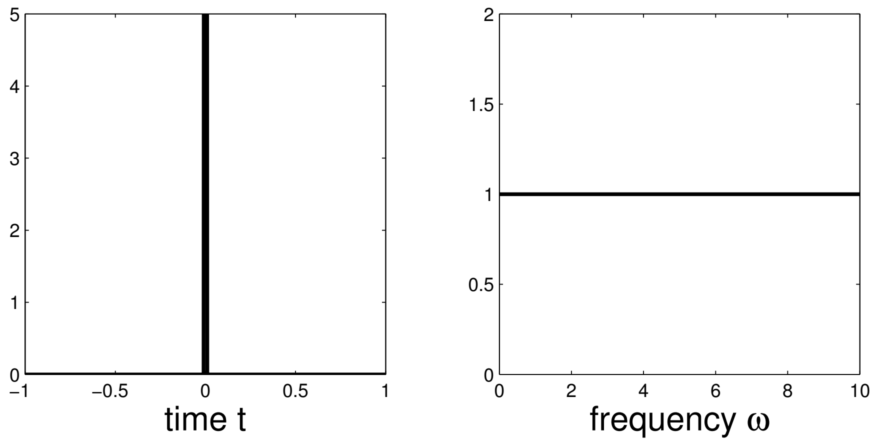

Figure 20: Left: Delta function. Right: Frequency composition.

Math aside: filtering property of delta function

The "filtering property" is an important feature of a delta function. For any continuous $f(x)$:

$$\begin{aligned}\int \delta(x-x_0)\, f(x)\, dx &= \lim_{\epsilon\to 0}\int_{x_0-\epsilon}^{x_0+\epsilon} \delta(x-x_0)\, f(x)\, dx = \\ &= f(x_0)\lim_{\epsilon\to 0}\int_{x_0-\epsilon}^{x_0+\epsilon} \delta(x-x_0)\, dx = \\ &= f(x_0)\end{aligned}$$

Harmonic functions

Consider a pure tone. Sound pressure increases and decreases regularly, following a sinus function.

$$S_{harmonic} = A\,\cos(\omega\, t + \phi)$$

Harmonic functions are parametrized by amplitude A, frequency ω and phase ϕ.

A **harmonic function** is the sum of exactly one sinus function with nonzero amplitude. Its Fourier decomposition is a delta function.

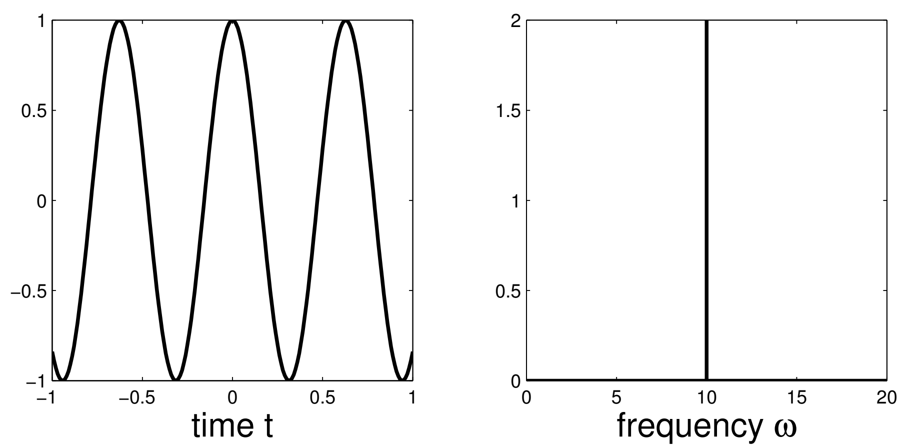

Figure 21: Harmonic function. Left: Harmonic function. Right: Frequency composition.

Gaussian functions

Consider a rapid (but not instantaneous) increase in pressure, followed by an equally rapid (but not instantaneous) decrease. Sound pressure follows a Gaussian functions:

$$S_{Gaussian} = \frac{A}{\sqrt{2\pi}\,\sigma}\,\exp\left(-\frac{(t-t_0)^2}{2\,\sigma^2}\right)$$

Gaussian functions are parameterized by an amplitude A, width σ, and peak time t_0.

A **Gaussian function** is a sum of sinus functions with comparably rapid modulation. Its Fourier decomposition is another Gaussian function.

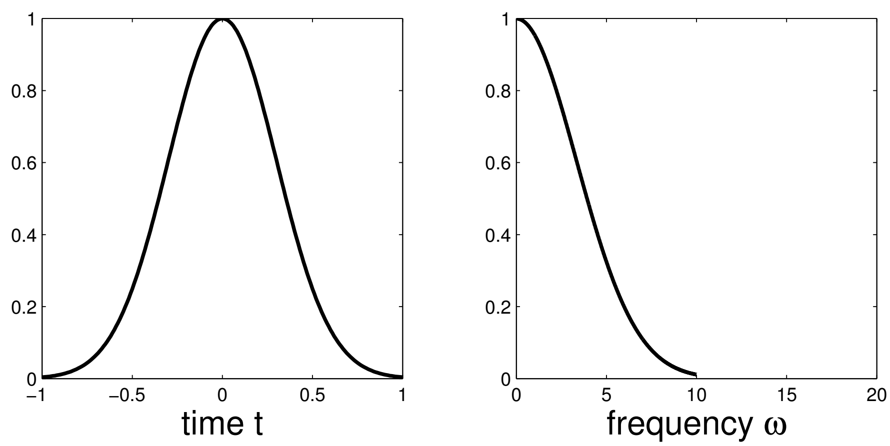

Figure 22: Gaussian function. Left: Harmonic function. Right: Frequency composition.

