Review of Linear System Theory

The following is a (*very*) brief review of linear system theory and Fourier analysis. I work primarily with discrete signals, but each result developed in this review has a parallel in terms of continuous signals and systems. I assume the reader is familiar with linear algebra (as reviewed in my handout *Geometric Review of Linear Algebra*), and least squares estimation (as reviewed in my handout *Least Squares Optimization*).

1 Linear shift-invariant systems

A system is **linear** if it obeys the principle of superposition: the response to a weighted sum of any two inputs is the (same) weighted sum of the responses to each individual input.

A system is **shift-invariant** (also called **translation-invariant** for spatial signals, or **time-invariant** for temporal signals) if the response to any input shifted by any amount Δ is equal to the response to the original input shifted by amount Δ.

These two properties are completely independent: a system can have one of them, both or neither [think of an example of each of the 4 possibilities].

The rest of this review is focused on systems that are *both* linear and shift-invariant (known as **LSI** systems). The diagram below decomposes the behavior of such an LSI system. Consider an arbitrary discrete input signal. We can rewrite it as a weighted sum of impulses (also called "delta functions"). Since the system is linear, the response to this weighted sum is just the weighed sum of responses to each individual impulse. Since the system is shift-invariant, the response to each impulse is just a shifted copy of the response to the first. The response to the impulse located at the origin (position 0) is known as the system's **impulse response**.

Putting it all together, the full system response is the weighted sum of shifted copies of the impulse response. Note that the system is fully characterized by the impulse response: This is all we need to know in order to compute the response of the system to any input!

To make this explicit, we write an equation that describes this computation:

$$y[n] = \sum_m x[m]r[n-m]$$

This operation, by which input x and impulse response r are combined to generate output signal y is called a **convolution**. It is often written using a more compact notation: $y = x * r$. Although we think of x and r playing very different roles, the operation of convolution is actually commutative: substituting $k = n - m$ gives:

$$y[n] = \sum_k x[n-k]r[k] = r * x[n]$$

• Author: Eero Simoncelli, Center for Neural Science, and Courant Institute of Mathematical Sciences, New York University.
• Thanks to Jonathan Pillow for generating some of the figures.
• Created: Fall 2001. Revised: Nov 2004.
• Send corrections or comments to eero.simoncelli@nyu.edu

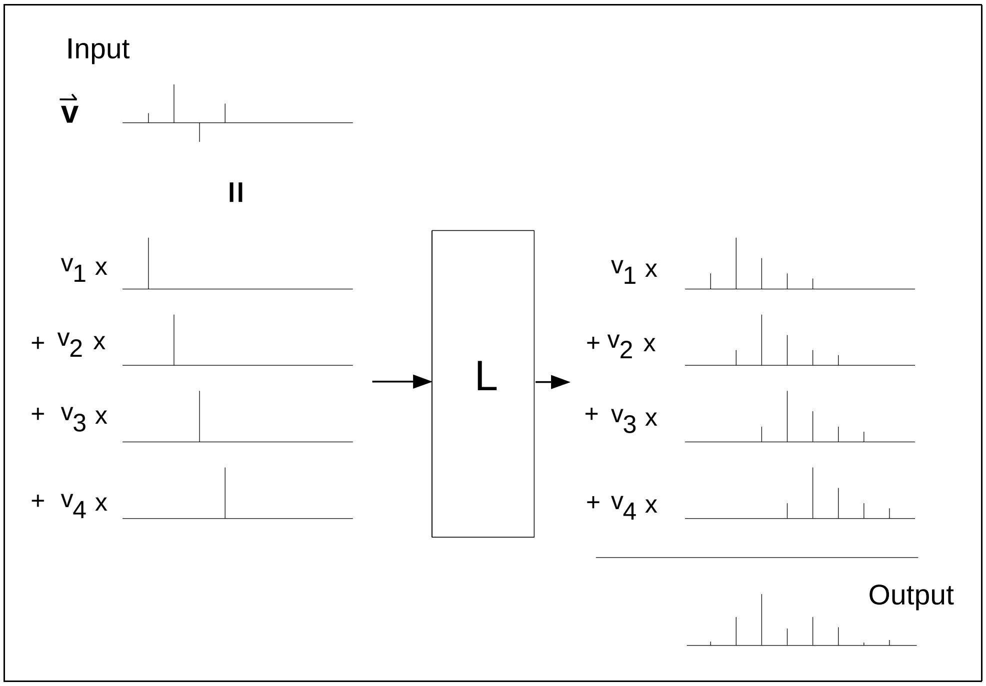

It is also easy to see that convolution is associative: $a*(b*c) = (a*b)*c$. And finally, convolution is distributive over addition: $a * (b + c) = a * b + a * c$.

