8.372 Quantum Information Science III **Fall 2024**

Lecture 15: October 31, 2024

Scribe: Daniel Z. Lee, Mahdi Hamad, Gary Yang Unitary k-designs and Representation Theory

In this lecture, we'll build the background for a more principled approach to these calculations.

15.0 Random Unitaries and Haar Measure

The uniform distribution over the unitary group $U(d)$ is given by the **Haar measure**, denoted μ_{Haar}. This measure is the analog of the uniform distribution for any compact group and is defined by the following properties:

- **Normalization**: $\mu_{\text{Haar}}(U(d)) = 1$.
- **Invariance**: For any subset $S \subseteq U(d)$ and any $U \in U(d)$, we have

$$\mu_{\text{Haar}}(S) = \mu_{\text{Haar}}(US) = \mu_{\text{Haar}}(SU).$$

Physicist's Perspective

This invariance implies that the measure is "uniform" in the sense that if we rotate any subset of unitaries, its measure does not change. Thus, the Haar measure provides a probability density that is independent of the choice of coordinates on $U(d)$.

The Haar measure is unique: if any probability distribution on $U(d)$ is left invariant, it must be the Haar measure. This left and right invariance essentially characterizes it, making it a fundamental tool in defining randomness for unitary operations.

15.0.1 Random Unit Vectors and Gaussian Sampling

For generating random unit vectors computationally, a common approach is to sample each component independently from a Gaussian distribution (which is rotationally invariant), and then normalize the vector. This produces a uniformly random vector on the unit sphere.

For random unitaries, one approach involves using the Gaussian Unitary Ensemble (GUE), a distribution over Hermitian matrices. To construct a GUE matrix:

- Diagonal entries are real, Gaussian-distributed.
- Off-diagonal entries are independent, with real and imaginary parts drawn from Gaussian distributions.

Given a Hermitian matrix X sampled from GUE, the unitary e^{iX} approximates a Haar-random unitary. This method is computationally feasible though not perfect, as it involves approximately $O(d^3)$ operations.

Practical Limitations

Sampling a truly Haar-random unitary is computationally challenging in high dimensions due to the large number of degrees of freedom. Approximate methods such as GUE provide feasible alternatives for many practical applications.

15.0.2 Introduction to Unitary k-designs

A **unitary k-design** is a distribution over unitaries that mimics the Haar measure up to the k-th moment. Formally, for any integer k, a distribution ν on $U(d)$ is a k-design if the following holds:

$$\mathbb{E}_{U\sim\nu}\left[U^{\otimes k,k}\right] = \mathbb{E}_{U\sim\text{Haar}}\left[U^{\otimes k,k}\right].$$

This means that, up to the k-th moment, the behavior of unitaries sampled from ν is indistinguishable from unitaries sampled according to the Haar measure.

Definition: k-design via Tensor Powers

The notation $U^{\otimes k,k}$ is shorthand for $U^{\otimes k} \otimes U^{*\otimes k}$, which represents k copies of U acting on the system along with k copies of the complex conjugate of U. A distribution ν on $U(d)$ is a k-design if choosing U from ν gives the same distribution for $U^{\otimes k,k}$ as choosing U from the Haar measure.

15.0.3 Level Repulsion and Eigenvalue Distribution

One notable feature of random unitary matrices is **level repulsion**, where eigenvalues tend to avoid being close to one another. The probability density of the eigenvalues $\{\lambda_i\}$ includes a term like

$$\prod_{i<j} |\lambda_i - \lambda_j|^2,$$

which vanishes when two eigenvalues coincide. This "repulsion" is similar to the behavior of charges repelling each other, leading to eigenvalues that spread out more evenly on the unit circle.

Visualization Exercise

Plotting the eigenvalues of a Haar-random unitary matrix on the complex unit circle reveals this level repulsion. Comparing this distribution with randomly chosen phases $e^{i\theta}$ illustrates the difference: in the Haar case, the eigenvalues push each other apart, while in the random phase case, they can cluster by chance.

This phenomenon is a well-known property in random matrix theory and also appears in GUE. When eigenvalues are distinct, each eigenvector has more degrees of freedom. For degenerate eigenvalues, the dimensionality of the associated subspace decreases.

15.0.4 Compactness and the Haar Measure

The Haar measure is defined for compact groups, such as $U(d)$, where the group has a finite total volume that can be normalized to 1. For non-compact groups, such as $SL(2,\mathbb{R})$, there is no normalized Haar measure due to infinite volume.

Fact: Invariant Subspace Projection

For a compact group G, averaging a representation R over G with the Haar measure yields a projection onto the invariant subspace V^G:

$$\int_G R(g)\, d\mu_{\text{Haar}}(g) = \text{Proj}_{V^G}.$$

This concept is crucial in representation theory and plays a foundational role in understanding k-designs.

15.0.5 Classical Analogue: k-wise Independent Hash Functions

In classical computing, k-designs have an analogy in **k-wise independent hash functions**. A hash function $h : U \to [m]$ is k-wise independent if, for any distinct inputs $x_1, \ldots, x_k$ and outputs $y_1, \ldots, y_k \in [m]$, we have:

$$\Pr_{h \in H} \left(h(x_1) = y_1 \wedge \cdots \wedge h(x_k) = y_k \right) = \frac{1}{m^k}.$$

This implies that the hash function behaves like a truly random function when viewed through any k inputs, although it is not completely random.

Construction

Degree-k polynomials are often used to construct k-wise independent hash functions. These provide a computationally efficient balance between determinism and the randomness required for various applications.

