Lecture Notes on Random Walks and Representation Theory

William Chuang

February 20, 2025

Contents

1 Introduction

These notes concern the interplay between group representation theory and random walks on finite groups. In particular, we will:

- Recall the notion of total variation distance for probability distributions on a finite group.
- Discuss how characters and irreducible representations help analyze mixing properties of random walks.
- Work through an example with the cycle group $\mathbb{Z}/n\mathbb{Z}$ (the "random walk on a cycle") and see how trigonometric expressions like $\cos\left(\frac{\pi k}{n}\right)$ appear in the analysis.
- Introduce (briefly) how these ideas extend to other examples, such as hypercubes or more general Hecke algebras.

We aim to fill in some standard definitions and classical propositions so that an undergraduate with a background in linear algebra and basic group theory can follow.

2 Preliminaries

2.1 Total Variation Distance

Definition 2.1 (Total Variation Distance)**.** *Let G be a finite set (in our case, a finite group), and let p and q be two probability distributions on G. The* total variation distance *between p and q is defined as*

$$\|p - q\|_{\mathrm{TV}} \;=\; \frac{1}{2}\sum_{x\in G}\bigl|p(x) - q(x)\bigr|.$$

Equivalently, it can be viewed as

$$\|p - q\|_{\mathrm{TV}} \;=\; \max_{A\subseteq G}\bigl|p(A) - q(A)\bigr|.$$

In the context of random walks, p^k often denotes the distribution of the walker's location after k steps, and u might denote the uniform distribution on G. A key question is how quickly $\|p^k - u\|_{\mathrm{TV}}$ goes to zero as $k \to \infty$.

2.2 Representation Theory of Finite Groups

Definition 2.2 (Representation)**.** *A* representation *of a finite group G over $\mathbb{C}$ is a group homomorphism $\rho : G \to GL(V)$, where V is a finite-dimensional vector space over $\mathbb{C}$.*

Definition 2.3 (Irreducible Representation)**.** *A representation ρ is said to be* irreducible *if it has no proper, nonzero, $\rho(G)$-invariant subspace.*

Every finite-dimensional representation of G decomposes as a direct sum of irreducible representations. The number of (isomorphism classes of) irreducible representations of G equals the number of conjugacy classes in G.

Definition 2.4 (Character)**.** *The* character *χ of a representation ρ is the function $\chi : G \to \mathbb{C}$ given by $\chi(g) = \mathrm{Trace}(\rho(g))$. If ρ is irreducible, χ is called an* irreducible character*.*

2.3 Group Algebra and Convolution

Definition 2.5 (Group Algebra over $\mathbb{C}[G]$)**.** *Given a finite group G, the group algebra $\mathbb{C}[G]$ is the vector space of formal linear combinations*

$$\left\{\sum_{g\in G} a_g\, g \;\middle|\; a_g \in \mathbb{C}\right\}$$

with the product (convolution) defined by extending

$$\left(\sum_{g\in G} a_g\, g\right)\left(\sum_{h\in G} b_h\, h\right) \;=\; \sum_{g,h\in G} a_g\, b_h\,(gh).$$

A probability distribution p on G can be viewed as an element of $\mathbb{C}[G]$ via

$$p \;=\; \sum_{g\in G} p(g)\, g,$$

and then convolution powers p^k correspond to k-fold convolution in the group algebra.

3 Mixing via Characters and Representations

One powerful technique to analyze the rate at which p^k approaches the uniform distribution is via characters. Roughly, each irreducible representation ρ gives an eigen-decomposition of the convolution operator associated to p. In many cases, bounding $\|p^k - u\|_{\text{TV}}$ is achieved by controlling the largest nontrivial eigenvalue in absolute value.

Theorem 3.1 (Upper Bound Lemma (Sketch)). *Let p be a probability distribution on G, and let $\{\rho_i\}$ be the irreducible representations of G. Suppose $\rho_i(p)$ denotes the operator*

$$\rho_i(p) \;=\; \sum_{g \in G} p(g)\, \rho_i(g).$$

Then

$$\|p^k - u\|_{\text{TV}} \;\leq\; \frac{1}{2} \sqrt{|G| \max_{i \neq \text{trivial}} \left|\lambda_i\big(\rho_i(p)\big)\right|^{2k}},$$

where $\lambda_i(\rho_i(p))$ are the eigenvalues of $\rho_i(p)$. In particular, if

$$\max_{i \neq \text{trivial}} \left|\lambda_i\big(\rho_i(p)\big)\right| \;=\; \lambda < 1,$$

then $\|p^k - u\|_{\text{TV}}$ decays on the order of λ^k.

(There are many variants of this bound; the above is just a typical statement.)

4 Example: Random Walk on the Cycle $\mathbb{Z}/n\mathbb{Z}$

Consider the cycle group $C_n = \mathbb{Z}/n\mathbb{Z}$. Let us define a simple random walk with the transition probability

$$p(x) \;=\; \begin{cases} \frac{1}{2}, & x = 1, \\ \frac{1}{2}, & x = -1 \equiv n-1, \\ 0, & \text{otherwise.} \end{cases}$$

In other words, from any point on the cycle, you step left or right with equal probability.