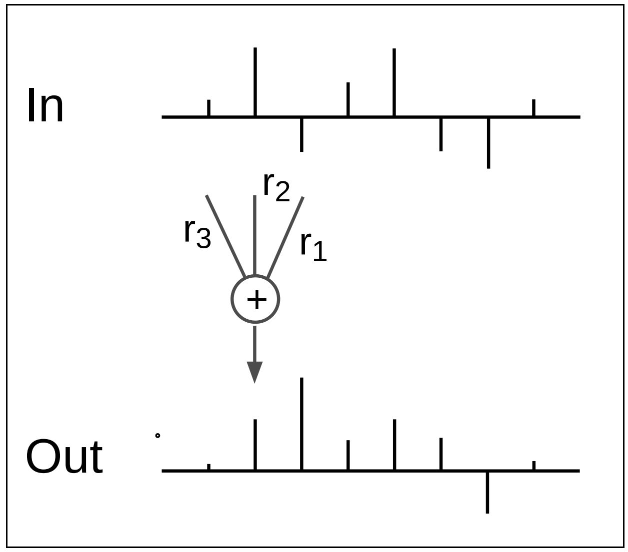

Back to LSI systems. The impulse response r is also known as a "convolution kernel" or "linear filter". Looking back at the definition, each component of the output y is computed as an inner product of a chunk of the input vector x with a *reverse-ordered* copy of r. As such, the convolution operation may be visualized as a weighted sum over a window that slides along the input signal.

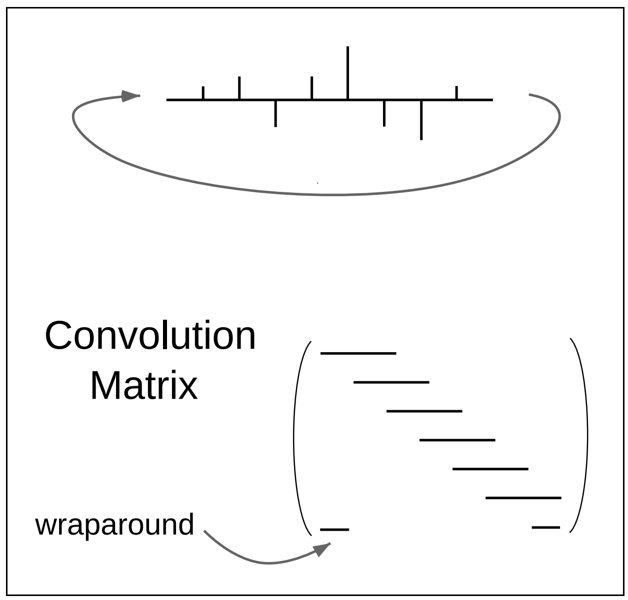

For finite-length discrete signals (i.e., vectors), one must specify how convolution is handled at the boundaries. The standard solution is to consider each vector to be one period of an infinitely periodic signal. Thus, for example, when the convolution operation would require one to access an element just beyond the last element of the vector, one need only "wrap around" and use the first element. There are other methods of handling boundaries. For example, one can pad the signal with zeros, or *reflect* it about the boundary.

The convolution operation is easily extended to multidimensional signals. For example, in 2D, both the signal and convolution kernal are two-dimensional arrays of numbers, and the operation corresponds to taking sums over a 2D window of the signal, with weights specified by the kernel.

2 Sinusoids and Convolution

The **sine** function, $\sin(\theta)$, gives the y-coordinate of the points on a unit circle, as a function of the angle θ. The **cosine** function $\cos(\theta)$, gives the x-coordinate. Thus, $\sin^2(\theta) + \cos^2(\theta) = 1$.

