Lecture 5: Discrete-Time Random Processes

Chapter 2.4

Charles A. Bouman

Purdue University

Wednesday, September 2, 2026

Outline

1. **Random processes and their specification**
   Finite-dimensional distributions, Kolmogorov, Gaussian processes.
2. **Stationarity**
   Strict-sense and wide-sense stationarity. Mean, autocovariance, and autocorrelation.
3. **The Gaussian simplification**
   Mean and autocovariance specify everything. The power spectral density and the Wiener–Khinchin theorem. Reversibility.
4. **Two dimensions**
   Random fields, 2-D stationarity, and the 2-D power spectral density.

Reading: Chapter 2.4, *Foundations of Computational Imaging*.

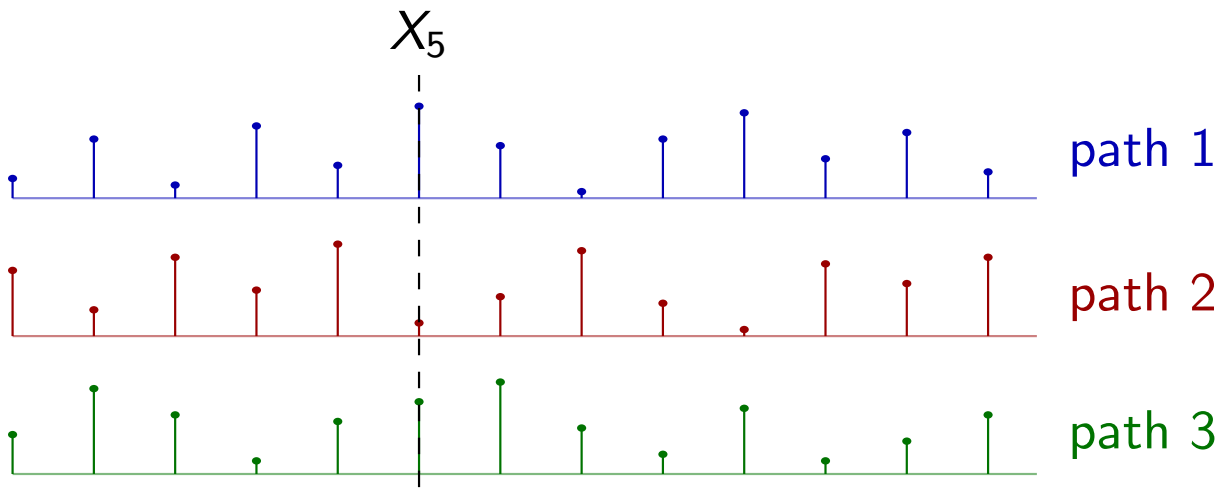
Today's Big Ideas

- A random process: infinitely many variables, one random object.
- Finite windows of the process specify its whole distribution.
- Stationarity: shifting in time does not change the distribution.
- A Gaussian process is completely specified by its mean and autocovariance. If it is also zero mean and stationary, the power spectral density alone specifies it.

In 2-D these same models describe images. We use them for the rest of the course.

Discrete-time random processes

A **discrete-time random process** X_n is a random variable defined for each integer $n \in \mathbb{Z}$. Each run of the experiment produces an entire sample path:



- At fixed n, X_n is one random variable that varies across paths.

The specification problem

A random process is an infinite collection of random variables. What does it take to specify its distribution?

Finite-dimensional distributions

The **finite-dimensional distributions** of a random process are the joint distributions of its finite windows,

$$F_{m,k}(t_m, \cdots, t_k) = P\{X_m \leq t_m,\ X_{m+1} \leq t_{m+1},\ \cdots,\ X_k \leq t_k\},$$

where m and k are any integers with $m \leq k$.

- These distributions are clearly necessary to specify the process.
- The next theorem says they are also sufficient.

The Kolmogorov extension theorem

Kolmogorov extension theorem

If the set of finite-dimensional distributions is consistent for all $m, k \in \mathbb{Z}$, then the distribution of the infinite-dimensional random process is well defined.

- Consistent finite windows determine the whole infinite process.
- See, for example, Wong and Hajek, *Stochastic Processes in Engineering Systems*, page 3.

Gaussian random processes

Definition

A discrete-time random process X_n is a **Gaussian random process** if for all $m \leq k$ the vector

$$[X_m,\ X_{m+1},\ \cdots,\ X_k]^T$$

is jointly Gaussian.

- This is a statement about every finite subset of the process, not about each X_n by itself. Recall from Lecture 3 that jointly Gaussian is stronger than Gaussian marginals.

Strict-sense stationary random process

It is often reasonable to assume that shifting a random process in time does not change its distribution. This form of time invariance is called **stationarity**.

