Fourier Epicycles

Yuchen Xin

Spring 2025

1 Introduction

In this paper, we will rigorously define the discrete Fourier transform, demonstrate how to invert it and interpret the inversion as a chain of epicycles, quantify how good the epicycle "approximation" is, use the epicycle representation to visualize various curve transformations (and in particular, demonstrate the convolution theorem), and finally explain the groundbreaking fast Fourier transform algorithm.

2 The Discrete Fourier Transform

Given a sequence of N points $\{x_n\}_{n=0}^{N-1}$ that all belong in $\mathbb{C}$, commonly referred to as an input "signal", we define the discrete Fourier transform (DFT) as

$$X_k := \frac{1}{N}\sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} \quad \text{for } k \in \mathbb{Z}, \text{creating a new sequence } \{X_k\}_{k=0}^{N-1} \text{ all in } \mathbb{C} \tag{1}$$

followed by the inverse discrete Fourier transform (IDFT)

$$\tilde{x}_n := \sum_{k=0}^{N-1} X_k e^{2\pi i \frac{n}{N} k}, \qquad \text{where } \tilde{x}_n = x_n \text{ as follows...} \tag{2}$$

Theorem 1 (Inversion)**.** *The IDFT of the DFT returns the original sequence.*

Proof. We first note a useful orthogonality property: let $\alpha \in \mathbb{Z}$ be an integer. Supposing $\alpha \not\equiv 0 \pmod N$, we know $\frac{\alpha}{N}$ is not an integer so $e^{2\pi i \frac{\alpha}{N}} \neq 1$, so

$$\sum_{k=0}^{N-1} e^{2\pi i \frac{\alpha}{N} k} = \frac{1 - e^{2\pi i \frac{\alpha}{N} N}}{1 - e^{2\pi i \frac{\alpha}{N}}} = \frac{1 - e^{2\pi i \alpha}}{1 - e^{2\pi i \frac{\alpha}{N}}} = \frac{1 - 1}{1 - e^{2\pi i \frac{\alpha}{N}}} = 0$$

using the finite geometric sum formula, where the denominator is nonzero as shown above. If $\alpha \equiv 0 \pmod N$ then $\frac{\alpha}{N}$ is an integer, so $e^{2\pi i \frac{\alpha}{N} k} = 1$ for all integers k, and

$$\sum_{k=0}^{N-1} e^{2\pi i \frac{\alpha}{N} k} = \sum_{k=0}^{N-1} (1) = N$$

Summarizing,

$$\sum_{k=0}^{N-1} e^{2\pi i \frac{\alpha}{N} k} = \begin{cases} 0 & \text{if } \alpha \not\equiv 0 \pmod N \\ N & \text{if } \alpha \equiv 0 \pmod N \end{cases} \tag{3}$$

Now,

$$\tilde{x}_n = \sum_{k=0}^{N-1} X_k e^{2\pi i \frac{n}{N} k} = \sum_{k=0}^{N-1} \left(\frac{1}{N} \sum_{m=0}^{N-1} x_m e^{-2\pi i \frac{k}{N} m} \right) e^{2\pi i \frac{n}{N} k} = \frac{1}{N} \sum_{m=0}^{N-1} \sum_{k=0}^{N-1} x_m e^{2\pi i \frac{n-m}{N} k} \tag{*}$$

$$= \frac{1}{N} \sum_{m=0}^{N-1} x_m \left(\sum_{k=0}^{N-1} e^{2\pi i \frac{n-m}{N} k} \right) = \frac{1}{N}(x_n \cdot N) = x_n \tag{**}$$

where the sums can be interchanged at (*) because they are finite, and where (**) comes from applying (3) to $\alpha = n - m$ such that $n - m \equiv 0 \pmod N$ if and only if $n = m$ (since $0 \leq n, m \leq N - 1$). □

We took our original sequence x_n and transformed it to X_k, just to get back to our original sequence again. What does that tell us? It means we can express each x_n in an arbitrary sequence of N complex numbers as a linear combination of the basis "vectors" $\omega_m = e^{2\pi i \frac{m}{N}}$ for $m \in [0, N-1]$. These ω_m are just the N roots of unity, and if we use Euler's formula

$$\omega_m = \cos(2\pi m/N) + i \sin(2\pi m/N)$$

we can interpret ω_m as having some "frequency" proportional to m, with respect to rotating about the unit disk $\partial\mathbb{D}$ on the complex plane. Then, the DFT as defined in (1) "projects" an input sequence onto the space generated by the ω_m, and the IDFT in (2) defines where the original sequence lies in this space.

In signal processing, we commonly treat x_n as a discrete sequence of signals over time — then, the DFT is said to tranform the *time domain* to the *frequency domain*. This interpretation is solidifed visually when we make epicycles in later sections.

One very useful property of transforming the "space" on which we think of sequences is how operations change — in particular, convolutions. We will come back to this in a later section. Another feature is that it allows us to bridge from discrete sequences to something continuous...

