

Signals, Linear Systems, and Convolution

Professor David Heeger

September 23, 1997

Characterizing the complete input-output properties of a system by exhaustive measurement is usually impossible. Instead, we must find some way of making a finite number of measurements that allow us to infer how the system will respond to other inputs that we have not yet measured. We can only do this for certain kinds of systems with certain properties. If we have the right kind of system, we can save a lot of time and energy by using the appropriate theory about the system's responsiveness. Linear systems theory is a good time-saving theory for *linear systems* which obey certain rules. Not all systems are linear, but many important ones are. When a system qualifies as a linear system, it is possible to use the responses to a small set of inputs to predict the response to any possible input. This can save the scientist enormous amounts of work, and makes it possible to characterize the system completely.

To get an idea of what linear systems theory is good for, consider some of the things in neuroscience that can be successfully modeled (at least, approximately) as shift-invariant, linear systems:

System	Input	Output
passive neural membrane	injected current	membrane potential
synapse	pre-synaptic action potentials	post-synaptic conductance
cochlea	sound	cochlear microphonic
optics of the eye	visual stimulus	retinal image
retinal ganglion cell	stimulus contrast	firing rate
human	pairs of color patches	color match settings

In addition, a number of neural systems can be approximated as linear systems coupled with simple nonlinearities (e.g., a spike threshold).

The aim of these notes is to clarify the meaning of the phrase: "The effect of any shift-invariant linear system on an arbitrary input signal is obtained by convolving the input signal with the system's impulse response function."

Most of the effort is simply definitional - you have to learn the meaning of technical terms such as "linear", "convolve", and so forth. We will also introduce some convenient mathematical notation, and we will describe two different approaches for measuring the system's impulse response function.

For more detailed introductions to the material covered in this handout, see Oppenheim, Wilsky, and Young (1983), and Oppenheim and Schaffer (1989).

Continuous-Time and Discrete-Time Signals

In each of the above examples there is an input and an output, each of which is a time-varying signal. We will treat a signal as a time-varying function, $x(t)$. For each time t , the signal has some value $x(t)$, usually called “ x of t .” Sometimes we will alternatively use $x(t)$ to refer to the entire signal x , thinking of t as a free variable.

In practice, $x(t)$ will usually be represented as a finite-length sequence of numbers, $x[n]$, in which n can take integer values between 0 and $N - 1$, and where N is the length of the sequence. This discrete-time sequence is indexed by integers, so we take $x[n]$ to mean “the n th number in sequence x ,” usually called “ x of n ” for short.

The individual numbers in a sequence $x[n]$ are called *samples* of the signal $x(t)$. The word “sample” comes from the fact that the sequence is a discretely-sampled version of the continuous signal. Imagine, for example, that you are measuring membrane potential (or just about anything else, for that matter) as it varies over time. You will obtain a sequence of measurements sampled at evenly spaced time intervals. Although the membrane potential varies continuously over time, you will work just with the sequence of discrete-time measurements.

It is often mathematically convenient to work with continuous-time signals. But in practice, you usually end up with discrete-time sequences because: (1) discrete-time samples are the only things that can be measured and recorded when doing a real experiment; and (2) finite-length, discrete-time sequences are the only things that can be stored and computed with computers.

In what follows, we will express most of the mathematics in the continuous-time domain. But the examples will, by necessity, use discrete-time sequences.

Pulse and impulse signals. The unit impulse signal, written $\delta(t)$, is one at $t = 0$, and zero everywhere else:

$$\delta(t) = \begin{cases} 1 & \text{if } t = 0 \\ 0 & \text{otherwise} \end{cases}$$

The impulse signal will play a very important role in what follows.

One very useful way to think of the impulse signal is as a limiting case of the pulse signal, $\delta_\Delta(t)$:

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

The impulse signal is equal to the pulse signal when the pulse gets infinitely short:

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

Unit step signal. The unit step signal, written $u(t)$, is zero for all times less than zero, and 1 for all times greater than or equal to zero:

$$u(t) = \begin{cases} 0 & \text{if } t < 0 \\ 1 & \text{if } t \geq 0 \end{cases}$$

Summation and integration. The Greek capital sigma, \sum , is used as a shorthand notation for adding up a set of numbers, typically having some variable take on a specified set of values. Thus:

$$\sum_{i=1}^5 i = 1 + 2 + 3 + 4 + 5$$

The \sum notation is particularly helpful in dealing with sums over discrete-time sequences:

$$\sum_{n=1}^3 x[n] = x[1] + x[2] + x[3].$$

An integral is the limiting case of a summation:

$$\int_{t=-\infty}^{\infty} x(t)dt = \lim_{\Delta \rightarrow 0} \sum_{k=-\infty}^{\infty} x(k\Delta)\Delta$$

For example, the step signal can be obtained as an integral of the impulse:

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

Up to $s < 0$ the sum will be 0 since all the values of $\delta(s)$ for negative s are 0. At $t = 0$ the cumulative sum jumps to 1 since $\delta(0) = 1$. And the cumulative sum stays at 1 for all values of t greater than 0 since all the rest of the values of $\delta(t)$ are 0 again.

This is not a particularly impressive use of an integral, but it should help to remind you that it is perfectly sensible to talk about infinite sums.

