Time Series Modelling and Forecasting with Applications in R

Prof. Sudeep Bapat

Shailesh J. Mehta School of Management

Indian Institute of Technology Bombay

Week 09

Lecture 41: Frequency Domain Analysis

Hello all, welcome to this course on Time Series Modelling and Forecasting using R. So, now we are starting with a new week, a fresh week right and then in front of you you can see the broad topic for this week which is spectral analysis or frequency domain analysis of a time series ok. So, we will spend again a whole lot on this broad idea about how can one convert a time series to a spectral component or how can one visualize a time series rather than a time series component to a frequency domain component. In order to analyze any time series, one can actually analyze a time series in two different ways. So, one would be the actual temporal.

So, temporal means over different time points, right? Let us say monthly data or weekly data, daily data, annual data, etc. At the same time, the other idea to visualize a time series is from a frequency domain. So, frequency domain means can we capture all the oscillations in the underlying time series? using some sinusoidal curves or let us say either sine curves, cosine curves or some combination of sine curves and cosine curves.

Make sense? So, just to again summarize as to what I have spoken till now is that a time series can be visualized using two different ideas. So, one would be a visualizing as to how we have been doing it so far right. So, through a time component or the other way to visualize the same series is through a frequency component.

So, frequency component means that you can sort of combine a bunch of sine curves and cosine curves and then create a linear combination of them. And superimpose them on the underlying time series. Because again, remember that any underlying practical time series contains all these oscillations, right, or all these frequencies. So, you have some seasonalities, you have some ups and downs, right. So, all these ups and downs or cyclicality in the underlying time series can be

superimposed by a bunch of sine curves and cosine curves, okay. So, this is the broad idea about how you convert the visual angle from a time component kind of framework to a frequency domain, okay. So, this is a broad idea, and again, today's session will be more like an introductory because the entire angle would probably be different to many of you, right. I mean, if you are sitting in any basic time series course, then this is not the predominant idea that people teach you, right. I mean, people teach you all the earlier models that we have covered so far.

So, AR, MA, ARMA, ARIMA, SARIMA, right, and probably, let us say, Volatility Modeling. So, we will cover Volatility Modeling soon in the next coming weeks. But this is one angle to visualize any time series slightly differently from a frequency domain approach. So, hopefully, I have cleared the air about what exactly you mean by frequency domain or spectral analysis and so on and so forth. All right.

Now, the very first thing is: What is spectral analysis? So, spectral analysis is nothing but decomposing a time series into, as mentioned before, different sine functions and cosine functions of different frequencies. Now, again, for this, you might require a very elementary sort of revision about trigonometry. So, sine curves, cosine curves—and again, do not worry. So, even as you progress through today's session, I will throw in some very introductory kind of relationships of

let us say, sine and then cosine, etcetera, OK? So, again, what exactly spectral analysis is, is nothing but decomposing an actual practical time series into a bunch of sine and cosine functions of different frequencies. So, again, remember that any of the sine and cosine functions can have different frequencies. So, what do you mean by that? So, I will give you an example.

Now, typically, this is exactly how a sine curve looks like, right? It starts at 0, then it oscillates around the x-axis like that, and similarly, any cosine curve might look like that. So, it will start—so at 0, it will be 1—and again, it will go down and then come back again, something like that, OK. Now, what do you mean by different frequency is that: Can you sort of translate the frequency or the amount or the length of the curve? So, for example, can I somehow tweak the sine by incorporating some constant here such that, let us say, probably the curves would be something like that?

Now, again, this would not be exactly a sine curve; this will be a combination of different sine and cosine curves, ok. So, the idea is to decompose a time series into sine and cosine functions coming from different frequencies. Then, the next thing is that cycle dynamics are predominant in many applications. So, firstly, one should understand where exactly one can apply the idea of spectral analysis. So, one can apply the idea of spectral analysis whenever one can see some cycle dynamics.

So, cycle dynamics mean nothing but oscillations. So, one example could be, let us say, infectious disease data. So, let us say environmental factors, periodicity of immunity, etc. So, if you are handling any disease data, then one can assume that the immunity would be periodic. There will be seasons when immunity is high, and there will be seasons when immunity is low, right?

Of course, depending on the outbreak of the disease or the severity of the disease, right. So, one can assume that all these applications would be some sort of oscillation around some constant, ok. Or, probably, the second example would be business cycles. So, let us say, how can one judge the magnitude and periodicity of cycles, right? I mean, how high and low are the periodicities, how high and low are the oscillations, right?

