Signal Processing Algorithms & Architecture

Dr. Anirban Dasgupta
Department of Electronics & Electrical Engineering
Indian Institute of Technology Guwahati

Lec 13: Introduction to Time-Frequency Analysis

Hello everyone, welcome to a fresh new lecture on the topic of introduction to time frequency analysis. This is Dr. Anirban Dasgupta and let us get started. So, we have already covered two modules, the first on time domain signal processing. and the second on frequency-domain signal processing and what have we learned? So, we can represent a signal like a sinusoid in the time domain, and in the time domain, we see how the signal amplitude varies with time.

From this, we can estimate or compute the frequency of the signal or of the sinusoid, and in one second, we see that there are about 50 complete oscillations. So, I can say that the frequency of this signal is 50 Hertz. The second modality is analysis in the frequency domain. So here we see that it has a frequency of $+50$ and -50.

So, both domains give an idea of what the signal is about of course, there can be a shift in phases. So, just the amplitude or the magnitude spectrum will not provide that information. So, you also need the phase spectrum but related to the frequency content, we can understand what the signal is all about.

And again, since this is a windowed version of the signal, this is just a part of the signal, so we do not get exactly one delta or one impulse at $+50$ and -50. But rather, we see that it is not very distinguishable; however, there is a finite spread because of this windowing effect. Now then, what is the problem? This looks very nice. So again, if we have a sum of sinusoids, for example, say we have a signal that is a sum of 10 hertz, 20 hertz, 30 hertz, 40 hertz, and 50 hertz. In the time domain, we can see that this signal is like this.

Now, of course, just looking at the time-domain signal, it is really difficult to say that this is composed of these five frequencies. But if we see the frequency content in the spectra, we can see that all five components are present: 10, 20, 30, 40, and 50. So, what is the problem again? What about this signal? Now, this is a very interesting case. So here in the first second, I have a sinusoidal signal of 10 hertz because we can see 10 complete oscillations and in the 2nd seconds, we have 20 oscillations, or we can say from 0 to 1 hertz the frequency is 10 hertz, from 1 to 2 hertz the frequency is 20 hertz, and from 2 to 3 hertz the frequency is 30 hertz.

Now you can calculate and count, and you will know that it is 30 hertz. From 3 to 4 seconds, the frequency is 40 hertz, and from 4 to 5 seconds, the frequency is 50 hertz.

And what about the spectra? Well, the spectra look very much similar to the waveform in the previous slide like this. Of course, because of this windowing effect, you see a lot of leakage because this window is short for each 10 and 20, and all the frequencies exist only for 1 second. So we get larger spectral leakage.

Whereas here the window is of course, for 1 second but this is combined for each so the leakage found here seems to be minimal. But what is the problem? The problem is that in the frequency analysis, we are able to detect all the frequencies present in the time-domain signal. But we do not know where that frequency is present. Like the frequencies can be a combination or superimposed version of all the sinusoids, or each sinusoid is occurring separately but at different time durations.

So this is a critical problem. Let us look at another problem. So here we see that we have a sinusoid of how many? 10 Hz and this is what we observe, except for the period of 2 to 3 Hertz, where we have a sign of 10 Hertz, multiplied or, say added to another sign of 50 Hertz. So, if we look at the spectra, we see that the 10 hertz component is dominant because it appears throughout.

However, we also get components at the 50-hertz location. So, we know that there is some 50 hertz present in the signal, and mostly it is a 10 hertz sinusoid. But where this is present, that information is not available in the frequency spectra. But if you see the time-domain signal, we know that the signal is oscillating at 10 hertz except for the interval from 2 to 3 hertz. So, in this duration, we know that this is not just a 10 hertz signal; there is something more to this 10 hertz signal.

But we do not know what frequencies are present. So, frequency spectra tell me what frequencies are present, but where they are, I do not know. But the time domain is telling me that there are some spectra, and while I can understand where they are, I do not know what frequencies they correspond to. So, just knowing one version of the signal may not be sufficient to get a clearer picture of the holistic thing and this is often what happens in the real world.

So real signals are not purely sinusoidal throughout, there can be places where high frequencies exist, and in some areas, there may be low frequencies. So, if we just look at the time, we can see where my faster oscillations are and where the slower oscillations are but we cannot exactly determine the frequency just by looking at the waveforms. Whereas in the frequency spectra or frequency domain we can locate what frequencies are present, we do not know where they are. So that brings up the concept of time frequency analysis.

