Lecture 19: Autocorrelation

Mark Hasegawa-Johnson

ECE 401: Signal and Image Analysis, Fall 2020

Outline

Review: Energy Spectrum and Parseval's Theorem

- The energy spectrum of a random noise signal has the DTFT form $|X(\omega)|^2$, or the DFT form $|X[k]|^2$.
- The easiest form of Parseval's theorem to memorize is the DTFT energy spectrum form:

$$\sum_{n=-\infty}^{\infty} x^2[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} |X(\omega)|^2 d\omega$$

- The DFT energy spectrum form is similar, but over a finite duration:

$$\sum_{n=0}^{N-1} x^2[n] = \frac{1}{N}\sum_{k=0}^{N-1} |X[k]|^2$$

Review: Power Spectrum and Parseval's Theorem

Energy of an infinite-length signal might be infinite. Wiener defined the **power spectrum** in order to solve that problem:

$$R_{xx}(\omega) = \lim_{N\to\infty} \frac{1}{N}|X(\omega)|^2$$

where $X(\omega)$ is computed from a window of length N samples. The DTFT power spectrum form of Parseval's theorem is

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=-(N-1)/2}^{(N-1)/2} x^2[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

White Noise

- White noise is a type of noise whose samples are uncorrelated ($E[x[n]x[m]] = E[x[n]]E[x[m]]$, unless $n = m$). If it is also zero mean and unit variance, then

$$E\left[x[n]x[m]\right] = \begin{cases} 1 & n = m \\ 0 & n \neq m \end{cases}$$

- The Fourier transform of any zero-mean random signal is, itself, a zero-mean random variable:

$$E\left[X(\omega)\right] = 0$$

- The power spectrum is also a random variable, but its expected value is not zero. The expected power spectrum of white noise is flat, like white light:

$$E\left[R_{xx}(\omega)\right] = E\left[\frac{1}{N}|X(\omega)|^2\right] = 1$$

Example: DTFT and Power Spectrum of White Noise

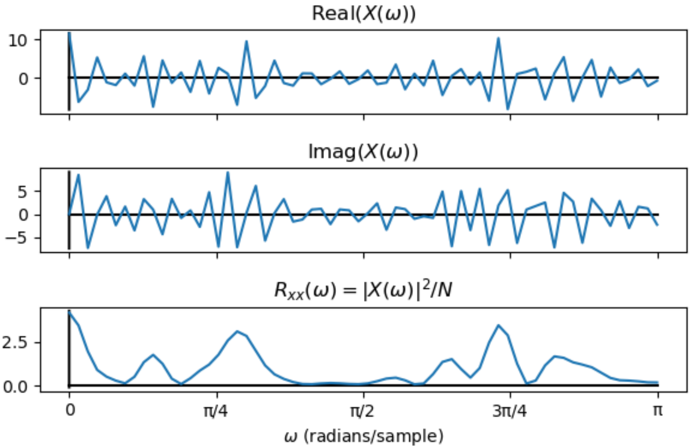

Example: Expected DTFT and Power Spectrum of White Noise

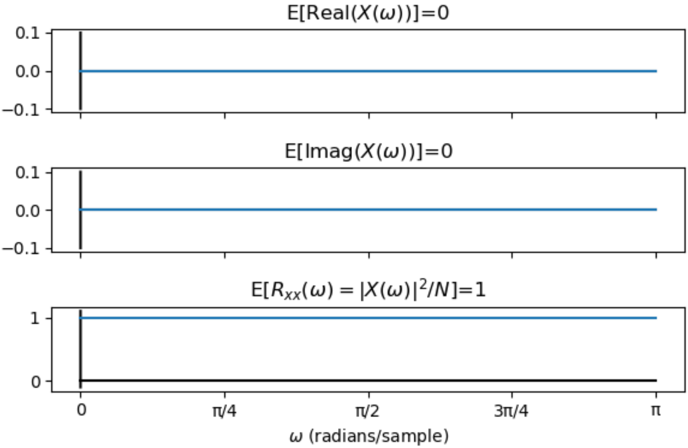

Colored Noise

- Most colored noise signals are well modeled as filtered white noise, i.e., $y[n] = h[n] * x[n]$. The filtering means that the samples of $y[n]$ are correlated with one another.
- If $x[n]$ is zero-mean, then so is $y[n]$, and so is $Y(\omega)$:

$$E\left[Y(\omega)\right] = 0$$

- The expected power spectrum is $|H(\omega)|^2$:

$$E\left[R_{yy}(\omega)\right] = E\left[\frac{1}{N}|Y(\omega)|^2\right] = |H(\omega)|^2$$