2.1 Generalizing to Continuous

Recall Fourier series: given some T-periodic function $f(t)$, we have

$$c_k = \frac{1}{T}\int_0^T f(t)e^{-2\pi i \frac{k}{T} t}dt, \qquad f(t) \sim \sum_{k=-\infty}^{\infty} c_k e^{2\pi i \frac{k}{T} t} \tag{4}$$

where (as a well-known theorem) the series converges pointwise to $f(t)$ when $f(t)$ is continuous and piecewise smooth.

Now let's try to relate (4) with the DFT/IDFT defined earlier. One way would be to convert (4) to the DFT/IDFT by forcing $f(t)$ into a discrete signal. Given a sequence $\{x_n\}_{n=0}^{N-1}$, we can define a function $f : [-0.5, N-0.5] \to \mathbb{C}$ to be N-periodic (the -0.5 is a minor technical detail) such that

$$f(t) = x_0\delta(t-0) + x_1\delta(t-1) + \cdots + x_{N-1}\delta(x-(N-1)) = \sum_{n=0}^{N-1} x_n \delta(t-n)$$

where $\delta(t-a)$ is the Dirac delta function, a generalized function that is 0 everywhere but at $t = a$, defined by the property that

$$\int_I g(t)\delta(t-a)dt = \begin{cases} g(a) & \text{if interval } I \text{ includes } a \text{ in its interior} \\ 0 & \text{otherwise} \end{cases}$$

for arbitrary continuous functions g. We write f in this way so then

$$c_k = \frac{1}{N}\int_0^N \left(\sum_{n=0}^{N-1} x_n\delta(t-n)\right) e^{-2\pi i \frac{k}{N} t}dt = \frac{1}{N}\sum_{n=0}^{N-1} x_n \int_0^N e^{-2\pi i \frac{k}{N} t}\delta(t-n)dt = \frac{1}{N}\sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} = X_k$$

where we exactly get the formula for the DFT X_k as in (1) for $0 \le k \le N-1$. If we truncate the series in (4) to just be these k values, we also recover exactly the IDFT.

This argument suggests that the DFT/IDFT are part of a more fundamental Fourier series phenomenon. But more relevant to our eventual epicycles is going *backwards*, from a discrete sequence to a Fourier series that converges to a nice curve that goes through the points of the sequence. More precisely, we take the IDFT in (2) and modify it as follows:

$$x_n = \sum_{k=0}^{N-1} X_k e^{2\pi i \frac{n}{N} k} \qquad \longrightarrow \qquad g(t) := \sum_{k=0}^{N-1} X_k e^{2\pi i \frac{k}{N} t}, \quad t \in [0, N] \tag{5}$$

This defines a C^∞ function $g(t)$ that is N-periodic (since each exponential in the sum is N-periodic with respect to t), such that $g(t)$ that still goes through all the original sequence points x_n whenever $t = n$ by Theorem 1. But how does $g(t)$ behave in between the sequence points? Does it nicely "interpolate" between the sequence points, whatever that means? We will explore this, but first, we must dive into epicycles.

3 Epicycle Drawings

To visually represent the DFT/IDFT, we treat $\mathbb{C} \cong \mathbb{R}^2$ such that the input sequence $\{x_n\}_{n=0}^{N-1}$ is a sequence of N points on a 2D plane. Afterwards, we compute the DFT $\{X_k\}_{k=0}^{N-1}$ as before, but then we take the "continuous" IDFT $g(t)$ as defined in (5).

Consider each term in $g(t)$. Rewrite $X_k = r_k e^{i\phi_k}$, where r_k, ϕ_k are fixed with respect to t. Then each term becomes

$$X_k e^{2\pi i \frac{k}{N} t} = r_k e^{i\left(2\pi \frac{k}{N} t + \phi_k\right)} = r_k \cos\left(2\pi \frac{k}{N} t + \phi_k\right) + i\ r_k \sin\left(2\pi \frac{k}{N} t + \phi_k\right)$$

Identifying the real and imaginary components as x and y components of a 2D Cartesian plane, we see that as t changes, $X_k e^{2\pi i \frac{k}{N} t}$ moves around in a circle of radius r, with angular frequency $2\pi \frac{k}{N}$ and an added phase ϕ. Call this circular path the k^{th} term traces out the k^{th} epicycle.

