Deriving the Discrete Fourier Transform via Representation Theory

Here, we show how the usual discrete Fourier transform is identical to the change-of-basis between the regular representation and the irreps of $\mathbb{Z}_N$.

We first need to prove certain facts about commuting matrices and irreps of abelian groups.

(A, B Diagonalizable and $AB = BA$) $\iff$ (A, B are Simultaneously Diagonalizable).

Proof of $\implies$:
Let A and B commute with each other: $AB = BA$. Let v be an eigenvector of A with eigenvalue λ, ie $v \in V_\lambda$. Then, $Av = \lambda v$. Now:

$$\begin{aligned} \Rightarrow A(Bv) &= (AB)v \\ &= (BA)v \\ &= B(Av) \\ &= B(\lambda v) \\ &= \lambda(Bv) \end{aligned}$$

Thus, $Bv \in V_\lambda$ as well. Thus, V_λ is stable under B. Thus, B has an eigenvector $w \in V_\lambda$ (with possibly a different eigenvalue μ):

$$Bw = \mu w$$

But, as $w \in V_\lambda$,

$$Aw = \lambda w$$

Thus, w is a simultaneous eigenvector of both A and B. In fact, B has an entire basis of eigenvectors in V_λ. Why? Because B is diagonalizable; there exists eigenvectors $\{w_\mu\}$ such that they span the entirety of $\mathbb{C}^n$ (where n is the dimension of the space of v). Let $\{W_\mu\}$ be the set of eigenspaces of B. Then:

$$\mathbb{C}^n = \oplus_\mu W_\mu$$

ie, we can write any vector $v \in V_\lambda$ uniquely as:

$$v = \sum_\mu w_\mu \text{ where } w_\mu \in W_\mu$$

Note that some of the w_μ may be 0. We need to show that the only non-zero w_μ must also be in V_λ, ie $w_\mu \in V_\lambda \cap W_\mu$. Then, by choosing a basis for each $V_\lambda \cap W_\mu$ and combining them, we obtain a basis of V_λ.

Now, by the same argument as above, but with A and B flipped: each W_μ is stable under A. Thus:

$$Av = \sum_\mu Aw_\mu$$
$$\Rightarrow \lambda v = \sum_\mu Aw_\mu$$
$$\Rightarrow \sum_\mu \lambda w_\mu = \sum_\mu Aw_\mu$$

Using the uniqueness of the direct sum decomposition: $Aw_\mu = \lambda w_\mu$ for each w_μ. Thus, each $w_\mu \in V_\lambda$ by definition, so $w_\mu \in V_\lambda \cap W_\mu$. Thus, we have shown:

$$V_\lambda = \oplus_\mu \left(V_\lambda \cap W_\mu\right)$$

Thus, for each V_λ, B has a corresponding basis of V_λ formed by choosing a basis for each $V_\lambda \cap W_\mu$ and combining them. As A is diagonalizable, we have:

$$\mathbb{C}^n = \oplus_\lambda V_\lambda$$

and hence, A and B are simultaneously diagonalized by $W = \left[w_\mu\right]$. ■

Proof of $\Longleftarrow$:

Let A, B be simultaneously diagonalizable. Then, there exists P:

$$A = P\Lambda_a P^{-1}$$
$$B = P\Lambda_b P^{-1}$$

ie, the columns of P are eigenvectors of A and B with eigenvalues given by the diagonal of diagonal matrix Λ_a and Λ_b. Let $v_i = P_{:,i}$ be one of these columns. Then:

$$Av_i = \Lambda_{ai} v_i$$
$$Bv_i = \Lambda_{bi} v_i$$

But then,

$$\begin{aligned}(AB)v_i &= A(Bv_i) \\ &= A(\Lambda_{bi} v_i) \\ &= \Lambda_{bi}(Av_i) \\ &= \Lambda_{bi}(\Lambda_{ai} v_i) \\ &= \Lambda_{bi}\Lambda_{ai} v_i \\ &= \Lambda_{ai}(\Lambda_{bi} v_i) \\ &= \Lambda_{ai}(Bv_i) \\ &= B(\Lambda_{ai} v_i) \\ &= B(Av_i) \\ &= (BA)v_i\end{aligned}$$

and as any v can be written in the basis of the v_i (as P^{-1} exists), we are done. ■