The angle, θ, is (by convention) assumed to be in units of radians - a full sweep around the unit circle corresponds to an angle of 2π radians.

For our purposes here, we will consider sinusoidal *signals*. A continuous sinusoidal signal in time might be written $A\sin(\Omega t - \phi)$, where A is the amplitude, Ω gives the frequency (in radians per unit time) and ϕ corresponds to a "phase shift", or a temporal delay by amount ϕ/Ω. A discretized sinusoid might be written as: $A\sin(2\pi\omega n - \phi)$. Here, n is the index into the vector, and the frequency (in cycles per sample) is determined by ω.

The usefulness of these "sinusoidal" functions comes from their unique behavior with respect to LSI systems. First, consider what happens when the function $\cos(\Omega t)$ is translated by an amount Δ. This corresponds to a change in phase $\phi = \Omega\Delta$, which may be written (using a basic trigonometric identity from high school mathematics):

$$A\cos(\Omega t - \phi) = A\cos(\phi)\cos(\Omega t) + A\sin(\phi)\sin(\Omega t)$$

That is, a shifted copy of a sinusoid is equivalent to a linear combination of the (non-shifted) $\sin$ and $\cos$ functions!

The left and right sides of this equation provide two different characterizations of a sinusoidal function: (a) in terms of its amplitude and phase, or (b) in terms of its decomposition into sinusoidal and cosinusoidal components. These two descriptions are just polar or rectangular coordinate system representations of the same information (see figure). The converse of this trigonometric identity is that a linear combination of sinusoids of the same frequency is a single sinusoid of that frequency, with a particular phase and amplitude.

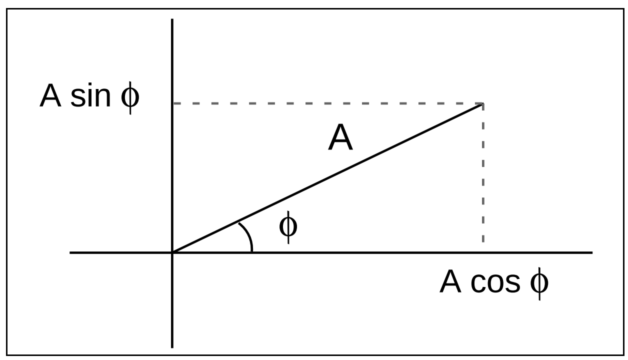

Now consider the response of an LSI system to a sinusoid:

$$\begin{aligned} y &= x * r \\ &= \sum_m r[m] A\sin(2\pi\omega(n-m) - \phi) \\ &= \sum_m r[m] A\sin\left(2\pi\omega n - (2\pi\omega m + \phi)\right) \end{aligned}$$

The response is a sum of sinusoids, all of the same frequency ($2\pi\omega$), but different phases, ($2\pi\omega m + \phi$). This sum is just a single sinusoid of the same frequency but (possibly) different phase and amplitude. That is, sinusoids preserve their frequency when they pass through LSI systems, even though their amplitude and phase may change.

Sinusoids as eigenfunctions of LSI systems. The relationship between LSI systems and sinusoidal functions may be expressed more compactly (and deeply) by bundling together a sine and cosine function into a single complex exponential:

$$e^{i\theta} \equiv \cos(\theta) + i\sin(\theta)$$

where $i = \sqrt{-1}$ is the imaginary number. This rather mysterious relationship can be derived by expanding the exponential in a Taylor series, and noting that the even (real) terms form the series expansion of a cosine and the odd (imaginary) terms form the expansion of a sine [See Feynman's beautiful derivation in his *Lectures on Physics*].

The notation may seem a bit unpleasant, but now changes in the amplitude and phase of a sinusoid may be expressed more compactly, and the relationship between LSI systems and sinusoids may be written as:

$$\begin{aligned} L\{e^{i\Omega t}\} &= A_L e^{i(\Omega t + \phi_L)} \\ &= A_L e^{i\phi_L} e^{i\Omega t} \end{aligned}$$

where A_L and ϕ_L represent the amplitude and phase change caused by the LSI system L in sinusoids with frequency Ω. Summarizing, the action of an LSI system on the complex exponential function is to simply multiply it by the complex number $A_L e^{i\phi_L}$. That is, the complex exponentials are *eigenfunctions* of LSI systems!

