Frontiers of Quantum Complexity and Cryptography Spring 2022

Lecture 3 - Haar-random states

Lecturer: Henry Yuen *Scribes: Rockwell Weiner*

1 Overview

We introduced the Haar distribution, which is a notion of "uniformly random quantum state", and discussed constructions and complexity.

2 Haar Measure

2.1 Definition

Informally, the Haar measure is a uniform distribution over pure quantum states. If we imagine the space of all pure states as a sphere, the Haar measure is equivalent to picking a random point on the sphere.

Suppose we are given an n-qubit state $|\psi\rangle$ which has been secretly and randomly sampled uniformly. If we let $d = 2^n$, we can describe this state as $|\psi\rangle \sim Haar(d)$. Now we apply an arbitrary unitary $U \in \mathbb{C}^{d\times d}$ to this state. If we consider the resulting state, $U|\psi\rangle$, we find that it is also described by the Haar distribution. The intuition for this is that we have just shifted all of the points on our sphere, but since they were all already equally likely, our new state is indistinguishable from the old. In fact, this is the formal definition of the Haar distribution.

Formal Definition: The Haar distribution is the unique unitary-invariant measure over $S(\mathbb{C}^d)$.

2.2 Construction

If we wish to actually constructively sample from the Haar distribution, we can use the following procedure:

- *Sample d complex Gaussians* – $\tilde{\alpha_1}, ..., \tilde{\alpha_d} \sim \mathbb{C}\mathcal{N}(0, 1)$
- *Construct a vector* – $v = \begin{pmatrix} \tilde{\alpha_1} \\ \vdots \\ \tilde{\alpha_d} \end{pmatrix}$
- *Normalize the vector* – $|\psi\rangle = \frac{v}{\|v\|} = \sum_{i=1}^{d} \alpha_i |i\rangle$

The complex Gaussian distribution $\mathbb{C}\mathcal{N}(0, 1)$ can be sampled as follows: sample two independent Gaussians $x, y \sim \mathcal{N}(0, \frac{1}{2})$ (i.e. mean 0, variance $\frac{1}{2}$), and set $z = x + iy$. This complex random variable has the property that it has mean 0 and its variance, defined to be $\mathbb{E}|z|^2$, is equal to 1.

2.3 Properties of the Haar Measure

2.3.1 Expected Amplitude

By symmetry, we can see that for any i, $\mathbb{E}[|\alpha_i|^2] = \frac{1}{d}$.

To prove this formally, first note that $\mathbb{E}[|\alpha_1|^2] = \mathbb{E}[|\langle 1|\psi\rangle|^2]$. Now consider fixing some unitary U such that $\langle 1|U = \langle 2|$. By the definition of the Haar distribution, we know that $|\psi\rangle = U|\psi\rangle$, so our expectation becomes $\mathbb{E}[|\langle 1|U|\psi\rangle|^2]$. From the way we've defined U, this expectation is equivalent to $\mathbb{E}[|\langle 2|\psi\rangle|^2] = \mathbb{E}[|\alpha_2|^2]$. So we've shown that $\mathbb{E}[|\alpha_1|^2] = \mathbb{E}[|\alpha_2|^2]$, and we can repeat this procedure for any two dimensions. Therefore all the expectations must be equal, and since they sum to one, we have $\mathbb{E}[|\alpha_i|^2] = \frac{1}{d}$.

2.3.2 Complexity

Before we discuss the complexity of the Haar measure, we need to define a concept of complexity (out of many possible definitions). For an n-qubit state $|\psi\rangle$, we will define the complexity, $\mathcal{C}_\epsilon(|\psi\rangle)$, to be the minimum size of a quantum circuit A which constructs this state, such that $A|0...0\rangle \approx_\epsilon |\psi\rangle$. Of course, this begs the question of how to measure the size of a quantum circuit.

In general, we will allow these circuits to use unitary gates, operating on either one or two qubits.

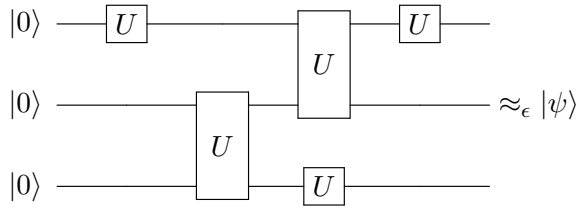

For the sake of our class, we will restrict ourselves to arbitrary single-qubit gates, and only the CNOT gate operating on two qubits. This is sufficient to act as a universal gate set. Going forward, we will measure complexity with respect to the number of gates in this model.

Note that this means our complexity measure can range from 0 (for the state $|0...0\rangle$) to exponential in n. The set of states which will be of particular interest to us are those where $\mathcal{C}(|\psi\rangle) \leq \text{poly}(n)$, since these are the states which we actually have a hope of constructing in a reasonable amount of time.

Claim: Haar states are typically very complex. Roughly: $\Pr_{|\psi\rangle\sim Haar}[\mathcal{C}(|\psi\rangle) \ll 2^n] < \exp(-d)$

Intuition: The see the intuition for this, we can use a counting argument. First, to estimate the number of d-dimensional states (if we temporarily suppose there are only finitely many states), imagine packing our ϵ-radius balls into a hypersphere, as before. The total number we can fit is on the order of $(\frac{1}{\epsilon})^d$.

Now we count how many states $|\psi\rangle$ exist such that $\mathcal{C}(|\psi\rangle) \leq S$. This is the number of quantum circuits of size at most S. For simplicity, suppose we have a fixed set of r different gates. Then we

have S opportunities to choose two qubits and apply one of r gates to them, giving us a count of $(\binom{n}{2}r)^S = O(n^S)$.

Finally, fix ϵ to some arbitrary value, say $\epsilon = \frac{1}{2}$. Then unless S is exponential in n, we have that $O(n^S)$
$(\frac{1}{\epsilon})^d = 2^{2^n}$.

To move from our somewhat simplified argument here to something more formal, we can argue that the infinite set of gates makes up a manifold in many dimensions, and when picking a random quantum state, we are extremely unlikely to encounter a point in this manifold.

Result: The key takeaway from this argument is that Haar states are rarely encountered in real life, given that most of them take an exponential number of gates to construct. Despite their rarity, they are still a mathematically useful concept.