15.0.6 Approximately Randomizing Maps

Instead of using d^2 number of unitaries to satisfy a unitary-1 design, we can settle for Approximate Randomizing Maps (AKA approximate 1-designs):

$$\mathcal{N}(\rho) = \frac{1}{m} \sum_{i=1}^{m} U_i \rho U_i^\dagger \tag{15.1}$$

We know for a fact that they exist and satisfy

$$\left\| \mathcal{N}(\rho) - \frac{I}{d} \right\|_\infty \leq \frac{\epsilon}{d}, \quad \forall \rho \tag{15.2}$$

where $m \lesssim \frac{d}{\epsilon^2}$. If ϵ is constant, then this is more economical than the exact randomizing maps.

One interesting consequence is data hiding, which we partially explore on our PSET. Examine the following state:

$$\left(I \otimes \mathcal{N} \right) \left(\Phi_d \right). \tag{15.3}$$

It has rank on the order $\frac{d}{\epsilon^2}$. This means that it is far from the maximally mixed state. However, $\left(I \otimes \mathcal{N} \right) \left(\Phi_d \right)$ and the maximally mixed state $\frac{I_{d^2}}{d^2}$ are almost indistinguishable using local operation and classical communication (LOCC). To distinguish them you need a global measurement. Say Bob measure on his basis (the second system), his measurement will got conjugated by a bunch of random unitary, which averages his measurement close to identity.

15.0.7 Summary of Designs

For 1−design, we need d^2 unitaries. For Approximate 1−design, we need $\frac{d}{\epsilon^2}$ unitaries. If we want discrete approximate for Haar randomness, we need $\exp\{d^2\}$ matrices. So we can see that (approximate) 1−designs provide significant savings compared to real Haar measures.

15.0.8 Approximate 1-Design on Teleportation

In regular teleportation, we use $2n$ cbits and n ebits to send n qubits. There is no asymptotic savings. However, if we use Remote State Preparation and Approximate Randomizing Maps (arXiv:quant-ph/0006044), where Alice knows classically what state is to be transmitted (Alice does not physically have a physical copy of the state), we can only use $1 + \delta$ cbits and n ebits to prepare n qubits. The intuitive understanding of the saving is that Alice has some extra information on the state that she is sending, which affects her action in the protocal.

15.1 2-Designs

The need for 2−Designs basically comes from where you which to apply U twice, or you wish to understand high moments of random unitaries. Just like state 1− versus 2− designs, for unitary 2− designs, we need entangling operations, which will take a product state to a state 2−designs.

$$T(X) \equiv \mathbb{E}_{\text{Haar}} \quad (U \otimes U)X(U \otimes U)^\dagger \tag{15.4}$$

If we just take U to be the Paulis, and $X = |00\rangle\langle 00|$, we would get

$$T(X) = \frac{|00\rangle\langle 00| + |11\rangle\langle 11|}{2}$$

where true random unitary will give you a combination of all the triplet states.

Let us examine what kind of X will be invariant under the map $T(X)$ to see what $T(X)$ does. One obvious answer is

$$X = I_{d^2}, \quad T(I_{d^2}) = I_{d^2} \tag{15.5}$$

because the unitaries would have canceled each other, and one less obvious answer is

$$X = F, \quad T(F) = F \tag{15.6}$$

in this case you can imagine that U and $U^\dagger$ will cancel on each end of the SWAP operator. Mathematicians have proven in our favor that these two maps are the only one invariant under this operation, so we can write

$$T = \text{proj Span}\{I, F\} \tag{15.7}$$

with respect to the Hilbert-Schmit inner product, where

$$\langle A, B\rangle = \operatorname{Tr} A^\dagger B. \tag{15.8}$$

One problem, however, is that I and F are not orthogonal. The better choices are

$$\Pi_\pm = \frac{I \pm F}{2} \tag{15.9}$$

where $\Pi_\pm$ projects to the ±1 eigenstates of F, which is the symmetric $(+1)$ and antisymmetric (-1) subspace. The symmetric/antisymmetric subspace, as we learned before, has dimension

$$\operatorname{Tr}\Pi_\pm = \frac{d(d\pm 1)}{2}$$

. This way we can rewrite

$$T(X) = \frac{\operatorname{Tr}\{X\Pi_+\}}{\operatorname{Tr}\Pi_+}Pi_+ + \frac{\operatorname{Tr}\{X\Pi_-\}}{\operatorname{Tr}\Pi_-}Pi_- \tag{15.10}$$

Note with $d = 2$, Π_+ gives you the projection onto the triplet state, and Π_- gives you projection to the singlet state.

15.1.1 Random Cliffords

In this subsection, we show that random Clifford gates are a unitary 2−design.

Because general matrix can be written as a combination of Paulis, we can just analyze how Cliffords would act on Paulis.

$$\mathbb{E}_{U\in\text{Clifford}} \quad (U\otimes U)\sigma_p\otimes\sigma_q(U\otimes U)^\dagger = \begin{cases} 0 & \text{if } p\neq q \\ I & \text{if } p=q=0^n \\ \sum_{r\neq 0^n}\frac{\sigma_r\otimes\sigma_r}{4^n-1} & \text{if } p=q\neq 0^n \end{cases} \tag{15.11}$$

Recall from last lecture, we know that

$$\sum_{r\neq 0^n}\sigma_r\otimes\sigma_r = 2^nF - I. \tag{15.12}$$

We can see here is that, whatever comes in, you just get a combination of I and F, basically what you want from a 2−design.