Arithmetic with signals. It is often useful to apply the ordinary operations of arithmetic to signals. Thus we can write the product of signals x and y as $z = xy$, meaning the signal made up of the products of the corresponding elements:

$$z(t) = x(t)y(t)$$

Likewise the sum of signals x and y can be written $z = x + y$. A signal x can be multiplied by a scalar α , meaning that each element of x is individually so multiplied. Finally, a signal may be shifted by any amount:

$$z(t) = x(t - s).$$




Figure 1: Staircase approximation to a continuous-time signal.

Representing signals with impulses. Any signal can be expressed as a sum of scaled and shifted unit impulses. We begin with the pulse or “staircase” approximation $\tilde{x}(t)$ to a continuous signal $x(t)$, as illustrated in Fig. 1. Conceptually, this is trivial: for each discrete sample of the original signal, we make a pulse signal. Then we add up all these pulse signals to make up the approximate signal. Each of these pulse signals can in turn be represented as a standard pulse scaled by the appropriate value and shifted to the appropriate place. In mathematical notation:

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

As we let Δ approach zero, the approximation $\tilde{x}(t)$ becomes better and better, and in the limit equals $x(t)$. Therefore,

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

Also, as $\Delta \rightarrow 0$, the summation approaches an integral, and the pulse approaches the unit impulse:

$$x(t) = \int_{-\infty}^{\infty} x(s) \delta(t - s) ds. \quad (1)$$

In other words, we can represent any signal as an infinite sum of shifted and scaled unit impulses. A digital compact disc, for example, stores whole complex pieces of music as lots of simple numbers representing very short impulses, and then the CD player adds all the impulses back together one after another to recreate the complex musical waveform.

This no doubt seems like a lot of trouble to go to, just to get back the same signal that we originally started with, but in fact, we will very shortly be able to use Eq. 1 to perform a marvelous trick.

Linear Systems

A *system* or *transform* maps an input signal $x(t)$ into an output signal $y(t)$:

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

where T denotes the transform, a function from input signals to output signals.

Systems come in a wide variety of types. One important class is known as *linear systems*. To see whether a system is linear, we need to test whether it obeys certain rules that all linear systems obey. The two basic tests of linearity are homogeneity and additivity.

Homogeneity. As we increase the strength of the input to a linear system, say we double it, then we predict that the output function will also be doubled. For example, if the current injected to a passive neural membrane is doubled, the resulting membrane potential fluctuations will double as well. This is called the *scalar rule* or sometimes the *homogeneity* of linear systems.

Additivity. Suppose we measure how the membrane potential fluctuates over time in response to a complicated time-series of injected current $x_1(t)$. Next, we present a second (different) complicated time-series $x_2(t)$. The second stimulus also generates fluctuations in the membrane potential which we measure and write down. Then, we present the sum of the two currents $x_1(t) + x_2(t)$ and see what happens. Since the system is linear, the measured membrane potential fluctuations will be just the sum of the fluctuations to each of the two currents presented separately.

Superposition. Systems that satisfy both homogeneity and additivity are considered to be linear systems. These two rules, taken together, are often referred to as the *principle of superposition*. Mathematically, the principle of superposition is expressed as:

$$T(\alpha x_1 + \beta x_2) = \alpha T(x_1) + \beta T(x_2) \quad (2)$$

Homogeneity is a special case in which one of the signals is absent. Additivity is a special case in which $\alpha = \beta = 1$.

Shift-invariance. Suppose that we inject a pulse of current and measure the membrane potential fluctuations. Then we stimulate again with a similar pulse at a different point in time, and again we measure the membrane potential fluctuations. If we haven't damaged the membrane with the first impulse then we should expect that the response to the second pulse will be the same as the response to the first pulse. The only difference between them will be that the second pulse has occurred later in time, that is, it is shifted in time. When the responses to the identical stimulus presented shifted in time are the same, except for the corresponding shift in time, then we have a special kind of linear system called a *shift-invariant* linear system. Just as not all systems are linear, not all linear systems are shift-invariant.

In mathematical language, a system T is shift-invariant if and only if:

$$y(t) = T[x(t)] \quad \text{implies} \quad y(t - s) = T[x(t - s)] \quad (3)$$

Convolution

Homogeneity, additivity, and shift invariance may, at first, sound a bit abstract but they are very useful. To characterize a shift-invariant linear system, we need to measure only one

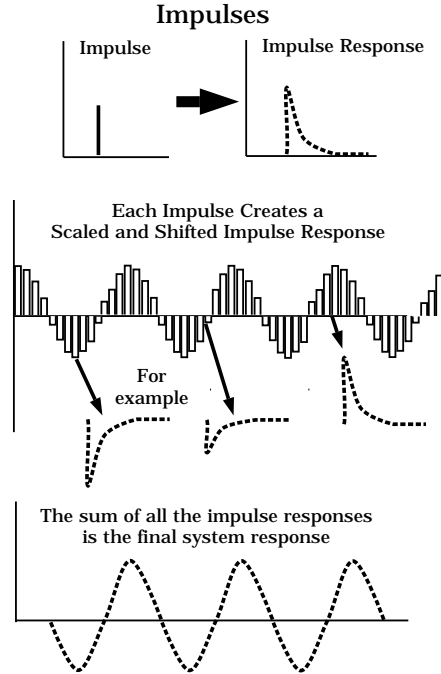


Figure 2: Characterizing a linear system using its impulse response.

thing: the way the system responds to a unit impulse. This response is called *the impulse response function* of the system. Once we've measured this function, we can (in principle) predict how the system will respond to any other possible stimulus.

