Random Signal Principles

Introduction

A physical process either natural or synthetic (human-made) can be viewed as a signal with time-varying features and properties that can be modeled as a random process. The term signal implies that some type of information is represented by the physical process, and random means that future outcomes of the process are unpredictable to some degree.

Natural random signals include:

- ultraviolet radiation impinging on a tree
- a planet revolving around a star
- a tornado moving across an open field

Synthetic (man made) examples of random signals:

- a microwave signal transmitted from a cell phone to a base station
- an automobile traveling from Los Angeles to San Francisco
- a baseball thrown by pitcher to a catcher.

These examples could be modeled using a function $g(x, y, z, t)$ that describes a trajectory in the three spatial dimensions $[x, y, z]$ as it varies over continuous time t. Obviously these "signals" have different physical mechanisms different levels of predictability, and contain different amounts of information. The figure below shows possible trajectories (called realizations) of a tornado in the x-direction as a function of time

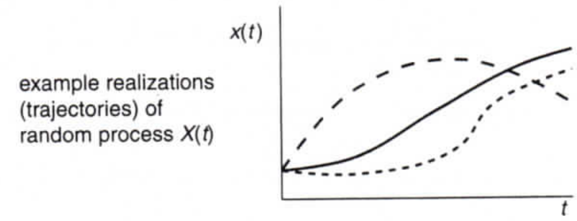

Three possible trajectories of a random process measured along the x-direction as a function of time t.

It is important to note that an observed trajectory of the tornado is one possible realization of its motion in the x-direction. Before observing the actual trajectory, many realizations (even an infinite number) could occur over some region. Alternatively, if we were able to "restart" the process and allow the tornado to proceed again (repeat the "experiment") and with different atmospheric conditions, then we would expect a different trajectory, which might be only slightly different from the first realization

We denote a random process by $X(t)$ (uppercase letter) and a realization of the process by $x(t)$ (lowercase letter). The outcome at a specific time t_o is also random: it is a random variable $X(t_o)$ and a particular outcome at that time is $x(t_o)$

A random process is a collection of random variables that are indexed by time as depicted in the figure below for one realization $x(t)$ and for one outcome $x(t_o)$ at $t = t_o$

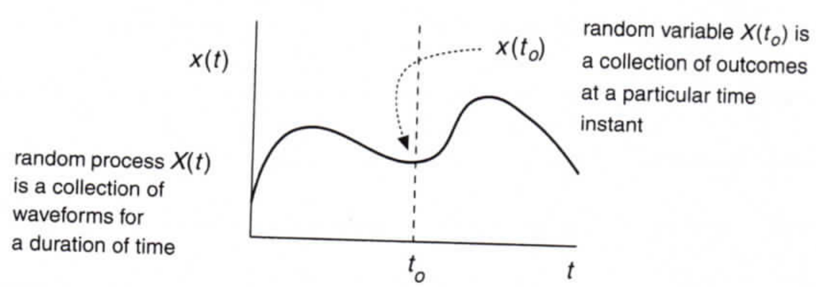

Realization $x(t)$ of random process $X(t)$ and outcome $x(t_o)$ of random variable $X(t_o)$ at time $t = t_o$.

Clearly, we cannot predict the trajectory of a tornado with certainty, even if we have numerous measurements about its velocity, the ground temperature, time of the day, geographic location, terrain, and so on. This is discussed further for the following simple example: a single toss of a fair coin as shown in the figure below

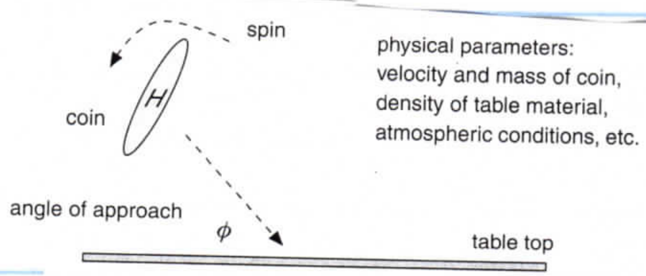

Physical approach to predicting whether a coin will land heads (H) or tails (T).