Remark: Extension to Arbitrary Families

The lemma extends to any finite family of diagonalizable, mutually commuting matrices $\{A_1, A_2, ..., A_k\}$ by induction on k.

Base case: The $k = 2$ case was proved above.

Inductive step: Suppose $\{A_1, ..., A_{k-1}\}$ are simultaneously diagonalizable by some P; ie in the basis of P, all of $A_1, ..., A_{k-1}$ are diagonal, with simultaneous eigenspaces:

$$\mathbb{C}^n = \oplus_{\lambda_1,...,\lambda_{k-1}} V_{\lambda_1,...,\lambda_{k-1}}$$

Now, A_k commutes with each A_i for $i < k$, so A_k maps each $V_{\lambda_1,...,\lambda_{k-1}}$ to itself (same argument as above). As A_k is diagonalizable, by the same argument as the lemma, A_k can be diagonalized within each $V_{\lambda_1,...,\lambda_{k-1}}$, refining the basis to simultaneously diagonalize all k matrices. ■

Irreps of Abelian Group are 1-dimensional

Let G be an abelian group (ie $ab = ba$ for every $a, b \in G$); then all irreps of G are 1-dimensional.

Proof:

Let ρ be an irrep of G. Since representations of finite groups over $\mathbb{C}$ can always be taken to be unitary, each $\rho(a)$ is diagonalizable. Further, $\rho(a)\rho(b) = \rho(b)\rho(a)$ for every $a, b \in G$. Thus, $\{\rho(a)\}_{a\in G}$ commute with each other; and hence, are all simultaneously diagonalizable. Hence, there exists P such that:

$$\rho(a) = P\Lambda(a)P^{-1} \text{ for all } a \in G$$

where $\Lambda(a)$ is diagonal. (Remember from our proof above that the columns of P are the simultaneous eigenvectors of $\{\rho(a)\}_{a\in G}$.) Note that each coordinate axis of $\Lambda(a)$ gives us an invariant subspace of ρ. But as ρ is an irrep, the only invariant subspaces of ρ can be the whole space or $\{0\}$. Thus, the whole space must be 1-dimensional, and hence $\rho(a)$ is 1-dimensional. ■

Irrep Multiplicity in Regular Representation

Let G be any group of order N. Let ρ_{reg} be the regular representation of G. Let $\rho_{\text{reg}} = \oplus_j n_j\rho_j$, where ρ_j is an irrep of G. We are interested in finding out the multiplicity n_j (ie how many times each irrep ρ_j occurs in ρ_{reg}). We claim:

$$n_j = \dim(\rho_j)$$

Proof:

From the Wonderful Orthogonality Theorem, we obtained the orthogonality relationship between characters χ of two irreps ρ_i and ρ_j:

$$\langle \chi_i, \chi_j \rangle = \frac{1}{N}\sum_{n=0}^{N-1} \chi_i(n)\overline{\chi_j(n)} = \delta_{ij}$$

As characters add linearly in a direct sum of representations:

$$\chi_{\text{reg}} = \sum_j n_j \chi_j$$
$$\Rightarrow \langle \chi_{\text{reg}}, \chi_i \rangle = \langle \sum_j n_j \chi_j, \chi_i \rangle = \sum_j n_j \langle \chi_j, \chi_i \rangle = \sum_j n_j \delta_{ij} = n_i$$

for each irrep ρ_i.