The way we use the impulse response function is illustrated in Fig. 2. We conceive of the input stimulus, in this case a sinusoid, as if it were the sum of a set of impulses (Eq. 1). We know the responses we would get if each impulse was presented separately (i.e., scaled and shifted copies of the impulse response). We simply add together all of the (scaled and shifted) impulse responses to predict how the system will respond to the complete stimulus.

Now we will repeat all this in mathematical notation. Our goal is to show that the response (e.g., membrane potential fluctuation) of a shift-invariant linear system (e.g., passive neural membrane) can be written as a sum of scaled and shifted copies of the system's impulse response function.

The convolution integral. Begin by using Eq. 1 to replace the input signal $x(t)$ by its representation in terms of impulses:

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

Using additivity,

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

Taking the limit,

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

Using homogeneity,

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

Now let $h(t)$ be the response of T to the unshifted unit impulse, i.e., $h(t) = T[\delta(t)]$. Then by using shift-invariance,

$$y(t) = \int_{-\infty}^{\infty} x(s) h(t - s) ds. \quad (4)$$

Notice what this last equation means. For any shift-invariant linear system T , once we know its impulse response $h(t)$ (that is, its response to a unit impulse), we can forget about T entirely, and just add up scaled and shifted copies of $h(t)$ to calculate the response of T to any input whatsoever. Thus any shift-invariant linear system is completely characterized by its impulse response $h(t)$.

The way of combining two signals specified by Eq. 4 is known as *convolution*. It is such a widespread and useful formula that it has its own shorthand notation, $*$. For any two signals x and y , there will be another signal z obtained by *convolving* x with y ,

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

Convolution as a series of weighted sums. While superposition and convolution may sound a little abstract, there is an equivalent statement that will make it concrete: a system is a shift-invariant, linear system if and only if the responses are *a weighted sum of the inputs*. Figure 3 shows an example: the output at each point in time is computed simply as a weighted sum of the inputs at recently past times. The choice of *weighting function* determines the behavior of the system. Not surprisingly, the weighting function is very closely related to the impulse response of the system. In particular, the impulse response and the weighting function are time-reversed copies of one another, as demonstrated in the top part of the figure.

Properties of convolution. The following things are true for convolution in general, as you should be able to verify for yourself with some algebraic manipulation:

$$\begin{array}{ll} x * y = y * x & \text{commutative} \\ (x * y) * z = x * (y * z) & \text{associative} \\ (x * z) + (y * z) = (x + y) * z & \text{distributive} \end{array}$$

Fourier Series Approximations

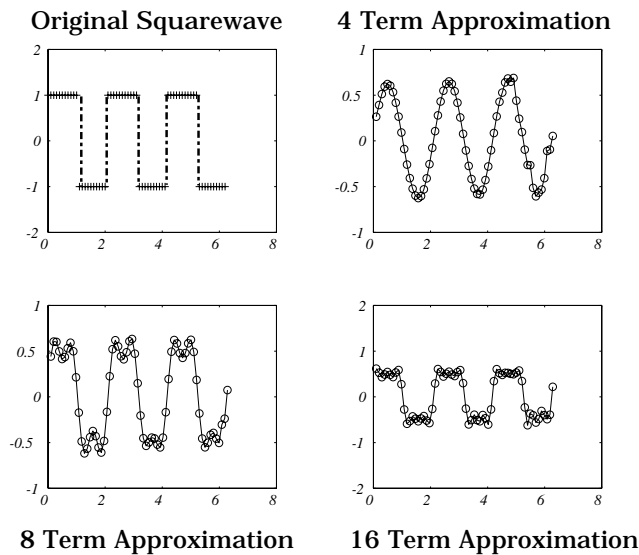


Figure 4: Fourier series approximation of a squarewave as the sum of sinusoids.

where ω is the frequency of each sinusoid, and A_ω and ϕ_ω are the amplitude and phase, respectively, of each sinusoid. You can go both ways. If you know the coefficients, A_ω and ϕ_ω , you can use this formula to reconstruct the original signal $x(t)$. If you know the signal, you can compute the coefficients by a method called the Fourier transform, a way of decomposing a complex stimulus into its component sinusoids (see Appendix II).

Example: Stereos as shift-invariant systems. Many people find the characterization in terms of frequency response to be intuitive. And most of you have seen graphs that describe performance this way. Stereo systems, for example, are pretty good shift-invariant linear systems. They can be evaluated by measuring the signal at different frequencies. And the stereo controls are designed around the frequency representation. Adjusting the bass alters the level of the low frequency components, while adjusting the treble alters the level of the high frequency components. Equalizers divide up the signal band into many frequencies and give you finer control.