Then, as $g(t)$ is a sum of N terms, visualizing complex addition as 2D vector addition, the effect is simply chaining up the epicycles, such that the k^{th} epicycle revolves around the $(k-1)^{\text{th}}$ epicycle at the $(k-1)^{\text{th}}$ frequency. Further, since addition is commutative, we can reorder the chain of epicycles in order of largest to smallest radius. Visually, in the demonstration I coded (based off of Shiffman [2019]), we get

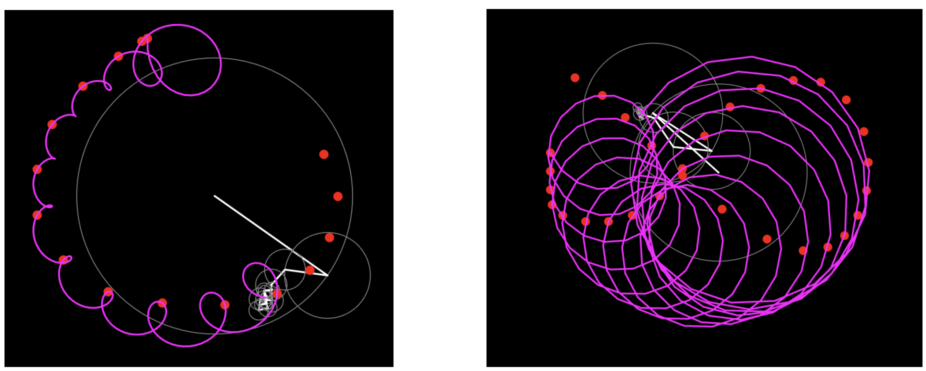

Figure 1: *Epicycle drawings of a circle-shaped and a lemniscate-shaped input sequence. The red dots constitute the input sequence, the white lines and circles represent the epicycles and their current rotation, and the magenta path traces out the curve $g(t)$. The origin is at the center.*

We can make a few observations from Figure 1. One is that the magenta path indeed passes through every input point, **in order** — as it should, since $g(n) = x_n$ by Theorem 1. Another is that $g(t)$ is indeed periodic — although it is not shown in Figure 1, the magenta path eventually traces back to where it started, and repeats the same path. This also directly comes from the formula for $g(t)$.

Most strikingly, though, the magenta path seems to take a loopy vacation between each x_n. This effect is most pronounced when the *orientation* of the input sequence is negative, i.e. clockwise (about the origin), as seen in the lemniscate-shaped input sequence in Figure 1. Why is this so?

Looking at the terms of $g(t)$, we note that the $(N-1)^{\text{th}}$ term has angular frequency $2\pi \frac{N-1}{N}$, which is very close to 2π. Hence, as t changes from one integer n to another integer $n+1$, where the curve must go from x_n to x_{n+1}, the $(N-1)^{\text{th}}$ epicycle almost completes a full circle (counterclockwise). Combined with the fact that **all other terms in $g(t)$ have nonnegative frequency**, we conclude that no other epicycle's motion can cancel that of the $(N-1)^{\text{th}}$ epicycle, so almost a full circle is made between each x_n.

To see why the effect is more exaggerated in negative orientation input sequences, we can show that the modulus of X_{N-1} (the radius of the $(N-1)^{\text{th}}$ epicycle) becomes large. Suppose, in the extreme case, that x_n defines a perfect clockwise circle about the origin, i.e. $x_n = Re^{-2\pi i \frac{n}{N} + i\varphi}$ for some radius R and phase φ. Looking at the DFT definition for X_{N-1}, we have

$$X_{N-1} = \frac{1}{N} \sum_{n=0}^{N-1} Re^{-2\pi i \frac{n}{N} + i\varphi} e^{-2\pi i \frac{N-1}{N} n} = \frac{Re^{i\varphi}}{N} \sum_{n=0}^{N-1} e^{-2\pi i \frac{N}{N} n} = \frac{Re^{i\varphi}}{N} \sum_{n=0}^{N-1} \cancelto{1}{e^{-2\pi i n}} = Re^{i\varphi} \frac{N}{N} = Re^{i\varphi}$$

Taking the modulus, we get that $|X_{N-1}| = R$, meaning the radius of the epicycle responsible for spinning (almost) a full circle between each term is the radius of the entire input sequence! The underlying mathematical mechanism is that in the projection of x_n onto its DFT, the frequency components of x_n add together *constructively* — such that even if the sequence $\{x_n\}$ was perturbed from the perfect clockwise circle, the product $x_n e^{-2\pi i \frac{n}{N}(N-1)}$ would still be close to 1 (after factoring out a constant phase offset term), thus producing a spin-ny epicycle with a large $|X_{N-1}|$ radius.

3.1 Can We Do Better?

What we want is for the magenta path to interpolate "nicely" between sequence points. We can state this problem in a more formal context: suppose the sequence $\{x_n\}$ are sampled points of an T-periodic smooth curve $f(t) : \mathbb{R} \to \mathbb{C}$ (say, at regular t intervals). We want some epicycle approximation $\tilde{g}(t)$ (similar to that defined in (5)) to be as close to $f(t)$ as possible, preferably in a uniform manner.