Now, from the definition of the regular representation: $\chi_{\text{reg}}(0) = N$ and $\chi_{\text{reg}}(n) = 0$ for any $n > 0$ (as 0 must be mapped to $\mathbb{I}_N$, and $(n + m = m) \Rightarrow (n = 0)$ so no other $\rho_{\text{reg}}(n)$ has 1 on its diagonal). Further, for any irrep ρ_j, $\chi_j(0) = \text{tr}\big(\rho_j(0)\big) = \dim\big(\rho_j\big)$ (as 0 must be mapped to $\mathbb{I}_{\dim(\rho_j)}$). Thus,

$$n_j = \langle \chi_{\text{reg}}, \chi_j \rangle = \frac{1}{N}\Big(\chi_{\text{reg}}(0)\overline{\chi_j(0)}\Big) = \frac{1}{N}\big(N \cdot \dim\big(\rho_j\big)\big) = \dim\big(\rho_j\big)$$

Thus, each irrep ρ_j occurs $\dim\big(\rho_j\big)$ times in the regular representation ρ_{reg}. Since the regular representation has dimension N, we have:

$$N = \dim\big(\rho_{\text{reg}}\big) = \chi_{\text{reg}}(0) = \sum_j n_j \chi_j(0) = \sum_j \dim\big(\rho_j\big)\dim\big(\rho_j\big) = \sum_j \dim\big(\rho_j\big)^2. \quad \blacksquare$$

Irreps of $\mathbb{Z}_N$

We are now ready to characterize the irreps of $\mathbb{Z}_N$. We state the main theorem here for clarity: The irreps $\mathbb{Z}_N$ are denoted by ρ_j for $j = 0$ to $N - 1$:

$$\rho_j(n) = \exp\left(\frac{2\pi i \cdot j \cdot n}{N}\right) = \omega^{jn}$$

where $\omega = \exp\big(\frac{2\pi i}{N}\big)$ is the N-th root of unity. Further, the change-of-basis matrix P from the regular representation ρ_{reg} to the direct sum of irreps $\rho_0 \oplus \rho_1 \oplus \cdots \oplus \rho_{N-1}$ is:

$$P_{nj} = \exp\left(-\frac{2\pi i \cdot j \cdot n}{N}\right) = \omega^{-jn}.$$

Proof:

Since $\mathbb{Z}_N$ is abelian, we know that all irreps ρ are 1-dimensional. Now, $\mathbb{Z}_N$ is generated by 1; note that the group operation is usual addition. Thus, the regular representation ρ_{reg} can simply be defined by:

$$\rho_{\text{reg}}(n) = \big(\rho_{\text{reg}}(1)\big)^n$$

where:

$$\rho_{\text{reg}}(1) = \begin{pmatrix} & & & & 1 \\ 1 & & & & \\ & 1 & & & \\ & & 1 & & \\ & & & \ddots & \\ & & & & 1 \end{pmatrix}$$

which clearly maps e_n to $e_{(n+1) \bmod N}$ for all n.

From our result above, the regular representation ρ_{reg} contains exactly $n_j = \dim(\rho_j) = 1$ copy of each irrep ρ_j. Thus, as the regular representation has dimension N, there are exactly N such irreps ρ_j.

To find the irreps, we need to find the simultaneous eigenvectors of $\{\rho_{\text{reg}}(n)\}_{n<N}$. Since $\rho_{\text{reg}}(n) = \big(\rho_{\text{reg}}(1)\big)^n$, we simply need to find the eigenvectors of $\rho_{\text{reg}}(1)$, and those eigenvectors will simultaneously be eigenvectors of each $\rho_{\text{reg}}(n)$:

$$\rho_{\text{reg}}(1) = P\Lambda(1)P^{-1} \quad \Longrightarrow \quad \rho_{\text{reg}}(n) = P(\Lambda(1))^n P^{-1}$$

