Linear Systems Theory Handout

David Heeger, Stanford University

Matteo Carandini, ETH/University of Zurich

Characterizing the complete input-output properties of a system by exhaustive measurement is usually impossible. 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.

These notes explain the following ideas related to linear systems theory:

- The challenge of characterizing a complex systems
- Simple linear systems
    - Homogeneity
    - Superposition
- Shift-invariance
    - Decomposing a signal into a set of shifted and scaled impulses
    - The impulse response function
    - Use of sinusoids in analyzing shift-invariant linear systems
    - Decomposing stimuli into sinusoids via Fourier Series
    - Characterizing a shift-invariant system using sinusoids

---

Linear systems theory can be applied to many systems. For example, it can be used to study the responses of a hi-fi system to sounds, or the response of the eye to color images. Here we describe the theory in general and in simple terms. The examples we give are applied to the responses of the retina to images. To simplify matters we imagine that the retina is a thin slit, and is looking at a single horizontal line of a monitor. This way the stimuli are one-dimensional (they live on a line). In real life they are two-dimensional, and that's a bit harder to visualize.

Systems, Inputs, and Responses

Step one is to understand how to represent possible inputs to systems. In our case, imagine a graph of the visual stimulus. Each point in space (abscissa) corresponds a light intensity (ordinate). If there is a single pixel on, and the rest are off, the signal is called an *impulse*. It looks like a single upwards blip on the graph. More complex images look like more complex graphs on this kind of plot. This sort of graph offers a general way to describe all of the possible stimuli.

One possible way to characterize the response of the eye to images might be to build a *look-up table*: a table that shows the exact neural response for every possible visual stimulus. Obviously, it would take an infinite amount of time to construct such a table, because the number of possible images is unlimited.

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 stimuli that we have not yet measured. We can only do this for certain kinds of systems with certain properties. If we have a good theory about the kind of system we are studying, 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.

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 superposition. Systems that satisfy both homogeneity and superposition are linear.

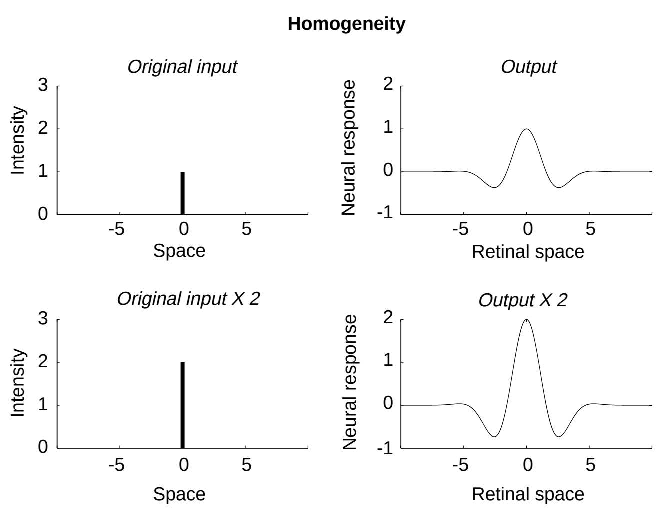

Homogeneity: As we increase the strength of a simple input to a linear system, say we double it, then we predict that the output function will also be doubled. For example, if the intensity of an image is doubled, the eye should respond twice as much if it's a linear system. This is called homogeneity.

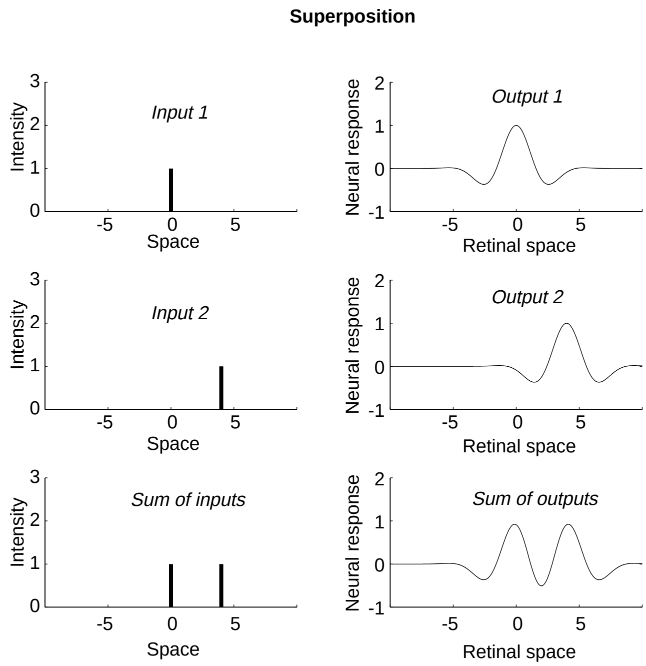

Additivity: Suppose we present a complex stimulus S1 such as the face of a person, and we measure the electrical responses of the nerve fibers coming from the retina. Next, we present a second stimulus S2 that is a little different: a different person's face. The second stimulus also generates a set of responses which we measure and write down. Then, we present the sum of the two stimuli S1 + S2: we present both faces together and see what happens. If the system is linear, then the measured response of each nerve fiber will be just the sum of its responses to each of the two stimuli presented separately.