Gabor functions

Consider a pure tone which rapidly (but not instantaneously) grows louder and equally rapidly (but not instantaneously) fades away. Sound pressure follows a sinus function within a Gaussian envelope:

$$S_{Gabor} = \frac{A}{\sqrt{2\pi}\,\sigma}\,\exp\left(-\frac{(t-t_0)^2}{2\,\sigma^2}\right)\,\cos(\omega\,t+\phi)$$

Gabor functions are parameterized by amplitude A, width σ, frequency ω, phase ϕ, and a peak time t_0.

A **Gabor function** is a sum of sinus functions with comparably rapid modulation. Its Fourier decomposition is another Gaussian function.

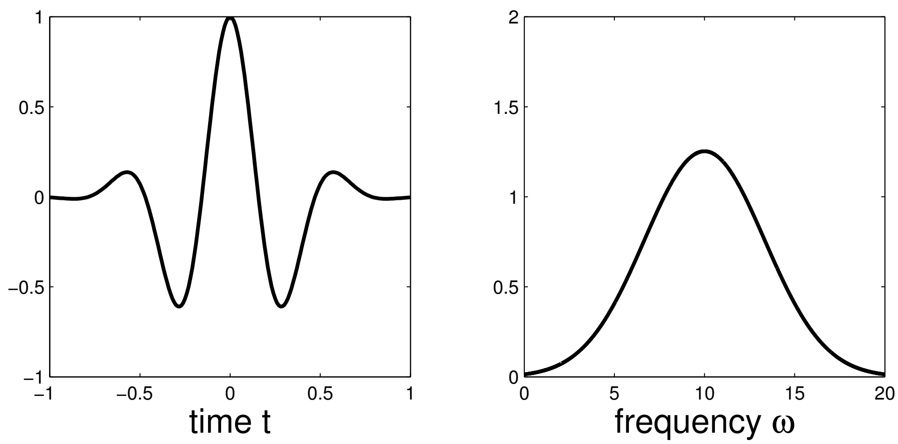

Figure 23: Gabor function. Left: Gabor functin. Right: Frequency composition.

Gabor functions in sensory biology

To understand the response of auditory neurons, we can expose the animal various kinds of test waveforms and observe the neuron's response in each case. Popular choices of test waveforms are harmonic functions (e.g., pure tones), short pulses of positive or negative pressure (e.g., narrow Gauss functions), and Gabor functions (also known as 'clicks' or 'pings'). Gabor functions are quite common in sensory biology (e.g., some bats and dolphins emit Gabor function waveforms in the context of echolocation).

Gabor-like clicks emitted by bats

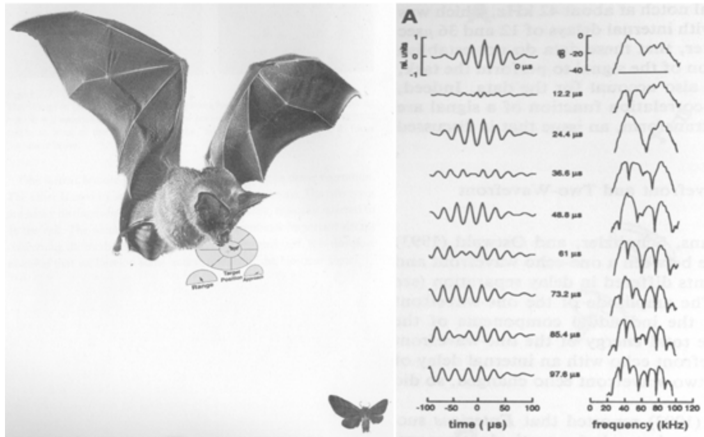

Figure 24: Gabor-like clicks emitted by bats[2]

Gabor-like clicks emitted by dolphins

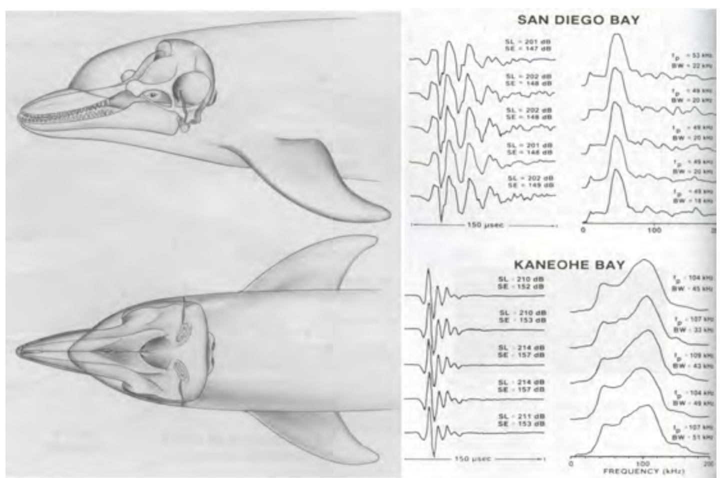

Figure 25: Gabor-like clicks emitted by dolphins[3]

White noise

Finally, we construct a **white noise** stimulus $s(t)$ from a series of discrete values s_m, each lasting for a brief interval Δt. We choose each s_m independently from a probability distribution with mean 0 and variance σ_s^2.