Definition: Strict-Sense Stationary Random Process

A 1-D discrete-time random process X_n is said to be *strict-sense stationary* if X_n and its shifted version $Y_n = X_{n-k}$ have the same finite-dimensional distributions for all integers k. This means that

$$F_{m,m+k}(t_0, t_1, \cdots, t_k) = F_{0,k}(t_0, t_1, \cdots, t_k)$$

for all integers m and $k \geq 1$ and for all values $t_0, t_1, \cdots, t_k \in \mathbb{R}$.

Second-order random processes

Strict-sense stationarity is a strong condition. In practice it is much easier to check the second-order properties of a random process.

Definition: Second-Order Random Process

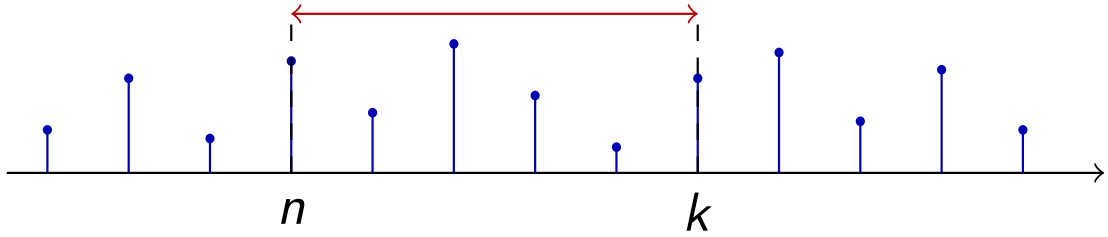
A 1-D discrete-time random process X_n is said to be a *second-order random process* if for all $n, k \in \mathbb{Z}$, $\mathbb{E}\left[|X_n X_k|\right] < \infty$, that is all second-order moments are finite.

Mean, autocovariance, and autocorrelation

If a random process is second order, then it has a well-defined mean μ_n and **time autocovariance** $R(n,k)$:

$$\begin{aligned} \mu_n &= \mathbb{E}[X_n] \ , \\ R(n,k) &= \mathbb{E}[(X_n - \mu_n)(X_k - \mu_k)] \ . \end{aligned}$$

$R(n,k)$: how strongly do these two values vary together?



We will also often talk about the **time autocorrelation** $C(n,k) = \mathbb{E}[X_n X_k]$.

- Autocorrelation equals autocovariance for a zero-mean process.

Wide-sense stationary random process

Definition: Wide-Sense Stationary Random Process

A 1-D second-order discrete-time random process X_n is said to be *wide-sense stationary* if for all $n, k \in \mathbb{Z}$,

$$\mathbb{E}[X_n] = \mu \qquad \text{and} \qquad \mathbb{E}[(X_n - \mu)(X_k - \mu)] = R(n-k) \ ,$$

where μ is a scalar constant and $R(n)$ is a scalar function of n.

- So a discrete-time random process is wide-sense stationary if its mean is constant and its time autocovariance is a function of only $n - k$, the difference between the two times.
- Strict-sense constrains all statistics; wide-sense constrains only two.

Gaussian processes are specified by mean and autocovariance

In general, specifying a random process requires every finite-dimensional distribution. For a Gaussian process, the mean and the autocovariance are enough. This is the main reason Gaussian models are used so widely.

Property: Autocovariance specification of Gaussian random processes

The distribution of a discrete-time Gaussian random process X_n is uniquely specified by its mean μ_n and time autocovariance $R(n, k)$.

- Every finite window of the process is jointly Gaussian, and a Gaussian vector is determined by its mean and covariance. Both are read from μ_n and $R(n, k)$.

Power spectral density: the Wiener–Khinchin theorem

- If a Gaussian random process is also stationary, then we can equivalently specify its distribution from its **power spectral density** $S_X(e^{j\omega})$.

Wiener–Khinchin theorem

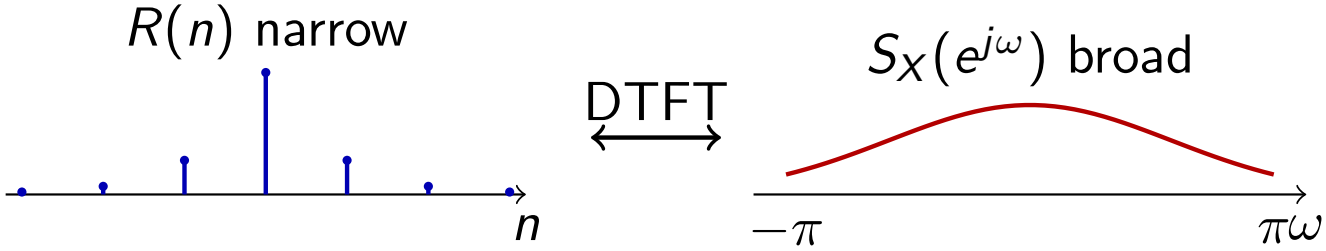
For a zero-mean wide-sense stationary random process, the power spectral density is given by

$$S_X(e^{j\omega}) = \sum_{n=-\infty}^{\infty} R(n)\, e^{-j\omega n} \ ,$$

the **discrete-time Fourier transform (DTFT)** of the time autocovariance $R(n)$.