Example: Filtered Noise

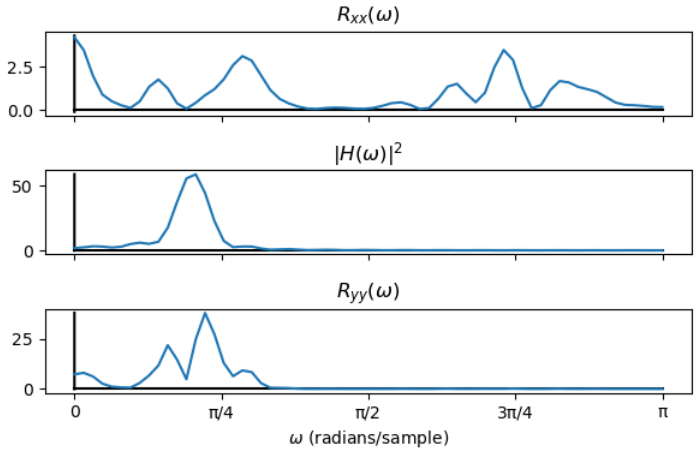

Outline

Finite-Duration Power Spectrum

In practice, we will very often compute the power spectrum from a finite-length window:

$$R_{xx}(\omega) = \frac{1}{N}|X(\omega)|^2, \quad R_{xx}[k] = \frac{1}{N}|X[k]|^2$$

where $X(\omega)$ is computed from a window of length N samples. The DTFT power spectrum form of Parseval's theorem is then

$$\frac{1}{N}\sum_{n=0}^{N-1} x^2[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega = \frac{1}{N}\sum_{k=0}^{N} R_{xx}[k]$$

Inverse DTFT of the Power Spectrum

Since the power spectrum of noise is MUCH more useful than the expected Fourier transform, let's see what the inverse Fourier transform of the power spectrum is. Let's call $R_{xx}(\omega)$ the power spectrum, and $r_{xx}[n]$ its inverse DTFT.

$$R_{xx}(\omega) = \frac{1}{N}|X(\omega)|^2 = \frac{1}{N}X(\omega)X^*(\omega)$$

where $X^*(\omega)$ means complex conjugate. Since multiplying the DTFT means convolution in the time domain, we know that

$$r_{xx}[n] = \frac{1}{N}x[n] * z[n]$$

where $z[n]$ is the inverse transform of $X^*(\omega)$ (we haven't figured out what that is, yet).

Inverse DTFT of the Power Spectrum

So what's the inverse DFT of $X^*(\omega)$? If we assume that $x[n]$ is real, we get that

$$X^*(\omega) = \left(\sum_{n=-\infty}^{\infty} x[n]e^{-j\omega n}\right)^*$$

$$= \sum_{n=-\infty}^{\infty} x[n]e^{j\omega n}$$

$$= \sum_{m=-\infty}^{\infty} x[-m]e^{-j\omega m}$$

So if $x[n]$ is real, then the inverse DTFT of $X^*(\omega)$ is $x[-n]$!

Autocorrelation

The power spectrum, of an N-sample finite-length signal, is

$$R_{xx}(\omega) = \frac{1}{N}|X(\omega)|^2$$

Its inverse Fourier transform is the autocorrelation,

$$r_{xx}[n] = \frac{1}{N}x[n] * x[-n] = \frac{1}{N}\sum_{m=-\infty}^{\infty} x[m]x[m-n]$$

This relationship, $r_{xx}[n] \leftrightarrow R_{xx}(\omega)$, is called Wiener's theorem, named after Norbert Wiener, the inventor of cybernetics.

Example: Autocorrelation of White Noise

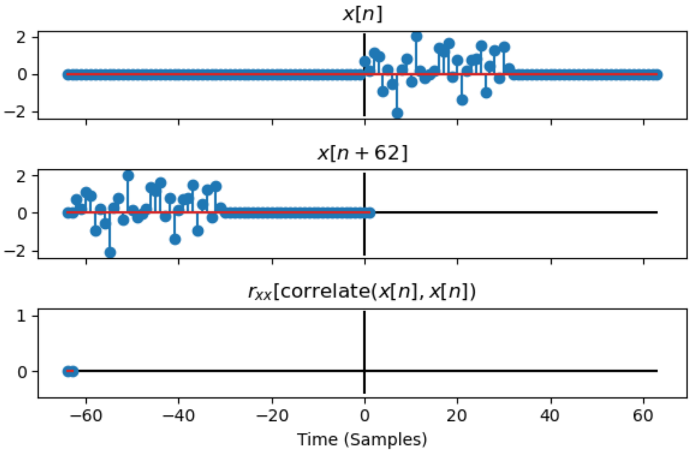

A warning about python

Notice, on the last slide, I defined autocorrelation as

$$r_{xx}[n] = \frac{1}{N}x[n] * x[-n] = \frac{1}{N}\sum_{m=-\infty}^{\infty} x[m]x[m-n]$$