3 Fourier Transform(s)

A collection of sinusoids may be used as a linear basis for representing (realistic) signals. The transformation from the standard representation of the signal (eg, as a function of time) to a set of coefficients representing the amount of each sinusoid needed to create the signal is called the **Fourier Transform**.

There are really four variants of Fourier transform, depending on whether the signal is continuous or discrete, and on whether it is periodic or aperiodic:

SIGNAL	continuous	discrete
aperiodic	**Fourier Transform** (continuous, aperiodic)	**Discrete-Time(Space) Fourier Transform** (continuous, periodic)
periodic	**Fourier Series** (discrete, aperiodic)	**Discrete Fourier Transform** (discrete, periodic)

For our purposes here, we'll focus on the Discrete Fourier Transform (DFT), defined for periodic discretized signals. We can write one period of such a signal as a vector of length (say) N. The following collection of N sinusoidal functions forms an orthonormal basis [verify]:

$$\begin{aligned} c_k[n] &\equiv \frac{1}{\sqrt{N}} \cos\left(\frac{2\pi k}{N} n\right), \qquad k = 0, 1, \ldots \frac{N}{2} \\ s_k[n] &\equiv \frac{1}{\sqrt{N}} \sin\left(\frac{2\pi k}{N} n\right), \qquad k = 1, 2, \ldots \frac{N}{2} - 1 \end{aligned}$$

Thus, we can write any vector $\vec{v}$ as a weighted sum of these basis functions:

$$v[n] = \sum_{k=0}^{N/2} a_k c_k[n] + \sum_{k=1}^{N/2-1} b_k s_k[n]$$

Since the basis is orthogonal, the Fourier coefficients $\{a_k, b_k\}$ may be computed by projecting the input vector $\vec{v}$ onto each basis function:

$$
\begin{aligned}
a_k &= \sum_{n=0}^{N-1} v[n] c_k[n] \\
b_k &= \sum_{n=0}^{N-1} v[n] s_k[n]
\end{aligned}
$$

Now, using the properties of sinusoids developed earlier, we can combine cosine and sine terms into a single phase-shifted sinusoid:

$$
v[n] = \sum_{k=0}^{N/2} A_k \sin\left(\frac{2\pi k}{N} n - \phi_k\right)
$$

with amplitudes $A_k = \sqrt{a_k^2 + b_k^2}$, and phases $\phi_k = \tan^{-1}(b_k/a_k)$. These are are referred to as the **Fourier amplitudes** and **Fourier phases** of the signal $v[n]$. Again, this is just a polar coordinate representation of the original amplitudes $\{a_k, b_k\}$.

At the right is an illustration of successive approximations of a triangle wave with sinusoids. The top panel shows the original signal. The next shows the approximation one gets by projecting onto a single zero-frequency sinusoid (i.e., the constant function). The next shows the result with three frequency components, and the last panel shows the result with six. [Try matlab's `xfourier` to see a live demo with square wave.]

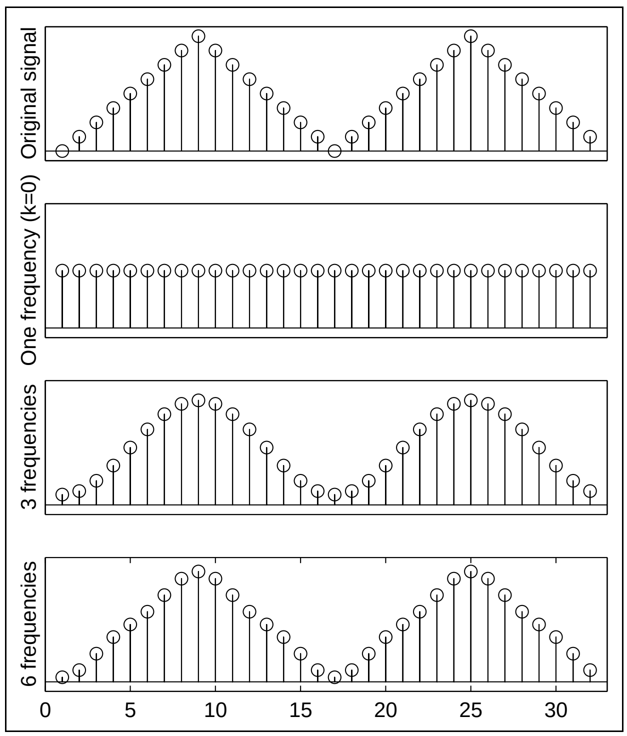

Alternatively, we can use a complex-valued number to represent the amplitude and phase of each frequency component, $A_k e^{i\phi_k}$. Now the Fourier amplitudes and phases correspond to the amplitude and phase of this complex number. This is the standard representation of the

Fourier coefficients.