So, what is time-frequency analysis? So, this is concerned with the simultaneous analysis of signals in time and frequency domains. Like, if you see this signal. So here I can see the signal in time and how the signal is behaving. I do not know what frequencies are present in this signal. But I can say that, okay, in the early stages of the signal, the signal has a very slow change with some kind of noise, and in the middle stage, there are a lot of fluctuations that I can observe over time.

But what frequencies are there, I do not know. Whereas in the frequency domain, by observing the energy spectral density, we know that these are the dominant frequencies. But where these dominant frequencies are occurring, we do not know. So, if we look at this kind of representation, which is called a time-frequency representation, often this is also called a spectrogram. So here we can see which frequencies are dominant at this time.

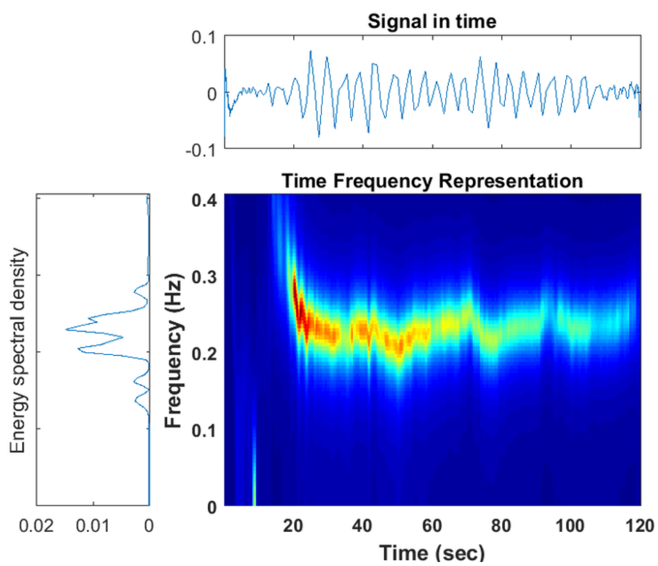

So this two-dimensional representation shows how the signal is behaving with respect to both time and frequency and the color map is giving the amplitude of that specific frequency at a specified time window. Why I said the window will be clear soon. So, this is a way to represent a signal in both time and frequency simultaneously. And this plot is often called a time-frequency representation (TFR) or time-frequency distribution (TFD).

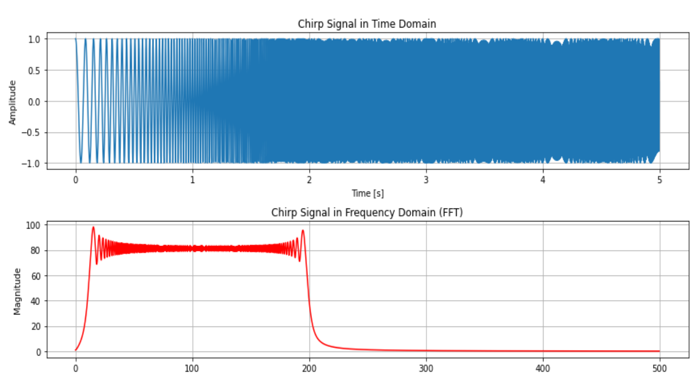

So, this is another real example of an EEG signal, and this kind of physiological signals like EEG and ECG are widely used in this time-frequency representation because here the practical challenges arise that I want to not only see the frequencies but also where those frequencies are present. So again, if you see in time, we see that the oscillations are faster here; they may be slightly slower here, and in the frequency, we see that near 10, you have a dominant oscillation or dominant frequency, but this is the dominant frequency throughout the signal. So, if you see the spectrogram or the time-frequency representation, here you can see that the red markings are showing that the intensity of that frequency is high. So, this location, this location, and this location, this 10 hertz frequency is very dominant, whereas in this location this frequency is absent because you get a blue line, which means low amplitude, and in these locations, these lower frequencies are present, which are below 5 hertz. So why do we need time-frequency analysis? So, one thing I have already explained, but let us still try to understand.

So, non-stationary signal analysis. Now, in time domain signal classification, we discussed that signals can be deterministic or stochastic, and in stochastic signals, they can be stationary or non-stationary. Stationary means that if I move throughout the signal, the signal properties like mean, variance, and its distribution or spectra will not change. While in non-stationary signals, these characteristics will change over time at different instances. Like speech signals, for example, different phonemes have different frequency patterns, and these patterns occur at different times.