By construction, the average is zero and the autocorrelation is a delta function:

$$\langle s \rangle = \frac{1}{T} \int_0^T s(t)\, dt = 0, \qquad \langle s^2 \rangle = \frac{1}{T} \int_0^T s^2(t)\, dt = \sigma_s^2$$

$$Q_{ss}(\tau) = \frac{1}{T} \int_0^T s(t)\, s(t+\tau)\, dt = \sigma_s^2\, \delta(\tau)$$

For all $\tau \neq 0$, the product $s(t)\, s(t+\tau)$ is equally likely to be positive and negative! Thus, the average product is zero.

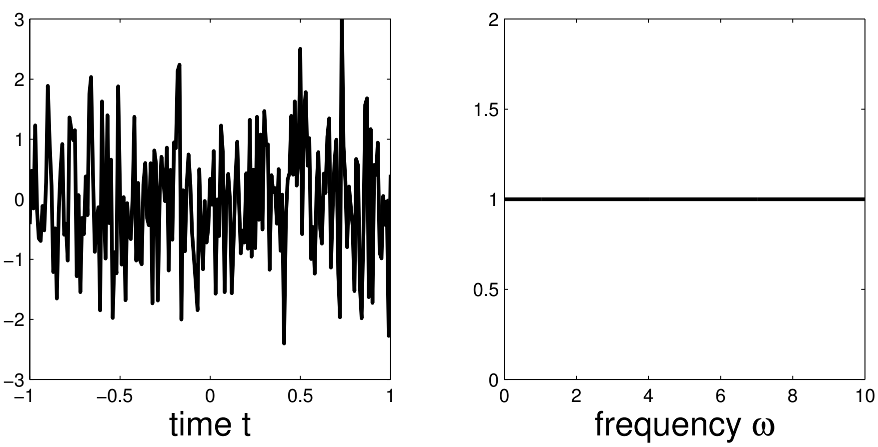

Figure 26: Left: White noise. Right: Frequency composition.

Just like a delta function, **white noise** contains sinus functions of all frequencies with equal amplitude. Its Fourier decomposition is therefore a constant (which depends on the noise variance σ^2).

Summary waveforms

- We introduce some useful waveforms, as well as their decomposition into sums of harmonics (Fourier series).

- Delta function, harmonic function, Gaussian function, Gabor function, white noise.

- Gabor functions are particularly common in sensory biology.

- Gabor functions are unique in that they are localized in both the time and the frequency domain.

3 Cross-correlation and convolution

Forward cross-correlation: Average integral of the product of stimulus $s(t)$ and response $r(t+\tau)$ as a function of time-shift τ:

$$Q_{rs}(\tau) = \frac{1}{T} \int_0^T r(t'+\tau)\, s(t')\, dt', \qquad \int_0^T s(t')\, dt' = 0$$

where τ is the delay between stimulus being presented and response being affected [1].

As long as stimuli are random (unpredictable), $Q_{rs}(\tau)$ can be non-zero only for positive τ, because responses may be affected only by past stimuli, never by future stimuli.

Choose delay τ and fix this value. Then, for all stimulus times t', multiply $s(t')$ with $r(t'+\tau)$. Finally, average all products.

Reverse cross-correlation:

Average integral of the product of stimulus $s(t-\tau)$ and response $r(t)$ as a function of time-shift τ:

$$Q_{rs}(\tau) = \frac{1}{T} \int_\tau^{T+\tau} r(t')\, s(t'-\tau)\, dt', \qquad \int_0^T s(t')\, dt' = 0$$

where τ is delay between stimulus being presented and response being affected. [1]

Choose delay τ and fix this value. Then, for all response times t', multiply $s(t'-\tau)$ with $r(t')$. Finally, average all products.

It should be clear that forward and reverse correlation compute the same function!

Convolution (dt. Faltungsintegral)

The convolution is defined for any two integrable functions $f(t)$ and $g(t)$. It is an *integral transform* in that it maps two functions onto a third:

$$f(t), g(t) \to C(t)$$

$$C(\tau) = f * g \equiv \int_{-\infty}^{+\infty} f(t)\, g(\tau - t)\, dt$$

The convolution is commutative

$$f * g = g * f = \int_{-\infty}^{+\infty} f(\tau - t')\, g(t')\, dt', \qquad t' = \tau - t$$

and can be understood as a 'weighted average'. It measures how well one function matches the time-reverse of the other.

Example: convolution with a 'balanced' kernel

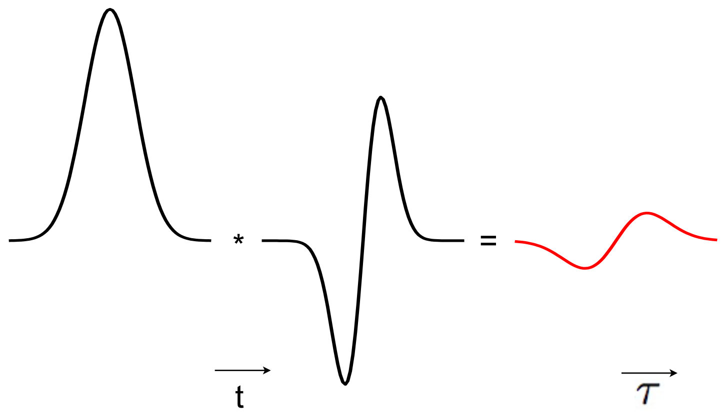

Figure 27: Convolution with a 'balanced' kernel 1

$$f(t) = \frac{1}{\sqrt{2\pi}}\, e^{-\frac{t^2}{2}}, \qquad s(t) = t\, e^{-t^2}, \qquad C(\tau) = f\, *\, s = \frac{\tau}{3\sqrt{3}}\, e^{-\frac{\tau^2}{3}}$$

