The Discrete Fourier Transform.

Recap: a bit of linear algebra, and a discussion on the basis.

A basis is any set of linearly independent vectors. It is useful because for a basis of R^n space, any n-dimensional vector can be exactly reproduced by a linear combination of the basis.

Are the following vector sets a basis?

$$\begin{Bmatrix}1\\1\\0\end{Bmatrix},\begin{Bmatrix}1\\-1\\0\end{Bmatrix},\begin{Bmatrix}0\\0\\1\end{Bmatrix}$$

Yes

Orthogonal

$$\begin{Bmatrix}0\\1\\0\end{Bmatrix},\begin{Bmatrix}1\\0\\0\end{Bmatrix},\begin{Bmatrix}0\\0\\1\end{Bmatrix}$$

Yes

Orthonormal

$$\begin{Bmatrix}1\\0\\0\end{Bmatrix},\begin{Bmatrix}1\\0\\1\end{Bmatrix},\begin{Bmatrix}0\\0\\1\end{Bmatrix}$$

No

$$\begin{Bmatrix}3\\2\\1\end{Bmatrix},\begin{Bmatrix}0\\3\\0\end{Bmatrix},\begin{Bmatrix}0\\0\\4\end{Bmatrix}$$

Yes

Non-orth

Useful basis traits:

Orthogonal basis : every vector in the basis is orthogonal to every other vector. IE $< \bar{v}_i, \bar{v}_j > = 0 \; ; \forall \, i \neq j$

we like this property because the inverse of the matrix formed by an orthogonal basis is of the form $B^T B = D$ where D is a diagonal matrix. So $B^T / D = B^{-1}$

Orthonormal basis: an orthogonal basis where each vector is of unit length, IE $\left|\left|\bar{v}\right|\right| = 1$. In this case $B^T B = I$, IE $B^T = B^{-1}$

For the bases given above, which are orthogonal? Which orthonormal?

Why is an orthonormal basis useful?

For transformation between two orthonormal bases, length and angles are maintained.

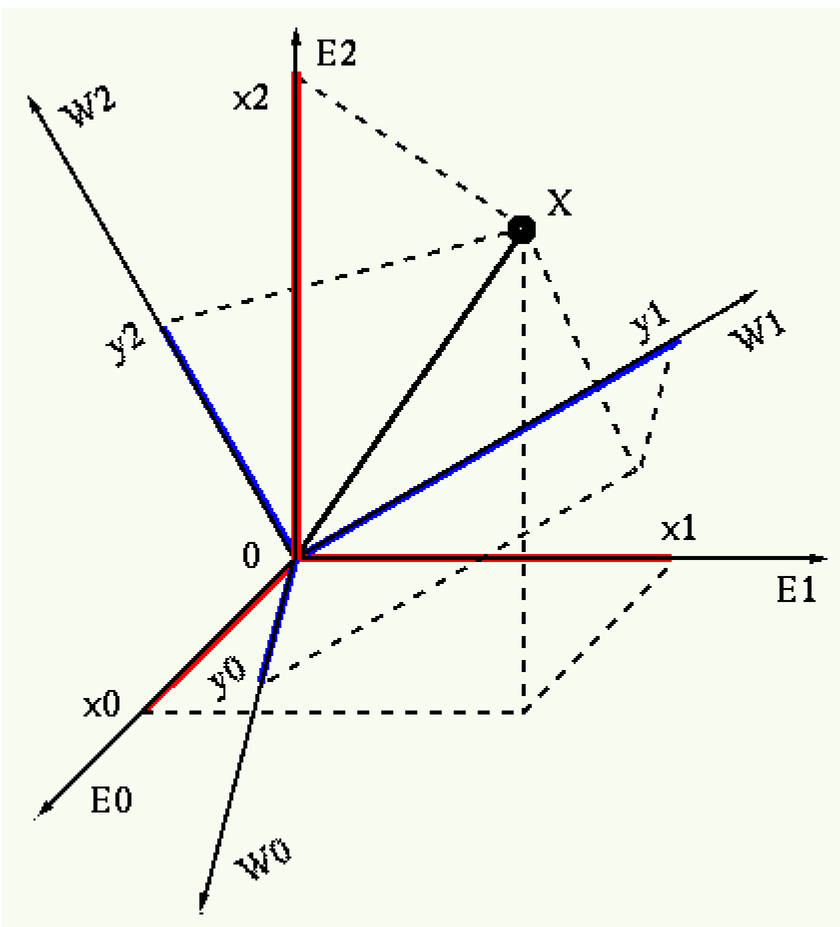

From a signal processing perspective, length is equivalent to signal energy, so for an orthonormal basis, signal energy is maintained. This is the basis of Parsevals theorem, which will be covered soon in lectures.

Estimation of vectors.

Let's consider the first basis vector, which I will call $\bar{v}$ here, that we saw today :

$$\bar{v} = \begin{Bmatrix} 1 \\ 1 \\ 0 \end{Bmatrix}$$

We want to try and approximate another vector, $\bar{x}$, by linearly scaling $\bar{v}$ by some constant. Formally we aim to optimize the following

$$\min_{k \in R}(||\bar{x} - k\bar{v}||)$$

Solving for k is straight forward, we have seen it briefly in lecture 3. Refer to lecture 3 pages 15 and 16.