Shift-invariant linear systems and sinusoids. The Fourier decomposition is important because if we know the response of the system to sinusoids at many different frequencies, then we can use the same kind of trick we used with impulses to predict the response via the impulse response function. First, we measure the system's response to sinusoids of all different frequencies. Next, we take our input (e.g., time-varying current) and use the Fourier transform to compute the values of the Fourier coefficients. At this point the input has been broken down as the sum of its component sinusoids. Finally, we can predict the system's response (e.g., membrane potential fluctuations) simply by adding the responses for all the

Shift-Invariant Linear Systems and Sinusoids

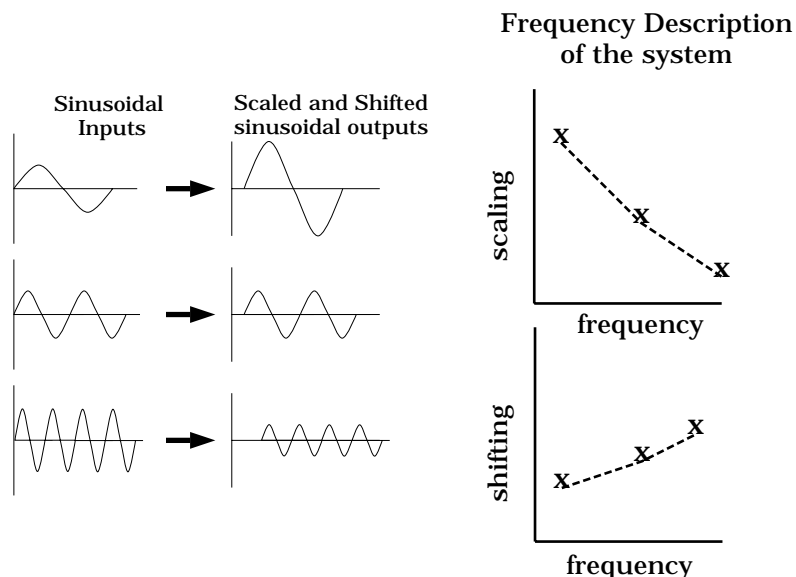


Figure 5: Characterizing a system using its frequency response.

component sinusoids.

Why bother with sinusoids when we were doing just fine with impulses? The reason is that sinusoids have a very special relationship to shift-invariant linear systems. When we use a sinusoids as the inputs to a shift-invariant linear system, the system's responses are always (shifted and scaled) copies of the input sinusoids. That is, when the input is $x(t) = \sin(2\pi\omega t)$ the output is always of the form $y(t) = A_\omega \sin(2\pi\omega t + \phi_\omega)$, with the same frequency as the input. Here, ϕ_ω determines the amount of shift and A_ω determines the amount of scaling. Thus, measuring the response to a sinusoid for a shift-invariant linear system entails measuring only two numbers: the shift and the scale. This makes the job of measuring the response to sinusoids at many different frequencies quite practical.

Often then, when scientists characterize the response of a system they will not tell you the impulse response. Rather, they will give you the *frequency response*, the values of the shift and scale for each of the possible input frequencies (Fig. 5). This frequency response representation of how the shift-invariant linear system behaves is equivalent to providing you with the impulse response function (in fact, the two are Fourier transforms of one another). We can use either to compute the response to any input. This is the main point of all this stuff: a simple, fast, economical way to measure the responsiveness of complex systems. If you know the coefficients of response for sine waves at all possible frequencies, then you can determine how the system will respond to any possible input.

Linear filters. Shift-invariant linear systems are often referred to as *linear filters* because they typically attenuate (filter out) some frequency components while keeping others

Linear Systems Logic

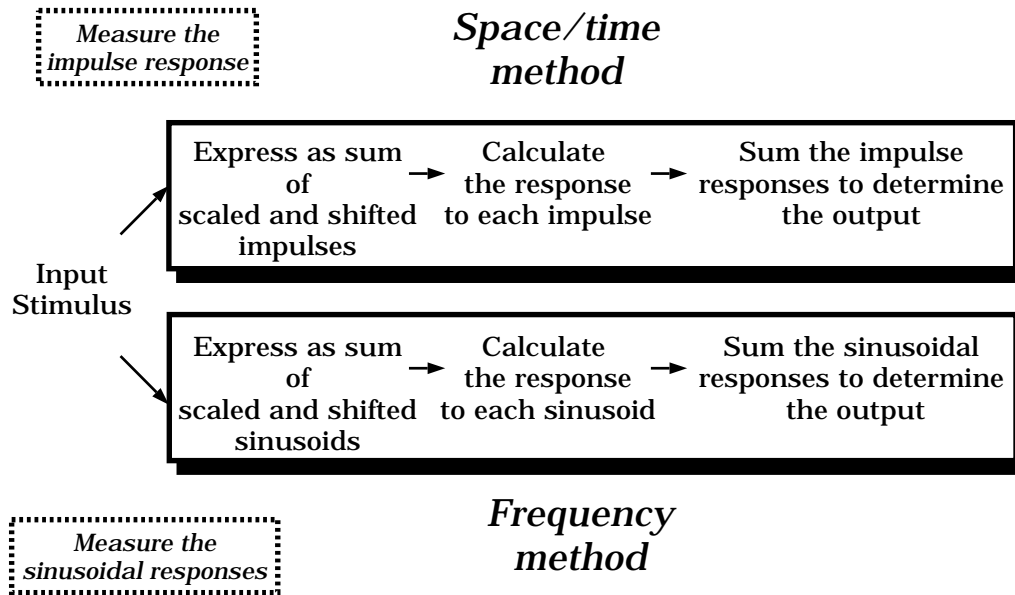


Figure 6: Alternative methods of characterizing a linear system.

intact.