Python defines an "energy version" of autocorrelation, instead of the "power version" shown above, i.e., `np.correlate` computes:

$$r_{\text{python}}[n] = \sum_{m=-\infty}^{\infty} x[m]x[m-n]$$

The difference is just a constant factor (N), so it usually isn't important. But sometimes you'll need to be aware of it.

Autocorrelation is also a random variable!

- Notice that, just as the power spectrum is a random variable, the autocorrelation is also a random variable.
- The autocorrelation is the average of N consecutive products, thus

$$E\left[r_{xx}[n]\right] = E\left[\frac{1}{N}\sum_{m=0}^{N-1} x[m]x[m-n]\right] = E\left[x[m]x[m-n]\right]$$

- The expected autocorrelation is related to the covariance and the mean:

$$E\left[r_{xx}[n]\right] = \text{Cov}\left(x[m], x[m-n]\right) + E\left[x[m]\right]E\left[x[m-n]\right]$$

- If $x[n]$ is zero-mean, that means

$$E\left[r[n]\right] = \text{Cov}\left(x[m], x[m-n]\right)$$

Autocorrelation of white noise

If $x[n]$ is zero-mean white noise, with a variance of σ^2, then

$$E\left[r_{xx}[n]\right] = E\left[x[m]x[m-n]\right] = \begin{cases} \sigma^2 & n = 0 \\ 0 & \text{otherwise} \end{cases}$$

We can write

$$E\left[r[n]\right] = \sigma^2 \delta[n]$$

Outline

Filtered Noise

What happens when we filter noise? Suppose that $x[n]$ is zero-mean white noise, and

$$y[n] = h[n] * x[n]$$

What is $y[n]$?

Example: Filtering of White Noise

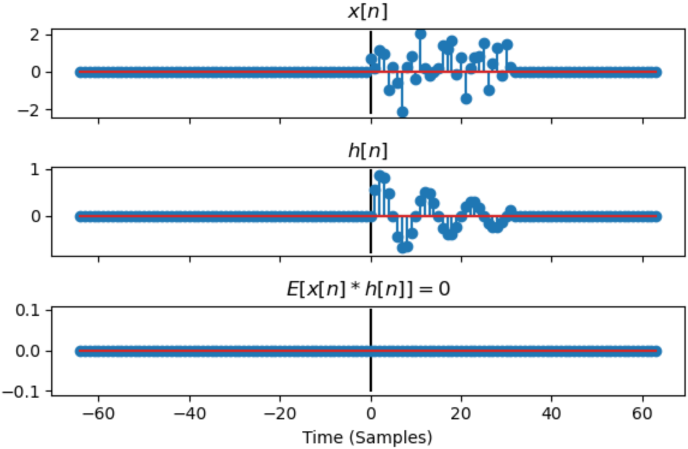

Filtered Noise

$$y[n] = h[n] * x[n] = \sum_{m=-\infty}^{\infty} h[m]x[n-m]$$

- $y[n]$ is the sum of zero-mean random variables, so it's also zero-mean.
- $y[n] = h[0]x[n] +$ other stuff, and $y[n+1] = h[1]x[n] +$ other stuff. So obviously, $y[n]$ and $y[n+1]$ are not uncorrelated. So $y[n]$ is not white noise.
- What kind of noise is it?

The variance of $y[n]$

First, let's find its variance. Since $x[n]$ and $x[n+1]$ are uncorrelated, we can write

$$\sigma_y^2 = \sum_{m=-\infty}^{\infty} h^2[m]\mathrm{Var}(x[n-m])$$
$$= \sigma_x^2 \sum_{m=-\infty}^{\infty} h^2[m]$$

The autocorrelation of $y[n]$

Second, let's find its autocorrelation. Let's define $r_{xx}[n] = \frac{1}{N}x[n] * x[-n]$. Then

$$
\begin{aligned}
r_{yy}[n] &= \frac{1}{N}y[n] * y[-n] \\
&= \frac{1}{N}(x[n] * h[n]) * (x[-n] * h[-n]) \\
&= \frac{1}{N}x[n] * x[-n] * h[n] * h[-n] \\
&= r_{xx}[n] * h[n] * h[-n]
\end{aligned}
$$