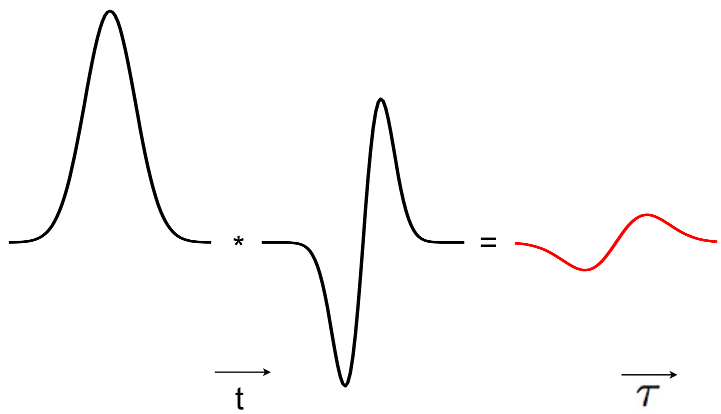

Figure 28: Convolution with a 'balanced' kernel 2

$$f(t) = \frac{1}{\sqrt{2\pi}}\, e^{-\frac{t^2}{2}} \qquad\qquad s(\tau - t) = (\tau - t)\, e^{-(\tau - t)^2}$$

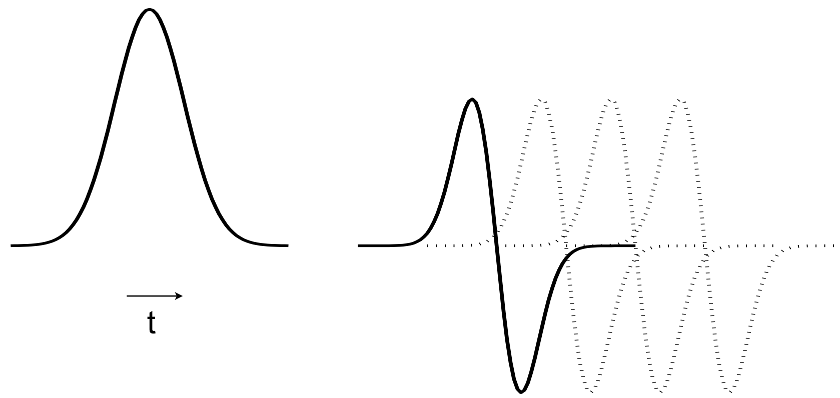

Figure 29: Convolution with a 'balanced' kernel 3

reversed *and* *shifted*

$\tau = -2$ $\tau = -1$ $\tau = 0$ $\tau = +1$ $\tau = +2$

$$f(t) = \frac{1}{\sqrt{2\pi}}\, e^{-\frac{t^2}{2}}, \qquad f(t)\, s(\tau - t) = \frac{1}{\sqrt{2\pi}}\, (\tau - t)\, e^{-(\tau - t)^2 - t^2/2}$$

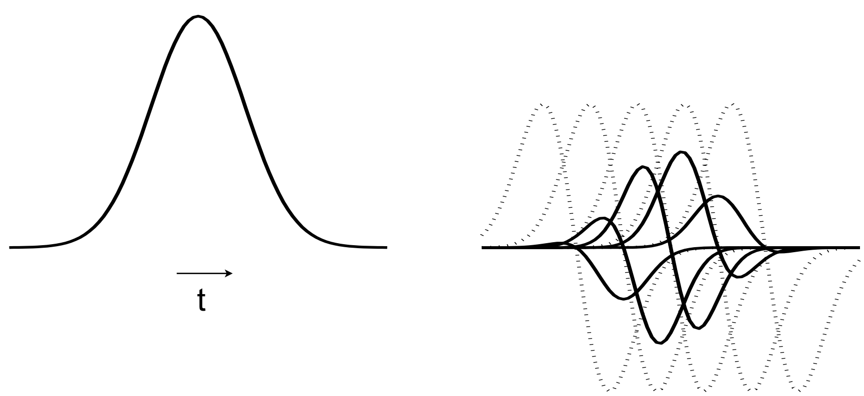

Figure 30: Convolution with a 'balanced' kernel 4

$\tau = -2$ $\tau = -1$ $\tau = 0$ $\tau = +1$ $\tau = +2$

$$C(\tau) = \int f(t)\, s(\tau - t)\, dt = \frac{\tau}{3\sqrt{3}}\, e^{-\frac{\tau^2}{3}}$$