To quantify this, we first define the L^2 norm over the interval $[0, T]$ of some function $h(t)$ as

$$\|h\| := \left(\frac{1}{T} \int_0^T |h(t)|^2 dt \right)^{1/2}$$

The L^2 norm satisfies the triangle inequality as a consequence of the Cauchy-Schwartz Inequality applied to the inner product $\langle f_1, f_2 \rangle = \int_0^T f_1(t)\overline{f_2(t)}dt$. If we apply the L^2 norm to the difference $\|f - \tilde{g}\|$, the result captures the "average" error between the desired curve $f(t)$ and the epicycle approximation curve $\tilde{g}(t)$. Stating error in terms of the L^2 norm allows us to use tools like Parceval's Identity to eventually place a bound.

There is one last thing to do before bounding the error. We have seen in the previous section that $g(t)$ as defined in (5) has great error compared to the underlying smooth curve (the circle and the lemniscate). We can improve the error (provably so, as we shall see) by instead considering

$$\tilde{g}(t) := \sum_{k=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} X_k e^{2\pi i \frac{k}{N} t}, \quad t \in [0, N] \tag{6}$$

The new curve $\tilde{g}(t)$ differs from $g(t)$ only in the summation indices, such that $\tilde{g}(t)$ now contains both positive and negative frequency epicycles. However, the fact that $\tilde{g}(n) = g(n) = x_n$ for all $n \in \{0, \dots, N-1\}$ still holds. To see why, we note the identity

$$X_k = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} e^{-2\pi i N} = \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k+N}{N} n} = X_{k+N}$$

holds, and the roots of unity $e^{2\pi i \frac{k}{N} t}$ in (6) for negative-k simply wrap around to the corresponding positive-k root of unity for $t = n \in \{0, \dots, N-1\}$, such that each term in the summation of (5) appears exactly once in the formula of (6), proving their equivalence. For real t in between the integers, $\tilde{g}(t)$ can differ greatly from $g(t)$.

Now we can finally bound the error between $\tilde{g}$ and f, depending on the smoothness of f:

Lemma 1. *If the Fourier series of the T-periodic function $f : [0, T] \to \mathbb{C}$ is finite, i.e. we can write*

$$f(t) = \sum_{k=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_k e^{2\pi i \frac{k}{T} t}$$

for some integer N, then if we sample N points $\{x_n\}_{n=0}^{N-1}$ at regular intervals $x_n = f\left(\frac{T}{N} n\right)$, the each item in the DFT satisfies $X_k = c_k$ (for $k \in \{0, \dots, N-1\}$). Thus the epicycle approximation $\tilde{g}(t)$ as defined in (6) matches $f(t)$ exactly, with zero error.

Proof. The proof of this is inspired from the proof of the Nyquist-Shannon Sampling Theorem, which is basically the Fourier transform version of Theorem 2.

We have

$$x_n = f\left(\frac{T}{N} n\right) = \sum_{k=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_k e^{2\pi i \frac{k}{T} \left(\frac{T}{N} n\right)} = \sum_{k=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_k e^{2\pi i \frac{k}{N} n}$$

Substituting this into the definition of the DFT, for k from $-\lceil N/2 \rceil + 1$ to $\lfloor N/2 \rfloor$, we have

$$\begin{aligned} X_k &= \frac{1}{N} \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} = \frac{1}{N} \sum_{n=0}^{N-1} \left(\sum_{m=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_m e^{2\pi i \frac{n}{N} m} \right) e^{-2\pi i \frac{k}{N} n} = \sum_{m=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_m \left(\frac{1}{N} \sum_{n=0}^{N-1} e^{2\pi i \frac{m-k}{N} n} \right) \\ &= \sum_{m=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_m \left(\begin{cases} 0 & \text{if } m \not\equiv k \pmod{N} \\ 1 & \text{if } m \equiv k \pmod{N} \end{cases} \right) = c_k \end{aligned}$$

□

Remark: In signal processing, if we try to reconstruct a signal $f(t)$ using only samples of $f(t)$ at different times, if the reconstruction differs from the original, we call it "aliasing". In Lemma 1, under the assumptions of f, we showed that there is no sampling or "aliasing" error.

Using this lemma, we are ready to prove the main theorem.

Theorem 2. *Given a T-periodic, continuous, and piecewise smooth $f : [0, T] \to \mathbb{C}$ such that $f \in C^m$ (with $m \geq 1$), if we sampled f at N points at regular intervals $x_n = f\left(\frac{T}{N} n\right)$, then after constructing $\tilde{g}(t)$ using these samples according to (6), we can bound*

$$\|f - \tilde{g}\| = O(N^{-m+1/2}) \qquad \text{as } N \to \infty$$

Proof. The assumptions of f imply it has a convergent Fourier series

$$f(t) = \sum_{k=-\infty}^{\infty} c_k e^{2\pi i \frac{k}{T} t}$$

But to apply our lemma, we must truncate this series. When a function has a finite number of Fourier frequency components we say it is "bandlimited" — in this case, we define the bandlimited function

$$\tilde{f}(t) := \sum_{k=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_k e^{2\pi i \frac{k}{T} t}$$