Excluding the "very unlikely" event that the coin lands on its edge, this experiment has, of course, only two outcomes: heads (H) and tails (T). Since we are interested only in how the coin lands (H or T) at a particular time (and not its trajectory) a single toss of a coin is modeled as a random variable and not as a random process. Given measurements such as the velocity of the coin just before impact, the angle of approach ϕ with respect to the table top, the mass of the coin, atmospheric conditions (temperature, humidity, and so on), and other relevant physical parameters, it might be possible to predict the outcome just before the coin lands.

We can never have sufficient physical measurements about most random events in order to predict an outcome without error (unless, of course, there is trivially only one outcome)

From the previous discussion, randomness can be viewed as a lack of complete information about a physical process, such that we cannot predict exactly an outcome or realization before it occurs. Thus it would be useful to develop a model of randomness that does not depend directly on physical attributes of a particular process, and can be applied generally to different types of signals and random experiments. For the coin toss example using a fair coin, we expect from intuition that with an increasing number of tosses, an equal number of H's and T's will occur. This viewpoint involving repeated experiments is known as the frequency interpretation of a probability model: the frequency with which an event occurs for a large number of repeated experiments determines its probability. If the coin is tossed N times and N_H heads are observed, then we can expect for M tosses of the coin that approximately $\left(\frac{N_H}{N}\right) M$ heads will occur. The ratio $\frac{N_H}{N}$ is the frequency of occurence for H. Likewise, the frequency of occurence for T is $\frac{N_T}{N}$ (since $N_T + N_H = N$) and we would expect to see approximately $\left(\frac{N_T}{N}\right) \cdot M$ tails for M tosses of the coin.

The frequency approach to developing a probability model is consistent with our notion of the likelihood of observing the various outcomes in repeated experiments

It turns out that this interpretation is the intuitive basis for developing a probability model based on three axioms. Observe that $0 \leq \frac{N_H}{N} \leq 1$: the frequency of occurence lies in the interval $[0,1]$. Let us denote the probability of some event E by $P(E)$. One axiom of probability is $P(E) \geq 0$. This lower bound is appealing not only from the frequency interpretation, but also as a mathematical representation. Another axiom states that the probability of "something happening" in an experiment is 1. Combining these two axioms gives $0 \leq P(E) \leq 1$, which are the same bounds for the frequency interpretation of a random experiment.

The third axiom is more complicated: it is concerned with combinations of events, and is crucial to assigning probabilities to any event of interest in an experiment. For example, in the coin-toss experiment, consider the trivial probability $P(E = T \text{ or } H)$. Obviously, these are the only outcomes (again excluding the possibility that the coin lands on its edge). Moreover they are mutually exclusive because either H or T occurs: both cannot happen on the same toss. We find that the frequency of ~~occurence~~ occurrence of either heads or tails is $\frac{(N_H + N_T)}{N}$; the frequencies simply add because the outcomes are "nonoverlapping" (they are disjoint). In addition, since no other outcomes are possible in this example, we ~~must~~

must have $N_H + N_T = N$ (the total number of tosses) and $\frac{N_H + N_T}{N} = 1$ as expected. This last result also implies that the probabilities of mutually exclusive events add. This is the motivation for the third axiom. Based on these three axioms, a probability model can be developed for a random variable, which can be extended to a random vector and then to a random process

Signals, Signal Processing, and Communications

Definition: Signal. A signal is a physical quantity propagating through space and time that "contains" information about an event. It propagates from one physical location to another through some medium.

There are many signals in the modern world such as those used in commercial radio and cable television. However, one can argue that nearly everything in the universe is a type of signal. For example, the light from an exploding star is certainly a type of signal that provides information about an extraordinary event. But a signal need not be restricted to electromagnetic radiation, which is perhaps the most familiar type of signal. A meteor passing through Earth's atmosphere could also be viewed as a signal, possibly providing information about the current location of Earth with respect to the Sun (such as the month of the year). An earthquake generates a signal that provides information about an event in the Earth's crust, and a volcanic eruption likewise imparts information caused by events far below the Earth's surface.

The figure below shows a block diagram of the three components of a signal model: (i) the source of the signal, (ii) the signal itself and the medium through which it propagates and, perhaps most importantly (iii) one or more sensors that perceive (observe or measure) the signal.