So joint time-frequency representation helps us see both which frequencies are present and where they are present. So, these two together help us get a very good idea that if a frequency is present in the signal, it does not mean that it is present throughout the signal in the time domain. It may be occurring in specific regions in time and that brings us to the concept of localized events, like, say, if I am observing a seismic signal and then there is a sudden vibration. So, that duration when there is a saying like this, like this, and suddenly this and this.

So, this is a high-frequency oscillation, and this is not present throughout; this is present in a specific window in time. So, if I can get this information about which area in time this frequency content is related to, that marks an event. Similarly, when I am recording the EEG signal and moving my eyes like this, there is a specific frequency content present at that location. I want to capture that specific moment in time when those specific frequencies representing the eye movements are present. And also improving signal processing.

If I know that the region is an artifact in the signal, I can easily filter it out or remove it.

Also, if I want to have some idea of noise, then denoising reconstruction will be very useful by doing a time-frequency analysis. So, these are applicable in areas like communication, radar, and biomedical signal processing, and the tasks are filtering and denoising; they are kind of the same, and this helps us better analyze and improve signal processing. So now we have a good introduction to what time frequency analysis is. Let us understand what the popular time frequency analysis methods are.

So, the first is the short-term Fourier transform. So, what is the short-term Fourier transform? As the name suggests, it is a Fourier transform that gives me frequency, but it is a short time; I will not do the Fourier transform of the whole signal together. I will take small windows to do the Fourier analysis, and in this manner, I can get the spectra of each window. But typically, this window size is kept the same in the short-time Fourier transform. What about the wavelet transform? Wavelet transform extends this idea to a non-sinusoidal basis.

Like if I have a signal that has a sharp change like this. So probably if I use a sinusoid, it might not be very convenient, and it may take a lot of coefficients or a lot of components to approximate this signal. But rather, if I analyze this kind of signal with a sharp, small module or basic signal, which I call a wavelet, then this will be easier and faster to analyze. Then the Wigner-Willett distribution is somewhat similar to the short-term Fourier transform, but it has a smoother representation and does not follow fixed window limitations like SDFD. And then finally, we have the Hilbert-Huang transform (HHT), which is also popularly known as empirical mode decomposition.

Now, one may argue that empirical mode decomposition is a purely time-domain analysis technique and that the concept of frequency, although not present, is addressed by the Hilbert-Huang transform, which decomposes any signal into two modes- One is an oscillatory mode, and one is a trend. So, a trend means that if the signal has some component that is going up or going down, that is a linear trend; the trend may not be linear, but this is the non-oscillatory component of the signal and then the oscillatory components can be components that are varying slowly; that is, you can say it is a band of frequencies, but the frequency term is not used in HHT, or it can be very fast-moving. Components are called IMFs, or intrinsic mode functions.

So, if we see an example of time-frequency representation, one of the best examples is to see the chirp signal. Now, where do we hear the word "chirp"? So, in the morning, we often hear the birds chirping. So, what is so special about a chirp? The chirp is a signal where the frequency slowly increases or decreases with time, and there are linear chirps where the frequency increases or decreases linearly with time; this is an example of a linear chirp. This is the spectrum of a chart, and from this, it is very difficult to identify that it is composed not only of many frequencies, but that these frequencies occurring separately at different time locations, such as this being a high-frequency region and this being a low-frequency region. So, in this kind of signal, time-frequency analysis is very useful because with just the time domain I do not know what the frequencies are, and in the frequency domain I do not know where it is occurring, because this is kind of a flat spectrum that denotes there is an impulse, but actually it is something totally different.

So, in the time-frequency representation in a spectrogram, which is typically used in STFT, this is what we are getting. So here the frequency axis is given in a logarithmic scale, and this is the time axis. Here we see that around 2 seconds, the frequency content is low. And as my time increases, my frequency component is also increasing. See, the yellow color is giving you the highest magnitude.

That means this is the dominant frequency, and this dominant frequency is increasing with time. So, here we get a very nice representation of what frequencies are present and where they are present. To explain the graph, take the chirp signal as an example; this is a frequency-modulated signal in which the frequency increases or decreases over time. Here, it is increasing, and the frequency starts at 10 Hertz. If we look closely, we can see that it starts at 10 Hertz.

$$s(t) = A \cdot \sin\,(2\pi(f_0\,t + \frac{f_1 - f_0}{2T}\,t^2))$$

$$s(t) = A \cdot \sin\left(2\pi\left(10\,t + \frac{200 - 10}{2 \times 5}\,t^2\right)\right)$$

I have programmed it that way, and it increases up to 200 Hertz; unfortunately, that is pretty difficult to see because it looks continuous here, and it spans over 5 seconds. That is this. So if we plug in this generic value, this is a general equation of a linear chirp signal. So A sine(2πft) is the starting frequency, which is 10 hertz. And then plus f_1 minus f naught by 2T, where F1 is the ending frequency of 200 hertz, f_0 is the starting frequency of 10 hertz, and 2T, where T is my duration of 10 seconds, and this T and this T square, this T square is giving me the instantaneous time T and t^2 .