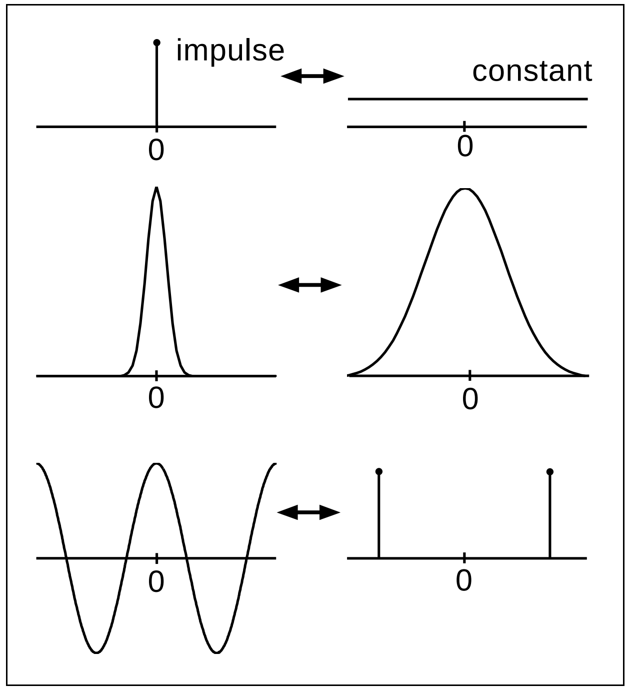

Shown are Fourier amplitude spectra (the amplitudes plotted as a function of frequency number k) for three simple signals: an impulse, a Gaussian, and a cosine function. These are shown here symmetrically arranged around the origin, but some authors plot only the positive half of the frequency axis. Note that the cosine function is constructed by adding a complex exponential with its frequency-negated cousin - this is why the Fourier transform shows two impulses.

Shifting property. When we shift an input signal, each sinusoid in the Fourier representation must be shifted. Specifically, shifting by m samples means that the phase of each sinusoid changes by $\frac{2\pi k}{N}m$. Note that the phase change is different for each frequency k, and also that the amplitude is unchanged.

Stretching property. If we stretch the input signal (i.e., rescale the x-axis by a factor α), the Fourier transform will be compressed by the same factor (i.e., rescale the frequency axis by $1/alpha$). Consider a Gaussian signal. The Fourier amplitude is also a Gaussian, with standard deviation *inversely* proportional to that of the original signal.

4 Convolution Theorem

The most important property of the Fourier representation is its relationship to LSI systems and convolution. To see this, we need to combine the eigenvector property of complex exponentials with the Fourier transform. The diagram below illustrates this. Consider applying an LSI system to an arbitrary signal. Use the Fourier transform to rewrite it as a weighted sum of sinusoids. The weights in the sum may be expressed as complex numbers, $A_k e^{i\phi_k}$, representing the amplitude and phase of each sinusoidal component. Since the system is linear, the response to this weighted sum is just the weighted sum of responses to each of the individual sinusoids. But the action of an LSI on a sinusoid with frequency number k will be to multiply the amplitude by a factor $A_L(k)$ and shift the phase by an amount $\phi_L(k)$. Finally, the system