What is the vector, $\mathrm{k}\bar{\mathrm{v}}$ *, that best approximates* $\bar{\mathrm{x}} = \begin{matrix} 4 \\ 5 \\ 1 \end{matrix}$ *?*

4.5

4.5

0

what is the value of k?

4.5

What is the mean squared error (note this is $\frac{||\bar{x}-k\bar{v}||}{N}$)?

0.5

The key point to this is that for an orthogonal basis our new coefficients are found by solving this problem for each basis vector. If we expand the equation

$$\bar{k} = \frac{B^T}{D}\ \bar{x}$$

We find that $k\ = \frac{<\bar{x},\bar{v}>}{||\bar{v}||}$ will fall out, and that the elements of $D_{i,i}$ are the energy of each basis vector $||\bar{v}_i||$.

For the orthonormal case, the energy $||\bar{v}||$ = 1, and so the result simplifies to $k\ =<\bar{x},\bar{v}>$. Fourier referred to this process as decomposing our vector $\bar{x}$ into a sum of linearly scaled basis vectors.

This is not the case for a basis that is not orthogonal. Because the basis vectors are non-orthogonal, the solutions for k_i are dependent on all the vectors in the set.

Clearly, there are an infinite number of choices for an orthonormal basis, and by extension for an arbitrary basis too. However there are choices of basis which help solve some important problems. Today, we will discuss the Fourier basis, which is useful for the case where our signals are periodic.

The Discrete Fourier transform:

Fourier provided the underlying math for decomposing a function into a set of sine waves, or more commonly a set of complex exponential functions.

In the discrete fourier transform we decompose a vector signal into a set of sampled complex exponential functions.

What do these basis vectors look like?

In the continuous case we have $f(t) = \cos(\omega t) + j\,\sin(\omega t)$ which alternatively is equal to the complex exponential $f(t) = e^{-2\pi j\omega t}$. To understand the FFT, we have to look at the problem in terms of the complex exponential, so we will continue with this from here on in.

When we sample f(t), we form a the following vector.

$$\bar{f}[n] = \{\, f(0) \quad f(T) \quad f(2T) \quad \cdots \quad f\big((n-1)\,T\,\big)\}$$

Assume that one of our fourier basis functions is $\mathrm{f(t)} = \mathrm{e}^{\mathrm{j}0.5\pi\mathrm{t}}$ *(IE* $\omega\ = \frac{1}{4}$*). Find the first 4 terms of the corresponding fourier basis vector* $\bar{\mathrm{f}}[\mathrm{n}]$ *where t = [0, 1, 2, 3].*

$$\begin{Bmatrix} 1 \\ i \\ -1 \\ -i \end{Bmatrix}$$

Find the first 4 terms of the fourier basis vector $\mathrm{f(t)} = \mathrm{e}^{\mathrm{j}0\mathrm{t}}$*, IE for* $\omega = 0$ *. What would we call this basis function?*

[1, 1, 1, 1], the DC term. It is also the *mean of the vector(once normalized)*

Are these vectors unit length (ie would a set of these vectors form an orthonormal basis)?

No, they need to be scaled by 1/N to be unit length, however, to formulate the FFT we apply this scaling AFTER solving the coefficents, so will continue using these terms. Remember that this scaling is consistent over all the basis functions

Sine functions are orthogonal, but what about when they are sampled?

While we won't prove this today, Sine functions of different frequencies are orthogonal functions, which as we showed earlier is very useful. However, when we sample them, we find that they aren't always orthogonal, they are only orthogonal if the sampling covers a whole unit period.

are the fourier basis vectors given above orthogonal? (use the inner product)

$yes\!: <\bar{\mathrm{v}}_1, \bar{\mathrm{v}}_2> = 0$

This is because for a bounded DFT, we have to assume a periodic extension. For that extension to not have a discontinuity, it has to have a whole number of periods captured.

Also we don't need to define our basis vectors in terms of ω, as the basis vector values will be invariant with changes in sampling period. Consider a sampling period of 0.01s, if we halve the samping period (doubling the frequency), we also halve the time required to capture N samples, and the basis vectors will remain the same (however the frequency they represent doubles)

so we find that we can describe the DFT basis functions for a vector of dimension N as follows

$$f_n(t) = e^{-2\pi jtn/N} \; for \; n = 0 \colon N - 1$$

A Discrete basis

By now you have calculated the first two basis vectors for the 4 point DFT. The other two vectors, correspond to n = 2, 3 . the following are the 4 basis functions

$$\begin{Bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{Bmatrix}, \begin{Bmatrix} 1 \\ i \\ -1 \\ -i \end{Bmatrix}, \begin{Bmatrix} 1 \\ -1 \\ 1 \\ -1 \end{Bmatrix}, \begin{Bmatrix} 1 \\ -i \\ -1 \\ i \end{Bmatrix}$$

Find the DFT of $\bar{f}[n] = \{\,0 \quad 1 \quad 0 \quad -1\,\}$ IE a sampled vector of $\sin(0.5\pi t)$ by whatever means you feel like.