For example, since a passive neural membrane is a shift invariant linear system, we know that injecting sinusoidally modulating current yields membrane potential fluctuations that are sinusoidal with the same frequency (sinusoid in, sinusoid out). The amplitude and phase of the output sinusoid depends on the choice of frequency relative to the properties of the membrane. The membrane essentially averages the input current over a period of time. For very low frequencies (slowly varying current), this averaging is irrelevant and the membrane potential fluctuations follow the injected current. For high frequencies, however, even a large amplitude sinusoidal current modulation will yield no membrane potential fluctuations. The membrane is called a *low-pass filter*: it lets low frequencies pass, but because of its time-averaging behavior, it attenuates high frequencies.

Figure 7 shows an example of a *band-pass filter*. When the frequency of a sinusoidal input matches the periodicity of the linear system's weighting function the output sinusoid has a large amplitude. When the frequency of the input is either too high or too low, the output sinusoid is attenuated.

Shift-invariant linear systems are often depicted with block diagrams, like those shown in Fig. 8. Fig. 8A depicts a simple linear filter with frequency response $\hat{h}(\omega)$. The equations that go with the block diagram are:

$$y(t) = h(t) * x(t)$$

$$\hat{y}(\omega) = \hat{h}(\omega)\hat{x}(\omega)$$

The first of these equations is the now familiar convolution formula in which $x(t)$ is the input

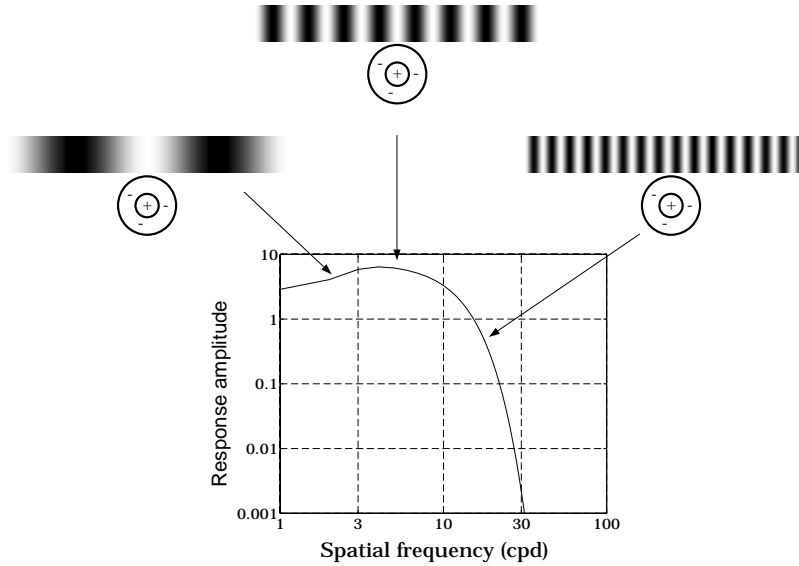


Figure 7: Illustration of an idealized, retinal ganglion-cell receptive field that acts like a bandpass filter (redrawn from Wandell, 1995). This linear on-center neuron responds best to an intermediate spatial frequency whose bright bars fall on-center and whose dark bars fall over the opposing surround. When the spatial frequency is low, the center and surround oppose one another because both are stimulated by a bright bar, thus diminishing the response. When the spatial frequency is high, bright and dark bars fall within and are averaged by the center (likewise in the surround), again diminishing the response.

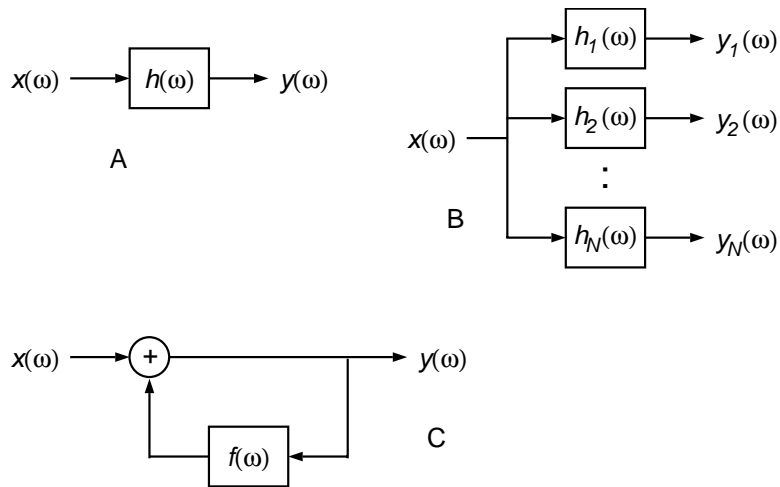


Figure 8: Block diagrams of linear filters. A: Linear filter with frequency response $\hat{h}(\omega)$. B: Bank of linear filters with different frequency responses. C: Feedback linear system.