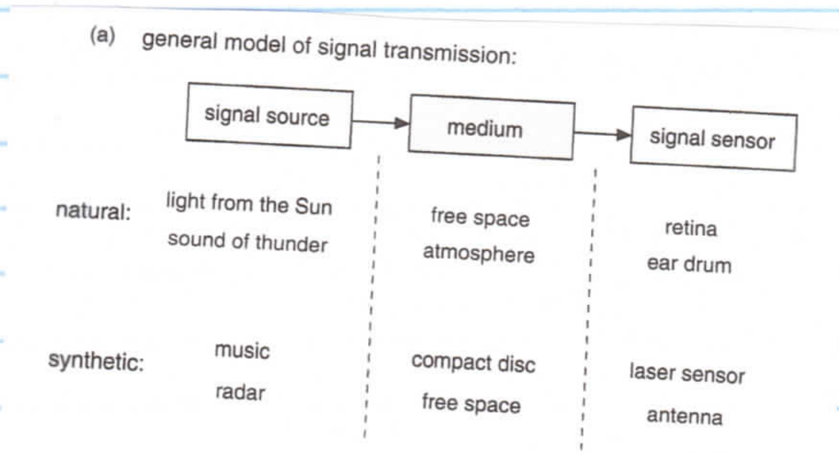

Natural signals are due to events in the physical world, such as the examples mentioned above. Other examples include the rise and fall of the tide, the movement of clouds, or a tree falling in a forest. The last example is often cited in the famous question "If a tree falls in the forest and no one is there to hear it, does it make a sound?" The falling tree and the disturbance it creates in the surrounding air can be viewed as a type of signal, and it may signify another physical event that just occurred (such as lightning striking the tree). The question above is obviously concerned with the sensor components of the signal model shown above. Sound is perceived vibrations in the ear drum, and thus if no one is present, then no sound will be heard. However it can be argued that the air has been disturbed by the falling tree, so a signal has been generated; moreover, additional signals (vibrations) are produced by the tree's impact with the ground.

Example

Consider the trajectory of a jet aircraft traversing an ocean. Its path contains "random" fluctuations due to turbulence that causes it to deviate from a "predictable" trajectory. If the trajectory is viewed as only along the x-axis, then we would see a realization of a one dimensional signal (simila to the tornado example mentioned earlier)

Thus, even an aircraft in flight can be viewed as a type of signal. If this "experiment" could be repeated, then the next realization would be different due to turbulence and course adjustments. Although random fluctuations of a signal itself might be small, we ~~also~~ would also need to take into account the fact that perfect measurements of a signal are not possible because sensors have limited accuracy and are generally exposed to noise.

This leads us to the sensor component of the signal model which is of particular interest because we intend to (i) develop a probabilistic models for signals based on observations from one or more sensors, and (ii) derive techniques for modifying signals to achieve some desired goal. A probabilistic model is required because it is not possible to obtain enough measurements to predict exactly how a signal will evolve. For the coin-toss experiment, although it might be possible to predict its outcome (H or T) reasonably well from numerous physical measurements, this approach is not practical. It is better that we accept the fact that an experiment/signal is random, develop a probabilistic model for the outcomes/realizations, and then exploit that model for a particular application. Randomness at the sensor end of a signal model can be viewed as a measure of the uncertainty in a received signal because it is not possible to identify all the underlying physical mechanisms responsible for generating and modifying the signal. We cannot quantify exactly all disturbances that interfere with the signal as it propagates through some medium.

The figure below shows a communication system model that is relevant to signals.

specific model of a communication system:

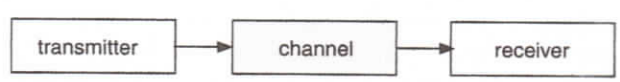

It consists of (i) a transmitter for generating and sending the signal, (ii) a channel that includes the medium of propagation as well as any impairments encountered before reception, and (iii) a receiver consisting of one or more sensors for detecting the signal. The communication model explicitly assumes that some type of information is transmitted and received. Consider a synthetic signal such as that transmitted by a cell phone. Obviously, the speaker knows the information contained in the transmitted signal. Randomness can be viewed as a property of the receiver because the listener does not know what the speaker intends to say, and also due to signal distortions caused by channel impairments. The receiver has incomplete information.