So, again, these are some elementary examples. Of course, we will see a lot of other examples down the line, but again, just to summarize in one line: wherever one can actually see some cycle dynamics that contain oscillations around a constant, one can break it apart and try to analyze it using a spectral analysis framework. So, now we can go into the intricacies of spectral analysis. So, again, decomposing a stationary time series—now, one assumption could be that the time series is stationary to start with. So, decomposing a stationary time series y_t into a combination of sinusoids—now, again, sinusoids is a collective term for sine and cosine functions, alright.

So, from now on, we will see sinusoids. So, sinusoids mean a combination of sines and cosines, okay. So, decomposing a stationary time series y_t into a combination of sinusoids with random and probably uncorrelated coefficients. So, it will be something like—just to give you an example. So, let us say the first coefficient could be, let us say, φ_1 applied on sine of a_1 plus φ_2 applied on cosine of probably a_2, and then so on and so forth.

So, this is just an example of what we are trying to express, right? And by the way, these two are the coefficients. So, φ_1 and φ_2 are the coefficients, right? One assumption is that these coefficients are random and uncorrelated. Makes sense so far? Now, the second point is to identify those frequencies that appear particularly strong or important, right? Now, the whole idea is, once you kind of

analyze or kind of put forward this frequency domain approach, then how can you say that which frequencies are kind of significant, right? So, those frequencies will be significant which appear to be pretty much strong or important. Make sense? And last point is can be viewed as a linear multiple regression problem. I mean, of course, this can be clearly viewed as a multiple regression problem, isn't it?

So, the only thing is you have bunch of independent variables which are nothing but sinusoids. And on the left hand side, you have the actual time series. So, y t equals that. And why exactly? Because we are trying to compare this structure underlying in y t with a bunch of sinusoids like that.

So, the last thing is one can actually view the spectral analysis as nothing but a linear multiple degradation problem. The dependent variable is y t as we have discussed so far and the independent variables are the sine cosine functions of all possible rather discrete frequencies. So, sine and then frequency is a 1, cosine frequency is a 2 etcetera. Make sense hopefully so far. Now, the next point is that time domain approach considers regression on past values of time series and shocks.

So, so far what we have been studying is the time domain approach, right. So, on one hand you have the time domain approach, and then all the models developed so far, right. So, ARMA, ARIMA, SARIMA, etc. All those were inside the time domain approach, right. So, the time domain approach considers regression on past values of the time series and shocks, right.

What do you mean by this is that yt depends on, let us say, yt minus 1, yt minus 2, plus probably some shocks. So, let us say epsilon t, plus something like that. So, if you consider the time domain approach, the current value of yt is regressed on its own past values plus some shocks. But when you transition from the time domain approach to the frequency domain approach, the frequency domain approach considers regression on sinusoids instead. So, rather than yt minus 1, yt minus 2, yt minus 3, one can actually create a regression on sinusoids.

So, sine functions, cosine functions, etc. Hence, the spectral density function is required for studying the frequency properties of a time series. So, we actually have a thing called the spectral density function, right? So, once you transition from the time domain to the frequency domain approach, one actually has to make use of a spectral density function, okay? So, a spectral density function is required for studying the frequency properties of a time series.

And then the next point is that the inferences regarding the spectral density are called the analysis in a frequency domain. So, whatever inferences one makes regarding the spectral density or rather the spectral density function are called the analysis in a frequency domain. So, if you want to analyze a time series in a frequency domain, one actually has to start with spectral analysis and by developing the spectral density function basically. Make sense so far? So, this is exactly what I was saying about earlier that just to throw you or just to tell you very very briefly as to what some preliminary trigonometric results are right.

So, some results to start with. So, let us say sine a plus or minus sine b. So, this would be 2 into sine of a plus and minus b by 2 and then cosine a minus plus b by 2. So, here so initially I am writing down plus and minus here, but then minus plus here this just means that the sine here is different from the sine here. So, if this is plus, this will be minus, if this is minus, this will be plus, depending on of course, whether this is plus or minus.

$$\sin A \pm \sin B = 2 \sin \frac{A \pm B}{2} \cos \frac{A \mp B}{2}$$

So, in a way, the first identity contains both the things. So, sine A plus sine B or sine A minus sine B. So, sine A plus sine B would be nothing but 2 sine A plus B by 2 cosine A minus B by 2. Similarly, sine A minus sine B would be 2 sine A minus B by 2 cosine A plus B by 2. Hopefully, this is clear. So, what exactly or rather how should one read this plus and minus collectively?