Example: Autocorrelation of Colored Noise

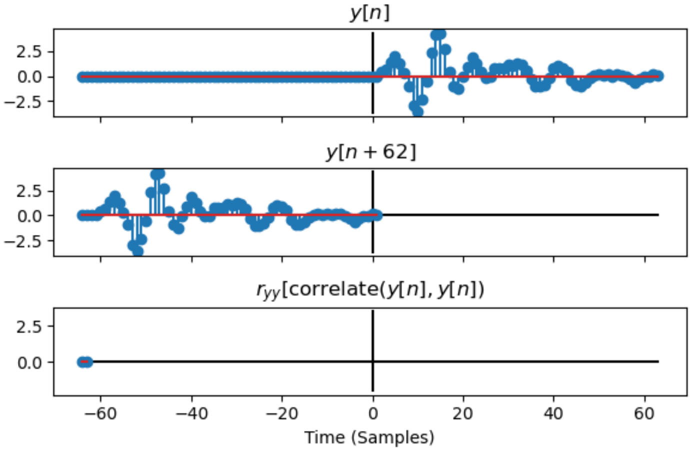

Expected autocorrelation of $y[n]$

$$r_{yy}[n] = r_{xx}[n] * h[n] * h[-n]$$

Expectation is linear, and convolution is linear, so

$$E\left[r_{yy}[n]\right] = E\left[r_{xx}[n]\right] * h[n] * h[-n]$$

Expected autocorrelation of $y[n]$

$x[n]$ is zero-mean white noise if and only if its autocorrelation is a delta function:

$$E\left[r_{xx}[n]\right] = \sigma_x^2 \delta[n]$$

If $y[n] = h[n] * x[n]$, and $x[n]$ is zero-mean white noise, then

$$E\left[r_{yy}[n]\right] = \sigma_x^2 \left(h[n] * h[-n]\right)$$

In other words, $x[n]$ contributes only its energy (σ_x^2). $h[n]$ contributes the correlation between neighboring samples.

Example: Expected Autocorrelation of Colored Noise

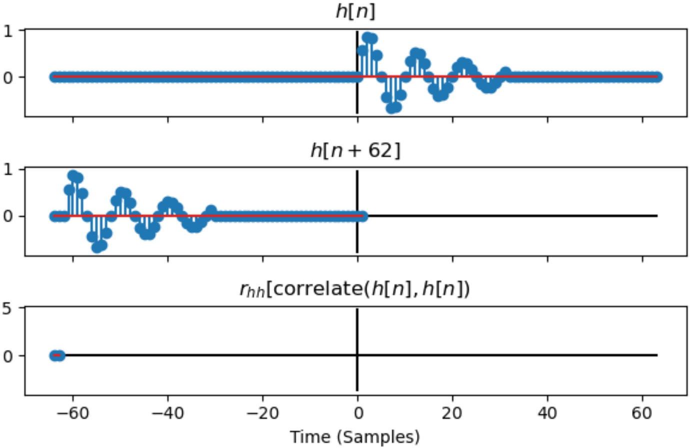

Example

Here's an example. The white noise signal on the top ($x[n]$) is convolved with the bandpass filter in the middle ($h[n]$) to produce the green-noise signal on the bottom ($y[n]$). Notice that $y[n]$ is random, but correlated.

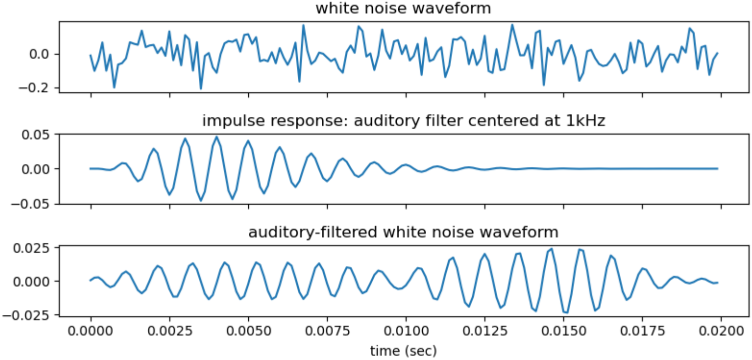

Example

Here's another example. The white noise signal on the left ($x[n]$) is convolved with an ideal lowpass filter, with a cutoff at $\pi/2$, to create the pink-noise signal on the right ($y[n]$). Notice that $y[n]$ is random, but correlated.