signal, $y(t)$ is the output signal, and $h(t)$ is the impulse response of the linear filter. The second equation (derived in Appendix II) says the same thing, but expressed in the Fourier domain: $\hat{x}(\omega)$ is the Fourier transform of the input signal, $\hat{y}(\omega)$ is the Fourier transform of the output signal, and $\hat{h}(\omega)$ is the frequency response of the linear filter (that is, the Fourier transform of the impulse response). At the risk of confusing you, it is important to note that $\hat{x}(\omega)$, $\hat{y}(\omega)$, and $\hat{h}(\omega)$ are complex-valued functions of frequency. The complex number notation makes it easy to denote both the amplitude/scale and phase/shift of each frequency component (see Appendix I for a quick review of complex numbers). We can also write out the amplitude and phase parts separately:

$$\begin{aligned}\text{amplitude}[\hat{y}(\omega)] &= \text{amplitude}[\hat{h}(\omega)] \text{amplitude}[\hat{x}(\omega)] \\ \text{phase}[\hat{y}(\omega)] &= \text{phase}[\hat{h}(\omega)] + \text{phase}[\hat{x}(\omega)]\end{aligned}$$

For an input sinusoid of frequency ω , the output is a sinusoid of the same frequency, scaled by $\text{amplitude}[\hat{x}(\omega)]$ and it is shifted by $\text{phase}[\hat{x}(\omega)]$.

Fig. 8B depicts a bank of linear filters that all receive the same input signal. For example, they might be spatially-oriented linear neurons (like V1 simple cells) with different orientation preferences.

Linear Feedback and IIR Filters

Fig. 8C depicts a linear feedback system. The equation corresponding to this diagram is:

$$y(t) = x(t) + f(t) * y(t). \quad (6)$$

Note that because of the feedback, the output $y(t)$ appears on both sides of the equation. The frequency response of the feedback filter is denoted by $\hat{f}(\omega)$, but the behavior of the entire linear system can be expressed as:

$$\hat{y}(\omega) = \hat{x}(\omega) + \hat{f}(\omega) \hat{y}(\omega).$$

Solving for $\hat{y}(\omega)$ in this expression gives:

$$\hat{y}(\omega) = \hat{h}(\omega) \hat{x}(\omega) = \frac{\hat{x}(\omega)}{1 - \hat{f}(\omega)},$$

where $\hat{h}(\omega) = 1/[1 - \hat{f}(\omega)]$ is the effective frequency response of the entire linear feedback system. Using a linear feedback filter with frequency response $\hat{f}(\omega)$ is equivalent to using a linear feedforward filter with frequency response $\hat{h}(\omega)$.

There is one additional subtle, but important, point about this linear feedback system. A system is called *causal* or *nonanticipatory* if the output at any time depends only on values of the input at the present time and in the past. For example, the systems $y(t) = x(t - 1)$ and $y(t) = x^2(t)$ are causal, but the system $y(t) = x(t + 1)$ is not causal. Note that not all causal systems are linear and that not all linear systems are causal (look for examples of each in the previous sentence).

For Eq. 6 to make sense, the filter $f(t)$ must be causal so that the output at time t depends on the input at time t plus a convolution with *past* outputs. For example, if $f(t) = \frac{1}{2}\delta(t-1)$ then:

$$y(t) = x(t) + \frac{1}{2}\delta(t-1) * y(t) = x(t) + \frac{1}{2}y(t-1).$$

The output $y(t)$ accumulates the scaled input values at 1 sec time intervals:

$$y(t) = x(t) + \frac{1}{2}x(t-1) + \frac{1}{4}x(t-2) + \dots$$

Linear feedback systems like this are often referred to as *infinite impulse response* or *IIR* linear filters, because the output depends on the full past history of the input. In practice, of course, the response of this system attenuates rather quickly over time; A unit impulse input from ten seconds in the past contributes only 2^{-10} to the present response.

Differential Equations as Linear Systems

Differentiation is a shift-invariant linear operation. Let $y(t) = \frac{d}{dt}x(t)$, and let's check the three conditions for a shift-invariant linear system:

- Shift-invariance: $\frac{d}{dt}[x(t-s)] = y(t-s)$.
- Homogeneity: $\frac{d}{dt}[\alpha x(t)] = \alpha y(t)$.
- Additivity: $\frac{d}{dt}[x_1(t) + x_2(t)] = y_1(t) + y_2(t)$.

In principle, then, we could try to express differentiation as convolution:

$$y(t) = d(t) * x(t),$$

where $d(t)$ is the “impulse response” of the differentiation operation. But this is weird because the derivative of an impulse is undefined.

On the other hand, the frequency response of differentiation makes perfect sense and it is often very useful to think about the differentiation operation in the Fourier domain. Differentiating a sinusoid,

$$\frac{d}{dt} \sin(2\pi\omega t) = 2\pi\omega \cos(2\pi\omega t),$$

produces another sinusoid that is scaled by $2\pi\omega$ and shifted by $\pi/2$. So the frequency response $\hat{d}(\omega)$ of differentiation is:

$$\begin{aligned} \text{amplitude}[\hat{d}(\omega)] &= 2\pi\omega \\ \text{phase}[\hat{d}(\omega)] &= \pi/2, \end{aligned}$$

And we can express differentiation as multiplication in the Fourier domain:

$$\hat{y}(\omega) = \hat{d}(\omega)\hat{x}(\omega).$$

Writing the amplitude and phase parts separately gives:

$$\begin{aligned}\text{amplitude}[\hat{y}(\omega)] &= \text{amplitude}[\hat{d}(\omega)] \text{amplitude}[\hat{x}(\omega)] = 2\pi\omega \text{amplitude}[\hat{x}(\omega)] \\ \text{phase}[\hat{y}(\omega)] &= \text{phase}[\hat{d}(\omega)] + \text{phase}[\hat{x}(\omega)] = (\pi/2) + \text{phase}[\hat{x}(\omega)].\end{aligned}$$

Now consider the behavior of the following first-order, linear, differential equation:

$$\frac{d}{dt}y(t) = -y(t) + x(t). \quad (7)$$

The equation for a passive neural membrane, for example, can be expressed in this form. There are two operations in this equation: differentiation and addition. Since both are linear operations, this is an equation for a linear system with input $x(t)$ and output $y(t)$. The output $y(t)$ appears on both sides of the equation, so it is a linear feedback system. Since the present output depends on the full past history, it is an infinite impulse response system (an IIR filter).