Shift-invariance: Suppose that we stimulate the left part of your eye with an impulse (a pixel turned on) and we measure the electrical response. Then we stimulate it again with a similar impulse on the right side, and again we measure the response. If the retina is symmetrical (which is true, with the exception of the blind spot), then we should expect that the response to the second impulse will be the same as the response to the first impulse. The only difference between them will be that the second impulse has occurred in a different location, that is, it is stimulating different cells that are shifted in retinal space. When the responses to the identical stimulus presented in different locations are the same, except for the corresponding shift in space, 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.

Why impulses are special: Homogeneity, superposition, and shift invariance may, at first, sound a bit abstract but they are very useful. They suggest that the system's response to an impulse can be the key measurement to make. The trick is to conceive of the complex stimuli we encounter (such as a person's face) as the combination of pixels. We can approximate any complex stimulus as if it were simply the sum of a number of impulses that are scaled copies of one another and shifted in space. Indeed, we can show any image we want on a computer monitor, which lights up different pixels. If we want a better approximation, we need a better computer, one that can show more pixels.

For shift-invariant linear systems, we can measure the system's response to an impulse and we will know how to predict the response to any stimulus (combinations of impulses) through the principle of superposition. To characterize shift-invariant linear systems, then, we need to measure only one thing: the way the system responds to an impulse of a particular intensity. This response is called *the impulse response function* of the system.

The problem of characterizing a complex system has become simpler now. For shift-invariant linear systems, there is only a single impulse response function to measure. Once we've measured this function, we can predict how the system will respond to any other possible stimulus.

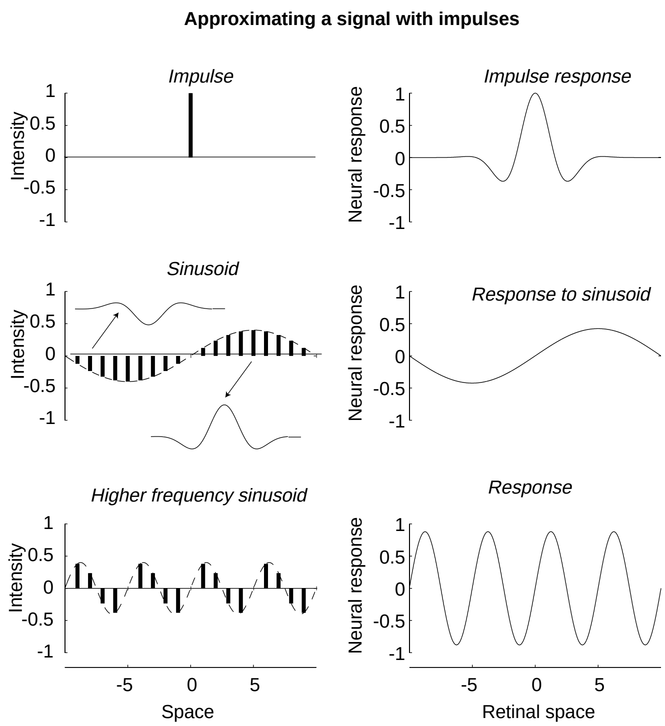

The way we use the impulse response function is illustrated in the above Figure. We conceive of the input stimulus, in this case a sinusoid, as if it were the sum of a set of impulses. We know the responses we would get if each impulse was presented separately (i.e., scaled and shifted copies of the 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.

Sinusoidal stimuli

Sinusoidal stimuli have a special relationship to shift-invariant linear systems. A sinusoid is a regular, repeating curve, that oscillates around a mean level. The sinusoid has a zero-value at time zero. The cosinusoid is a shifted version of the sinusoid; it has a value of one at time zero.

The sine wave repeats itself regularly. The distance from one peak of the wave to the next peak is called the *wavelength* or *period* of the sinusoid and it is generally indicated by the greek letter lambda. The inverse of wavelength is frequency: the number of peaks in the stimulus that arrive per second at the ear. The longer the wavelength, the lower the frequency. Apart from frequency, sinusoids also have various amplitudes, which represent how high they get at the peak of the wave and how low they get at the trough. Thus, we can describe a sine wave completely by its frequency and by its amplitude.

When we write the mathematical expression of a sine-wave, the two mathematical variables that correspond to the frequency and the amplitude are A and f:

$$A \sin(2 \pi f t)$$

The height of the peaks increase as the value of the amplitude, A, increases. The spacing between the peaks becomes smaller as the frequency, f, increases.

The response of shift-invariant systems to sine waves: Just as we can express any stimulus as the sum of a series of shifted and scaled impulses, so too we can express any periodic stimulus (a stimulus that repeats itself over time) as the sum of a series of (shifted and scaled) sinusoids at different frequencies. This is called the *Fourier Series* expansion of the stimulus. The equation describing this expansion works as follows. Suppose that $s(t)$ is a periodic stimulus. Then we can always express $s(t)$ as a sum of sinusoids:

$$s(t) = A0 + A1 \sin(2 \pi f1\ t + p1) + A2 \sin(2 \pi f2\ t + p2) + A3 \sin(2 \pi f3\ t + p3) + \ldots$$

(Do not memorize this equation!)