According to Information theory, the amount of information contained in a signal is related to its degree of randomness: signals with "greater randomness" contain more information. For example if a small planet suddenly has extreme orbital fluctuations due to an impact with a comet, that randomness provides more information than when the orbit is close to being perfectly elliptical with small variation. The same can be said for ~~the~~ synthetic signals shown in figures below: a sinusoidal signal with fixed amplitude a, frequency f_0, and phase ϕ is at one extreme: it is deterministic and thus perfectly predictable. As such, it contains essentially no useful information. A passband binary phase shift keying (BPSK) signal is obtained by multiplying the sinusoidal carrier in the figure with a random sequence of ± 1 called pulse amplitude modulation (PAM) The BPSK is "more random"

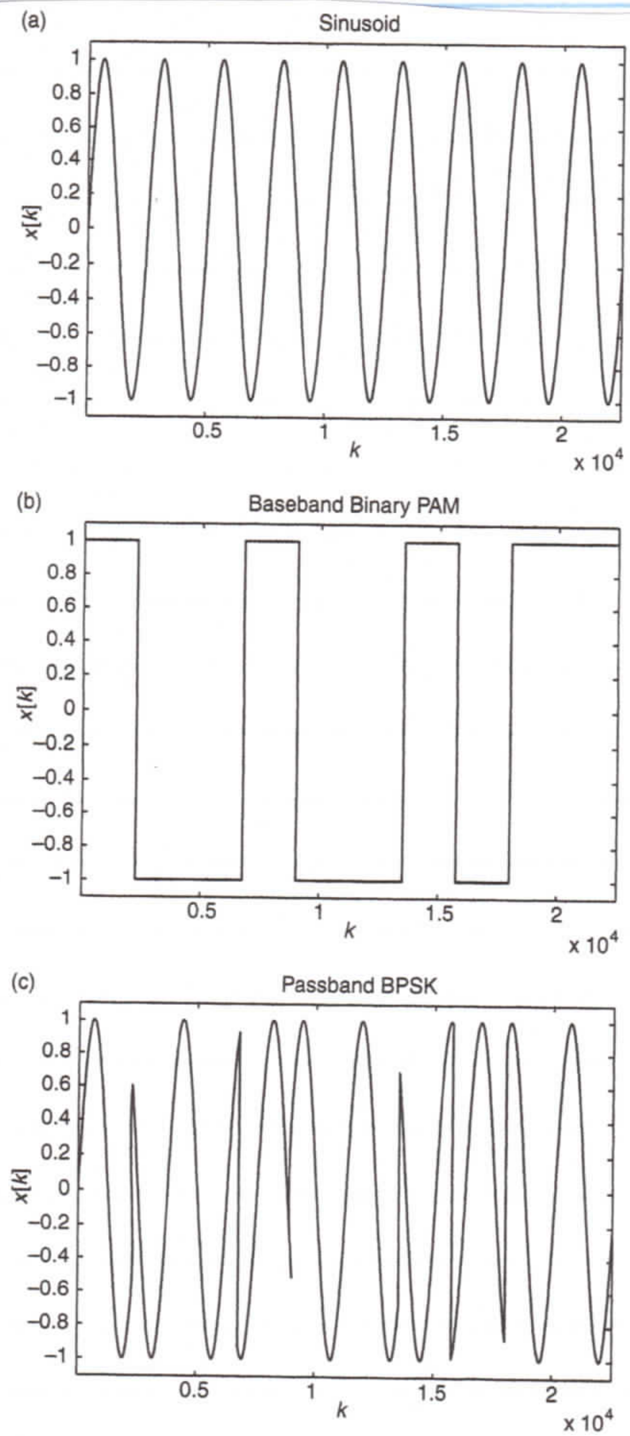

The speech signal shown below exhibits even more randomness and thus contains greater information than the previous synthetic signals. The figure shows also a realization of a sound from wind and a "white noise"