Comparing with f under the L^2 norm and applying the triangle inequality, we have

$$\|f - \tilde{g}\| = \|(f - \tilde{f}) + (\tilde{f} - \tilde{g})\| \leq \underbrace{\|f - \tilde{f}\|}_{\text{truncation error}} + \underbrace{\|\tilde{f} - \tilde{g}\|}_{\text{sampling error}}$$

By Lemma 1, we get zero sampling error: $\|\tilde{f} - \tilde{g}\| = 0$. For the truncation error, we first explicitly write out $f - \tilde{f}$, calling the difference a new function $h : [0, T] \to \mathbb{C}$:

$$h(t) := f(t) - \tilde{f}(t) = \sum_{k=-\infty}^{\infty} c_k e^{2\pi i \frac{k}{T} t} - \sum_{k=-\lceil N/2 \rceil + 1}^{\lfloor N/2 \rfloor} c_k e^{2\pi i \frac{k}{T} t} = \sum_{k \leq -\lceil N/2 \rceil, k > \lfloor N/2 \rfloor} c_k e^{2\pi i \frac{k}{T} t} \tag{7}$$

By the linearity of Fourier series, since h is a linear combination of f and $\tilde{f}$, we know the Fourier series of h is a linear combination of the Fourier series of f and $\hat{f}$, which is given by the right-hand-side in (7). Thus h represents all the higher-magnitude frequencies that constitute f — the Fourier coefficients of h are 0 for a finite range of frequencies near 0, and c_k outside this range (as given in (7)). Using Parceval's Identity, we know the L^2 norm of h over $[0, T]$ satisfies

$$\|h\|^2 = \sum_{k \leq -\lceil N/2 \rceil, k > \lfloor N/2 \rfloor} |c_k|^2$$

Now $f \in C^m$, and $\hat{f} \in C^\infty$, so $h \in C^m$, meaning the Fourier coefficients of h, which are c_k (for large $|k|$) decay on the order of k^{-m} as $k \to \infty$. In other words, there is a constant C such that

$$|c_k| \leq \frac{C}{|k|^m} \qquad \Longrightarrow \qquad \sum_{k \leq -\lceil N/2 \rceil, k > \lfloor N/2 \rfloor} |c_k|^2 \leq 2C^2 \sum_{k \geq \lfloor N/2 \rfloor} k^{-2m}$$

The series on the right is nonnegative and decreasing (with $m \geq 1$), so it is comparable to the integral

$$2C^2 \sum_{k \geq \lfloor N/2 \rfloor} k^{-2m} \leq C' \int_{N/2}^{\infty} x^{-2m} dx = C' \frac{x^{-2m+1}}{-2m+1}\bigg|_{N/2}^{\infty} = -C' \frac{(N/2)^{-2m+1}}{-2m+1}$$

$$= \frac{C'}{(2m+1) \cdot 2^{-2m+1}} N^{-2m+1} = O(N^{-2m+1})$$

Thus

$$\|f - \tilde{g}\| \leq \|h\| \leq [O(N^{-2m+1})]^{1/2} = O(N^{-m+1/2})$$

□

We proved the theorem, but what does it all mean? First, it suggests that we should draw epicycles based on $\tilde{g}(t)$ instead of $g(t)$, including negative and positive frequency epicycles. This makes intuitive sense — the epicycle approximation is more suited to handle positively and negatively oriented input sequences, and epicycles that revolve in opposite directions (with positive and negative frequencies) can cancel out erratic spin-ny behavior.

If we use $\tilde{g}(t)$ to make epicycle drawings, then Theorem 2 tells us that the smoother the underlying curve the input sequence is sampled from, the closer the epicycle drawing will match. Theorem 2 also justifies the intuition that the more points are sampled (the greater the N), the better the reconstructed curve is.

In the demonstration I coded, we can clearly see Theorem 2 in action:

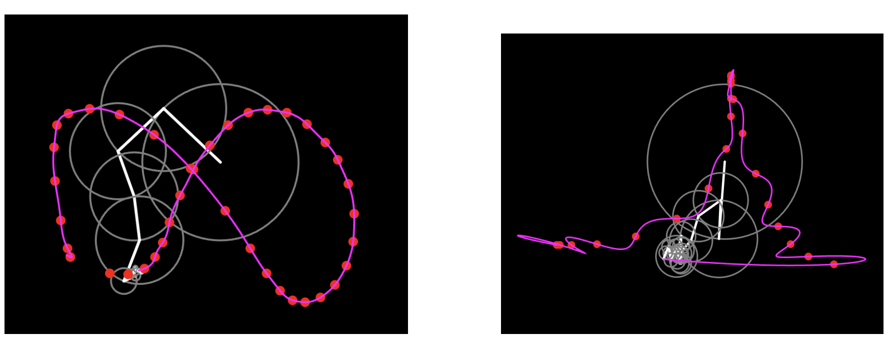

Figure 2: *On the left, the lemniscate is drawn with smooth interpolation in between the points, in contrast to Figure 1. On the right, the epicycle curve is still relatively smooth, but fails to interpolate nicely near the bottom, where the input sequence is irregular and has a "jump discontinuity".*