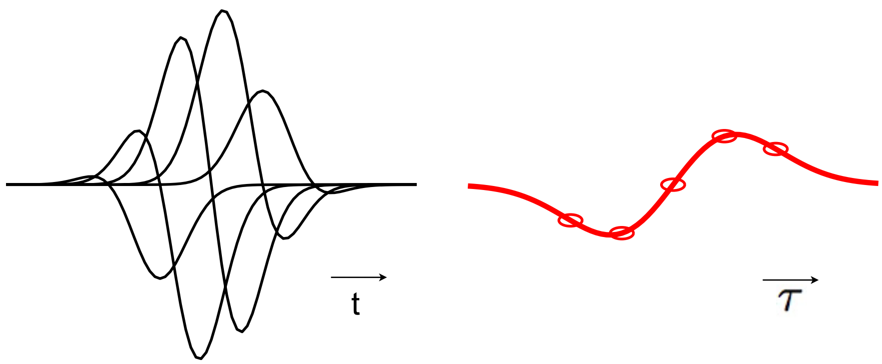

Figure 31: Convolution with a 'balanced' kernel 5

$\tau = -2 \qquad \tau = -1 \qquad \tau = 0 \qquad \tau = +1 \qquad \tau = +2$

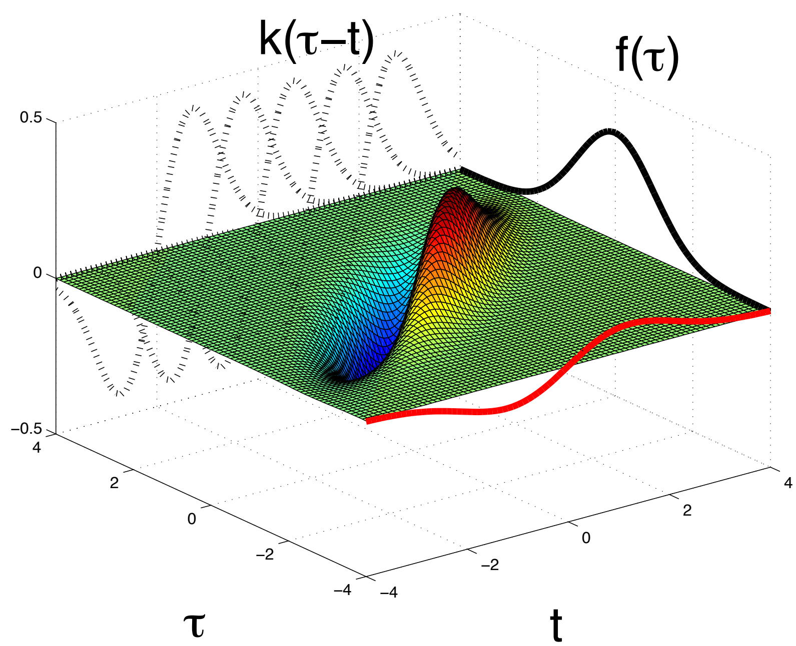

Figure 32: Convolution with a 'balanced' kernel 6

Summary cross-correlation and convolution

- The cross-correlation of time-varying stimuli & responses describes their covariance as a function of delay.
- It measures the effect of earlier stimuli on later responses.

- A convolution maps two integrable functions $f(t)$ and $g(t)$ onto a third function $C(\tau)$, where τ is a time-shift.
- It measures how well (or poorly) the two functions match, as a function of time-shift.

- Linear-filter models of neural responses are based on convolutions.

4 Linear-filter model of neural responses (advanced)

We use one-dimensional stimuli to introduce 'linear filter' models of neural responses. Examples of such a stimuli are soundwaves (audition), tactile vibrations (somatosensation), or time-varying potentials (electrosensation). The next lecture will extend this concept to two- and three-dimensional stimuli (vision).

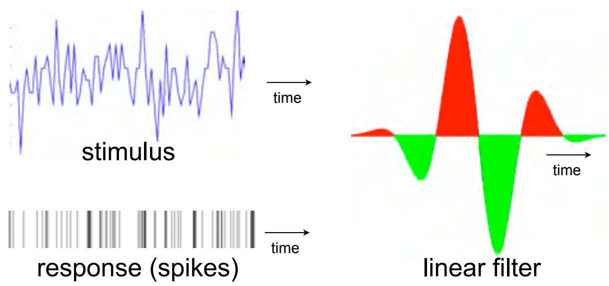

Figure 33: Linear filter

We will consider stimuli $s(t)$ with zero-mean (to account for adaptation).

Linear kernel $D(\tau)$

We model a neuronal response as a linear filtering operation. Response $L(t)$ is defined as the *convolution* of the stimulus $s(t)$ with a *linear filter* or *linear kernel* $D(\tau)$, where τ is 'stimulus-response latency'. $D(\tau)$ describes how strongly a stimulus at time $t - \tau$ affects the response at time t:

$$L(t) = \int_0^\infty s(t - \tau)\, D(\tau)\, d\tau$$

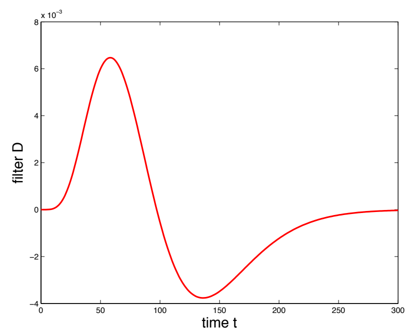

Figure 34: Convolution of the stimulus with a linear filter

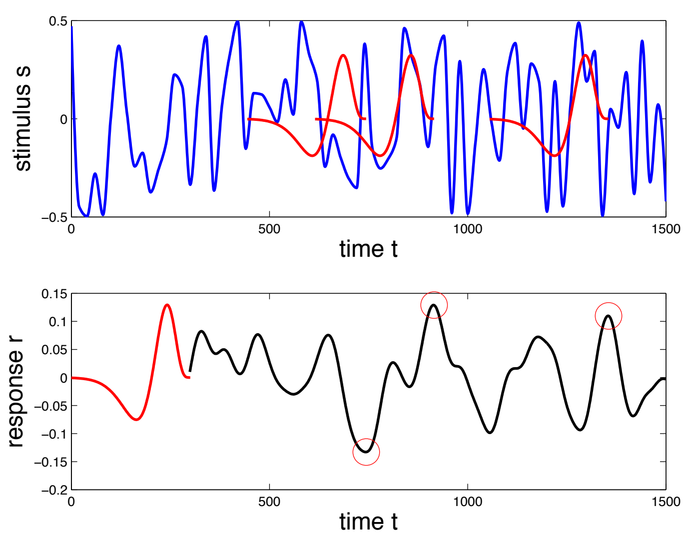

Figure 35: Large and positive model response

The model response is large and positive when the *time-reversed kernel* matches the *positive stimulus* well.

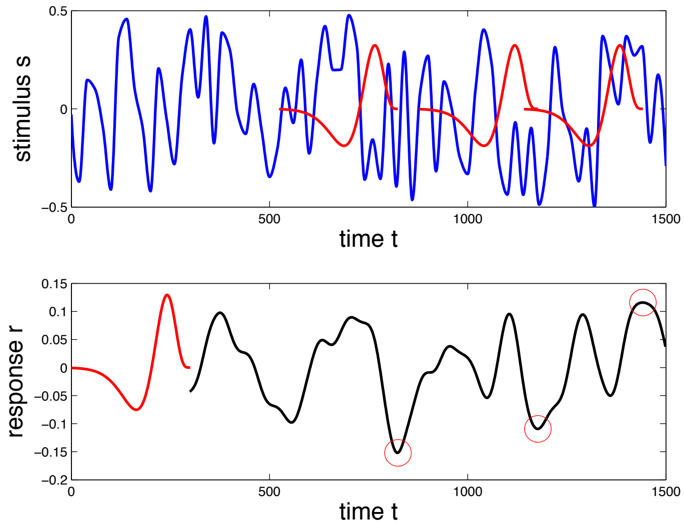

Figure 36: Large and negative model response

It is large and negative when the *time-reversed kernel* matches the *negative stimulus* well.

Optimal linear filter

A linear-filter model provides an estimate $r_{est}(t)$ of a neuron's actual response $r(t)$.

$$r_{est}(t) = r_0 + \int_0^\infty s(t-\tau)\, D(\tau)\, d\tau$$

The instantaneous error $E(t)$ and the mean square error MSE of this estimate are:

$$E(t) \;= r_{est}(t) - r(t)$$

$$MSE \;= \tfrac{1}{T} \int E^2(t)\, dt$$

Which filter $D(\tau)$ is **optimal** in the sense that it minimizes the MSE? The optimal filter $D(\tau)$ satisfies ('functional derivative')

$$\frac{\delta MSE}{\delta D} = 0$$

This leads to an integral equation (D&A, Chapter 2, Appendix A) with the stimulus-response correlation Q_{rs} and the stimulus autocorrelation

Q_{ss}:

$$\int_0^\infty Q_{ss}(\tau - \tau')\, D(\tau')\, d\tau = Q_{rs}(-\tau)$$

$$\begin{aligned} Q_{rs}(\tau) &= \frac{1}{T} \int r(t)\, s(t+\tau)\, dt \\ Q_{ss}(\tau) &= \frac{1}{T} \int s(t)\, s(t+\tau)\, dt \end{aligned}$$

Note that $Q_{rs}(-\tau)$ is evaluated at $-\tau$ and a *reverse* correlation.
The integral equation is easy to solve for white noise stimuli:

$$Q_{ss}(\tau) = \sigma_s^2\, \delta(\tau)$$

$$\begin{aligned} \int_0^\infty Q_{ss}(\tau - \tau')\, D(\tau')\, d\tau' &= Q_{rs}(-\tau) \\ \sigma_s^2 \int_0^\infty \delta(\tau - \tau')\, D(\tau')\, d\tau' &= Q_{rs}(-\tau) \\ \sigma_s^2\, D(\tau) &= Q_{rs}(-\tau) \\ D(\tau) &= \frac{Q_{rs}(-\tau)}{\sigma_s^2} \end{aligned}$$

The **optimal linear filter** is the **reverse correlation** with a white noise stimulus!

Optimal kernel and preferred stimulus

- Reverse correlation between response and preceding stimulus yields *response-weighted average stimulus*.

- If linearity assumption holds, this is also the *optimal linear kernel* (best fit to actual response).

- Additionally, it is also the *preferred stimulus* (stimulus evoking maximal response)!

- To see this, compare stimuli of identical power and position ($t = 0$)!

$$\int_{-\infty}^{\infty} s^2(\tau)\, d\tau = \int_{-\infty}^{\infty} D^2(\tau)\, d\tau$$

Stimuli with different phase

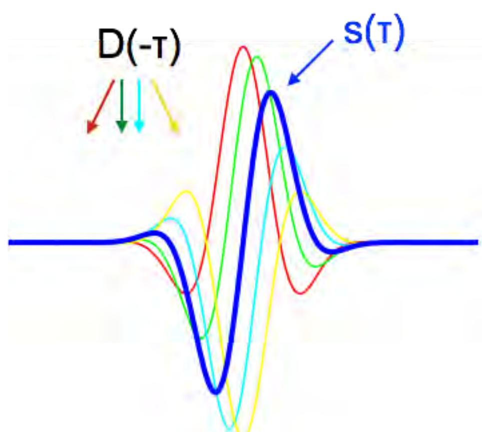

Figure 37: Stimuli with different phase

$\int_{-\infty}^{\infty} s(-\tau)\, D(\tau)\, d\tau$ is maximal for $s(\tau) = D(-\tau)$.