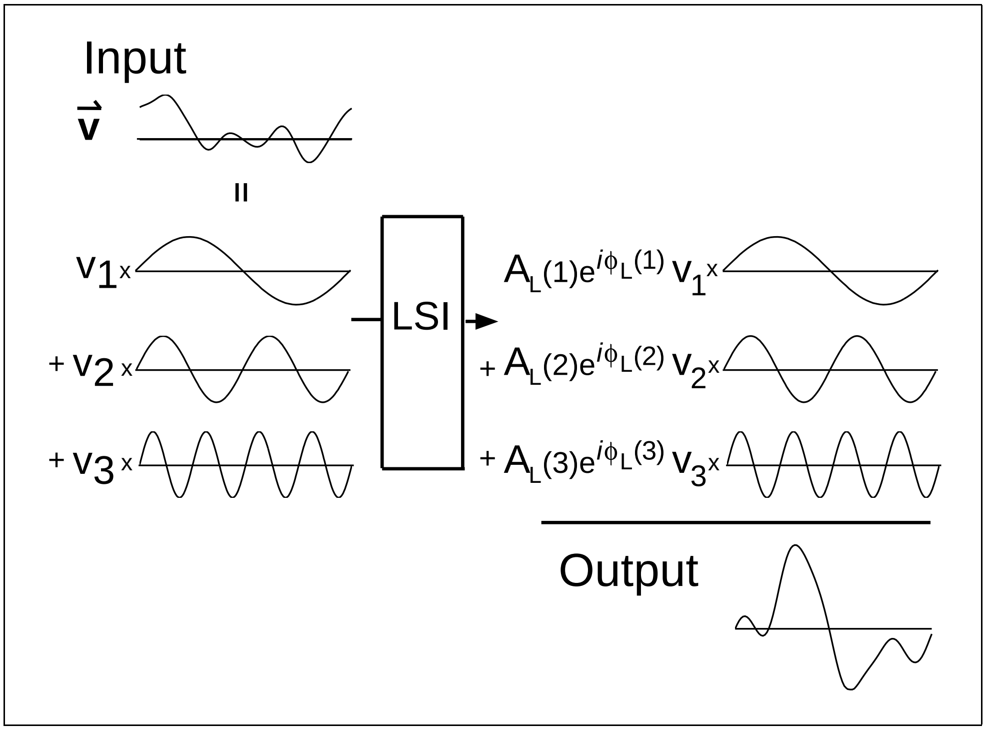

response is a sum of sinusoids with amplitudes/phases corresponding to

$$(A_L(k)A_k)e^{i(\phi_L(k)+\phi_k)} = (A_L(k)e^{i\phi_L(k)})(A_k e^{i\phi_k}).$$

Earlier, using a similar sort of diagram, we explained that LSI systems can be characterized by their impulse response, $r[n]$. Now we have seen that the complex numbers $A_L(k)e^{i\phi_L(k)}$ provide an alternative characterization. We now want to find the relationship between these two characterizations. First, we write an expression for the convolution (response of the LSI system):

$$y[n] = \sum_m r[m]x[n-m]$$

Now take the DFT of both sides of the equation:

$$\begin{aligned} Y[k] &= \sum_n y[n]e^{i2\pi nk/N} \\ &= \sum_n \sum_m r[m]x[n-m]e^{i2\pi nk/N} \\ &= \sum_m r[m] \sum_n x[n-m]e^{i2\pi nk/N} \\ &= \sum_m r[m]e^{i2\pi mk/N} \\ &= X[k]R[k] \end{aligned}$$

This is quite amazing: the DFT of the LSI system response, y[n], is just the product of the DFT of the input and the DFT of the impulse response! That is, the complex numbers $A_L(k)e^{i\phi_L(k)}$ correspond to the Fourier transform of the function $r[n]$.

Summarizing, the response of the LSI system may be computed by a) Fourier-transforming the input signal, b) multiplying each Fourier coefficient by the associated Fourier coefficient of the impulse response, and c) Inverse Fourier-transforming. A more colloquial statement of this **Convolution theorem** is: "convolution in the signal domain corresponds to multiplication in the Fourier domain". Reversing the roles of the two domains (since the inverse transformation is essentially the same as the forward transformation) means that "multiplication in the signal domain corresponds to convolution in the Fourier domain".