The power spectral density specifies the distribution

- Since the DTFT is invertible, $S_X(e^{j\omega})$ is equivalent to $R(n)$.



- A narrow autocovariance describes a fast-changing process, and it spreads power across all frequencies. A broad autocovariance describes a smooth process, and it concentrates power at low frequencies.

Property: Power spectral density specification

The distribution of a zero-mean wide-sense stationary discrete-time Gaussian random process X_n is uniquely specified by its power spectral density $S_X\left(e^{j\omega}\right)$.

Reversible random process

A random process is **reversible** if it has the same distribution as its time reverse X_{-n}. Reversibility is a second invariance property, alongside stationarity.

Definition: Reversible Random Process

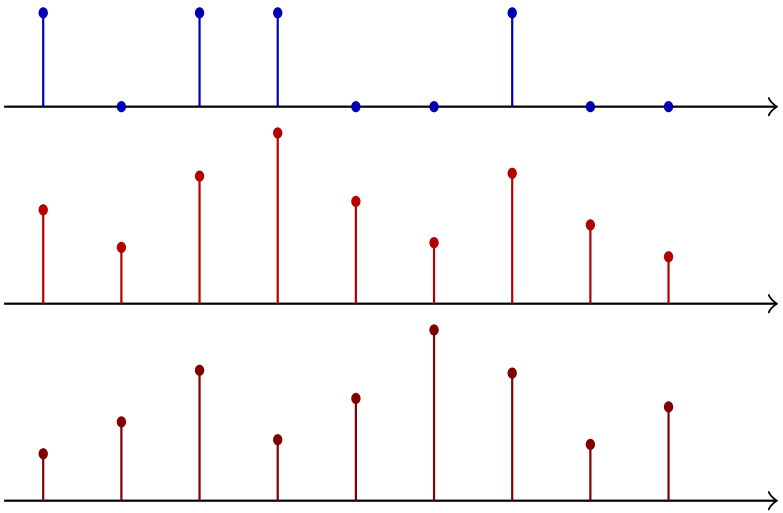
A 1-D discrete-time random process X_n is said to be *reversible* if its finite-dimensional distributions have the property that

$$F_{0,k}(t_0, t_1, \cdots, t_k) = F_{m,m+k}(t_k, t_{k-1}, \cdots, t_0)$$

for all integers m and $k \geq 1$ and for all values $t_0, t_1, \cdots, t_k \in \mathbb{R}$.

Not all random processes are reversible

Let X_n be an i.i.d. Bernoulli process, which is a random binary sequence. It is stationary and reversible, because it has the same distribution run forward and run backward. Now filter it with a causal IIR filter.



X_n: reversible

Y_n: sharp rise, slow decay

Y_{-n}: slow rise, sharp fall

- Filtered Y_n is stationary, but not reversible: its time reverse differs.

Reversibility of Gaussian random processes

Let X_n be a wide-sense stationary Gaussian random process, and let $Y_n = X_{-n}$ be its time reverse. By the definition of time autocovariance,

$$\begin{aligned} R_Y(n) &= \mathbb{E}\left[(Y_k - \mu)(Y_{k+n} - \mu)\right] \\ &= \mathbb{E}\left[(X_{-k} - \mu)(X_{-k-n} - \mu)\right] \\ &= \mathbb{E}\left[(X_{l+n} - \mu)(X_l - \mu)\right] \\ &= R_X(n) \ , \end{aligned}$$

where $l = -k - n$, and μ is the mean of both X_n and Y_n. So X_n and Y_n have the same autocovariance, and by the autocovariance specification property they have the same distribution.

Property: Reversibility of stationary Gaussian random processes

Any wide-sense stationary Gaussian random process is also strict-sense stationary and reversible.

Generalization to 2-D

Strict-sense stationarity, wide-sense stationarity, autocovariance, power spectral density, and reversibility all generalize naturally to 2-D and higher dimensions.