Taking the Fourier transform of both sides, the differential equation can be expressed using multiplication in the Fourier domain:

$$\hat{d}(\omega)\hat{y}(\omega) = -\hat{y}(\omega) + \hat{x}(\omega).$$

Solving for $\hat{y}(\omega)$ in this expression gives:

$$\hat{y}(\omega) = \hat{h}(\omega)\hat{x}(\omega) = \frac{\hat{x}(\omega)}{1 + \hat{d}(\omega)},$$

where $\hat{h}(\omega) = 1/[1 + \hat{d}(\omega)]$ is the effective frequency response of this linear feedback system.

Appendix I: A Quick Review of Complex Numbers

Fourier transforms involve complex numbers, so we need to do a quick review. A complex number $z = a + jb$ has two parts, a real part x and an imaginary part jb , where j is the square-root of -1. A complex number can also be expressed using complex exponential notation and Euler's equation:

$$z = a + jb = Ae^{j\phi} = A[\cos(\phi) + j \sin(\phi)],$$

where A is called the amplitude and ϕ is called the phase. We can express the complex number either in terms of its real and imaginary parts or in terms of its amplitude and phase, and we can go back and forth between the two:

$$\begin{aligned}a &= A \cos(\phi), & b &= A \sin(\phi) \\ A &= \sqrt{a^2 + b^2}, & \phi &= \tan^{-1}(b/a)\end{aligned}$$

Euler's equation, $e^{j\phi} = \cos(\phi) + j \sin(\phi)$, is one of the wonders of mathematics. It relates the magical number e and the exponential function e^ϕ with the trigonometric functions,

$\sin(\phi)$ and $\cos(\phi)$. It is most easily derived by comparing the Taylor series expansions of the three functions, and it has to do fundamentally with the fact that the exponential function is its own derivative: $\frac{d}{d\phi}e^{\phi} = e^{\phi}$.

Although it may seem a bit abstract, complex exponential notation, $e^{j\phi}$, is very convenient. For example, let's say that you wanted to multiply two complex numbers. Using complex exponential notation,

$$(A_1 e^{j\phi_1}) (A_2 e^{j\phi_2}) = A_1 A_2 e^{j(\phi_1 + \phi_2)},$$

so that the amplitudes multiply and the phases add. If you were instead to do the multiplication using real and imaginary, $a + jb$, notation you would get four terms that you could write using sin and cos notation, but in order to simplify it you would have to use all those trig identities that you forgot after graduating from high school. That is why complex exponential notation is so widespread.

Appendix II: The Fourier Transform

Any signal can be written as a sum of shifted and scaled sinusoids, as was expressed in Eq. 5. That equation is usually written using complex exponential notation:

$$x(t) = \int_{-\infty}^{\infty} \hat{x}(f) e^{j\omega t} dt. \quad (8)$$

The complex exponential notation, remember, is just a shorthand for sinusoids and cosinusoids, but it is mathematically more convenient. The $\hat{x}(\omega)$ are the Fourier transform coefficients for each frequency component ω . These coefficients are complex numbers and can be expressed either in terms of their real (cosine) and imaginary (sine) parts or in terms of their amplitude and phase.

A second equation tells you how to compute the Fourier transform coefficients, $\hat{x}(\omega)$, from from the input signal:

$$\hat{x}(\omega) = \mathcal{F}\{x(t)\} = \frac{1}{2\pi} \int_{-\infty}^{\infty} x(t) e^{-j\omega t} dt. \quad (9)$$

These two equations are inverses of one another. Eq. 9 is used to compute the Fourier transform coefficients from the input signal, and then Eq. 8 is used to reconstruct the input signal from the Fourier coefficients.

The equations for the Fourier transform are rather complex (no pun intended). The best way to get an intuition for the frequency domain is to look at a few examples. Figure 9 plots sinusoidal signals of two different frequencies, along with their Fourier transform amplitudes. A sinusoidal signal contains only one frequency component, hence the frequency plots contain impulses. Both sinusoids are modulated between plus and minus one, so the impulses in the frequency plots have unit amplitude. The only difference between the two sinusoids is that




Figure 9: Fourier transforms of sinusoidal signals of two different frequencies.

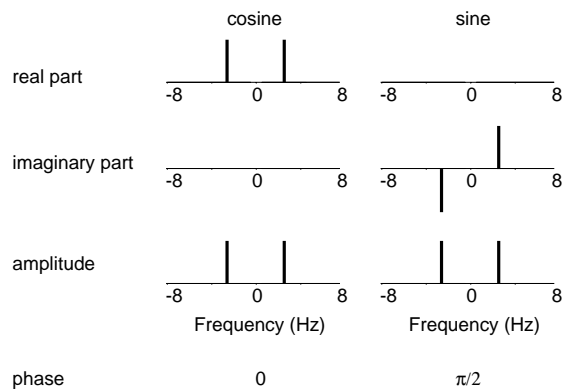


Figure 10: Fourier transforms of sine and cosine signals. The amplitudes are the same, but the phases are different.

one has 4 cycles per second and the other has 8 cycles per second. Hence the impulses in the frequency plots are located at 4 Hz and 8 Hz, respectively.