So, plugging in the values, we get this relation that I have plotted in the time domain. So, this type of signal is commonly used in radar and communication systems to observe how frequencies change over time. So, in this spectrogram, we see that the time-frequency plot is generated, and the colors are based on the magnitude of the Fourier coefficients for each time segment and frequency. We have also provided the color bar. If you can see, this is the color bar, and this yellow region shows the high value of frequency amount or magnitude, while black means low value, which indicates that in these regions, the frequencies are not present.

So, if you take just a section that is the Fourier series or Fourier transform of that specific window, and we have used a logarithmic scale because a linear scale would make it harder to see, as the frequency is increasing very fast, like we have 10 squared and 10 cubed. So, in a logarithmic scale 10, 10 squared, and 10 cubed will be equally spaced. So, we see that this STFT provides a dynamic view of how the frequency contents are changing over time. Unlike our simple Fourier transform, which will give a single spectrum of what frequencies are present. So, what are the challenges in time-frequency analysis? The first is the time-frequency trade-off.

Now, what does it mean? So, for example, if I draw a complete sine wave and this is my time t, then I can know the frequency is 1 by t at least if I draw half, I am still able to know what this frequency is. But if I draw one fourth of the wave, it is probably that frequency, but it may also increase. So, if I localize in time very small, like if this is my curve, what does it mean? I do not know.

It may mean a high frequency or a low frequency. It is very difficult to say. And if I just give a dot, like at a specific instant, if I say that at t equals t naught, tell me what frequencies there are. Can I say this point can represent any frequencies? So, if I localize in time very accurately at this exact moment, then I am very uncertain about what frequencies are present. This is looking like an impulse; everything is there, and the higher I visit, because this is again a sign multiplied by a rectangular window, the more cycles I plot of the sign, the better. Like this will have a frequency at something like 1 by T.

But the more cycles I plot, the narrower the spectra will be. And if I go with an infinite cycle, then I get a precisely like spectrum at that point. So, the question is whether if I localize too much in time, I do not get the exact frequency. I am very uncertain about the frequency. But if I increase the time plot, I am very certain about the frequency.

But then I am uncertain about time because I am getting a wider window. Right? So, if I am very precise in frequency and can exactly tell that 50 hertz is present, then I have to go the entire time from minus infinity to infinity. And if I want to say that at t equals 5 seconds, what the frequency is, I do not know. I am absolutely clueless.

So, there is a trade-off between time and frequency. Like if I say there is a delta t, I am observing the signal. So, in this $\delta(t)$ time, tell me what the frequencies $\delta(f)$ present in the signal are and if this $\delta(t)$ decreases, the uncertainty in $\delta(f)$ will increase and similarly, if the certainty in frequency increases, the uncertainty in time will decrease, and vice versa.

So that is the uncertainty principle. This makes the selection of window size very critical. Because what constitutes a good window for achieving optimal resolution in time and frequency? Another factor is complexity. Because calculating 1D vectors in the time domain or frequency domain is easier than computing or dealing with matrices.

And this will give high-dimensional data. The spectrograms. So, this is a real challenge, and people are still working on it because, in real-time applications like radar, telecommunications, or even medical signal processing, this is a critical challenge in using time-frequency analysis methods. Also, noise is already a challenge, be it in time domain or frequency domain analysis, and in time-frequency analysis, because you are using windowed data, you may sometimes not be sure whether a specific pattern is due to noise or a signal component, which makes filtering and denoising extremely difficult and complex. So, to start with these challenges, one of the basics is to use adaptive resolution. And what is adaptive resolution? Seeing the signal at different rates is the concept of multirate signal processing. So, here the signal can be downsampled for analyzing low frequencies and upsampled when focusing on high frequencies.

So, this method is the fundamental basis for studying wavelet transforms, and by using multi-rate or multi-resolution, which is sometimes referred to in the case of images, this helps to analyze the signal very nicely. So, we will understand what multirate signal processing is in our next topic. Thank you so much. Have a nice day.