3.2 Epicycle Transformations and Convolution

With a working epicycle drawer, we can explore key properties of the DFT/IDFT by visually representing them through epicycles. For example, the $k = 0$ epicycle never turns since its angular frequency is 0, and acts as the central pivot of the entire epicycle drawing. Hence if we translated the entire input sequence by some constant $re^{i\theta}$, we expect that only the $k = 0$ epicycle moves by $re^{i\theta}$, while all the other epicycles remain the same. This is easily proved: compared to the old DFT $\{X_k\}$, the new DFT $\{\tilde{X}_k\}$ is

$$\tilde{X}_k = \frac{1}{N}\sum_{n=0}^{N-1}(x_n + re^{i\theta})e^{-2\pi i \frac{k}{N} n} = X_k + re^{i\theta}\frac{1}{N}\sum_{n=0}^{N-1} e^{-2\pi i \frac{k}{N} n} = X_k + re^{i\theta}\delta_{k0}$$

where δ_{k0} denotes the Kronecker delta that is 1 if $k = 0$ and 0 otherwise. But this illuminates a more general fact about the linearity of the DFT: if the input sequence is a linear combination $x_n = \alpha a_n + \beta b_n$, the DFT is

$$X_k = \frac{1}{N}\sum_{n=0}^{N-1}(\alpha a_n + \beta b_n)e^{-2\pi i \frac{k}{N} n} = \alpha\frac{1}{N}\sum_{n=0}^{N-1} a_n e^{-2\pi i \frac{k}{N} n} + \beta\frac{1}{N}\sum_{n=0}^{N-1} b_n e^{-2\pi i \frac{k}{N} n} = \alpha A_k + \beta B_k \tag{8}$$

where $\{A_k\}, \{B_k\}$ are the DFT of $\{a_n\}, \{b_n\}$, respectively.

Remark: as long as we show the DFT of two sequences are equal for k from 0 to $N - 1$, then their epicycle drawings must be equivalent even though we use negative k in the definition of $\tilde{g}(t)$ in (6), since $X_k = X_{k+N}$ as shown earlier.

Another transformation we could perform is scaling. If we doubled the modulus of every x_n, we expect the radii of every epicycle to double — a quick look at the DFT formula confirms this. But what if we wanted to perform a more general "scaling" operation: given input sequences $\{a_n\}_{n=0}^{N-1}$ and $\{b_n\}_{n=0}^{N-1}$, we define

$$x_n := (a * b)_n = \sum_{m=0}^{N-1} a_m b_{n-m \bmod N} \tag{9}$$

We call (9) a convolution of the sequences $\{a_n\}$ and $\{b_n\}$. To construct x_n, we take b_n and scale it by a_0 — then add on b_{n-1} scaled by a_1, then b_{n-2} scaled by b_0, and so on. This can be interpreted as a weighted average of b_n using weights a_m, and it generalizes many important operations used in signal processing. Audio effects like gain and delay are specially designed convolutions of audio signals; image filters (like blurring) is just a weighted average of the colors around each pixel, i.e. a 2D convolution. Some neural network architectures are even specifically designed to train a convolution filter. Convolution is even at the heart of polynomial multiplication:

$$\left(\sum_{n=0}^{N-1} a_m x^m\right)\left(\sum_{m=0}^{N-1} b_n x^n\right) = \sum_{n=0}^{N-1}\sum_{m=0}^{N-1} a_m b_n x^{m+n} = \sum_{n=0}^{N-1}\sum_{m=0}^{N-1} a_m b_{(m+n)-m} x^{m+n}$$

If we define a new index $n' = m + n$ and rearrange the sum, we can identify the new inner sum as just a convolution of the coefficients of the original polynomial. The details are not important in this paper.

What we are actually interested in is how epicycles can represent a convolution, as to gain a new perspective to what is actually happening in a convolution. For that, we need an important theorem, based off of McFee [2023]:

Theorem 3 (Convolution)**.** *Let $\{a_n\}, \{b_n\}$ both be sequences of N complex numbers. Define $x = a * b$. Then*

$$X_k = A_k \cdot B_k$$

Proof. We first compute

$$X_k = \frac{1}{N}\sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} = \frac{1}{N}\sum_{n=0}^{N-1}\left(\sum_{m=0}^{N-1} a_m b_{n-m \bmod N}\right) e^{-2\pi i \frac{k}{N} n} = \sum_{m=0}^{N-1} a_m \left(\frac{1}{N}\sum_{n=0}^{N-1} b_{n-m \bmod N} e^{-2\pi i \frac{k}{N} n}\right)$$