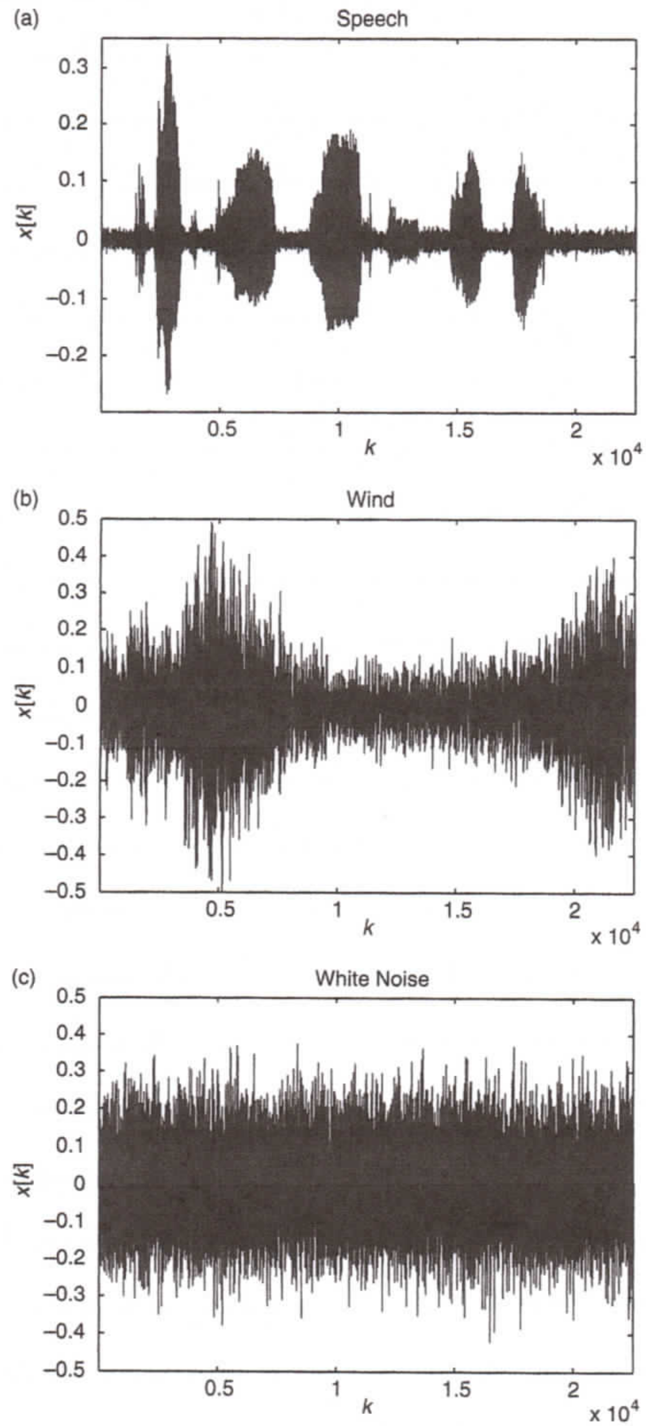

There is a difference between ~~rar~~ information and meaning of a signal

information : ~~measured by~~ content measured by its randomness

meaning : depends on the receiver, language used: meaning is an interpretation of what is received. Here we are interested in information as defined

In the following figure, two basic types of signal processing are shown. The first is estimation of the parameters of some underlying model of the signal. The received signal is manipulated to derive estimators of the parameters. Perhaps the most well-known estimator is the sample mean

$$\hat{X} = \frac{1}{N} \sum_{k=1}^{N} X[k]$$

where N samples of $X(t)$ (denoted by random sequence $X[k]$) are averaged to estimate the true mean μ_X of the signal. The second type of processing is filtering where the goal is to transform the signal usually such that it has "better" properties. This is achieved by convolving the signal with filter coefficients $[h[k]]$

$$Y[k] = \sum_{n=0}^{\infty} h[n] X[k-n]$$

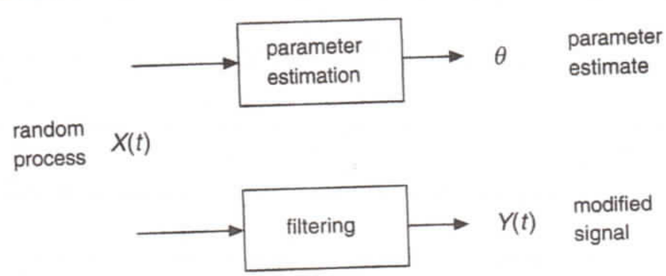

Two basic types of signal processing.