The frequencies of the sinusoids are very simple: $f1 = 1$ (one single cycle over the stimulus duration), $f2 = 2$ (two cycles over the stimulus duration), $f3 = 3$, etc.

You can go either way: if you know the coefficients (the *A*'s and *p*'s), you can reconstruct the original stimulus $s(t)$; if you know the stimulus, you can compute the coefficients by a method called the Fourier Transform (a way of decomposing complex stimuli into its component sinusoids).

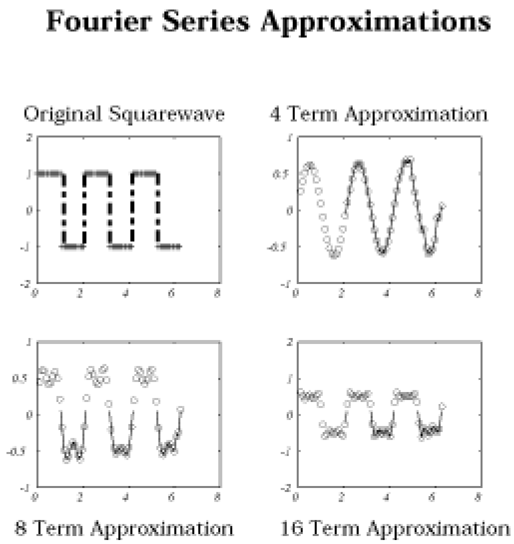

This decomposition is important because if we know the response of the system to sinusoids at many different frequencies, then we can use the same kind of trick we used with impulses to predict the response via the impulse response function. First, we measure the system's response to sinusoids of all different frequencies. Next, we take our input stimulus (a complex sound) and use the Fourier Transform to compute the values of the coefficients in the Fourier Series expansion. At this point the stimulus has been broken down as the sum of its component sinusoids. Finally, we can predict the system's response to the (complex) stimulus simply by adding the responses for all the 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 sinusoidal stimulus as input to a shift-invariant linear system, the system's responses is always a (shifted and scaled) copy of the input, *at the same frequency as the input*. That is, when the input is sin(2 pi f t) the output is always of the form A sin(2 pi f t + p). Here, p determines the amount of shift and A 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.

Shift-Invariant Linear Systems and Sinusoids

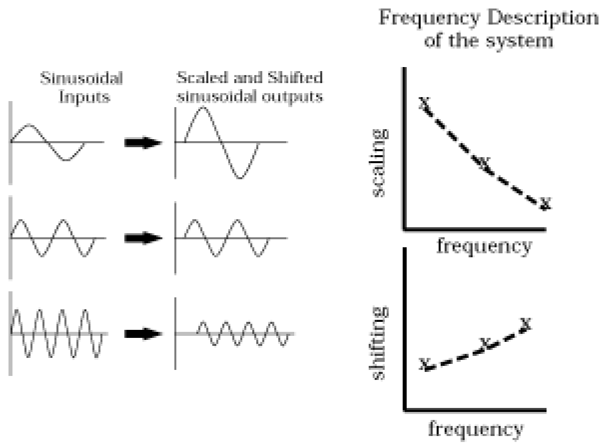

Often, then, when scientists characterize the response of a shift-invariant linear system they will not tell you the impulse response. Rather, they will give you plots that tell you about the values of the shift and scale for each of the possible input frequencies. This representation of how the shift-invariant linear system behaves is equivalent to providing you with the impulse response function. We can use these numbers to compute the response to any stimulus. 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 periodic stimulus.

Linear Systems Logic

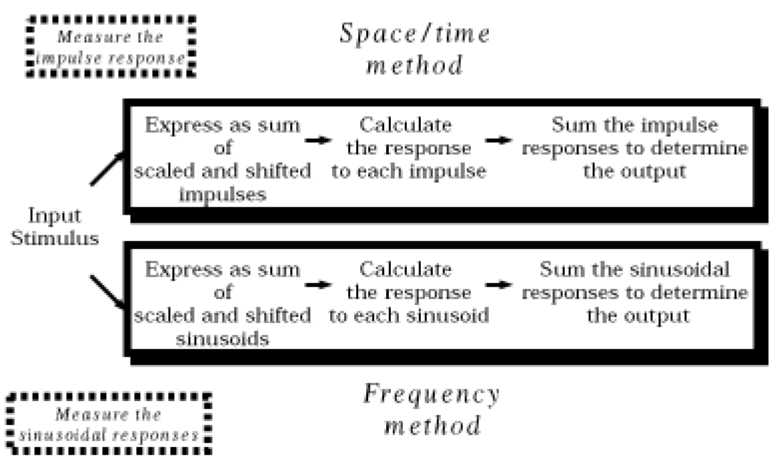

Modified 1999 by Matteo Carandini, ETH/University of Zurich.