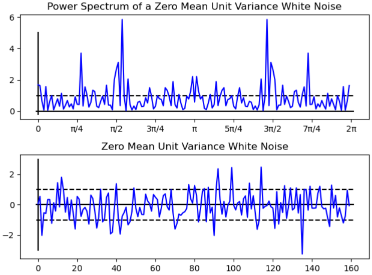

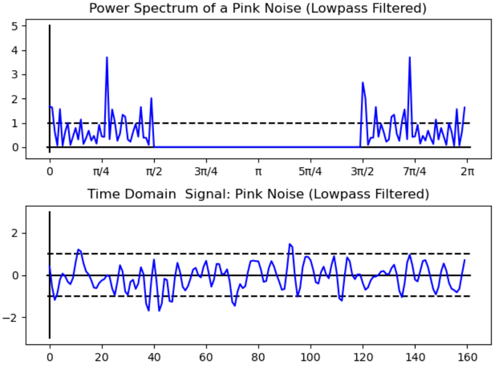

Example

Here's a third example. The white noise signal on the left ($x[n]$) is convolved with an ideal highpass filter, with a cutoff at $\pi/2$, to create the blue-noise signal on the right ($y[n]$). Here, it's a lot less obvious that the samples of $y[n]$ are correlated with one another, but they are: in fact, they are **negatively correlated**. If $y[n] > 0$, then $y[n+1] < 0$ with a probability greater than 50%.

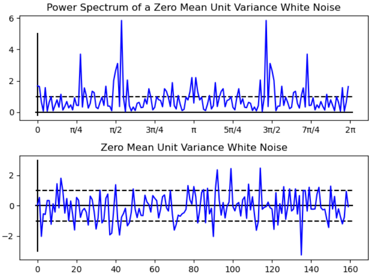

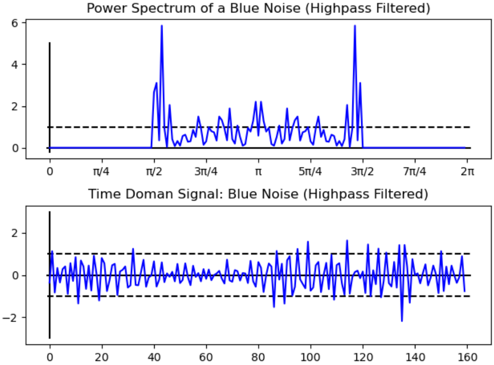

Outline

Power Spectrum of Filtered Noise

So we have $r_{yy}[n] = r_{xx}[n] * h[n] * h[-n]$. What about the power spectrum?

$$\begin{aligned} R_{yy}(\omega) &= \mathcal{F}\left\{r_{yy}[n]\right\} \\ &= \mathcal{F}\left\{r_{xx}[n] * h[n] * h[-n]\right\} \\ &= R_{xx}(\omega)|H(\omega)|^2 \end{aligned}$$

Example

Here's an example. The white noise signal on the top ($|X[k]|^2$) is multiplied by the bandpass filter in the middle ($|H[k]|^2$) to produce the green-noise signal on the bottom ($|Y[k]|^2 = |X[k]|^2|H[k]|^2$).

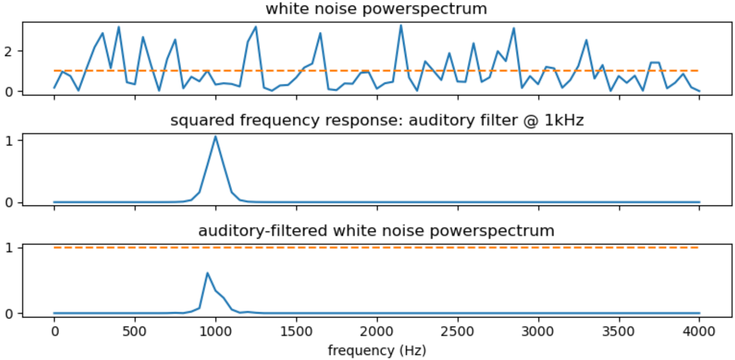

Units Conversion

The DTFT version of Parseval's theorem, assuming a finite window of length N samples, is

$$\frac{1}{N}\sum_n x^2[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

Let's consider converting units to Hertz. Remember that $\omega = \frac{2\pi f}{F_s}$, where F_s is the sampling frequency, so $d\omega = \frac{2\pi}{F_s}df$, and we get that

$$\frac{1}{N}\sum_n x^2[n] = \frac{1}{F_s}\int_{-F_s/2}^{F_s/2} R_{xx}\left(\frac{2\pi f}{F_s}\right) df$$

So we can use $R_{xx}\left(\frac{2\pi f}{F_s}\right)$ as if it were a power spectrum in continuous time, at least for $-\frac{F_s}{2} < f < \frac{F_s}{2}$.