Stimuli with different width

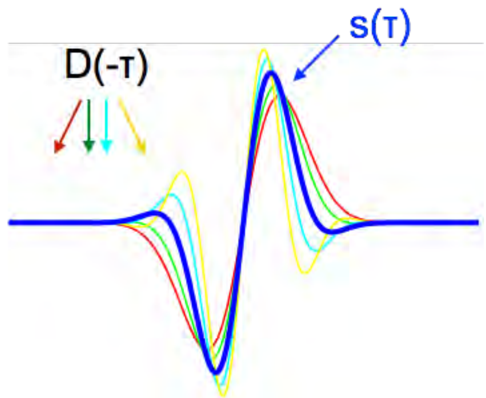

Figure 38: Stimuli with different width

$\int_{-\infty}^{\infty} s(-\tau)\, D(\tau)\, d\tau$ is maximal for $s(\tau) = D(-\tau)$.

Stimuli with different periodicity

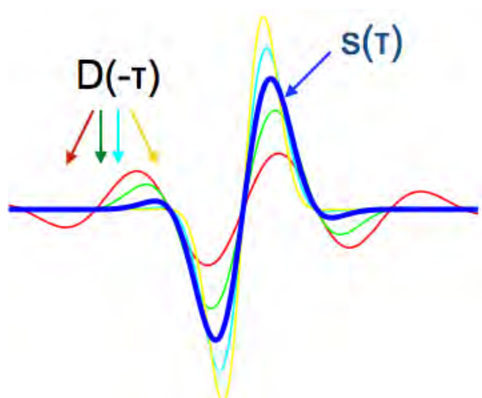

Figure 39: Stimuli with different periodicity

$\int_{-\infty}^{\infty} s(-\tau)\, D(\tau)\, d\tau$ is maximal for $s(\tau) = D(-\tau)$.

- We have compared Gabor stimuli of different phase, width, and periodicity (but identical power and position).
- The convolution between simulus and *optimal linear kernel* is maximal when the stimulus matches the time-reversed kernel!
- *Preferred stimulus* is best possible match of the *optimal kernel* and, thus, identical to the time-reversed kernel.

Summary linear filter model (LFM)

- A linear-filter model estimates the response to a stimulus $s(t)$ as its convolution with a *linear kernel* $D(\tau)$:

$$r_{est}(t) = r_0 + \int_0^{\infty} s(t-\tau)\, D(\tau)\, d\tau$$

- A LFM predicts a maximally positive response when the stimulus matches the time-reversed linear kernel.
- A LFM predicts a maximally negative response when the stimulus matches the time-reversed and *inverted* linear kernel.
- The optimal linear kernel is the reverse correlation between response and white noise stimulus.
- The preferred stimulus matches the optimal linear kernel.
- Theoretical justification for reverse correlation methods!

5 Non-linear parts of LFM: 'static non-linearities'

Linear-filter models of neuronal responses have evident problems:

- LFM predicts both positive and negative responses. Neuronal responses are only positive.
- LFM predicts that response amplitude grows linearly with stimulus intensity/amplitude. Neuronal responses saturate.

Both problems are easy to fix!

Static non-linearity
We introduce a non-linear function $F()$ to link the linear response $L(t)$ to the predicted response $r_{est}(t)$:

$$r_{est}(t) = F\left[L(t)\right] \qquad \text{where} \qquad L(t) = \int s(t-\tau)\, D(\tau)\, d\tau$$

The non-linear function $F()$ is called a *static non-linearity* because it is evaluated instantaneously at each time t, with no explicit dependence on time. Thus $F()$ is 'memoryless'.

Pattern matching between $s(t)$ and $D(\tau)$ remains *linear* !

Scaling of response with stimulus amplitude is *not linear* !

Fitting functional form

The functional form of a static nonlinearity is typically determined empirically, by comparing linear prediction with actual response.

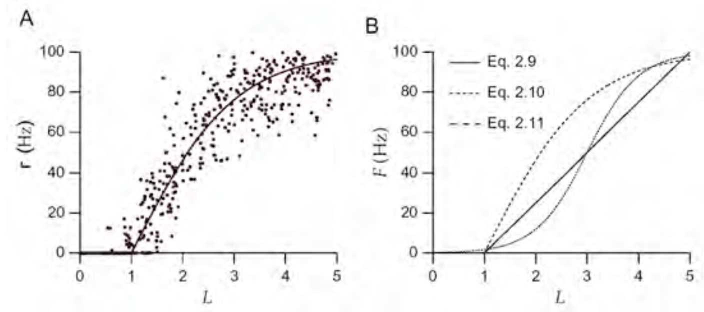

Figure 40: Comparing linear prediction with actual response

Various functions are used to approximate static nonlinearities. All are *rectifying* functions, in that they suppress negative responses!

$$F(L) = k\ [L]_+ \qquad \text{rectifying}$$

$$F(L) = k\ [L - L_0]_+ \qquad +\text{threshold}$$

$$F(L) = \frac{r_{max}}{1 + \exp\left[-k(L - L_{1/2})\right]} \qquad \text{logistic}$$

$$F(L) = r_{max}\ [\tanh\left(k(L - L_0)\right)]_+ \qquad +\text{threshold}$$

$$F(L) = \frac{r_{max}\, k\ [L]_+^2}{A_{1/2}^2 + k\ [L]_+^2} \qquad \text{Naka} - \text{Rushton}$$