Let X_{s_1,s_2} be a 2-D discrete-space random process that is wide-sense stationary. Its 2-D autocovariance function is

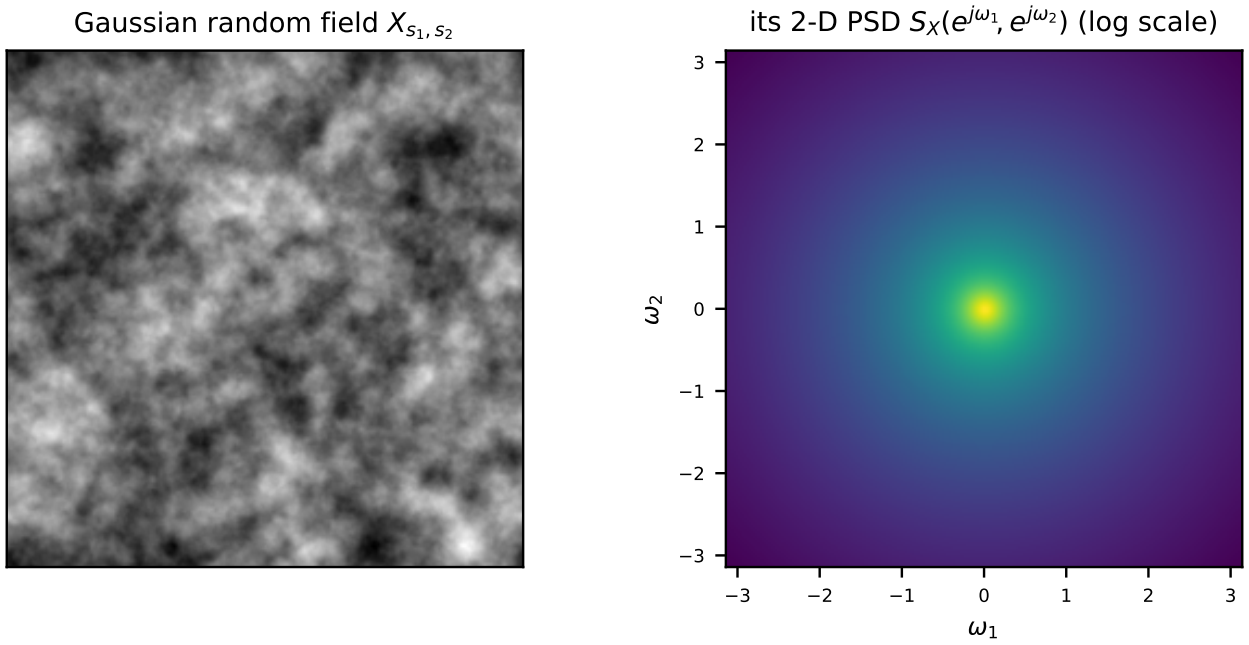
$$R(s_1 - r_1,\, s_2 - r_2) = \mathbb{E}\left[(X_{s_1,s_2} - \mu)(X_{r_1,r_2} - \mu)\right] \;,$$

where $\mu = \mathbb{E}\left[X_{s_1,s_2}\right]$. Its 2-D power spectral density is

$$S_X(e^{j\omega_1}, e^{j\omega_2}) = \sum_{s_1=-\infty}^{\infty} \sum_{s_2=-\infty}^{\infty} R(s_1, s_2)\, e^{-j\omega_1 s_1 - j\omega_2 s_2} \;,$$

the **discrete-space Fourier transform (DSFT)** of the 2-D autocovariance.

A stationary Gaussian random field



- Left: one realization of a 2-D stationary Gaussian random process, called a **random field**. Any shifted window of the image has the same distribution, which is what 2-D stationarity means.
- Right: its power spectral density; low frequencies give smooth texture.
- Our first image-like random object; later image models build on it.

Stationarity and reversibility in 2-D

- Stationary: X_{s_1,s_2} and its shift by (k_1, k_2) have the same distribution.
- Reversible: X_{s_1,s_2} and $X_{-s_1,-s_2}$ have the same distribution.

Again, it can be shown that 2-D zero-mean wide-sense stationary Gaussian random processes are always reversible and strict-sense stationary.

Summary

- A random process is an infinite collection of random variables X_n. Its distribution is specified by its finite-dimensional distributions $F_{m,k}(t_m, \cdots, t_k) = P\{X_m \leq t_m, \cdots, X_k \leq t_k\}$, and by the Kolmogorov extension theorem those are enough.
- Strict-sense stationarity means that no finite-dimensional distribution changes under a shift, $F_{m,m+k} = F_{0,k}$.
- Wide-sense stationarity constrains only the first two moments, $\mathbb{E}[X_n] = \mu$ and $\mathbb{E}[(X_n - \mu)(X_k - \mu)] = R(n-k)$.
- Gaussian: specified by mean μ_n and autocovariance $R(n,k)$.
- If it is also zero mean and wide-sense stationary, the power spectral density $S_X(e^{j\omega}) = \sum_n R(n)e^{-j\omega n}$ specifies it. This is the Wiener–Khinchin theorem.
- Such a process is also strict-sense stationary and reversible. In 2-D the same results hold, with the DSFT of the 2-D autocovariance.

Reading: Chapter 2.4. **Next lecture:** Chapter 3.1–3.2.