Example: Power Spectrum of Colored Noise

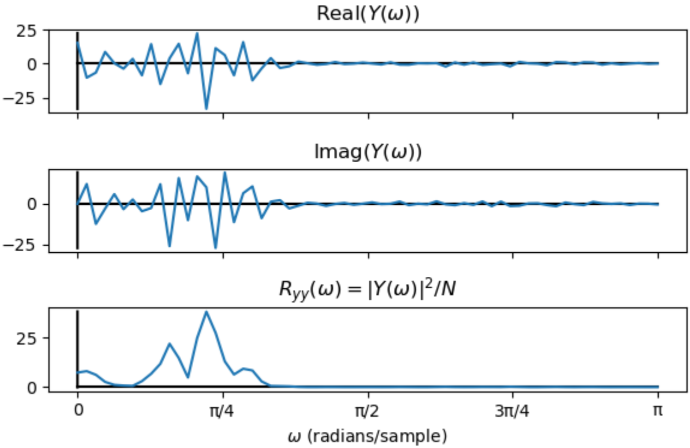

Example: Expected Power Spectrum of Colored Noise

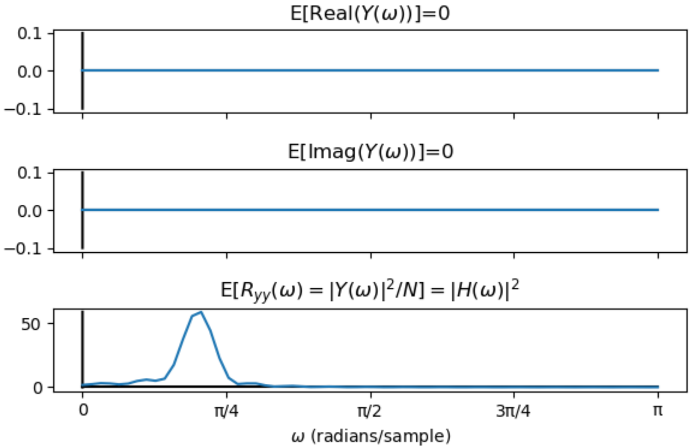

Outline

Parseval's Theorem

Now we have everything we need to prove Parseval's theorem. Let's prove the DTFT power form of the theorem, for a finite-length signal:

$$\frac{1}{N}\sum_{n=0}^{N-1} x^2[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

where

$$R_{xx}(\omega) = \frac{1}{N}|X(\omega)|^2$$

Parseval's Theorem

$$\frac{1}{N}\sum_{n=0}^{N-1} x^2[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

Notice that the left-hand side is the autocorrelation, with a lag of 0:

$$r_{xx}[m] = \frac{1}{N}\sum_{n=0}^{N-1} x[n]x[n-m]$$

So Parseval's theorem is just saying that

$$r_{xx}[0] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

Wiener's Theorem

Wiener's theorem says that the power spectrum is the Fourier transform of the autocorrelation:

$$r_{xx}[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)e^{j\omega n}d\omega$$

But notice what happens if we plug in $n = 0$:

$$r_{xx}[0] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

Q.E.D.