4.1 Characters of $\mathbb{Z}/n\mathbb{Z}$

All irreducible representations of $\mathbb{Z}/n\mathbb{Z}$ are one-dimensional. Let $\omega = e^{2\pi i/n}$. Then the irreps are given by

$$\rho_k(x) \;=\; \omega^{kx}, \quad x \in \mathbb{Z}/n\mathbb{Z},$$

for $k = 0, 1, \ldots, n-1$. The corresponding character is simply

$$\chi_k(x) \;=\; \text{Trace}(\rho_k(x)) \;=\; \omega^{kx}.$$

4.2 Convolution Operator

The probability distribution p can be seen as

$$p = \frac{1}{2}\,1 + \frac{1}{2}\,(-1) \quad \text{in } \mathbb{Z}/n\mathbb{Z},$$

where we identify $-1 \equiv n-1$. Thus, for each irreducible representation ρ_k,

$$\rho_k(p) = \frac{1}{2}\,\rho_k(1) + \frac{1}{2}\,\rho_k(-1) = \frac{1}{2}\,\omega^k + \frac{1}{2}\,\omega^{-k} = \cos\Big(\frac{2\pi k}{n}\Big).$$

Hence, the eigenvalue of the convolution operator $\rho_k(p)$ is exactly $\cos\big(\frac{2\pi k}{n}\big)$.

4.3 Distribution After k Steps

The distribution after k steps is p^{*k} (the k-fold convolution of p with itself). In representation-theoretic terms, the eigenvalue corresponding to ρ_k becomes

$$\big(\cos(\tfrac{2\pi k}{n})\big)^k.$$

In particular, for large k, the dominant term is given by the largest (in absolute value) of these cosines when $k \neq 0$. (When $k = 0$, ρ_0 is the trivial representation with eigenvalue 1, corresponding to the uniform distribution.)

Proposition 4.1 (Mixing Time on a Cycle)**.** *For the simple random walk on $\mathbb{Z}/n\mathbb{Z}$, the total variation distance to the uniform distribution satisfies*

$$\|p^{*k} - u\|_{\mathrm{TV}} \le \frac{1}{2}\sum_{j=1}^{n-1}\big|\cos(\tfrac{2\pi j}{n})\big|^k.$$

A more refined analysis (using e.g. local central limit theorems or direct Fourier estimates) shows that $\|p^{*k} - u\|_{\mathrm{TV}}$ is typically on the order of

$$\exp\big(-c\,k/n^2\big)$$

for some constant $c > 0$, once k is large compared to n^2. (Exact forms can vary, but n^2 is the standard cutoff for simple random walk on a cycle.)

Remark 4.2. *Sometimes you see $\cos\big(\frac{\pi k}{n}\big)$ in certain variations of the cycle walk or in expansions involving $\omega^{kx} = e^{2\pi i kx/n}$. The exact angles depend on whether you label group elements from 0 to $n-1$ or from $-\frac{n-1}{2}$ to $\frac{n-1}{2}$, and so forth. The key idea is that each irreducible representation is one-dimensional, and so the eigenvalue is always a complex root of unity average, which simplifies to a cosine for the real part.*

5 Hypercube and Other Examples

5.1 Hypercube $\{0,1\}^n$

Another classical example is the hypercube graph on $\{0,1\}^n$. The group structure can be taken as $(\mathbb{Z}/2\mathbb{Z})^n$, and a simple random walk flips one coordinate (chosen uniformly at random) at each step. Its eigenvalues can be analyzed via the Walsh–Fourier basis, leading to expansions in terms of ± 1 characters. One finds that the random walk has mixing time on the order of $n \log n$.

5.2 Hecke Algebras and Graphs

Hecke algebras arise in more advanced settings (e.g. association schemes, double coset algebras). In many such cases, one again uses representation theory to decompose the algebra into irreducibles, obtaining mixing bounds or spectral gap estimates. For instance, in the study of Ramanujan graphs or certain p-adic groups, Hecke operators have eigenvalues whose absolute values control expansion or mixing properties.

6 Appendix: Additional Propositions and Lemmas

Proposition 6.1 (Orthogonality of Characters)**.** *Let $\{\chi_i\}$ be the irreducible characters of a finite group G. Then*

$$\sum_{g \in G} \chi_i(g)\, \overline{\chi_j(g)} \;=\; \frac{|G|}{d_i}\, \delta_{ij},$$

where $d_i = \chi_i(e)$ is the dimension of the i-th representation and δ_{ij} is the Kronecker delta.

Lemma 6.2 (Lower Bound Lemma (Typical Form))**.** *If $\lambda = \max_{i \neq \text{trivial}} |\lambda_i(\rho_i(p))|$ is close to 1, then for certain small k, one can show*

$$\|p^k - u\|_{\text{TV}} \;\geq\; \textit{(some constant)} \cdot (1-\lambda)^k.$$

Hence, we can also get lower bounds on the mixing time by controlling how λ approaches 1.

Remark 6.3. *Precise statements depend on details of the group, the support of p, and so forth. For random walks with small spectral gap (i.e. λ very close to 1), one often needs more refined tools to get exact cutoff phenomena or sharp lower bounds.*

Conclusion

These notes have illustrated how tools from representation theory (characters, irreps, and the group algebra) are used to understand random walks on finite groups. The example of the cycle group $\mathbb{Z}/n\mathbb{Z}$ and references to the hypercube and Hecke algebras demonstrate the generality of the approach.

In a more advanced course, one would study in detail the role of character tables, the Plancherel formula, and explicit expansions of Markov transition operators in terms of irreps. One would also delve deeper into the analysis of mixing times and cutoff phenomena.