Okay. So, now cosine A plus cosine B is 2 cosine A plus B by 2 cosine A minus B by 2. And the next is cosine A minus cosine B is nothing but minus 2 of sine A plus B by 2 sine of A minus B by 2. Okay. So, unfortunately, one cannot combine these two identities in a collective manner.

$$\cos A + \cos B = 2 \cos \frac{A + B}{2} \cos \frac{A - B}{2}$$

$$\cos A - \cos B = -2 \sin \frac{A + B}{2} \sin \frac{A - B}{2}$$

So, I have kept them separate. But the first one, one can do that. Okay. And now, if I am not wrong, the next slide also tells you some other brief elementary trigonometric results. So, for example, sine of any k pi happens to be 0, of course, for any k which is let us say plus or minus 1, plus or minus 2, plus or minus 3, etc.

$$\sin k\pi = 0,\ \forall\ k = \pm 1, \pm 2, \ldots$$

$$\cos k\pi = (-1)^{|k-1|},\ \forall\ k = \pm 1, \pm 2, \ldots$$

$$\sin(-A) = -\sin A$$

$$\cos(-A) = \cos A$$

And why is that? Because this is a property of the sine function. So, if you have something like sine pi or let us say sine of 2 pi or sine of 3 pi. So, all these are 0. So, this is more like a repetition.

So, again if you go with the sine curve also right. So, at 0 it is 0 and every extension or every transition of pi. So, pi, 2 pi, 3 pi, 4 pi, 5 pi value of sine happens to be 0 everywhere. And what happens to cos then? So, cosine of k pi happens to be this number.

So, minus 1 to the power of k minus 1. minus 1 to the power of k minus 1. For example, cosine, so if you plug in again the same values for k. So, for all values of k ranging from plus and minus 1, plus and minus 2, etc. Just one single example here is that let us say if k is 1, what would happen? So, if cosine of pi because k is 1 happens to be minus 1 raised to the power 1 minus 1.

Now, 1 minus 1 is 0 and this would be 1 basically. Now, if k is 2, what would happen? So, if k is 2, this would be nothing but cosine of 2 pi and then this would be minus 1 raised to the power 2 minus 1 because k is 2. 2 minus 1 is 1, this is minus 1. So, in a way the cosine function oscillates between every transition of pi.

So, cosine pi is 1, cosine 2 pi is minus 1, right, etc. Then the next thing is sine of minus a is minus sine of a, and then cosine of minus a is cosine a, ok. So, again, just to summarize some elementary results of trigonometry so that you can start somewhere, basically, ok. All right. So, now the next thing we will study is, again, this is technical.

So, do not get bogged down with all the notations here. So, I will try to explain them in a bit clearer fashion. But the heading that you see is called Fourier and inverse Fourier transforms. Now, again, this is so I have put this initially because this slide just tells you what exactly is the definition of a Fourier transform. So, the discrete Fourier transform, or in short, one can call this as a DFT.

So, many textbooks use the acronym DFT. So, DFT stands for discrete Fourier transform of any function, let us say HT, ok. So, if you have a function HT, then one can actually write down or create the DFT or the discrete Fourier transform of that function HT. And obviously, the values that T can take can be any integer value. So, let us say minus 1, 0, 1, 2, 3, 4, even on the negative side, ok.

$$H(\omega) = \sum_{t=-=\infty}^{\infty} h(t)e^{-i\omega t} \ (-\pi \leq \omega \leq \pi)$$

So, if you start with the function ht, then what exactly the dft of that function ht is given by this value. So, capital H of omega. So, omega is a parameter. This is nothing but an infinite sum. So, t goes from minus infinity to infinity.

Of course, this t stands for this t here. So, the index goes from all the indices that t can take from minus infinity to infinity. Then you have this function here which is h t and then e to the power minus i into omega t. Now, again this i is nothing but the under root of minus 1, this is the imaginary number. So, the Fourier transform of any function is that you take that function h t multiplied by e to the power minus i omega t. So, omega is a parameter and by the way omega is a parameter lying between minus pi to pi. All right.