Further, since the irreps are all 1-dimensional, we can read out the irreps by looking at the corresponding eigenvalues:

$$\rho_{\text{reg}}(1) = P \begin{pmatrix} \rho_1(1) & & & & \\ & \rho_2(1) & & & \\ & & \rho_3(1) & & \\ & & & \ddots & \\ & & & & \rho_N(1) \end{pmatrix} P^{-1} \qquad \rho_{\text{reg}}(n) = P \begin{pmatrix} \rho_1(n) & & & & \\ & \rho_2(n) & & & \\ & & \rho_3(n) & & \\ & & & \ddots & \\ & & & & \rho_N(n) \end{pmatrix} P^{-1}$$

so $\rho_j(n) = \big(\rho_j(1)\big)^n$. Thus, we only need to look at the eigenvectors and eigenvalues of $\rho_{\text{reg}}(1)$ to define all of the irreps.

Let v be an eigenvector of $\rho_{\text{reg}}(1)$ with eigenvalue λ. Then, we can expand v in the (standard) basis of the e_i:

$$v = \begin{pmatrix} v_0 \\ v_1 \\ \vdots \\ v_{N-1} \end{pmatrix} \quad \Longrightarrow \quad v = \sum_{n=0}^{N-1} v_n e_n$$

From the definition of $\rho_{\text{reg}}(1)$, we have:

$$\rho_{\text{reg}}(1)e_n = \begin{cases} e_{n+1} & \text{for } n < N-1 \\ e_0 & \text{for } n = N-1 \end{cases}$$

Hence,

$$\begin{aligned} \rho_{\text{reg}}(1)v = \sum_{n=0}^{N-1} v_n\big(\rho_{\text{reg}}(1)e_n\big) &= \sum_{n=0}^{N-2} v_n\big(\rho_{\text{reg}}(1)e_n\big) + v_{N-1}\big(\rho_{\text{reg}}(1)e_{N-1}\big) \\ &= \sum_{n=0}^{N-2} v_n e_{n+1} + v_{N-1}e_0 \\ &= \sum_{n=1}^{N-1} v_{n-1} e_n + v_{N-1}e_0 \end{aligned}$$

But, as v is an eigenvector:

$$\rho_{\text{reg}}(1)v = \lambda v$$
$$\Rightarrow \sum_{n=1}^{N-1} v_{n-1}e_n + v_{N-1}e_0 = \sum_{n=0}^{N-1} \lambda v_n e_n$$
$$\Rightarrow (v_{N-1} - \lambda v_0)e_0 + \sum_{n=1}^{N-1}(v_{n-1} - \lambda v_n)e_n = 0$$

As the e_i are linearly independent, each of the coefficients must be 0. Hence:

$$v_{N-1} = \lambda v_0$$
$$v_{n-1} = \lambda v_n \text{ for } n = 1, ..., N-1$$

Thus, we have:

$$v_{N-2} = \lambda v_{N-1} = \lambda^2 v_0$$
$$v_{N-3} = \lambda v_{N-2} = \lambda^3 v_0$$
$$v_{N-4} = \lambda v_{N-3} = \lambda^4 v_0$$
$$\cdots$$

Continuing, the general formula is $v_{N-k} = \lambda^k v_0$. Setting $k = N$, we see:

$$v_{N-N} = v_0 = \lambda^N v_0$$
$$\Rightarrow \lambda^N = 1$$

Thus, λ is an N-th root of unity. The N-th roots of unity are $\exp\left(\frac{2\pi i j}{N}\right)$ for $j = 0$ to $(N-1)$, which are exactly the N eigenvalues of $\rho_{\text{reg}}(1)$. Let $\omega = \exp\left(\frac{2\pi i}{N}\right)$. Then, $\lambda = \omega^j$ for each choice of $j = 0$ to $(N-1)$. The simplest choice of ordering the eigenvalues λ (and hence irreps ρ_j) is to simply take $\lambda \in [\omega^0, \omega^1, ..., \omega^{N-1}]$:

$$\rho_0(1) = \exp\left(\frac{2\pi i \cdot 0}{N}\right) = \omega^0$$
$$\rho_1(1) = \exp\left(\frac{2\pi i \cdot 1}{N}\right) = \omega^1$$
$$\rho_2(1) = \exp\left(\frac{2\pi i \cdot 2}{N}\right) = \omega^2$$
$$\cdots$$
$$\rho_{N-1}(1) = \exp\left(\frac{2\pi i \cdot (N-1)}{N}\right) = \omega^{N-1}$$

and see that $\rho_j(1) = \exp\left(\frac{2\pi i \cdot j}{N}\right) = \omega^j$ in general. Thus, extending to all n:

$$\rho_j(n) = \big(\rho_j(1)\big)^n = \exp\left(\frac{2\pi i \cdot j \cdot n}{N}\right) = \omega^{jn}$$

This characterizes all of the irreps ρ_j of $\mathbb{Z}_N$. Finally, the change-of-basis matrix is given by the orthonormal eigenvectors. Let $v(\lambda)$ be the eigenvector corresponding to each λ. We set $v_0(\lambda) = \frac{1}{\sqrt{N}}$ (since we can scale eigenvectors arbitrarily, we show that this works below), and get from the recurrence relations above:

$$v(\lambda) = \begin{pmatrix} v_0(\lambda) \\ v_1(\lambda) \\ v_2(\lambda) \\ \vdots \\ v_{N-2}(\lambda) \\ v_{N-1}(\lambda) \end{pmatrix} = \frac{1}{\sqrt{N}} \begin{pmatrix} 1 \\ \lambda^{N-1} \\ \lambda^{N-2} \\ \vdots \\ \lambda^2 \\ \lambda \end{pmatrix} = \frac{1}{\sqrt{N}} \begin{pmatrix} \lambda^0 \\ \lambda^{-1} \\ \lambda^{-2} \\ \vdots \\ \lambda^{-(N-2)} \\ \lambda^{-(N-1)} \end{pmatrix}$$

for each eigenvalue λ after using the fact that $\lambda^N = 1$. This gives us $v_n = \lambda^{-n}$ for the eigenvector v corresponding to λ. Let $\lambda_1 = \omega^{j_1}, \lambda_2 = \omega^{j_2}$. Then, the corresponding eigenvectors $v(\lambda_1)$ and $v(\lambda_2)$ as constructed above are indeed orthonormal:

$$\begin{aligned} \langle v(\lambda_1), v(\lambda_2) \rangle &= \sum_{n=0}^{N-1} \left(\frac{1}{\sqrt{N}} \lambda_1^{-n} \right) \overline{\left(\frac{1}{\sqrt{N}} \lambda_2^{-n} \right)} \\ &= \frac{1}{N} \sum_{n=0}^{N-1} \omega^{-j_1 n} \omega^{j_2 n} \\ &= \frac{1}{N} \sum_{n=0}^{N-1} \omega^{(j_2 - j_1) n} \\ &= \frac{1}{N} \sum_{n=0}^{N-1} \left(\omega^{(j_2 - j_1)} \right)^n = \delta_{j_1 j_2} \end{aligned}$$

using the well-known property of roots of unity:

$$\sum_{n=0}^{N-1} \omega^n = \begin{cases} N & \text{if } \omega = 1 \\ 0 & \text{if } \omega \neq 1 \end{cases}$$

which can be derived by factorizing:

$$0 = \omega^N - 1 = (\omega - 1) \left(\sum_{n=0}^{N-1} \omega^n \right)$$

Thus, writing the corresponding eigenvectors for $\lambda \in [\omega^0, \omega^1, ..., \omega^{N-1}]$, we get the change-of-basis matrix as:

$$P = \frac{1}{\sqrt{N}} \begin{pmatrix} \left(\omega^0\right)^0 & \left(\omega^1\right)^0 & \left(\omega^2\right)^0 & ... & \left(\omega^{N-1}\right)^0 \\ \left(\omega^0\right)^{-1} & \left(\omega^1\right)^{-1} & \left(\omega^2\right)^{-1} & ... & \left(\omega^{N-1}\right)^{-1} \\ \left(\omega^0\right)^{-2} & \left(\omega^1\right)^{-2} & \left(\omega^2\right)^{-2} & ... & \left(\omega^{N-1}\right)^{-2} \\ \vdots & & & & \\ \left(\omega^0\right)^{-(N-2)} & \left(\omega^1\right)^{-(N-2)} & \left(\omega^2\right)^{-(N-2)} & ... & \left(\omega^{N-1}\right)^{-(N-2)} \\ \left(\omega^0\right)^{-(N-1)} & \left(\omega^1\right)^{-(N-1)} & \left(\omega^2\right)^{-(N-1)} & ... & \left(\omega^{N-1}\right)^{-(N-1)} \end{pmatrix}$$

which has the element $P_{nj} = \frac{1}{\sqrt{N}} \left(\omega^j\right)^{-n} = \frac{1}{\sqrt{N}} \left(\omega^{-jn}\right) = \frac{1}{\sqrt{N}} \exp\left(-\frac{2\pi i \cdot j \cdot n}{N}\right)$ for each n, j from 0 to $(N-1)$. P is exactly the discrete Fourier transform matrix! Note that P is symmetric: $P^T = P$ but also orthonormal: $P^{-1} = P^\dagger = \overline{P}$ by construction. ■

The Discrete Fourier Transform

The discrete Fourier transform converts a function $f : \mathbb{Z}_N \to \mathbb{R}$ to a new function $\hat{f} : \mathbb{Z}_N \to \mathbb{R}$, where:

$$\hat{f}(j) = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} f(n) \exp\left(-\frac{2\pi i \cdot j \cdot n}{N}\right)$$

(This $\frac{1}{\sqrt{N}}$ factor differs between conventions.)

If we write f as a vector:

$$f = \begin{pmatrix} f(1) \\ f(2) \\ \vdots \\ f(n) \end{pmatrix}$$

Then, we see that those factors map exactly to our symmetric change-of-basis matrix P from above:

$$\hat{f}(j) = \sum_n P_{jn} f(n) = \sum_n P_{nj} f(n) = (Pf)(j)$$
$$\implies \hat{f} = Pf$$

Thus, performing the discrete Fourier transform is simply multiplying by the change-of-basis P. To do the inverse discrete Fourier transform, we just have to multiply by $P^{-1} = \overline{P}$:

$$f = \overline{P}\hat{f}$$

Recall that:

$$\exp\left(-\frac{2\pi i \cdot j \cdot n}{N}\right) = \overline{\rho_j(n)}$$

So, we have:

$$\hat{f}(j) = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} f(n)\overline{\rho_j(n)} = \sqrt{N} \cdot \left(\frac{1}{N} \sum_{n=0}^{N-1} f(n)\overline{\rho_j(n)}\right) = \sqrt{N}\langle f, \rho_j \rangle$$

Thus, using the orthogonality of the characters (irreps here, since they are one-dimensional), we see that:

$$f = \sum_{j=0}^{N-1} \langle f, \rho_j \rangle \rho_j$$
$$\Rightarrow f(n) = \sum_{j=0}^{N-1} \langle f, \rho_j \rangle \rho_j(n)$$
$$\Rightarrow f(n) = \sum_{j=0}^{N-1} \left(\frac{1}{\sqrt{N}} \hat{f}(j)\right) \rho_j(n)$$

ie, we have decomposed f in the basis of the irreps ρ_j, with coefficients given exactly by $\frac{1}{\sqrt{N}}\hat{f}$.

It turns out that one can generalize this construction to any group G; we can decompose a function on G in terms of the irreps ρ_j of G. This is called the Peter-Weyl Theorem.