

Big Idea (Week 4): Time-Frequency Representations and Conversion

There are many different ways to represent information. For example, we might represent time as (1) (year, date, month, and hours) in Universal Time, (2) elapsed seconds since midnight (UT), January 1, 1970, (3) years into the reign of the current monarch, days since the last equinox, and hours since sunrise or sunset, or (4) days, hours, and minutes to graduation. Which one is most convenient depends on what we want to do with the information. Some representations may be more natural for data collection, others may be more convenient for certain kinds of processing, and some representations may expose more structure in the data allowing greater compression or understanding.

In the case of sound, samples in time are a very direct way to capture and reproduce acoustic pressure waves. We can take these samples directly by converting analog voltages from a microphone to digital samples, and we can use them directly to drive a speaker by converting the digital samples back to voltages. However, time samples are not the most efficient way to represent the common structure that often occurs in the patterns of acoustic pressure waves that humans call “sound” because of the manner in which those patterns are perceived. For example, arrays of time samples do not offer the most directly compressible representation. For many operations it is more convenient to describe, work with, and store sounds using a representation based upon frequency content rather than time varying amplitudes. Musical notes, for example, can be represented directly by their frequencies, with one or a few frequencies often defining the sound wave exactly for a period of a very large number of time samples.

As an alternative to representing a sound sequence as a discrete set of amplitude samples in time, we can also represent sound as a weighted sum of “harmonics”— discrete sine and cosine waves whose frequencies are multiples of a base frequency, the “fundamental”. That is, rather than representing a sound wave as a vector of time samples:

$\{\dots, \text{Sample}[-3T], \text{Sample}[-2T], \text{Sample}[-T], \text{Sample}[0], \text{Sample}[T], \text{Sample}[2T], \text{Sample}[3T], \dots\}$,

we can represent it as a vector of coefficients: $\{a_{c0}, a_{s1}, a_{c1}, a_{s2}, a_{c2}, \dots\}$ chosen such that:

$$\text{Sample}[t] \approx a_{c0} \cos(0t) + a_{s1} \sin(t) + a_{c1} \cos(t) + a_{s2} \sin(2t) + a_{c2} \cos(2t) + \dots$$

for all $t \in \{0, T, -T, 2T, -2T, \dots\}$.

The relationship between time and frequency representations for sampled signals can be geometrically interpreted as a linear change of basis between different frames of reference in the (typically high dimensional) vector space comprised of all possible time samples. This brings the powerful and highly efficient computational machinery of linear algebra to bear on the problem of audio signal processing.