Figure 10 shows the Fourier transforms of a sinusoid and a cosinusoid. We can express the Fourier transform coefficients either in terms of their real and imaginary parts, or in terms of their amplitude and phase. Both representations are plotted in the figure. Sines and cosines of the same frequency have identical amplitude plots, but the phases are different.

Do not be put off by the negative frequencies in the plots. The equations for the Fourier transform and its inverse include both positive and negative frequencies. This is really just a mathematical convenience. The information in the negative frequencies is redundant with that in the positive frequencies. Since $\cos(-f) = \cos(f)$, the negative frequency components in the real part of the frequency domain will always be the same as the corresponding positive frequency components. Since $\sin(-f) = -\sin(f)$ the negative frequency components in the imaginary part of the frequency domain will always be minus one times the corresponding positive frequency components. Often, people plot only the positive frequency components, as was done in Fig. 9, since the negative frequency components provide no additional information. Sometimes, people plot only the amplitude. In this case, however, there is information missing.

There are a few facts about the Fourier transform that often come in handy. The first of the properties is that the Fourier transform is itself a linear system, which you can check for yourself by making sure that Eq. 9 obeys both homogeneity and additivity. This is important because it makes it easy for us to write the Fourier transforms of lots of things. For example, the Fourier transform of the sum of two signals is the sum of the two Fourier transforms:

$$\mathcal{F}\{x(t) + y(t)\} = \mathcal{F}\{x(t)\} + \mathcal{F}\{y(t)\} = \hat{x}(\omega) + \hat{y}(\omega),$$

where I have used $\mathcal{F}\{\cdot\}$ as a shorthand notation for “the Fourier transform of”. The linearity of the Fourier transform was one of the tricks that made it easy to write the transforms of both sides of Eq. 6.

A second fact, known as the convolution property of the Fourier transform, is that the Fourier transform of a convolution equals the product of the two Fourier transforms:

$$\mathcal{F}\{h(t) * x(t)\} = \mathcal{F}\{h(t)\}\mathcal{F}\{x(t)\} = \hat{h}(\omega)\hat{x}(\omega).$$

This property was also used to write the Fourier transform of Eq. 6. Indeed this property is central to much of the discussion in this handout. Above I emphasized that for a shift-invariant linear system (i.e., convolution), the system’s responses are always given by shifting and scaling the frequency components of the input signal. This fact is expressed mathematically by the convolution property above, where $\hat{x}(\omega)$ are the frequency components of the input and $\hat{h}(\omega)$ is the frequency response, the (complex-valued) scale factors that shift and scale each frequency component.

The convolution property of the Fourier transform is so important that I feel compelled to write out a derivation of it. We start with the definition of convolution:

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

By the definition of the Fourier transform,

$$\hat{y}(\omega) = \int_{-\infty}^{\infty} \left[\int_{-\infty}^{\infty} x(s)h(t-s) ds \right] e^{-j\omega t} dt.$$

Switching the order of integration,

$$\hat{y}(\omega) = \int_{-\infty}^{\infty} x(s) \left[\int_{-\infty}^{\infty} h(t-s)e^{-j\omega t} dt \right] ds.$$

Letting $\sigma = t - s$,

$$\begin{aligned} \hat{y}(\omega) &= \int_{-\infty}^{\infty} x(s) \left[\int_{-\infty}^{\infty} h(\sigma)e^{-j\omega(\sigma+s)} d\sigma \right] ds \\ &= \int_{-\infty}^{\infty} x(s)e^{-j\omega s} \left[\int_{-\infty}^{\infty} h(\sigma)e^{-j\omega\sigma} d\sigma \right] ds \end{aligned}$$

Then by the definition of the Fourier transform,

$$\begin{aligned} \hat{y}(\omega) &= \int_{-\infty}^{\infty} x(s)e^{-j\omega s} \hat{h}(\omega) ds \\ &= \hat{h}(\omega) \int_{-\infty}^{\infty} x(s)e^{-j\omega s} ds \\ &= \hat{h}(\omega)\hat{x}(\omega) \end{aligned}$$

A third property of the Fourier transform, known as the differentiation property, is expressed as:

$$\mathcal{F} \left\{ \frac{dx}{dt} \right\} = 2\pi j\omega \mathcal{F}\{x\}.$$

This property was used to write the Fourier transform of Eq. 7. It is also very important and I would feel compelled to write a derivation of it as well, but I am running out of energy, so you will have to do it yourself.

References

- [1] A V Oppenheim and R W Schaffer. *Discrete-Time Signal Processing*. Prentice-Hall, Englewood Cliffs, New Jersey, 1989.
- [2] A V Oppenheim, A S Willsky, and I T Young. *Signals and Systems*. Prentice-Hall, Englewood Cliffs, New Jersey, 1983.
- [3] B A Wandell. *Foundations of Vision*. Sinauer, Sunderland, MA, 1995.