For the inner sum, we can get rid of the mod by "shifting": multiply by $e^{2\pi i \frac{k}{N} m}$ and break up the summation to get

$$\begin{aligned} e^{2\pi i \frac{k}{N} m} \cdot \frac{1}{N}\sum_{n=0}^{N-1} b_{n-m \bmod N} e^{-2\pi i \frac{k}{N} n} &= \frac{1}{N}\left(\sum_{n=0}^{m-1} + \sum_{n=m}^{N-1}\right) b_{n-m \bmod N} e^{-2\pi i \frac{k}{N}(n-m)} \\ &= \frac{1}{N}\left(\sum_{n=0}^{m-1} b_{n-m+N} e^{-2\pi i \frac{k}{N}(n-m+N)} + \sum_{n'=0}^{N-m-1} b_{n'} e^{-2\pi i \frac{k}{N} n'}\right) \qquad (*) \\ &= \frac{1}{N}\left(\sum_{n'=N-m}^{N-1} b_{n'} e^{-2\pi i \frac{k}{N} n'} + \sum_{n'=0}^{N-m-1} b_{n'} e^{-2\pi i \frac{k}{N} n'}\right) = \frac{1}{N}\sum_{n'=0}^{N-1} b_{n'} e^{-2\pi i \frac{k}{N} n'} = B_k \qquad (**) \end{aligned}$$

where we changed indices to $n' = n - m$ in (*) and to $n' = n - m + N$ along with the $2\pi i$-periodicity of the exponential $e^{-2\pi i \frac{k}{N}}$ in (**). Returning to X_k, we have

$$X_k = \sum_{m=0}^{N-1} a_m \left(e^{-2\pi i \frac{k}{N} m} B_k\right) = B_k \sum_{m=0}^{N-1} a_m e^{-2\pi i \frac{k}{N} m} = B_k \cdot A_k$$

□

We can see convolution in action using epicycles:

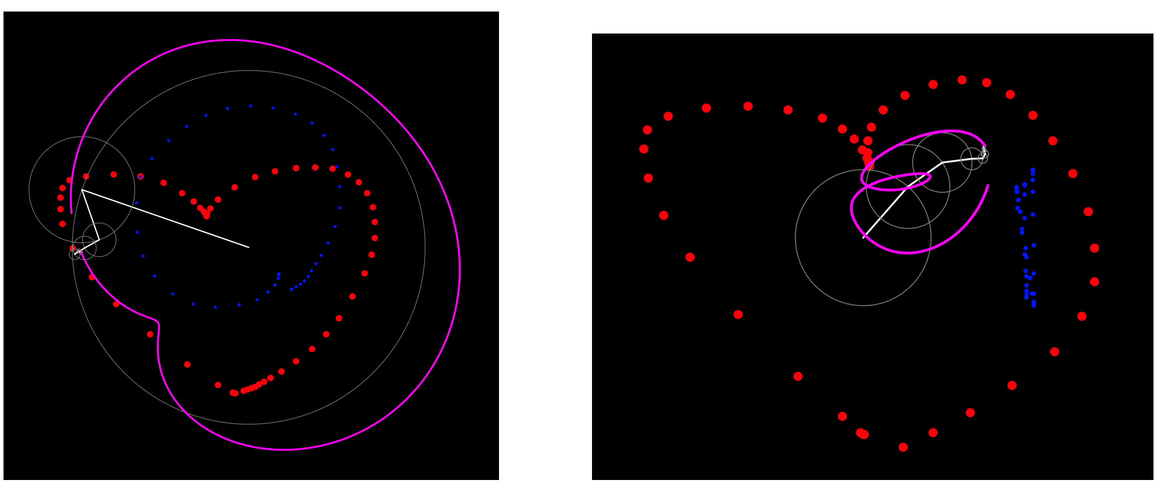

Figure 3: *Convolutions of a heart-shaped sequence (red) with the "filter" sequence (blue), which is a circle around the origin on the left and a vertical line on the right. The result is the magenta curve, a transformation of the original heart. It is left as an exercise to the reader to figure out why the convolutions produce the illustrated results :)*

4 Fast Fourier Transform

One final important detail with regards to the DFT is how efficiently we can compute it. Looking at the definition in (1), to compute each X_k it takes N "operations" (an operation here just means a single addition or multiplication of two complex numbers), so to compute the entire $\{X_k\}$ sequence it takes N^2 operations, which is far too slow by computer science standards.