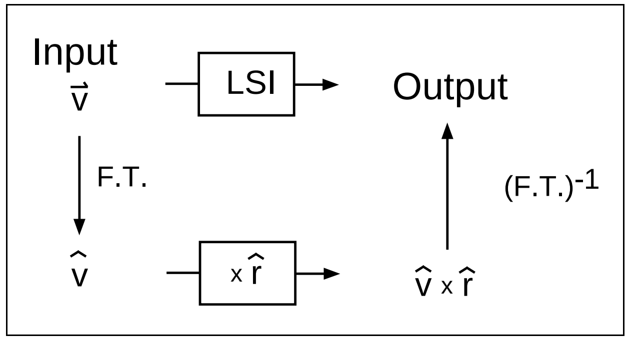

Why would we want to bother going through three sequential operations in order to compute a convolution? Conceptually, multiplication is easier to understand than convolution, and thus we can often gain a better understanding of an LSI by thinking about it in terms of its effect in the Fourier domain. More practically, there are very efficient algorithms for the Discrete Fourier Transform, known as the *Fast Fourier Transform (DFT)*, such that this three-step computation may be more computationally efficient than direct convolution.

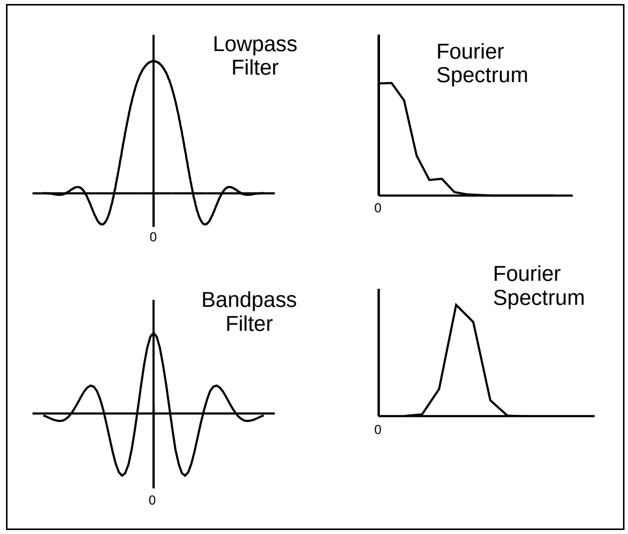

As an example of conceptual simplification, consider two impulse responses, along with their Fourier amplitude spectra. It is often difficult to anticipate the behavior of these systems solely from their impulse responses. But their Fourier transforms are quite revealing. The first is a **lowpass** filter meaning that it discards high frequency sinusoidal components (by multiplying them by zero). The second is a **bandpass** filter - it allows a central band of frequencies to pass, discarding the lower and higher frequency components.

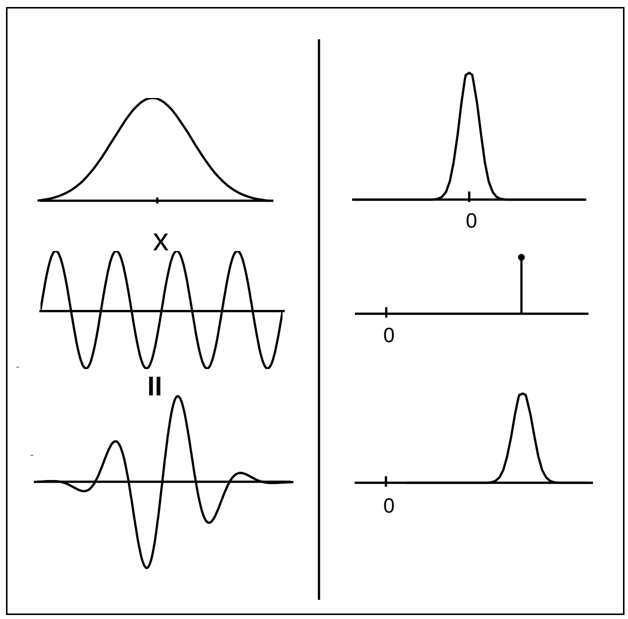

As another example of conceptual simplification, consider an impulse response formed by the product of a Gaussian function, and a sinusoid (known as a **Gabor function**). How can we visualize the Fourier transform of this product? We need only compute the convolution of the Fourier transforms of the Gaussian and the sinusoid. The Fourier transform of a Gaussian is a Gaussian. The Fourier transform of a sinusoid is an impulse at the corresponding frequency. Thus, the Fourier transform of the Gabor function is a Gaussian centered at the frequency of the sinusoid.