And similarly, the inverse Fourier transform of H omega. So, once you write down the Fourier transform, which is H omega, I can write down the inverse Fourier transform of H omega. So, inverse Fourier transform of H omega is given by small h of t, which is given by this function. So, 1 by 2 pi integral from minus pi to pi of the Fourier transform, which is H omega and then multiplied by the same constant, which is e to the power minus i omega t and then d omega. Now, as in this is not a very, very really difficult kind of a statement here.

$$h(t) = \frac{1}{2\pi}\int_{-\pi}^{\pi} H(\omega)\, e^{-i\omega t} d\omega$$

So, if you understand each and every part here, right. So, let us see if you take any function h(t), you simply plug that function here, you have an infinite sum here on the outside, and you multiply by this constant, which is e to the power minus i omega t. Now, again, if you are very good in math, let us say in the school days or even in college days, then you might remember that you have some resemblance between this constant, which is e to the power minus i omega t, and sine function and cosine functions. I can actually write down sine of theta as a function of e to the power i omega t plus cosine of—I mean, you know—so there is some resemblance between this guy, this constant here, and the sine functions and cosine functions. So, in a way, all these things are kind of tied together, right.

So, Fourier transform, inverse Fourier transform, spectral density, right, that involves sines and cosines, etc., okay. Now, a couple of properties of the inverse Fourier transform is that, let us say, if h(t) happens to be the same as h of minus t—so h of t happens to be h of minus t—so the function is symmetric in a way. Then, I can actually write down the Fourier transform as this. So, h of omega happens to be small h of 0. Plus, the summation would now only run from 1 to infinity because all the negative signs are the same as the positive signs, right?

If $h(t) = h(-t)$, then,

$$H(\omega) = h(0) + \sum_{t=1}^{\infty} h(t)(e^{-i\omega t} + e^{i\omega t})$$

Or,

$$H(\omega) = h(0) + 2\sum_{t=1}^{\infty} h(t)\cos\omega t, \ \ (-\pi \leq \omega \leq \pi)$$

And then, t going from 1 to infinity, h of t, and then e to the power minus i omega t plus e to the power i omega t, or in other terms, h of omega is nothing but small h of 0 plus 2 into summation t going from 1 to infinity h of t and then cosine omega t. So, this is exactly what I was mentioning just a short while back, that you have some relation between this combination. So, e to the power minus i omega t plus e to the power i omega t, and then 2 into cosine of omega t. So, from here, can you—can you say that, can you tell me what exactly is the relationship now? So, in a way, 2 into cosine of omega t would be nothing but this thing. So, e to the power minus i omega t plus e to the power i omega t, or in other words, cosine omega t would be nothing but one half of that.

So, if you bring this 2 here, it will be one half, is it not? And then, similarly, one can actually establish a relationship between sine of omega t and these two individual entities, OK. So, further, since you already know that cosine of minus omega t happens to be cosine of omega t, I can actually replace or reduce the inverse Fourier transform as follows. So, small h of t, which is nothing but the inverse Fourier transform, happens to be 1 by 2 pi integral from minus pi to pi of the Fourier transform h of omega e to the power i omega t d omega, which reduces to this guy. So, 1 by 2 pi integral from 0 to pi.

$$h(t) = \frac{1}{2\pi} \int_{-\pi}^{\pi} H(\omega)\, e^{i\omega t} d\omega = \frac{1}{2\pi} \int_{0}^{\pi} H(\omega)\, (e^{i\omega t} + e^{-i\omega t}) d\omega$$

$$= \frac{1}{\pi} \int_{0}^{\pi} H(\omega) \cos(\omega t)\, d\omega$$

So, instead of minus pi to pi, it happens to be 0 to pi again because of this relation, right? h of omega and then the same thing. So, e to the power i omega t plus e to the power minus i omega t d omega, which again further reduces to this because of the same relation. So, cosine omega t has a relationship between this and the constant one half. So, actually, you have to spend some more time if you are really interested or keen in kind of exploring more about the intricacies of, let us say, trigonometric functions and e to the power i omega t or e to the power minus i omega t and then all these constants. So, further, suppose the time series contains a periodic component at a known frequency.

Suppose the time series contains a periodic component at a known frequency; then, I can actually write down the time series as follows. So, the time series y t can be written down as capital R cosine of omega t plus some nu plus u t. And what exactly are all these individual entities? So, omega is nothing but the frequency of the periodic variation. Omega lies between 0 and 2 pi.

$$y_t = R \cos(\omega t + \vartheta) + u_t$$

ω: frequency of the periodic variation ($0 \le \omega \le 2\pi$)

R: amplitude of variation

ϑ: phase

$\{u_t\}$: purely random process

So, omega is the frequency. Capital R is the amplitude. So, amplitude of variation. So, amplitude means the height. Of any of the frequencies, right?