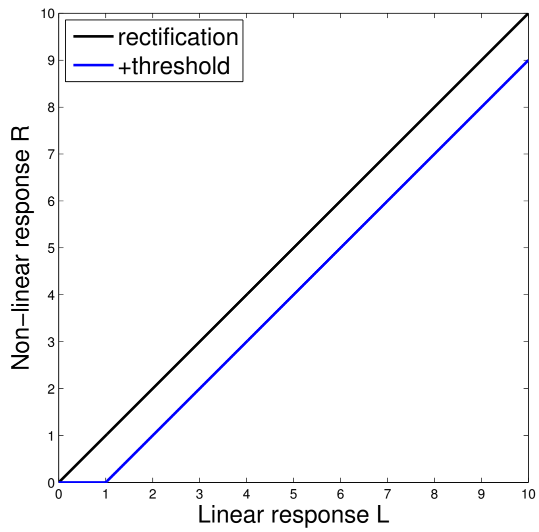

Figure 41: Fitting functional form 1

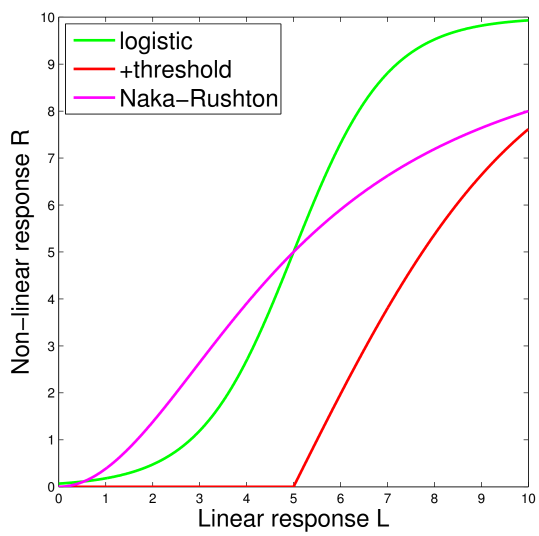

Figure 42: Fitting functional form 2

Example: lateral geniculate neuron

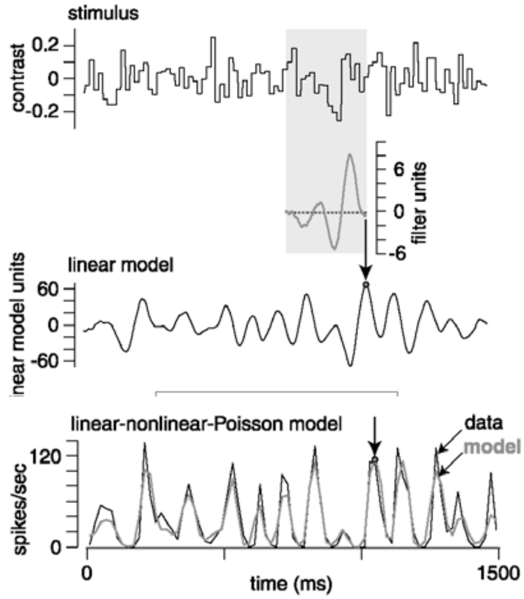

Figure 43: Lateral geniculate neuron 1

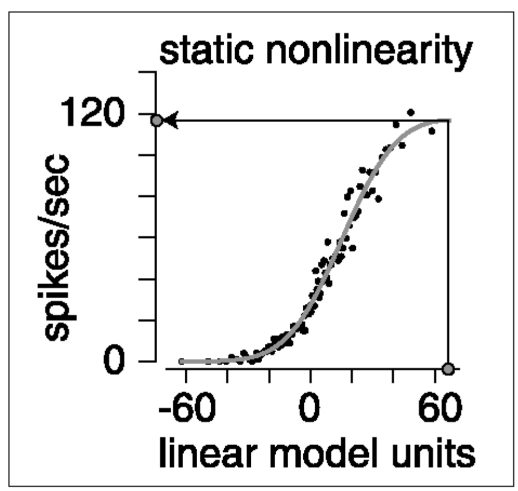

Figure 44: Static nonlinearity

Summary static non-linearities

- A complete LFM consists of one (or more) linear filter(s) followed by a static-nonlinearity.
- The static non-linearity converts the linear convolution result (positive or negative, non-saturating) into a neuronal response (positive and saturating).
- LFMs capture many aspects of the response properties of sensory neurons (e.g., “simple cells” in visual cortex).
- Extending LFMs to include multiple linear filters captures the response properties of even more sensory neurons (e.g., “complex cells” in visual cortex).
- Further extensions are needed to model responses to all stimuli (including natural stimuli).

6 Bibliography

1. Is it possible to heal an eye injury/ defect/ or Blindness by doing surgery on a specific brain region(s)? Quora. Brandon Mclaughlin, 2014. Figure 1. Ref: https://www.quora.com/Is-it-possible-to-heal-an-eye-injury-defect-or-Blindness-by-doing-surgery-on-a-specific-brain-region-s

2. Howard C. Hughes. Sensory exotica. MIT Press 1999

3. Howard C. Hughes. Sensory exotica. MIT Press 1999