Outline

Example: Autocorrelation of White Noise

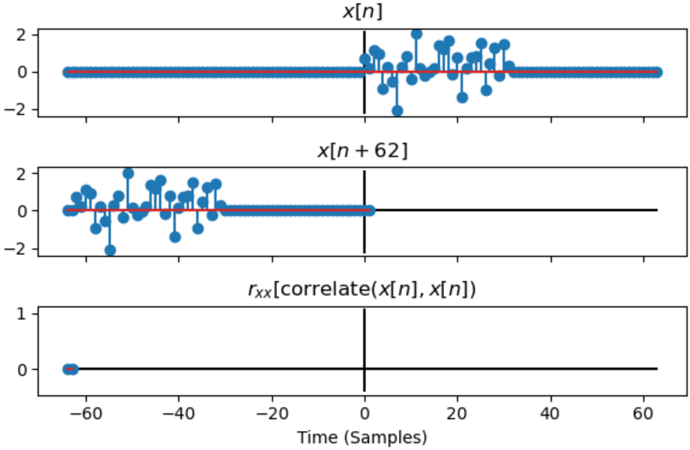

Example: Power Spectrum of White Noise

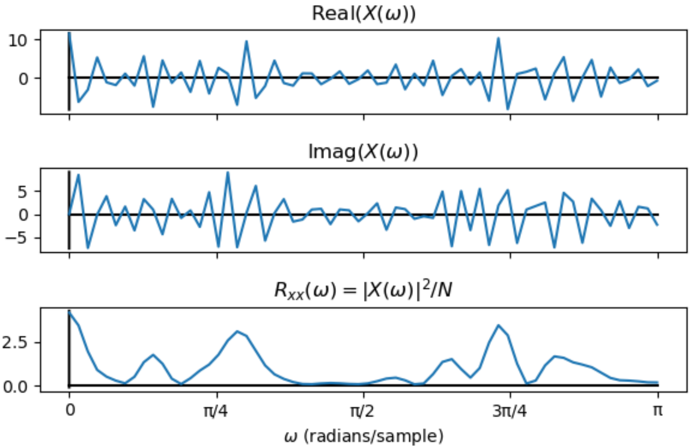

Example: Expected Power Spectrum of White Noise

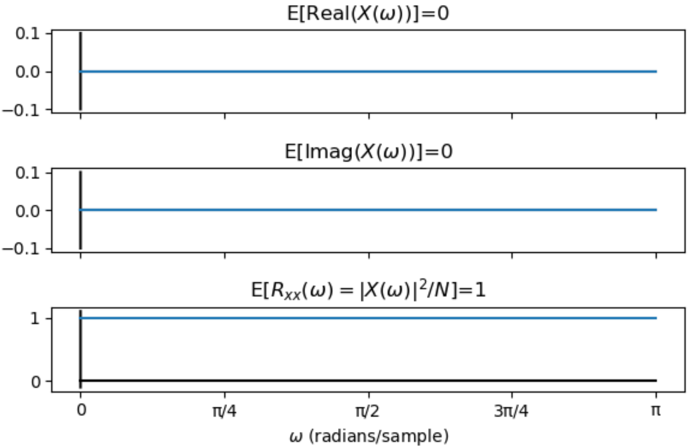

Example: Filtering of White Noise

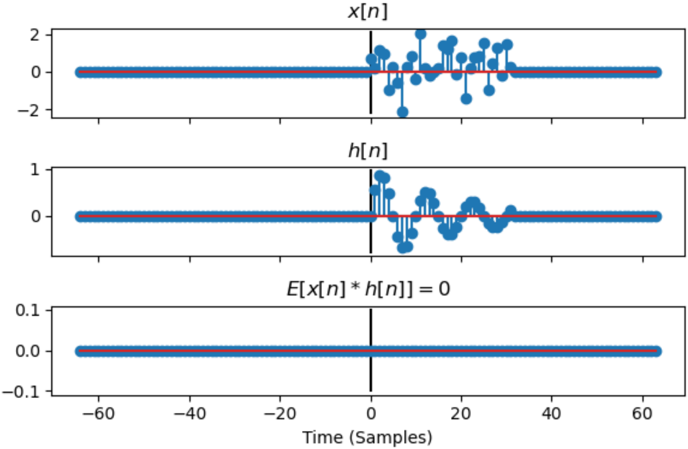

Example: Power Spectra of White and Colored Noises

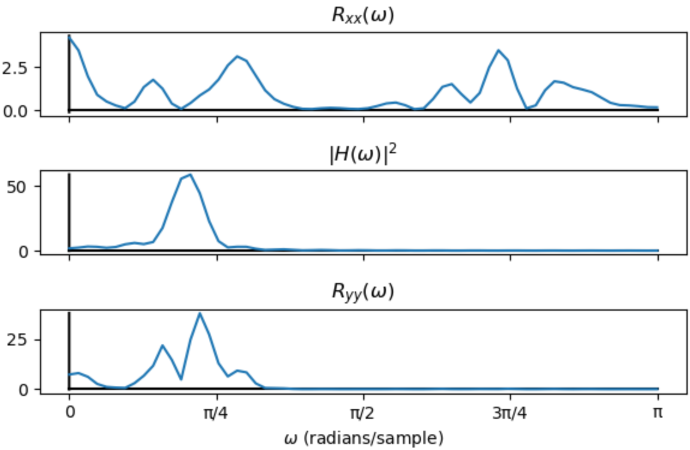

Example: Autocorrelation of Colored Noise

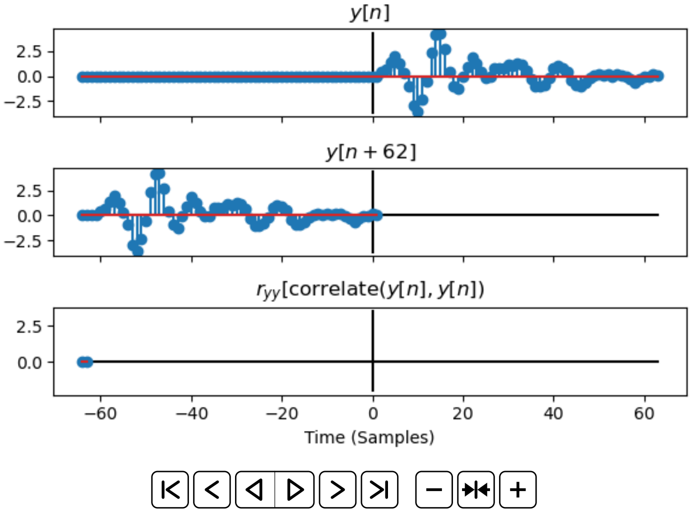

Example: Power Spectrum of Colored Noise

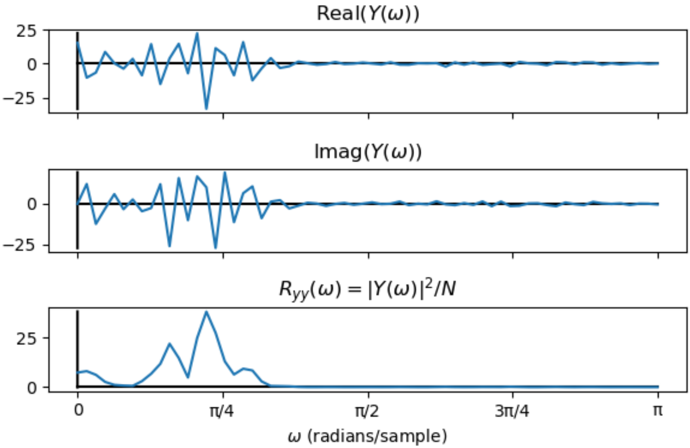

Example: Expected Power Spectrum of Colored Noise

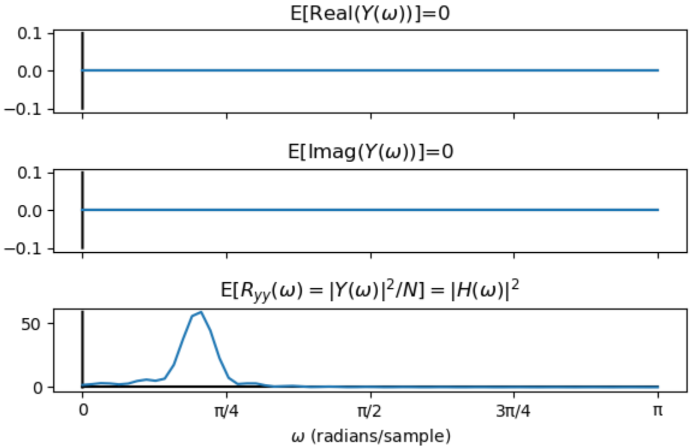

Outline

Wiener's Theorem and Parseval's Theorem

- Wiener's theorem says that the power spectrum is the DTFT of autocorrelation:

$$r_{xx}[n] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)e^{j\omega n}d\omega$$

- Parseval's theorem says that average power in the time domain is the same as average power in the frequeny domain:

$$r_{xx}[0] = \frac{1}{2\pi}\int_{-\pi}^{\pi} R_{xx}(\omega)d\omega$$

Filtered Noise

If $y[n] = h[n] * x[n]$, $x[n]$ is any noise signal, then

$$
\begin{aligned}
r_{yy}[n] &= r_{xx}[n] * h[n] * h[-n] \\
R_{yy}(\omega) &= R_{xx}(\omega)|H(\omega)|^2
\end{aligned}
$$

White Noise and Colored Noise

If $x[n]$ is zero-mean unit variance white noise, and $y[n] = h[n] * x[n]$, then

$$\begin{aligned} E\left[r_{xx}[n]\right] &= \delta[n] \\ E\left[R_{xx}(\omega)\right] &= 1 \\ E\left[r_{yy}[n]\right] &= h[n] * h[-n] \\ E\left[R_{yy}(\omega)\right] &= |H(\omega)|^2 \end{aligned}$$