However, we can do better, using an algorithm introduced by Cooley and Turkey. Without loss of generality, take $N = 2^K$ to be some power of 2 (otherwise add 0s to the input sequence — this will not change the time complexity, as we will see later). Then, decompose the DFT formula into a sum over even indices and a sum over odd indices:

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi i \frac{k}{N} n} = \underbrace{\sum_{n'=0}^{N/2-1} x_{2n'} e^{-2\pi i \frac{k}{N}(2n')}}_{\text{Even indices}} + \underbrace{\sum_{n'=0}^{N/2-1} x_{2n'+1} e^{-2\pi i \frac{k}{N}(2n'+1)}}_{\text{Odd indices}}$$

Now let $x^E_{n'} = x_{2n'}$ be the even subsequence of $\{x_n\}$ and $x^O_{n'} = x_{2n'+1}$ the odd subsequence, and X^E_k, X^O_k be their DFTs. Examining the even part of X_k, we have

$$\sum_{n'=0}^{N/2-1} x_{2n'} e^{-2\pi i \frac{k}{N}(2n')} = \sum_{n'=0}^{N/2-1} x^E_{n'} e^{-2\pi i \frac{k}{N/2} n'} = X^E_k$$

while the odd part is

$$\sum_{n'=0}^{N/2-1} x_{2n'+1} e^{-2\pi i \frac{k}{N}(2n'+1)} = e^{-2\pi i \frac{k}{N}} \sum_{n'=0}^{N/2-1} x^O_{n'} e^{-2\pi i \frac{k}{N/2} n'} = e^{-2\pi i \frac{k}{N}} X^O_k$$

Combined, they give

$$X_k = X^E_k + e^{-2\pi i \frac{k}{N}} X^O_k \tag{10}$$

Remark: For $k \geq N/2$, the DFT components X^E_k, X^O_k wrap around modulo $N/2$ by periodicity of the definition of the DFT. It would be more accurate to write $k \bmod N/2$ in (10) but we skip that technical detail.

How does (10) speed up computation? If we wanted to compute X_k for a fixed k, (10) would not change the number of operations. However, if we used (10) to *recursively* compute the entire DFT sequence $\{X_k\}$, we will get a speed-up. To see why, let $T(N)$ be the number of operations it takes to compute a DFT sequence from some length N input sequence. Exploiting (10), we first use $T(N/2)$ operations to compute $\{X^E_k\}$ and another $T(N/2) + 1$ operations to compute $\{X^O_k\}$ (the extra 1 for the $e^{-2\pi i \frac{k}{N}}$ factor). Then, iterating through all k (N additional computations), we can complete the computation for $\{X_k\}$. Thus we have the recurrence

$$T(N) = T\left(\frac{N}{2}\right) + 1 + T\left(\frac{N}{2}\right) + N \leq 2T\left(\frac{N}{2}\right) + CN$$

for a constant C. We also have the base case $T(1) = 1$ since the DFT of a single number is just itself. It remains to find an explicit formula for $T(n)$ — we do so by counting the total number of operations per recursive "level". At the first level, we start with N, complete CN operations, and then use $T(N/2)$ twice. At the second level, we complete $C(N/2)$ operations (twice), totaling CN operations, and call $T(N/4)$ four times. Repeating, we have $\log_2 N$ recursive levels, each with CN total operations, so

$$T(n) \leq \sum_{n=1}^{\log_2 N} CN = CN \log_2 N = O(N \log N)$$

which performs significantly better than $O(N^2)$, asymptotically.

The same logic applies here to give a fast $O(N \log N)$ IDFT algorithm. These constitute the "fast Fourier transform", which not only speeds up epicycle drawing computations, but also (by the convolution theorem) dramatically speeds up computations of convolutions of sequences. This makes applying filters to audio and images, as well as training convolutional neural networks, a lot faster than doing it the naive way — making the fast Fourier transform a groundbreaking discovery.

5 Conclusion

Overall, epicycles provide a great way to visualize the discrete Fourier transform of an input sequence, and they are especially effective at approximating the underlying smooth curve (as proven). The epicycle representation also provides a new perspective of understanding how sequences of points transform under operations like convolutions.

Of course, the real star of the show is the discrete Fourier transform itself, allowing us to change from the time domain to the frequency domain. The main sidekick is the fast Fourier transform algorithm, that makes everything in this paper very practical from a computational perspective, readily used in all sorts of industries.

There are also many extensions. Epicycles can be used to visualize audio filters, image filters, polynomial multiplication, and pretty much anything to do with convolution. Or, epicycles can be bumped up to higher dimensions. I even read a crazy paper from Birmanns [2022] where quaternions were used in place of complex numbers to define a 4-dimensional DFT/IDFT, that allows representing a sequence of points in 3D/4D using quaternion-based epicycles.

References

Jan Philipp Birmanns. Creating multidimensional drawings with epicycles, 2022. URL `https://www.maturitaetsarbeiten.ch/cms/images/2022/Birmanns_Jan/MP_JanBirmanns_final.pdf`.

Brian McFee. *Digital Signals Theory*. CRC Press, 2023. URL `https://brianmcfee.net/dstbook-site`.

Daniel Shiffman. Challenge #130: Drawing with fourier transform and epicycles, 2019. URL `https://thecodingtrain.com/challenges/130-drawing-with-fourier-transform-and-epicycles`.

Appendix

The demonstration I coded is available at `https://editor.p5js.org/XYuchen/sketches/BKDlBvRLK`.