Then, this nu stands for phase. And again, do not worry. So, we will try to explain each and every entity, each and every quantity in some other lecture which will follow this week, okay? So, omega—again, just to summarize—omega is the frequency, capital R is the amplitude, nu is the phase, and then the last thing, ut is a purely random process. So, ut is a random variable.

Make sense? So, suppose the time series contains a periodic component or an oscillating factor; I can actually write down the time series in this manner. So, the amplitude multiplied by cosine of the frequency into t plus the phase plus ut. So, ut is a purely random process. And further, if R and nu are constants, then I can actually derive that the expectation of the underlying time series yt is nothing but R multiplied by cosine of omega t plus nu, which varies with t, and hence the process becomes non-stationary.

$$E(y_t) = R\cos(\omega t + \vartheta)$$

Further, if you assume that capital R is a random variable with 0 mean and finite variance or this new quantity follows a uniform distribution between 0 to 2 pi, then my expected value of y t would be 0 and the process becomes stationary. So, I think this slide just tells you that under some restrictions or under some specific choices, one can actually get the expected value of y t to be 0, ok. Now, the last thing we will focus on in today's lecture is a few practical applications of the Fourier transform. So, where exactly can one apply this entire idea about Fourier transform or spectral density or converting from a time domain to a frequency domain etc. So, predominantly a lot of applications could be found in signal processing.

So, the first example is signal processing, let us say audio processing. And even within audio processing, where exactly? So let's say noise reduction. So isolating or removing noise frequencies from an audio signals. This could be one application.

Second one could be equalization. So equalization means what? So adjust the balance of frequency components for a better sound quality. So, again if you if you kind of think in your heads that if you have a sound wave right. So, a sound wave can actually be split into the same frequency domain right.

I mean any sound wave has some frequencies has some amplitudes right. So, can you superimpose let us say some time series from a frequency domain approach to sound waves for let us say audio processing or in general some signal processing. So, if you transmit a signal right. Any signal also contains all these intricate, let us say, frequencies and oscillations, amplitude, phases, etc. So, cannot we sort of superimpose a particular frequency domain time series on a, let us say, some audio signal or in general some signal?

The answer is yes, of course. Then probably in the same feature, the same domain, the third one could be compression. So, represent audio signals more compactly by prioritizing significant frequency components. For example, MP3 encoding, etc. Then the next one within signal processing could be image processing.

So the first one is edge detection. So one can actually use frequency domain analysis to detect sharp transitions in images. So one can actually use frequency domain analysis to detect some sharp transitions in images. So if you have any deteriorations in let's say images or if you see any transitions in any of the images. then those could be potential outliers in the images and one can actually detect those using a frequency domain analysis, all right.

Or the next one could be, let us say, image compression. So, the JPEG format uses discrete cosine transform. So, in short, DCT, which is a Fourier-related transform, compresses the images, OK. So, JPEG, which is a very famous kind of image format, uses the discrete cosine transform, in short, DCT. which is a Fourier-related transform to compress the images.

Okay. Okay. And then, the next thing that one can actually apply the Fourier transform to is in communication. So, the first one is modulation. So, let us say, if you want to convert some signals to higher frequencies for transmission.

So, let us say, AM, FM radio, TV, cell phones, etc. Right. So, it can either be used in modulation or demodulation. So, demodulation means extracting information from received signals by analyzing the frequency content in them. So, either modulation or demodulation.

The last thing could be spectrum analysis. So, identify interference, optimize bandwidth usage, and detect some channel characteristics. And the last slide here is, let us say, you have some applications in physics and engineering as well. So, let us say, waveform analysis. So, analyze all the vibrations, sound waves, electromagnetic waves, etcetera, or optics.

Design and analyze lenses using Fourier optics, where light fields are treated as frequency-domain signals. Or the last one is structural analysis. So, study stresses and vibrations in mechanical structures. So, in short, wherever you see any kind of frequencies or, let us say, oscillations which are kind of inbuilt in practical applications, one can actually go ahead and then apply some Fourier transform and try to superimpose some time series from a frequency-domain sort of approach to model and analyze that. So, now in the coming lectures this week, we will try to again extend the idea of spectral density.

What do you mean by spectral density? How do you connect spectral density to a Fourier transform? And all the intricacies of this entire idea about moving from a time-domain approach to a frequency-domain approach. Thank you.