0, -2i, 0, +2i

You should find only two terms are significant, why is this?

We considered the DFT of a sine wave with a frequency which exactly matches the two basis vectors (remember there are two coefficents for each frequency, except the Nyquist frequency and the DC term) it is orthogonal to all the other basis vectors, so they have no contribution.

You should also find that the terms are complex, why is this? If we consider this complex number in polar form, what does it correspond to?

The term is complex because the basis vector is $f(t) = \cos(\omega t) + j\sin(\omega t)$ and the inner product is only significant for the sine part of the equation, which is purely complex. when we look at the number in polar form we find a magnitude of 2 and a phase of -90 degrees. A cosine function is simply a sine function shifted by 90degrees to the left. This is why we say the complex terms capture the phase.

The FFT

The Fast Fourier Transform is an efficient way to obtain the Discrete Fourier Transform. It gives an identical result to the DFT you found above, but does so more efficiently.

How much more efficient? Well unless you were using a scheme akin to the FFT, when you calculated the DFT you probably used the following

$$\overline{fft} = B\bar{x}$$

$$\overline{fft} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \\ -1 \end{bmatrix}$$

This maxtrix multiplication will involve 4^2 multiplications, in computing we would say it is an $O(n^2)$ algorithm. Scale this up to a million coefficents, and your looking at 1 trillion multiplications. That will take ~250 seconds to compute on a quadcore at 1ghz assuming theres little overhead.

The FFT reduces the cost to $O(n\ln(n))$, here your looking at something closer to 20 million operations, and a execution time of ~5ms quadcore at 1ghz (this analysis is a little basic, as there is some overhead to the FFT-remember we have to scale AFTER the fact, and there is also an accumulate in the dot product)

The Plan

If we look at the basis, we can see that there are a number of redundant multiplications

$$\begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & i & -1 & -i \\ 1 & -1 & 1 & -1 \\ 1 & -i & -1 & i \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \\ -1 \end{bmatrix}$$

The values for the terms calculated by the coloured squares will be exactly the same. They also appear on a regular spacing. Lets see how they appear

Polar form, and why the FFT is done on an non-orthonormal basis

Recall that for a complex multiplication, if we do the multiplication in polar form the magnitudes multiply, and the angles add. If the two complex numbers are unit length, then the magnitude of the result is unit length as well.

This leads to the trick of the FFT. In the DFT, we talked about our basis vectors being regularly sampled points of a complex exponential:

$$\bar{f}[n] = \{\, f(0) \quad f(T) \quad f(2T) \quad \cdots \quad f\big((n-1)\,T\,\big)\,\}$$
$$\bar{f}[n] = \ \{\, e^{0} \quad e^{-jk1} \quad e^{-jk2} \quad \ldots \quad e^{-jk(n-1)}\,\}$$

Where $k = \ 2\pi n/N$, and is fixed for a given DFT we intend to take. What's interesting is that $\left\|e^{-jx}\right\| = 1$ for any given value of x. We also find that the samples are evenly spaced around the unit circle. You should have seen this last Thursday in the discussion of the Z-Transform

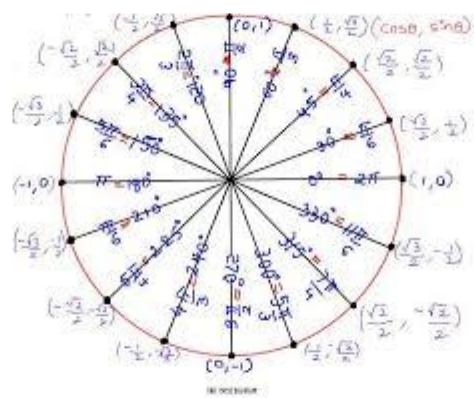

Lets define a new function, $w = \ e^{-jk1}$, which is the value of our basis vector at sample 2. We can find the basis vector at sample 1 as follows:
$w^0 = e^{-jk1\times 0} = e^0 = 1$

And sample 3:

$$w^2 = e^{-jk1\times 2} = e^{-jk2}$$

And so on. Now lets rewrite the basis matrix in terms of this new function, w

$$\begin{bmatrix} w^0 w^0 w^0 w^0 \\ w^0 w^1 w^2 w^3 \\ w^0 w^2 w^4 w^6 \\ w^0 w^3 w^6 w^9 \end{bmatrix} = \begin{bmatrix} 1 \;\; 1 \;\;\; 1 \;\;\; 1 \\ 1 w^1 w^2 w^3 \\ 1 w^2 w^4 w^6 \\ 1 w^3 w^6 w^9 \end{bmatrix}$$

If we prescaled the basis functions, the magnitude of we would have $w = \ k_{norm}e^{-jk1}$. If we look at the next coefficient w^2 we see the following:

$w^2 = {k_{norm}}^2 e^{-jk2} \neq k_{norm} e^{-jk2}$

Hopefully its clear that this unit circle expansion will only work if the terms are unit length, and that if we are going to factorize that matrix in terms of these w^n functions (as the FFT does), we will have to apply the scaling later.

Factorising the DFT matrix